## DISCLAIMER:

This document does not meet the current format guidelines of the Graduate School at The University of Texas at Austin.

It has been published for informational use only.

## Copyright

by
Zhao Song
2019

The Dissertation Committee for Zhao Song certifies that this is the approved version of the following dissertation:

# Matrix Theory: <br> Optimization, Concentration, and Algorithms 

Committee:

Eric Price, Supervisor

Georgios-Alex Dimakis

Adam Klivans

Yin Tat Lee
C. Greg Plaxton

# Matrix Theory: <br> Optimization, Concentration, and Algorithms 

by

## Zhao Song

## DISSERTATION

Presented to the Faculty of the Graduate School of The University of Texas at Austin in Partial Fulfillment of the Requirements for the Degree of DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN
August 2019

Dedicated to my parents Li Song and Mei Sun, also my grandparents Kexin Song, Weifeng Guan, Yongzhe Sun, and Fenglin Yang.

Thank Wei Ye for accompanying me over a half of my Ph.D. time.

## Acknowledgments

First, I would like to thank my advisor, Professor Eric Price, for all the inspirations and support that he has given me throughout my Ph.D. career. He has been, and will always be, a great source of wisdom and encouragement for me. I genuinely want to thank Eric for giving me so many freedoms during the Ph.D. I thank Eric for allowing me to work on topics from algorithm to complexity, from theoretical computer science to machine learning, and from applications to applied Mathematics. I thank Eric for allowing me to travel around the country to pursue the best scholarship in the field. I also wanted to thank Eric for introducing David and Yin Tat to me. I could never imagine a Ph.D life without Eric as my advisor. It is an honor of life for me to be the first student of Eric.

I am very thankful for my spritual advisor, Professor David P. Woodruff, who had been teaching and training me during the first half of my Ph.D. time.

I really want to thank my informal advisor, Professor Yin Tat Lee, who had been helping and guiding me during the second half of my Ph.D. time.

I would appreciate the contribution from David and Yin Tat on my thesis.

I wish to thank the multitudes of people who helped me. Time would fail me to tell of ...

Coauthors The main purpose of writing this thesis is to express my gratitude to all coauthors and collaborators. I want to thank all of you $x$ times, where $x$ is the number of pages in this thesis. The list includes, but is not limited to: Alexandr Andoni, Josh Alman, Sanjeev Arora, Zeyuan Allen-Zhu, Maria-Florina Balcan, Duane Boning, Peter L. Bartlett, Aritra Banik, Binay K. Bhattacharya, Dana H. Ballard, Paul Beame, Jan van den Brand, Mitali Bafna, Hongge Chen, Michael Cohen, Xue Chen, Sitan Chen, Constantine Caramanis, Timothy Chu, Chi-Ning Chou, Inderjit S. Dhillon, Luca Daniel, Huaian Diao, Sandip Das, Alexandros G. Dimakis, Minati De, Simon S. Du, Martin Ester, Ankit Garg, Anna Gál, Cho-Jui Hsieh, Wei Hu, Bo Hu, Prateek Jain, Ziheng Jiang, Rajesh Jayaram, Chi Jin, Haotian Jiang, Daniel Kane, Tsunehiko Kameda, Weihao Kong, Rasmus Kyng, Ka Yui Lee, Yin Tat Lee, Yingyu Liang, Yibo Lin, David Liau, Jerry Li, Meng Li, Xingguo Li, Yuanzhi Li, Daogao Liu, Kyle Luh, Dana Moshkovitz, Vasileios Nakos, Michael Orshansky, Eric Price, Aviad Rubinstein, Milan Rubinstein, Ilya P. Razenshteyn, Ankit Singh Rawat, Sasanka Roy, Victor Reis, Xiaorui Sun, Aaron Schild, Aaron Sidford, Nikhil Srivastava, Wen Sun, Vladyslav Sokol, Seyed Abbas Sadat, Saeed Seddighin, Ruoqi Shen, Clifford Stein, Ewin Tang, Richard T. Vaughan, Santosh S. Vempala, Ye Wang, Mengdi Wang, Ruosong Wang, Zhengyu Wang, Tsui-Wei Weng, David P. Woodruff, Lin F. Yang, Xin Yang, Ger Yang, Xinyang Yi, Huacheng Yu, Yi Zhang, Qiuyi Zhang, Huan Zhang, Jiong Zhang, Hongyang Zhang, Ruohan Zhang, Ruizhe Zhang, Kai Zhong, Peilin Zhong, Yuke Zhu.

Useful Discussions I want to thank the researchers who had done super useful discussions with me, and they probably even didn't notice our conversation lead to a paper finally. More surprisingly, perhaps most of the researchers mentioned below couldn't remember they've ever talked to me. I'd like to thank Udit Agarwal, Alexandr Andoni, Josh Alman, Zeyuan Allen-Zhu, Amir Abboud, Megasthenis Asteris, Scott Aaronson, Arturs Backurs, Saugata Basu, Boaz Barak, Jarosław Błasiok, Saugata Basu, Sebastien Bubeck, Peter L. Bartlett, Michael Borokhovich, Mitali Bafna, Paul Beame, Abhishek Bhowmick, Ainesh Bakshi, Eshan Chattopadhyay, Lijie Chen, Xue Chen, Kenneth Clarkson, Hongge Chen, Danqi Chen, Sitan Chen, Chi-ning Chou, Thomas Dillig, Isil Dillig, Alexandros G. Dimakis, Ethan Elenberg, Ronen Eldan, Yu Feng, Vitaly Feldman, Uriel Feige, Ronald Fagin, Rong Ge, Anna Gál, Ankit Garg, Mika Göös, Surbhi Goel, Badih Ghazi, Daniel Hsu, Changyong Hu, Yige Hu, Cho-Jui Hsieh, Wei Hu, Samuel Hopkins, William Hoza, Bernhard Haeupler, Russell Impagliazzo, Chi Jin, Prateek Jain, T.S. Jayram, Ravindran Kannan, Michael Kapralov, Anna Karlin, Rasmus Kyng, Daniel Kane, Janardhan Kulkarni, Murat Kocaoglu, Pritish Kamath, Mrinal Kumar, J. M. Landsberg, Qi Lei, Fu Li, Yixuan Li, Yin Tat Lee, Yibo Lin, Jason D. Lee, James Lee, Yuanzhi Li, Zhiyuan Li, Xin Li, Yingyu Liang, Kyle Luh, Meng Li, Wuxi Li, Daogao Liu, Zhixian Lei, David Liau, Kasper Green Larsen, Syed Mohammad Meesum, Ankur Moitra, Dana Moshkovitz, Cameron Musco, Christopher Musco, Raghu Meka, Jelani Nelson, Vasileios Nakos, Preetum Nakkiran, Richard Peng, Eric Price, Greg Plaxton, Daniel Perrucci, Dimitris Papailiopoulos, Govind Ramnarayan, Ilya Razenshteyn, Anup Rao, James Renegar, Aviad Rubinstein, Rocco Servedio, Tselil Schramm, Clifford Stein, Wen Sun, Madhu Sudan, Aaron Sidford, Aaron Schild, Nikhil Srivastava, Xiaorui Sun, Tianxiao Shen, Saeed Seddighin, Thomas Steinke, Karthikeyan Shanmugam, Xiaofei Shi, Ewin Tang, Elias

Tsigaridas, Rashish Tandon, Avishay Tal, Santosh S. Vempala, Salil Vadhan, Jan Vondrák, Yuhao Wan, Yining Wang, Zhengyu Wang, Zhaoran Wang, Xinyu Wang, Ye Wang, David P. Woodruff, Ryan R. Williams, John Wright, Omri Weinstein, Mary Wootters, Shanshan Wu, Xiaoxia Wu, Bo Xiong, Yuanzhong Xu, Huacheng Yu, Bei Yu, Huacheng Yu, Xin Yang, Lin F. Yang, Ger Yang, Huan Zhang, Hongyang Zhang, Yi Zhang, Jiong Zhang, Ruizhe Zhang, Kai Zhong, Yuke Zhu, Danyang Zhuo, and David Zuckerman for useful discussions.

Visiting My research has benefited enormously from a number of visits to research institutions over the past few years. I would like to thank all my hosts, with whom I have established invaluable friendships: David P. Woodruff at IBM Almaden, Jelani Nelson at Harvard University, Yin Tat Lee at the University of Washington, Zeyuan Allen-Zhu, Jerry Li, and Ilya Razenshteyn at MSR Redmond, Peter L. Bartlett, David P. Woodruff, and Santosh Vempala at Simons Institute at the University of California, Berkeley, Nikhil Srivastava at the University of California, Berkeley, Omri Weinstein, Alex Andoni and Peilin Zhong at Columbia University, and Sanjeev Arora at Princeton University.

I would like to thank all my office mates at the institutions that I have stayed in over the last few years: Pravesh Kothari during my first year at UT-Austin; Mika Göös and Ilya Razenshteyn at IBM in 2015; Cameron Musco, Badih Ghazi, and Xingguo Li at IBM in 2016; Chen Shao, Lin F. Yang, Peilin Zhong, Hongyang Zhang at IBM in 2017. I had a productive year during my visit to Harvard University in 2017 and 2018, where I shared the office with Aviad Rubinstein, Mika Göös, Kyle Luh, Tselil Schramm, Rasmus Kyng, Tobias Christiani, and Huacheng Yu. I enjoyed my leisure time with Josh Alman, Chi-Ning Chou, Tobias Christiani, Zhixian Lei, Zhengyu Wang and Huacheng Yu while I was at Harvard. I had a productive summer when I was interning MSR Redmond in 2018 while sharing the office with Mark Sellke, and Yuanzhi Li. During my visit to UC Berkeley in 2019 Fall, I shared an office with Lijie Chen, Gautam Kamath, Jerry Li, Samuel Hopkins, Ainesh Bakshi, Simon S. Du, Themis Gouleakis, Rajesh Jayaram, Ruosong Wang, Xiaoxia Wu, Xiaofei Shi, and Hongyang Zhang. I was struggling on my thesis while I was interning MSR Redmond in 2019 Summer while sharing the office with Ronen Eldan, Mark Sellke, Sitan Chen, Lixin Huang, Kevin Tian. I enjoyed my time at the University of Washington where I shared the
office with Xin Yang, Haotian Jiang, Swati Padmanabhan, Victor Reis, Ruoqi Shen, Ewin Tang, and Robbie Weber in 2018 and 2019. I enjoyed my leisure time with Lequn Chen, Zihou Gao, Ziheng Jiang, Weihao Kong, Jialin Li, Hanchuan Li, Daogao Liu, Xuhai Xu, Tianyi Zhou, and Danyang Zhuo while I was at the University of Washington. I thank my office mates Fu Li, Kai Zhong, Qi Lei, Jiong Zhang, and Ruizhe Zhang at UT-Austin for their support during the latter part of my Ph.D. career. I enjoy my leisure time with Ruizhe Zhang during the last a few days of my Ph.D. program at UT-Austin.

Traveling As a frequent traveler, I was homeless for the majority of my Ph.D. years. I would like to thank all the friends who have kindly provided me a place to stay: Feng Xiao at IBM; Xindi Hu and Aviad Rubinstein at Harvard; Xin Yang and Jerry Li during my visit to UW and MSR; and Lijie Chen, Themis Gouleakis, Chi Jin, Gautam Kamath, and Santosh Vempala who shared their house with me while I was visiting UC Berkeley.

Over my entire Ph.D., I don't know how to drive. My life would be very inconvenient if I don't know the following friends who had given me a ride many times. I want to thank for Beidi Chen, Qiren Chen, Lequn Chen, Yu Feng, Chi Jin, Weihao Kong, Meng Li, Xingguo Li, Yuanzhi Li, Yibo Lin, Jialin Li, Jinsong Liu, Aviad Rubinstein, Ilya Razenshteyn, Ruoqi Shen, Aaron Sidford, Xinyu Wang, Ye Wang, David P. Woodruff, Yun Wu, Biying Xu, Feng Xiao, Xin Yang, Lin F. Yang, Bei Yu, Zhuoran Zhao, Ruohan Zhang, Kai Zhong, Danyang Zhuo.

RA and TA Over the past few years, I have received fundings from many generous professors, including Alexandros Dimakis, Inderjit Dhillon, Adam Klivans, Rasmus Kyng, Yin Tat Lee and Eric Price.

Over the past few years, I have been doing TA so many times. I'd like to thank Rasmus Kyng and Bill Young for working with at Harvard and UT-Austin.

Undergraduate The decision to pursue a Ph.D. is already a brave one for me. I really want to thank for my undergraduate advisors Richard Vaughan, Tsunehiko Kameda, Binay K. Bhattacharya, Ke Wang, Martin Ester, Ze-Nian Li, Hao Zhang (at Simon Fraser University) for kindly training and guiding a young researcher. I would like to thank for my undergraduate friends Song Chen, Tiaoxiang Gao, Junsong Hu, Wen Sun, Chi Jin, Yatao Wang, Mingzhou Yang, Huan Zhang, Zhuoran Zhao, Yuke Zhu who encouraged me to pursue Ph.D.

Struggling Time Over the entire Ph.D. time, I have been thinking about quitting the Ph.D. program seriously thousands of times. I really want to thank all the friends who had suggested me not to give up. I had a very struggling time during my first year at Austin. I truly appreciate Adam Klivans and Pravesh Kothari's support without which I might have quitted Ph.D. I really like to thank Alexandros G. Dimakis, Anna Gál, and Ankit Singh Rawat for helping me publishing the first theory paper during my Ph.D. This was extremely important to me, as it gave me the confidence to continue my Ph.D. research.

During the second year, I met my advisor Eric. I want to thank him for giving me a set of toys (Fox, etc.). Those inspire my research for many years. I finally collected six before graduation.

In the summer of my second year, I was an intern at IBM Almaden. This was the first time that I did research outside my home institution. I would like to thank my intern advisor David Woodruff for his help on many things, especially on giving me the first stock in my life. When I was planning to graduate, that stock is five times than before. Since then, I started to collaborate with David heavily, and he became my pseudo-advisor.

While I was at the University of Washington, I was struggling with many things. I want to thank Yin Tat Lee for chatting with me about many non-research kinds of stuff, sharing so many funny stories with me, and giving many bits of help in my life.

Even during the last year of my Ph.D., I was still considering quitting the Ph.D. due to a variety of reasons. I really appreciate the support from so many professors (Scott Aaronson, Greg Plaxton, Adam Klivans, Eric Price, David P. Woodruff) and friends (Zeyuan Allen-Zhu, Josh Alman, Lijie Chen, Gautam Kamath Rasmus Kyng, Yin Tat Lee, Xingguo

Li, Jerry Li, Ilya Razenshteyn, Aviad Rubinstein, Ruosong Wang, Zhengyu Wang, Lin Yang, Xin Yang, Peilin Zhong) who convinced me to finish the Ph.D. program.
??? Friends I wanted to thank Lijie Chen, Zhengyu Wang, and Peilin Zhong for their enourmous help on my research. Their contributions range from proposing new research ideas to paper proofreading. Many times they kindly declined the authorship they deserved. I am so grateful for their help and friendship that I could not find a proper category for them. This is why I put "???" above.

I wanted to thank Ziheng Jiang for his generous help at the University of Washington, without which I would not be able to write this thesis.

I really wanted to thank Billy (Vasileios) and Zhengyu for writing my only student papers with me.

At the end of the third year, I started to get interested in Machine Learning and Deep Learning. I would like to thank my friend Kai Zhong for introducing this new field for me.

Thesis I want to thank Zeyuan Allen-Zhu and Yin Tat Lee for the suggestion of the title of the thesis. I want to thank Jerry Li for choosing the final title. I want to thank Jerry Li (Sana Minatozaki) for his generous help on the abstract, introduction, and the open problem section of the thesis. I want to thank Lijie Chen for his help on acknowledgment. I want to thank Ruoqi Shen for her generous advice on the abstract of the thesis. I appreciate Zhengyu Wang's generous help on the acknowledgment, abstract, introduction and open problem section of the thesis. I want to thank Yin Tat Lee, Victor Reis, Xiaorui Sun, Xin Yang, Peilin Zhong for their generous help on the open problem section.

Defense and Graduation I want to thank Ruizhe Zhang for his generous help during my defense day. I want to thank all my committees and Vijaya Ramachandran for coming to my defense at UT-Austin. I want to thank Yin Tat Lee and Xin Yang for creating two Skype call lives of my defense. One is at MSR Redmond, and the other is at the University of Washington. I want to thank Weihao Kong, Jialin Li, Daogao Liu, Swati Padmanabhan, Ruoqi Shen, Xin Yang, and Danyang Zhuo for joining my defense remotely. I want to thank Beidi Chen, Lixin Huang, Gautam Kamath, Yin Tat Lee, Jerry Li, Chen Luo, Aaron Sidford, and Sam Wong for attending my defense remotely.

I want to thank Katie Dahm (who is the CS Graduate Program Coordinator at The University of Texas at Austin) for giving all kinds of help for letting me graduate on time.

To Junior Ph.D. Student To be honest, I never thought I was a successful Ph.D. student. At the end of my thesis defense, someone asked a question: "do you have any suggestions for junior students?". I am still not sure if I have a satisfactory answer to that question today. Instead, I wanted to write down how I do research, and what I feel when something happened, for the reference of future readers.

1. After every meeting, I will write down the notes as formal as possible via latex. Even if the idea proposed in the meeting is completely wrong, I still write down the details of all the calculations and mark why it doesn't work. To me, the definition of the research progress is not just solving the problem. Sometimes, ruling out some wrong directions also counts as remarkable progress.
2. For the first three months of my Ph.D., I almost spent all the time on solving one problem. But, it didn't go anywhere. I felt truly sad and wanted to quit the Ph.D. program. Then someone suggested me to go to some classes to learn more things and read more papers. I took the suggestion. Luckily, after that class, I got my first theory paper during my Ph.D. This gives me very strong confidence to continue my Ph.D.
3. After that first three months of bad experience, I usually spent more time reading papers than thinking about how to solve problems. I am reading ArXiv papers every day. Usually, once a conference released the accepted paper list, I will google every title of the paper and try to take a quick look.
4. I am in general happy to talk about research with anyone in any field. I am also open to working on any problems in any field. As long as my coauthors want to work on something with me, I am happy to work on it. As far as I remember, I never say "no" to
anyone or any problem.
To clarify, I am just sharing my experience with everyone. I am NOT suggesting anyone try to follow my way.

Parents Finally, it is time to say something for my parents.
Although I am neither a lousy Ph.D. student nor a stupid coauthor, I have to admit that I am a horrible son. Over my entire Ph.D. career, I have only met my parents twice for a total of one month. I still remember the shock that I had during our first meeting. I could barely recognize their faces, and they became so much older compared to the time I left them.

My parents do not know much about Computer Science or graduate school, and I was not patient enough to explain to them what was going on in my life. For all of these, I want to say "sorry" to my parents. I truly appreciate their support in my life. Mom and dad, I love you!

It is difficult to forget the pleasant time with my grandparents during my childhood. I felt really sorry that I did not give enough time to them as I grew older. That is the reason why I have mentioned all of my grandparents' names in this thesis.

## Previous Publications

This thesis is based on several previous publications. We list the chapters together with the corresponding papers.

## Optimization

## Chapter 2 is based upon the following previous publication

- Michael B. Cohen, Yin Tat Lee, Zhao Song

Solving Linear Programs in the Current Matrix Multiplication Time.
STOC 2019 [CLS19]

## Chapter 3 is based upon the following previous publication

- Yin Tat Lee, Zhao Song, Qiuyi Zhang

Solving Empirical Risk Minimization in the Current Matrix Multiplication. COLT 2019 [LSZ19]

## Chapter 4 is based upon the following previous publication

- Yin Tat Lee, Zhao Song, Santosh S. Vempala

Algorithmic Theory of ODEs and Sampling from Well-conditioned Logconcave Densi-
xxiv
ties.
Manuscript 2018 [LSV18]

Chapter 5 is based upon the following previous publication

- Zeyuan Allen-Zhu, Yuanzhi Li, Zhao Song

A Convergence Theory for Deep Learning via Over-Parameterization.
ICML 2019 [AZLS19]

Chapter 6 is based upon the following previous publication

- Zeyuan Allen-Zhu, Yuanzhi Li, Zhao Song

On the convergence rate of training recurrent neural networks.
Manuscript 2018 [AZLS18]

## Chapter 7 is based upon the following previous publication

- Ankit Garg, Yin Tat Lee, Zhao Song, Nikhil Srivastava

A Matrix Expander Chernoff Bound.
STOC 2018 [GLSS18]

## Chapter 8 is based upon the following previous publication

- Rasmus Kyng, Zhao Song

A Matrix Chernoff Bound for Strongly Rayleigh Distributions and Spectral Sparsifiers
from a few Random Spanning Trees.
FOCS 2018 [KS18]

Chapter 9 is based upon the following previous publication

- Rasmus Kyng, Kyle Luh, Zhao Song

Four Deviations Suffice for Rank 1 Matrices.
Manuscript 2019 [KLS19]

Chapter 10 is based upon the following previous publication

- Vasileios Nakos, Zhao Song

Stronger L2/L2 Compressed Sensing; Without Iterating. STOC 2019 [NS19]

Chapter 11 is based upon the following previous publication

- Eric Price, Zhao Song

A Robust Sparse Fourier Transform in the Continuous Setting. FOCS 2015 [PS15]

Chapter 12 is based upon the following previous publication

- Xue Chen, Eric Price, Daniel Kane, Zhao Song.

Fourier-sparse interpolation without a frequency gap.
FOCS 2016 [CKPS16]

Chapter 13 is based upon the following previous publication

- Vasileios Nakos, Zhao Song, Zhengyu Wang
(Nearly) Sample-Optimal Sparse Fourier Transform in Any Dimension; RIPless and Filterless.

FOCS 2019 [NSW19a]

## Chapter 14 is based upon the following previous publication

- Vasileios Nakos, Zhao Song, Zhengyu Wang

The Power of (Careful) Iterative Loop Analysis in Compressed Sensing. Manuscript 2019 [NSW18]

## Chapter 16 is based upon the following previous publication

- Ilya Razenshteyn, Zhao Song, David Woodruff

Weighted low rank approximations with provable guarantees.
STOC 2016 [RSW16]

Chapter 17 is based upon the following previous publication

- Zhao Song, David Woodruff, Peilin Zhong

Low Rank Approximation with Entrywise $\ell_{1}$-Norm Error.
STOC 2017 [SWZ17]

Chapter 19 and 20 are based upon the following previous publication

- Zhao Song, David Woodruff, Peilin Zhong

Towards a Zero-One Law for Entrywise Low Rank Approximation.
Manuscript 2019 [SWZ18]

Chapter 21 is based upon the following previous publication

- Zhao Song, David Woodruff, Peilin Zhong

Relative Error Tensor Low Rank Approximation.
SODA 2019 [SWZ19b]

## Chapter 22 is based upon the following previous publication

- Zhao Song, David Woodruff, Huan Zhang

Sublinear Time Orthogonal Tensor Decomposition.
NeurIPS 2016 [SWZ16]

Chapter 23 is an open problem section, which is based on the all the publications corresponding to Chapters $2 \sim 22$ and also the following previous publications

- Alexandr Andoni, Zhao Song, Clifford Stein, Zhengyu Wang, Peilin Zhong

Parallel graph connectivity in log diameter rounds.
FOCS 2018 [ASS $\left.{ }^{+} 18\right]$

- Aviad Rubinstein, Saeed Seddighin, Zhao Song, Xiaorui Sun Approximation Algorithms for LCS and LIS with Truly Improved Running Times. FOCS 2019 [RSSS19]
- Aviad Rubinstein, Zhao Song

Reducing approximate Longest Common Subsequence to approximate Edit Distance. Manuscript 2019 [RS19b]

We remark that the papers covered by this thesis are only about part of the author's gradschool-publications. Topics not covered in this thesis are deep learning theory [ZSJ+17, ZSD17, WZC ${ }^{+}$18, LSY19, SY19], recurrent neural network [ZLSD18], finding adversarial examples [WZC $\left.{ }^{+} 18, \mathrm{ZCS}^{+} 19\right]$, reinforcement learning [ZSB15, SS19], coding theory [DGRS15], unique games [BCS18], optimization [LPS19, vdBLSS19], learning mixture distributions [CLS20], multi-arm bandits [LPSY18], matrix completion [ZSJD18, BLS ${ }^{+}$19], Fourier transform with restricted samples [LSS19], regression [PSW17, DSSW18, DSWY19, $\mathrm{DJS}^{+} 19, \mathrm{SWY}^{+}$19], Laplacian system [ASS19], low-rank approximation [LSW ${ }^{+}$19], EDA design $\left[W L Y^{+} 15\right]$, fast multipole method [ASS19], geometry algorithms $\left[\mathrm{BBD}^{+} 18, \operatorname{ASS19]}\right.$, robotics [ZS16], quantum algorithm [TS19] and discrepancy [RS19a].

The author also had done a few research when he was an undergraduate student. Since the author never wrote a undergraduate thesis, we like to list those papers here. The topics are robotics [SV12b, SSV12, SV13, SS14], graph and geometry algorithms [BKS12b, BKS12c, BKS14, $\left.\mathrm{BDK}^{+} 14, \mathrm{BKS15b}, \mathrm{BBD}^{+} 16\right]$, data-mining [HSE12, SZ13].

# Matrix Theory: Optimization, Concentration, and Algorithms 

Publication No.<br>$\qquad$<br>Zhao Song, Ph.D.<br>The University of Texas at Austin, 2019

Supervisor: Eric Price

The matrix plays an essential role in many theoretical computer science and machine learning problems. In this thesis, we develop a better understanding of matrices with a view towards these applications. Our insights yield improvements for a number of old, well-studied algorithmic problems.

In this thesis, we study matrices from three perspectives. We first consider their role in optimization. We study a number of matrix optimization problems, and propose new solvers and results for linear programs, empirical risk minimization, ordinary differential equations, deep neural networks. We next consider how random matrices concentrate. Specifically, we generalize a number of scalar Chernoff-type concentration inequalities and the Spencer-type discrepancy theorems to matrices. Finally, we develop new algorithms for problems on matrices. These fall roughly into two sub-categories, namely matrix factorization problems and structured recovery problems. In the first category, we propose a number of new algorithms
for a variety of low-rank matrix factorization problems. In the second category, we give new algorithms for some recovery tasks with structured matrices. We design matrices and corresponding algorithms for compressed sensing tasks, and we give fast algorithms for the sparse Fourier transform problem, which can be thought of as a sparse recovery problem where one cannot freely choose the matrix. We now describe our contributions in more detail.

Linear programming is one of the fundamental problems in computer science. In both theory and practice, many problems can be solved via linear programming, and doing so often yields the best-known runtime. We present an algorithm that runs in the current matrix multiplication time, which breaks a thirty-year-old barrier. Furthermore, our technique can be generalized to speed up a large family of convex optimization problems, i.e., empirical risk minimization.

Sampling logconcave functions which arise in statistics and machine learning has been the subject of intensive study. This problem can be thought of as a special case of solving multivariate ordinary differential equations (ODEs). Under sufficiently strong smoothness conditions, discrete algorithms of Hamiltonian Monte Carlo (HMC) have runtime and number of function evaluations growing with the dimension. We give new algorithms that obtain a nearly linear implementation of HMC for a broad class of smooth, strongly logconcave densities, with the number of iterations (parallel depth) and gradient evaluations being polylogarithmic in the dimension (rather than polynomial, as in previous work).

Deep neural networks (DNNs) have demonstrated dominating performance in many fields; since AlexNet, networks used in practice are going wider and deeper. We prove why stochastic gradient descent (SGD) can find global minima on the training objective of DNNs in polynomial time. We only make two assumptions: the inputs are non-degenerate and
the network is over-parameterized. Furthermore, our results also hold for Recurrent Neural Networks (RNNs) which are multi-layer networks widely used in natural language processing. They are harder to analyze than feedforward neural networks, because the same recurrent unit is repeatedly applied across the entire time horizon.

The Chernoff bound for the concentration of scalar random variables is a fundamental tool in the analysis of randomized algorithms. Over the past decade, a matrix generalization of the Chernoff bound has also found widespread application, but this generalization is fairly restrictive. For a number of decades, it has been open whether one can remove some of these restrictions. We answer this question in the affirmative by giving a number of new matrix Chernoff bounds under more relaxed independence assumptions than before.

Spencer's theorem is a famous result in discrepancy theory, and it is an important open question how to generalize this result to the matrix setting. We make progress on this, and prove a matrix generalization of this result, in some restricted settings. Our result also generalizes the famous Kadison-Singer conjecture.

The classical low rank approximation problem is : given a matrix $A$, find a rank- $k$ matrix $B$ such that the Frobenius norm of $A-B$ is minimized. It can be solved in polynomialtime using, for instance, singular value decomposition. If one allows randomization and approximation, it can be solved in time proportional to the number of non-zero entries of $A$ with high probability. We consider a number of natural generalizations of this important problem. We generalize the problem to consider a number of different norms, including weighted norms, entry-wise L1, and more general loss functions including Huber and Tukey loss. We also consider the natural tensor version of the problem. For all these settings, we give new state-of-the-art algorithms, including a number of new fixed parameter tractable
algorithms.
Compressed Sensing, or sparse recovery, is a powerful mathematical framework whose goal is to reconstruct an approximately sparse signal from linear measurements. We consider the extensively studied problem of L2/L2 compressed sensing. Our algorithm has faster decoding time and significantly smaller column sparsity, answering two open questions of prior work. Previous work on sublinear-time compressed sensing employed an iterative procedure, recovering the heavy coordinates in phases. We completely depart from that framework, and give the first sublinear-time algorithm which achieves the optimal number of measurements without iterating.

The Fourier transform is ubiquitous in computer science, mathematics, and beyond. In recent years, a number of works have studied methods for computing the Fourier transform in sublinear time if the output is sparse. We consider two important settings in which this occurs: namely, the sparse high-dimensional setting, and the sparse continuous setting. In the high dimensional setting, we consider the extensively studied problem of computing a $k$-sparse approximation to the $d$-dimensional Fourier transform of a length $n$ signal. Our algorithm achieves a nearly optimal sample complexity and runs in time comparable to the Fast Fourier Transform. All previous algorithms proceed either via the Restricted Isometry Property or via filter functions. Our approach totally departs from the aforementioned techniques, and we believe is a new approach to the sparse Fourier transform problem.

While many of the works on sparse Fourier transforms have focused on the discrete setting, in many applications the input signal is continuous and naive discretization significantly worsens the sparsity level. Assuming the frequency gap is known, we present an algorithm for robustly computing sparse Fourier transforms in the continuous setting. Know-
ing the frequency gap is necessary for robustly identifying individual frequencies. However, was unknown whether interpolating the signal also requires a frequency gap. We resolve this problem by giving an algorithm which shows that such a gap not necessary for estimating the signal.

## Table of Contents

Acknowledgments ..... vi
Abstract ..... xxx
List of Tables ..... lxvi
List of Figures ..... lxxi
Chapter 1. Introduction ..... 1
1.1 Introduction ..... 1
1.2 Optimization ..... 10
1.3 Concentration ..... 25
1.4 Compressed Sensing and Sparse Fourier Transform ..... 33
1.5 Factorization ..... 48
1.6 More Algorithms ..... 61
Part I Optimization ..... 1
Chapter 2. Linear Programs ..... 2
2.1 Introduction ..... 3
2.1.1 Related Work ..... 6
2.2 Results and Techniques ..... 7
2.2.1 Central Path Method ..... 9
2.2.1.1 Short Step Central Path Method ..... 10
2.2.1.2 Stochastic Central Path Method ..... 11
2.2.2 Projection Maintenance via Lazy Update ..... 14
2.3 Notations ..... 17
2.4 Stochastic Central Path Method ..... 18
2.4.1 Proof Outline ..... 18
2.4.2 Bounding each quantity of stochastic step ..... 23
2.4.2.1 Bounding $\widetilde{\delta}_{s}, \widetilde{\delta}_{x}$ and $\widetilde{\delta}_{\mu}$ ..... 25
2.4.2.2 Bounding $\mu^{\text {new }}-\mu$ ..... 29
2.4.3 Stochastic central path ..... 34
2.4.4 Analysis of cost per iteration ..... 41
2.4.5 Main result ..... 44
2.5 Projection Maintenance ..... 46
2.5.1 Proof outline ..... 48
2.5.2 Proof of Theorem 2.5.1 ..... 51
2.5.3 Initialization time, update time, query time ..... 54
2.5.4 Bounding $w$ move ..... 58
2.5.5 Bounding $v$ move ..... 62
2.5.6 Potential function $\psi$ ..... 65
2.6 Omitted Proofs ..... 66
Chapter 3. Empirical Risk Minimization ..... 72
3.1 Introduction ..... 74
3.1.1 Related Work ..... 77
3.2 Overview of Techniques ..... 79
3.2.1 Central Path Method ..... 81
3.2.2 Robust Central Path ..... 83
3.2.3 Speeding up via Sketching ..... 85
3.2.4 Maintaining the Sketch ..... 87
3.2.5 Fast rectangular matrix multiplication ..... 89
3.3 Preliminaries ..... 90
3.4 Robust Central Path ..... 91
3.4.1 Newton Step ..... 92
3.4.2 Robust Central Path Method ..... 94
3.5 Robust Central Path ..... 99
3.5.1 Outline of Analysis ..... 100
3.5.2 Changes in $\mu$ and $\gamma$ ..... 103
3.5.3 Movement from $(x, x, s)$ to $\left(x^{\text {new }}, x, s^{\text {new }}\right)$ ..... 107
3.5.4 Movement from ( $\left.x^{\text {new }}, x, s^{\text {new }}\right)$ to $\left(x^{\text {new }}, x^{\text {new }}, s^{\text {new }}\right)$ ..... 112
3.5.5 Movement of $t$ ..... 114
3.5.6 Potential Maintenance ..... 116
3.6 Central Path Maintenance ..... 119
3.6.1 Proof of Theorem 3.6.1 ..... 123
3.6.2 Initialization time, update time, query time, move time, multiply time ..... 129
3.6.3 Bounding $W$ move ..... 135
3.6.4 Bounding $V$ move ..... 139
3.6.5 Potential function $\psi$ ..... 144
3.6.6 $x$ and $\bar{x}$ are close ..... 146
3.6.7 $s$ and $\bar{s}$ are close ..... 149
3.6.8 Data structure is maintaining ( $x, s$ ) implicitly over all the iterations ..... 150
3.7 Combining Robust Central Path with Data Structure ..... 153
3.7.1 Guarantee for $W$ matrices ..... 154
3.7.2 Main result ..... 156
3.8 Initial Point and Termination Condition ..... 160
3.9 Basic Properties of Subsampled Hadamard Transform Matrix ..... 164
3.9.1 Properties obtained by random projection ..... 165
Chapter 4. Algorithmic Theory of ODEs ..... 169
4.1 Introduction ..... 171
4.1.1 Results ..... 174
4.1.2 HMC and improved contraction rate ..... 178
4.1.3 Techniques ..... 180
4.2 ODE Solver for any basis ..... 182
4.2.1 Picard-Lindelöf theorem ..... 182
4.2.2 Intuition of collocation method ..... 184
4.2.3 Collocation method ..... 186
4.2.4 Proof of first order ODE ..... 191
4.2.5 A basis for piece-wise polynomials ..... 196
4.3 Improved Contraction Bound for HMC ..... 199
4.4 Strongly Convex functions with Lipschitz Gradient ..... 207
4.4.1 Bounding the ODE solution ..... 207
4.4.2 Sampling ..... 212
4.5 Sampling from Incoherent Logistic Loss Functions and More ..... 215
4.5.1 $\ell_{\infty}$ bound of the dynamic ..... 215
4.5.2 Lipschitz constant of $F$ ..... 219
4.5.3 Existence of low-degree solutions ..... 220
4.5.4 Main result ..... 223
4.6 Preliminaries ..... 229
4.6.1 Notation ..... 229
4.6.2 Simple ODEs ..... 230
4.6.3 Cauchy Estimates and Method of Majorants ..... 232
4.7 Deferred Proof for ODE (Section 4.2) ..... 235
4.7.1 Proof of general $k$-th order ODE ..... 235
4.8 Deferred Proof for Cauchy Estimates (Section 4.6.3) ..... 241
4.9 Cauchy Estimates of Some Functions ..... 245
4.9.1 Logistic loss function ..... 245
4.9.2 Pseudo-Huber loss function ..... 246
Chapter 5. Convergence result of Deep Neural Network ..... 248
5.1 Introduction ..... 250
5.1.1 Our Result ..... 253
5.1.2 Other Related Works ..... 257
5.2 Preliminaries ..... 259
5.2.1 Objective and Gradient ..... 262
5.3 Our Results and Techniques ..... 263
5.4 Conceptual Messages and Technical Theorems ..... 265
5.4.1 Objective is Almost Convex and Semi-Smooth ..... 267
5.4.2 Equivalence to Neural Tangent Kernel ..... 269
5.5 Proof Overview ..... 271
5.6 Notable Extensions ..... 275
5.7 Detailed Proofs ..... 278
5.8 Properties at Random Initialization ..... 279
5.8.1 Forward Propagation ..... 280
5.8.2 Intermediate Layers ..... 284
5.8.3 Backward Propagation ..... 288
5.8.4 $\delta$-Separateness ..... 289
5.8.4.1 Auxiliary Claim ..... 290
5.9 Stability against Adversarial Weight Perturbations ..... 293
5.9.1 Forward Perturbation ..... 294
5.9.1.1 Auxiliary Claim ..... 297
5.9.2 Intermediate Layers ..... 300
5.9.3 Backward ..... 303
5.10 Gradient Bound at Random Initialization ..... 304
5.10.1 Proof of Lemma 5.10.1: Upper Bound ..... 305
5.10.2 Proof of Lemma 5.10.1: Lower Bound ..... 306
5.10.2.1 Proof of Claim 5.10.4 ..... 309
5.11 Theorem 5.11.1: Gradient Bound at After Perturbation ..... 311
5.12 Theorem 5.12.1: Objective Semi-Smoothness ..... 313
5.12.1 Proof of Claim 5.12.2 ..... 316
5.13 Theorem 5.13.1: Convergence Rate of GD ..... 319
5.14 Theorem 5.14.1: Convergence Rate of SGD ..... 322
5.15 Theorem 5.15.1: Equivalence to Neural Tangent Kernel ..... 328
5.16 Extension to Other Loss Functions ..... 331
5.17 Extension to Convolutional Neural Networks ..... 334
5.17.1 Changes in the Proofs ..... 336
5.18 Extension to Residual Neural Networks ..... 339
5.18.1 Changes in the Proofs ..... 340
Chapter 6. Convergence result of Recurrent Neural Network ..... 350
6.1 Introduction ..... 351
6.1.1 Our Question ..... 354
6.1.2 Other Related Works ..... 358
6.2 Notations and Preliminaries ..... 359
6.2.1 Elman Recurrent Neural Network ..... 360
6.2.2 Objective and Gradient ..... 362
6.3 Our Results ..... 363
6.3.1 Conclusion ..... 366
6.4 Proof Sketch ..... 367
6.5 Basic Properties at Random Initialization ..... 368
6.6 Stability After Adversarial Perturbation ..... 372
6.7 Proof Sketch of Theorem 6.15.2: Polyak-Łojasiewicz Condition ..... 375
6.7.1 Indicator and Backward Coordinate Bounds ..... 377
6.7.2 Thought Experiment: Adding Small Rank-One Perturbation ..... 379
6.7.3 Real Proof: Randomness Decomposition and McDiarmid's Inequality ..... 381
6.8 Proof Sketch of Theorem 6.16.1: Objective Semi-Smoothness ..... 384
6.9 Roadmap of Later Proof Details ..... 386
6.10 Preliminaries on Probability Theory ..... 388
6.10.1 Swapping Randomness ..... 389
6.10.2 Concentration of Chi-Square Distribution ..... 390
6.10.3 Concentration of Sum of Squares of ReLU of Gaussians ..... 392
6.10.4 Gaussian Vector Percentile: Center ..... 396
6.10.5 Gaussian Vector Percentile: Tail ..... 399
6.10.6 McDiarmid's Inequality and An Extension ..... 401
6.11 Basic Properties at Random Initialization ..... 404
6.11.1 Forward Propagation ..... 406
6.11.2 Forward Correlation ..... 410
6.11.2.1 Tools ..... 412
6.11.3 Forward $\delta$-Separateness ..... 416
6.11.4 Intermediate Layers: Spectral Norm ..... 421
6.11.4.1 Tools ..... 423
6.11.5 Intermediate Layers: Sparse Spectral Norm ..... 425
6.11.6 Backward Propagation ..... 427
6.12 Stability After Adversarial Perturbation ..... 430
6.12.1 Forward ..... 432
6.12.1.1 Tools ..... 435
6.12.2 Intermediate Layers ..... 440
6.12.3 Backward ..... 442
6.12.4 Special Rank-One Perturbation ..... 445
6.13 Indicator and Backward Coordinate Bounds ..... 451
6.13.1 Indicator Coordinate Bound ..... 452
6.13.2 Backward Coordinate Bound ..... 457
6.14 Gradient Bound at Random Initialization (Theorem 6.14.2) ..... 462
6.14.1 Randomness Decomposition ..... 465
6.14.2 Gradient Lower Bound in Expectation ..... 469
6.14.2.1 Proof of Lemma 6.14.5 ..... 469
6.14.2.2 Core Lemma A ..... 474
6.14.3 Gradient Stability ..... 477
6.14.3.1 Proof of Lemma 6.14.7 ..... 477
6.14.3.2 Core Lemma B ..... 480
6.14.4 Main Theorem: Proof of Theorem 6.14.2 ..... 482
6.15 Gradient Bound After Perturbation (Theorem 6.15.2) ..... 485
6.16 Objective Semi-Smoothness (Theorem 6.16.1) ..... 489
6.16.1 Tool ..... 494
6.17 Convergence Rate of Gradient Descent (Theorem 6.17.1) ..... 495
6.18 Convergence Rate of Stochastic Gradient Descent (Theorem 6.18.1) ..... 497
Part II Matrix Concentrations ..... 500
Chapter 7. Matrix Chernoff Bound on Expander ..... 501
7.1 Introduction ..... 502
7.1.1 A Matrix Expander Chernoff Bound ..... 504
7.1.2 Martingale Approximation of Expander Walks ..... 509
7.2 Preliminaries ..... 512
7.2.1 Linear Algebra ..... 512
7.2.2 Complex analysis ..... 515
7.3 New Golden-Thompson inequality ..... 517
7.3.1 Complex estimate on the half disk ..... 518
7.3.2 Bounded Multimatrix Golden-Thompson type inequality ..... 522
7.4 Proof of Theorem 7.1.2 ..... 526
7.5 Proof of Theorem 7.1.6 ..... 538
7.6 Elementary calculations ..... 540
Chapter 8. Matrix Chernoff Bound for Random Spanning Trees ..... 546
8.1 Introduction ..... 548
8.1.1 Previous work ..... 553
8.1.2 Our results and techniques ..... 556
8.2 Notation ..... 564
8.3 Preliminaries ..... 565
8.3.1 Useful facts and tools ..... 565
8.3.2 Strongly Rayleigh distributions ..... 566
8.3.3 Random spanning trees ..... 567
8.4 A Matrix Chernoff Bound for Strongly Rayleigh Distributions ..... 569
8.4.1 Main result ..... 576
8.5 Applications to Random Spanning Trees ..... 578
8.6 Lower bounds ..... 580
8.6.1 Single spanning tree, low probability ..... 580
8.6.2 Single spanning tree, high probability ..... 583
8.6.3 Sum of a batch of spanning trees ..... 584
8.7 Shrinking Marginals Lemma ..... 588
Chapter 9. Four Derivation Suffice for Rank 1 Matrix ..... 589
9.1 Introduction ..... 590
9.1.1 Related work and open questions ..... 594
9.1.2 Our techniques ..... 597
9.2 Preliminaries ..... 598
9.2.1 Linear Algebra ..... 598
9.2.2 Real Stability ..... 598
9.2.3 Interlacing Families ..... 599
9.3 Expected Characteristic Polynomial ..... 601
9.4 Defining the Interlacing Family ..... 605
9.5 Largest Root of the Expected Characteristic Polynomial ..... 607
9.6 Omitted Proofs ..... 612
9.6.1 Proof of Lemma 9.5.1 ..... 612
Part III Compressive Sensing and Sparse Fourier Transform ..... 615
Chapter 10. Compressive Sensing ..... 616
10.1 Introduction ..... 617
10.1.1 Previous work ..... 619
10.1.2 Our result ..... 623
10.1.3 Notation ..... 623
10.1.4 Technical statements ..... 624
10.1.5 Overview of techniques and difference with previous work ..... 626
10.1.5.1 Summary of [GLPS10]. ..... 626
10.1.5.2 Our approach ..... 628
10.1.6 Possible applications of our approach to other problems ..... 634
10.2 The Identification Linear Sketch ..... 638
10.2.1 Design of the sketch and decoding algorithm ..... 638
10.2.2 Concentration of estimation ..... 641
10.2.3 Analysis of running time ..... 645
10.2.4 Guarantees of the algorithm ..... 648
10.2.5 Bounding the number of measurements ..... 650
10.3 The Pruning Linear Sketch ..... 651
10.3.1 Design of the sketching matrix, and helpful definitions ..... 651
10.3.2 Analyzing head and badly-estimated coordinates ..... 655
10.3.3 Guarantees of the algorithm ..... 658
10.3.4 Time, measurements, column sparsity, and probability ..... 664
10.4 Tail Estimation ..... 665
10.4.1 Random walks ..... 665
10.4.2 $p$-stable distributions ..... 666
10.4.3 $\ell_{p}$-tail estimation algorithm ..... 668
10.5 Putting It All Together ..... 674
Chapter 11. Continuous Fourier Transform I ..... 675
11.1 Introduction ..... 676
11.1.1 Comparison to Naive Methods ..... 683
11.1.2 Previous Work In Similar Settings ..... 686
11.2 Algorithm Overview ..... 688
11.3 Proof outline ..... 693
11.4 Notation and Definitions: Permutation, Hashing, Filters ..... 700
11.5 Proofs of basic hashing-related lemmas ..... 709
11.6 Proofs for one stage of recovery ..... 715
11.7 Proofs for combining multiple stages ..... 733
11.8 Proofs for converting (11.2) into (11.3) ..... 743
11.9 Lower Bound ..... 752
Chapter 12. Continuous Fourier Transform II ..... 757
12.1 Introduction ..... 758
12.1.1 Related work ..... 762
12.1.2 Our techniques ..... 766
12.1.3 Organization ..... 772
12.2 Proof Sketch ..... 773
12.3 Preliminaries ..... 779
12.3.1 Notation ..... 780
12.3.2 Facts about the Fourier transform ..... 781
12.3.3 Tools and inequalities ..... 784
12.3.4 Legendre polynomials ..... 785
12.3.5 Gram matrix and its determinant ..... 786
12.4 Robust Polynomial Interpolation Algorithm ..... 787
12.4.1 Constant success probability ..... 788
12.4.2 Boosting success probability ..... 793
12.5 Bounding the Magnitude of a Fourier-sparse Signal in Terms of Its Average Norm ..... 796
12.5.1 Bounding the maximum inside the interval ..... 797
12.5.2 Bounding growth outside the interval ..... 800
12.6 Hash Functions and Filter Functions ..... 802
12.6.1 Permutation function and hash function ..... 803
12.6.2 Filter function ..... 804
12.6.3 HashToBins ..... 806
12.7 Frequency Recovery ..... 808
12.7.1 Overview ..... 809
12.7.2 Analysis of GetLegal1Sample and GetEmpirical1Energy ..... 813
12.7.3 A cluster of frequencies, times $H$, is a one-cluster signal per Defini- tion 12.7.1 ..... 820
12.7.4 Frequency recovery of one-cluster signals ..... 822
12.7.5 The full signal, after multiplying by $H$ and convolving with $G$, is one- clustered. ..... 828
12.7.6 Frequency recovery of $k$-clustered signals ..... 835
12.7.7 Time and sample complexity of frequency recovery of $k$-clustered signals ..... 837
12.8 One-cluster Signal Recovery ..... 840
12.8.1 Overview ..... 840
12.8.2 Bounding the Gram matrix determinant ..... 843
12.8.3 Perturbing the frequencies does not change the subspace much ..... 845
12.8.4 Existence of nearby $k$-Fourier-sparse signal with frequency gap bounded away from zero ..... 848
12.8.5 Approximating $k$-Fourier-sparse signals by polynomials ..... 850
12.8.6 Transferring degree- $d$ polynomial to $(d+1)$-Fourier-sparse signal ..... 852
$12.9 k$-cluster Signal Recovery ..... 855
12.9.1 Overview ..... 855
12.9.2 Heavy clusters separation ..... 856
12.9.3 Approximating clusters by polynomials ..... 858
12.9.4 Main result, with constant success probability ..... 861
12.9.5 Boosting the success probability ..... 865
12.10Technical Proofs ..... 866
12.10.1Proof of Theorem 12.8.2 ..... 866
12.10.2Proofs of Lemma 12.5.3 and Lemma 12.5.4 ..... 871
12.10.3Proof of Lemma 12.4.3 ..... 876
12.10.4Proof of Lemma 12.6 .1 ..... 877
12.10.5Proof of Lemma 12.3.5 ..... 878
12.10.6Proof of Lemma 12.3.9 ..... 879
12.11Known Facts ..... 881
12.11. Linear regression ..... 881
12.11.2Multipoint evaluation of a polynomial ..... 882
12.12Analysis of Hash Functions and Filter Functions ..... 883
12.12.1Analysis of filter function $(H(t), \widehat{H}(f))$ ..... 883
12.12.2Analysis of filter function $(G(t), \widehat{G}(f))$ ..... 899
12.12.3Parameters setting for filters ..... 901
12.12.4Analysis of HashToBins ..... 903
12.13Algorithm ..... 910
Chapter 13. High Dimensional Fourier Transform ..... 916
13.1 Introduction ..... 917
13.1.1 Preliminaries ..... 921
13.1.2 Our result ..... 921
13.1.3 Summary of previous Filter function based technique ..... 924
13.1.4 RIP property-based algorithms: a quick overview ..... 925
13.1.5 Overview of our technique ..... 926
13.1.5.1 $O(k \log n)$ measurements for $k=O(\log n)$ ..... 929
13.1.5.2 $O(k \log k \log n)$ measurements for arbitrary $k$ ..... 932
13.2 Algorithm for $d$-dimensional Sparse Fourier Transform ..... 937
13.2.1 Notations ..... 937
13.2.2 Algorithm ..... 940
13.2.3 Analysis ..... 943
Chapter 14. Set Query ..... 963
14.1 Introduction ..... 964
14.2 Classical set query ..... 966
14.2.1 Construction of sketching matrix ..... 966
14.2.2 Algorithm ..... 967
14.2.3 Iterative loop analysis ..... 968
14.2.4 Main result ..... 973
14.3 Fourier set query ..... 976
14.3.1 Definitions and some backgrounds of Fourier transform ..... 976
14.3.2 Permutation and filter function ..... 977
14.3.3 Collision event, large offset event, and large noise event ..... 978
14.3.4 Iterative loop analysis ..... 979
14.3.5 Main result ..... 984
Chapter 15. Robust Sparse Recovery via M-estimators ..... 987
15.1 Introduction ..... 988
15.1.1 Our techniques ..... 991
15.1.2 Our results ..... 992
15.1.2.1 Set-query ..... 993
15.1.2.2 Sparse Recovery ..... 996
15.2 Algorithm and Analysis ..... 996
15.2.1 Set query algorithm for Huber ..... 997
15.2.2 Proof of Theorem 15.1.3 ..... 1000
15.3 Conclusion ..... 1006
15.4 Preliminaries ..... 1007
15.4.1 Huber estimator ..... 1007
15.4.2 Tukey estimator ..... 1007
15.5 Set Query ..... 1009
15.5.1 $\epsilon$-approximation for reverse Huber function ..... 1011
15.5.2 $\epsilon$-approximation for Tukey function ..... 1013
15.6 Fourier set query ..... 1015
15.7 Sparse recovery ..... 1020
15.7.1 $k$-sparse recovery for Huber function ..... 1020
15.7.2 $k$-sparse recovery for Tukey function ..... 1028
Part IV Matrix Factorizations ..... 1033
Chapter 16. Weighted Low Rank Approximation ..... 1034
16.1 Introduction ..... 1036
16.1.1 Our Results ..... 1040
16.1.2 Other Results ..... 1044
16.1.3 Our Techniques ..... 1045
16.2 Preliminaries ..... 1051
16.3 Multiple Regression Sketch ..... 1055
16.4 Additive Approximation ..... 1058
16.4.1 Handling Weight Matrices With Zero Entries ..... 1062
16.5 Multiplicative Approximation ..... 1063
16.6 Recovering the Solution Itself ..... 1068
16.7 Adversarial Matrix Completion ..... 1069
16.8 Few distinct columns ..... 1072
16.8.1 Few Distinct Rows and Columns ..... 1073
16.8.2 Few Distinct Columns When OPT $=0$ ..... 1075
16.8.3 Few Distinct Columns When OPT $\neq 0$ ..... 1078
16.9 Hardness ..... 1081
16.10Bicriteria Approximation ..... 1083
16.10.1Rank $r$ and Binary Weights $W$ ..... 1083
16.10.2Modulo $p$ Integer Weights $W$ ..... 1085
16.10.3r Distinct Columns Weights W ..... 1087
16.11Weighted Nonnegative Matrix Factorization Problem ..... 1088
16.11.1Combining with Our Techniques ..... 1092
16.11.2Few Entries Determine $B$ and $C$ ..... 1094
16.11.3A Semi-Algebraic Set with Double Exponential Depence in $r$ ..... 1099
16.11.4A Semi-Algebraic Set with Single Exponential Dependence in $r$ ..... 1102
Chapter 17. Entry-wise L1 Low Rank Approximation ..... 1108
17.1 Introduction ..... 1109
17.1.1 Our Results ..... 1113
17.1.2 Technical Overview ..... 1116
17.1.3 Several Theorem Statements, an Algorithm, and a Roadmap ..... 1128
17.2 Notation ..... 1130
17.3 Preliminaries ..... 1131
17.3.1 Polynomial system verifier ..... 1131
17.3.2 Cauchy and $p$-stable transform ..... 1132
17.3.3 Lewis weights ..... 1133
17.3.4 Frobenius norm and $\ell_{2}$ relaxation ..... 1135
17.3.5 Converting entry-wise $\ell_{1}$ and $\ell_{p}$ objective functions into polynomials ..... 1137
17.3.6 Converting entry-wise $\ell_{1}$ objective function into a linear program ..... 1138
17.4 $\ell_{1}$-Low Rank Approximation ..... 1139
17.4.1 Existence results via dense Cauchy transforms, sparse Cauchy trans- forms, Lewis weights ..... 1140
17.4.2 Input sparsity time, $\operatorname{poly}(k, \log n, \log d)$-approximation for an arbitrary matrix $A$ ..... 1145
17.4.3 $\operatorname{poly}(k, \log d)$-approximation for an arbitrary matrix $A$ ..... 1150
17.4.4 $\widetilde{O}(k)$-approximation for an arbitrary matrix $A$ ..... 1153
17.4.5 Rank- $3 k$ and $O(1)$-approximation algorithm for an arbitrary matrix $A$ ..... 1159
17.4.6 CUR decomposition for an arbitrary matrix $A$ ..... 1164
17.4.7 Rank- $r$ matrix $B$ ..... 1168
17.4.7.1 Properties ..... 1168
17.4.7.2 poly $(k, r)$-approximation for rank- $r$ matrix $B$ ..... 1171
17.5 Contraction and Dilation Bound for $\ell_{1}$ ..... 1176
17.5.1 Definitions ..... 1177
17.5.2 Properties ..... 1179
17.5.3 Cauchy embeddings, no dilation ..... 1186
17.5.4 Cauchy embeddings, no contraction ..... 1189
17.5.5 Cauchy embeddings, $k$-dimensional subspace ..... 1190
17.5.6 Sparse Cauchy transform ..... 1195
17.5.7 $\ell_{1}$-Lewis weights ..... 1199
17.6 Hardness Results for Cauchy Matrices, Row Subset Selection, OSE ..... 1202
17.6.1 Hard instance for Cauchy matrices ..... 1203
17.6.2 Hard instance for row subset selection ..... 1211
17.6.3 Hard instance for oblivious subspace embedding and more row subset selection ..... 1219
17.6.3.1 Definitions ..... 1219
17.6.3.2 Main results ..... 1221
17.7 $\ell_{p}$-Low Rank Approximation ..... 1244
17.7.1 Definitions ..... 1245
17.7.2 Properties ..... 1247
17.7.3 Tools and inequalities ..... 1254
17.7.4 Dense $p$-stable transform ..... 1256
17.7.5 Sparse $p$-stable transform ..... 1258
17.7.6 $\ell_{p}$-Lewis weights ..... 1261
17.8 EMD-Low Rank Approximation ..... 1262
17.8.1 Definitions ..... 1263
17.8.2 Analysis of no contraction and no dilation bound ..... 1264
17.9 Hardness ..... 1267
17.9.1 Previous results ..... 1268
17.9.2 Extension to multiplicative error $\ell_{1}$-low rank approximation ..... 1269
17.9.3 Using the ETH assumption ..... 1273
17.9.4 Extension to the rank- $k$ case ..... 1279
17.10Experiments and Discussions ..... 1285
17.10.1Setup ..... 1286
17.10.2Counterexample for [DZHZ06] ..... 1287
17.10.3Counterexample for [BDB13] ..... 1289
17.10.4Counterexample for [Kwa08] ..... 1291
17.10.5Counterexample for [KK05] ..... 1295
17.10.6Counterexample for all ..... 1298
17.10.7Discussion for Robust PCA [CLMW11] ..... 1300
17.11Limited Independent Cauchy Random Variables ..... 1302
17.11.1Notations and tools ..... 1303
17.11.2Analysis of limited independent random Cauchy variables ..... 1304
17.12Streaming Setting ..... 1310
17.12.1Definitions ..... 1311
17.12.2Turnstile model, $\operatorname{poly}(k, \log (d), \log (n))$ approximation ..... 1312
17.12.3Row-update model, poly $(k) \log d$ approximation ..... 1315
17.13Distributed Setting ..... 1318
17.13.1Definitions ..... 1319
17.13.2Arbitrary-partition model, subspace, $\operatorname{poly}(k, \log (d), \log (n))$ approxima- tion ..... 1320
17.13.3Arbitrary-partition model, decomposition, $\operatorname{poly}(k, \log (d), \log (n))$ approx- imation ..... 1322
17.13.4Row-partition model, subspace, $\operatorname{poly}(k) \log d$ approximation ..... 1324
17.13.5Row-partition model, decomposition, $\operatorname{poly}(k) \log d$ approximation ..... 1326
Chapter 18. L1 Low Rank Approximation with Regularization ..... 1331
18.1 Introduction ..... 1333
18.1.1 Our results ..... 1335
18.1.2 Technique overview ..... 1337
18.2 Preliminaries ..... 1343
18.2.1 Frobenius norm ridge low-rank approximation and $\ell_{2}$ relaxation ..... 1343
18.2.2 Cauchy transforms ..... 1344
18.3 The Algorithm ..... 1347
18.3.1 Sketching Lemmas ..... 1348
18.3.2 Reducing to small problems ..... 1351
18.3.3 Solving small problems ..... 1354
18.4 Proof of Lemma 18.3.1 ..... 1358
18.5 Proof of Lemma 18.3.3 ..... 1363
18.6 Proof of Theorem 18.3.7 ..... 1365
18.7 Proof of Theorem 18.3.8 ..... 1367
Chapter 19. Average Case L1 Low Rank Approximation ..... 1369
19.1 Introduction ..... 1370
19.1.1 Notation ..... 1371
19.1.2 Our Results ..... 1372
19.1.3 Our Techniques ..... 1373
$19.2 \ell_{1}$-Norm Column Subset Selection ..... 1377
19.2.1 Properties of the Noise Matrix ..... 1378
19.2.2 Definition of Tuples and Cores ..... 1381
19.2.3 Properties of a Good Tuple and a Coefficients Tuple ..... 1383
19.2.4 Main Result ..... 1385
19.3 Missing Proofs in Section 19.2 ..... 1388
19.3.1 Proof of Lemma 19.2.1 ..... 1388
19.3.2 Proof of Lemma 19.2.2 ..... 1389
19.3.3 Proof of Lemma 19.2.3 ..... 1392
19.3.4 Proof of Lemma 19.2.4 ..... 1393
19.3.5 Proof of Lemma 19.2.5 ..... 1394
19.3.6 Proof of Lemma 19.2.6 ..... 1396
19.3.7 Proof of Lemma 19.2.7 ..... 1397
19.3.8 Proof of Lemma 19.2.8 ..... 1397
19.4 Hardness Result ..... 1399
19.4.1 A Useful Fact ..... 1402
19.4.2 One-Sided Error Concentration Bound for a Random Cauchy Matrix ..... 1403
19.4.3 "For Each" Guarantee ..... 1405
19.4.4 From "For Each" to "For All" via an $\epsilon$-Net ..... 1406
19.4.5 Bounding the Cost from the Large-Entry Part via "Bad" Regions ..... 1408
19.4.6 Cost from the Sign-Agreement Part of the Small-Entry Part ..... 1413
19.4.7 Cost from the Sign-Disagreement Part of the Small-Entry Part ..... 1416
19.4.8 Overall Cost of the Small-Entry Part ..... 1418
19.4.9 Main result ..... 1420
Chapter 20. Towards a Zero-One Law for Low Rank Approximation ..... 1424
20.1 Introduction ..... 1426
20.1.1 Our Results ..... 1429
20.1.1.1 A Zero-One Law ..... 1430
20.1.1.2 Lower Bound on the Number of Columns ..... 1432
20.1.2 Overview of our Approach and Related Work ..... 1433
20.2 Algorithm for General Loss Low Rank Approximation ..... 1437
20.2.1 Properties of Uniform Column Sampling ..... 1438
20.2.2 Correctness of the Algorithm ..... 1439
20.3 Missing Proofs in Section 20.2 ..... 1441
20.3.1 Proof of Lemma 20.2.1 ..... 1441
20.3.2 Proof of Lemma 20.2.2 ..... 1442
20.3.3 Proof of Claim 20.2.3 ..... 1444
20.3.4 Proof of Claim 20.2.4 ..... 1445
20.3.5 Proof of Lemma 20.2.5 ..... 1446
20.3.6 Proof of Theorem 20.1.2 ..... 1447
20.4 Necessity of the Properties of $g$ ..... 1450
20.4.1 Functions without Approximate Triangle Inequality ..... 1451
20.4.2 ReLU Function Low Rank Approximation ..... 1453
20.5 Regression Solvers ..... 1455
20.5.1 Regression for Convex $g$ ..... 1455
20.5.2 $\ell_{p}$ Regression ..... 1456
20.5.3 $\ell_{0}$ Regression ..... 1457
20.6 Hardness ..... 1460
20.6.1 Column Subset Selection for the Huber Function ..... 1460
20.6.2 Column Subset Selection for the Reverse Huber Function ..... 1465
Chapter 21. Tensor Low Rank Approximation ..... 1466
21.1 Introduction ..... 1468
21.1.1 Our Results ..... 1473
21.1.2 Our Techniques ..... 1479
21.1.3 Other Low Rank Approximation Algorithms Following Our Framework. ..... 1484
21.1.4 Comparison to [BCV14] ..... 1491
21.1.5 A Roadmap ..... 1492
21.2 Notation ..... 1493
21.3 Preliminaries ..... 1498
21.3.1 Subspace Embeddings and Approximate Matrix Product ..... 1499
21.3.2 Tensor CURT decomposition ..... 1501
21.3.3 Polynomial system verifier ..... 1506
21.3.4 Lower bound on the cost of a polynomial system ..... 1507
21.3.5 Frobenius norm and $\ell_{2}$ relaxation ..... 1508
21.3.6 CountSketch and Gaussian transforms ..... 1510
21.3.7 Cauchy and $p$-stable transforms ..... 1513
21.3.8 Leverage scores ..... 1514
21.3.9 Lewis weights ..... 1515
21.3.10TEnsorSketch ..... 1518
21.4 Frobenius Norm for Arbitrary Tensors ..... 1520
21.4.1 $(1+\epsilon)$-approximate low-rank approximation ..... 1522
21.4.2 Input sparsity reduction ..... 1529
21.4.3 Tensor multiple regression ..... 1533
21.4.4 Bicriteria algorithms ..... 1535
21.4.4.1 Solving a small regression problem ..... 1535
21.4.4.2 Algorithm I ..... 1538
21.4.4.3 poly $(k)$-approximation to multiple regression ..... 1545
21.4.4.4 Algorithm II ..... 1547
21.4.5 Generalized matrix row subset selection ..... 1550
21.4.6 Column, row, and tube subset selection, $(1+\epsilon)$-approximation ..... 1556
21.4.7 CURT decomposition, $(1+\epsilon)$-approximation ..... 1560
21.4.7.1 Properties of leverage score sampling and BSS sampling ..... 1560
21.4.7.2 Row sampling for linear regression ..... 1562
21.4.7.3 Leverage scores for multiple regression ..... 1566
21.4.7.4 Sampling columns according to leverage scores implicitly, im- proving polynomial running time to nearly linear running time ..... 1568
21.4.7.5 Input sparsity time algorithm ..... 1575
21.4.7.6 Optimal sample complexity algorithm ..... 1579
21.4.8 Face-based selection and decomposition ..... 1580
21.4.8.1 Column-row, column-tube, row-tube face subset selection ..... 1581
21.4.8.2 CURT decomposition ..... 1586
21.4.9 Solving small problems ..... 1590
21.5 Extension to general $q$-th order tensors ..... 1592
21.5.1 Fast sampling of columns according to leverage scores, implicitly ..... 1593
21.5.2 General iterative existential proof ..... 1598
21.5.3 General input sparsity reduction ..... 1600
21.5.4 Bicriteria algorithm ..... 1601
21.5.5 CURT decomposition ..... 1602
21.6 Matrix CUR decomposition ..... 1604
21.6.1 Algorithm ..... 1605
21.6.2 Stronger property achieved by leverage scores ..... 1609
21.7 Entry-wise $\ell_{1}$ Norm for Arbitrary Tensors ..... 1616
21.7.1 Facts ..... 1617
21.7.2 Existence results ..... 1619
21.7.3 Polynomial in $k$ size reduction ..... 1624
21.7.4 Solving small problems ..... 1631
21.7.5 Bicriteria algorithms ..... 1633
21.7.5.1 Input sparsity time ..... 1633
21.7.5.2 Improving cubic rank to quadratic rank ..... 1636
21.7.6 Algorithms ..... 1640
21.7.6.1 Input sparsity time algorithm ..... 1640
21.7.6.2 $\widetilde{O}\left(k^{3 / 2}\right)$-approximation algorithm ..... 1642
21.7.7 CURT decomposition ..... 1644
21.8 Weighted Frobenius Norm for Arbitrary Tensors ..... 1649
21.8.1 Definitions and Facts ..... 1650
21.8.2 $r$ distinct faces in each dimension ..... 1652
21.8.3 $r$ distinct columns, rows and tubes ..... 1659
21.8.4 $r$ distinct columns and rows ..... 1662
21.9 Hardness ..... 1669
21.9.1 Definitions ..... 1670
21.9.2 Symmetric tensor eigenvalue ..... 1673
21.9.3 Symmetric tensor singular value, spectral norm and rank-1 approximation1676
21.9.4 Tensor rank is hard to approximate ..... 1681
21.9.4.1 Cover number ..... 1681
21.9.4.2 Properties of 3SAT instances ..... 1683
21.9.4.3 Reduction ..... 1686
21.10Extension to Other Tensor Ranks ..... 1705
21.10.1Tensor Tucker rank ..... 1706
21.10.1.1Definitions ..... 1706
21.10.1.2Algorithm ..... 1706
21.10.2Tensor Train rank ..... 1711
21.10.2.1Definitions ..... 1711
21.10.2.2Algorithm ..... 1711
21.11Entry-wise $\ell_{p}$ Norm for Arbitrary Tensors, $1<p<2$ ..... 1717
21.11. 1 Existence results for matrix case ..... 1718
21.11.2Existence results ..... 1721
21.11.3Polynomial in $k$ size reduction ..... 1725
21.11.4Solving small problems ..... 1729
21.11.5Bicriteria algorithm ..... 1730
21.11.6Algorithms ..... 1734
21.11.7CURT decomposition ..... 1736
21.12Robust Subspace Approximation (Asymmetric Norms for Arbitrary Tensors) ..... 1740
21.12.1Preliminaries ..... 1741
21.12. $2_{1}$-Frobenius (a.k.a $\ell_{1}-\ell_{2}-\ell_{2}$ ) norm ..... 1742
21.12.2.1Definitions ..... 1742
21.12.2.2Sampling and rescaling sketches ..... 1743
21.12.2.3No dilation and no contraction ..... 1744
21.12.2.4Oblivious sketches, MSketch ..... 1747
21.12.2.5Running time analysis ..... 1750
21.12.2.6Algorithms ..... 1751
21.12.3 $\ell_{1}-\ell_{1}-\ell_{2}$ norm ..... 1764
21.12.3.1Definitions ..... 1764
21.12.3.2Projection via Gaussians ..... 1765
21.12.3.3Reduction, projection to high dimension ..... 1767
21.12.3.4Existence results ..... 1769
21.12.3.5Running time analysis ..... 1772
21.12.3.6Algorithms ..... 1775
21.13Streaming Setting ..... 1779
21.14Distributed Setting ..... 1784
21.15Distributed Setting ..... 1788
Chapter 22. Orthogonal Tensor Decomposition ..... 1796
22.1 Introduction ..... 1797
22.1.1 Our Contributions ..... 1799
22.2 Notation. ..... 1805
22.3 Main Results ..... 1806
22.4 Experiments ..... 1813
22.4.1 Experiment Setup and Datasets ..... 1813
22.4.2 Results ..... 1815
22.5 Generate importance sampling indices ..... 1817
22.6 Importance sampling lemmas ..... 1821
22.7 Approximate Tensor Contractions ..... 1824
22.7.1 Case 1: sampling with known per-slice Frobenius norm ..... 1827
22.7.2 Case 2: sampling without known per-slice Frobenius norm ..... 1828
22.7.3 Case 3: improve the bounds in case 2 ..... 1829
22.8 Estimating Slice Norms in Sublinear Time ..... 1830
22.8.1 Definitions and Facts ..... 1830
22.8.2 Order $p=3$, rank $k=1$ tensor ..... 1835
22.8.3 Order $p=3$, rank $k=2$ tensor ..... 1837
22.8.4 Even order tensor ..... 1845
22.9 Robust Tensor Power Method Analysis for General Order $p \geq 3$ ..... 1850
22.9.1 Useful facts ..... 1850
22.9.2 Convergence guarantee and deflation ..... 1855
22.9.3 Main result ..... 1871
22.10Algorithm ..... 1877
22.11Additional Experiment Results ..... 1879
22.11.1Synthetic tensors with inverse decaying eigenvalues ..... 1880
22.11.2Synthetic tensors with inverse square decaying eigenvalues ..... 1883
22.11.3Synthetic tensors with linearly decaying eigenvalues ..... 1885
22.11.4Results from real-life datasets ..... 1887
Part V Open Problems ..... 1893
Chapter 23. Open Problems ..... 1894
23.1 Linear Programs ..... 1895
23.2 Matrix Multiplication ..... 1896
23.3 Ordinary Differential Equations ..... 1897
23.4 Matrix Factorization ..... 1898
23.4.1 Weighted Factorization ..... 1898
23.4.2 Nonnegative Matrix Factorization ..... 1899
23.4.3 Tensor factorization ..... 1899
23.5 Continuous Fourier Transform ..... 1901
23.5.1 One dimensional ..... 1901
23.5.2 High dimensional ..... 1901
23.6 Discrete Fourier Transform ..... 1902
23.6.1 High dimensional ..... 1902
23.7 Discrete time continuous frequency Fourier transform ..... 1904
23.8 Applications of Fourier transform ..... 1907
23.8.1 Speeding up linear programs ..... 1907
23.8.2 Proving convergence result for deep neural network ..... 1908
23.9 An Over-parameterization Theory of Neural Networks ..... 1909
23.10Matrix Chernoff and Discrepancy ..... 1911
23.10.1Matrix Chernoff ..... 1911
23.10.2Discrepancy ..... 1911
23.11Edit Distance and Longest Common Subsequence ..... 1914
23.12Map-Reduce ..... 1916
Part VI More Algorithms for Fundamental Problems ..... 1918
Chapter 24. $\quad \ell_{\infty}$ Regression ..... 1919
24.1 Introduction ..... 1920
24.1.1 Our Contributions ..... 1923
24.1.2 Notation ..... 1930
24.2 Warmup: Gaussians OSEs ..... 1931
24.3 SRHT Matrices ..... 1933
24.4 Proof of Lemma 24.3.2 ..... 1937
24.5 Lower bound for $\ell_{2}$ and $\ell_{\infty}$ guarantee ..... 1939
24.6 Proof for Gaussian case ..... 1945
24.7 Combining Different Matrices ..... 1947
24.7.1 Removing dependence on $n$ via Count-Sketch ..... 1948
24.7.2 Combining Gaussians and SRHT ..... 1948
24.7.3 Combining all three ..... 1948
24.8 Count-Sketch does not obey the $\ell_{\infty}$ guarantee ..... 1948
24.9 Leverage score sampling does not obey the $\ell_{\infty}$ guarantee ..... 1953
24.10Bounding $\mathbb{E}\left[\|Z\|_{F}^{2}\right]$ ..... 1956
Chapter 25. Symmetric Norm Regression ..... 1964
25.1 Introduction ..... 1966
25.1.1 Our Contributions ..... 1969
25.1.2 Technical Overview ..... 1972
25.2 Linear Regression for Orlicz Norms ..... 1979
25.3 Linear Regression for Symmetric Norms ..... 1983
25.4 Conclusion ..... 1986
25.5 Preliminaries ..... 1987
25.6 Missing Proofs in Section 25.2 ..... 1988
25.6.1 Proof of Lemma 25.2.1 ..... 1988
25.6.2 Proof of Lemma 25.2.2 ..... 1989
25.6.3 Proof of Lemma 25.2.3 ..... 1992
25.6.4 Proof of Lemma 25.2.4 ..... 1992
25.6.5 Proof of Lemma 25.2.5 ..... 1992
25.6.6 Proof of Theorem 25.2.6 ..... 1997
25.7 Missing Proofs in Section 25.3 ..... 1999
25.7.1 Background ..... 1999
25.7.1.1 Known $\ell_{2}$ Oblivious Subspace Embeddings ..... 1999
25.7.1.2 Properties of Symmetric Norms ..... 2000
25.7.2 Proof of Lemma 25.3.1 ..... 2006
25.7.3 Proof of Lemma 25.3.2 ..... 2006
25.7.4 Contraction Bound of SymSketch ..... 2010
25.7.5 Dilation Bound of SymSketch ..... 2015
25.7.6 Proof of Theorem 25.3.3 ..... 2020
25.8 Missing Proofs of Main Theorems ..... 2021
25.8.1 Proof of Theorem 25.1.3 ..... 2021
25.8.2 Proof of Theorem 25.1.4 ..... 2021
Chapter 26. Kronecker Product Regression ..... 2023
26.1 Introduction ..... 2024
26.1.1 Our Contributions ..... 2027
26.2 Preliminaries ..... 2030
26.3 Kronecker Product Regression ..... 2032
26.3.1 Kronecker Product $\ell_{p}$ Regression ..... 2034
26.4 All-Pairs Regression ..... 2035
26.5 Low Rank Approximation of Kronecker Product Matrices ..... 2037
26.6 Numerical Simulations ..... 2038
26.7 Missing Proofs from Section 26.3 ..... 2041
26.8 Missing Proofs from Section 26.3.1 ..... 2044
26.8.1 Sampling from an $\ell_{p}$-Well-Conditioned Base ..... 2048
26.8.2 $\ell_{p}$ Sampling From the residual of a $O(1)$-factor approximation ..... 2052
26.9 Missing Proofs from Section 26.4 ..... 2064
26.9.1 Proof of Fast Sampling Lemma 26.4.2 ..... 2068
26.10Missing Proofs from Section 26.5 ..... 2074
26.11Entry-wise Norm Low T-rank Approximation ..... 2080
26.12Properties for General Loss Functions for Low trank Approximation ..... 2083
Chapter 27. Dynamic $k$-means Clustering ..... 2087
27.1 Introduction ..... 2088
27.1.1 Our Result ..... 2089
27.1.2 Our Techniques ..... 2089
27.1.3 Related Work ..... 2093
27.1.4 Roadmap ..... 2095
27.2 Preliminaries ..... 2096
27.3 An Offline Sensitivity Sampling Procedure ..... 2100
27.3.1 Randomly Shifted Grids ..... 2100
27.3.2 Sensitivity Estimation and Coreset Construction ..... 2101
27.3.3 Analysis ..... 2102
27.4 Conclusion ..... 2107
27.5 Missing Details in Section 27.3 ..... 2108
27.6 Coreset Construction over a Dynamic Stream ..... 2116
27.6.1 The Dynamic Point-Cell Storing Data Structure ..... 2117
27.6.2 Estimating the Number of Points in Each Cell ..... 2118
27.6.3 Sensitivity Sampling over a Dynamic Stream ..... 2124
27.6.4 The Final Algorithm ..... 2132
27.7 Why Do Previous Techniques Fail? ..... 2135
Chapter 28. Convergence Result of one-hidden Layer Neural Network ..... 2139
28.1 Introduction ..... 2140
28.1.1 Our Result ..... 2142
28.1.2 Technical Overview ..... 2146
28.2 Preliminaries ..... 2149
28.2.1 Notation ..... 2149
28.3 Problem Formulation ..... 2150
28.4 Quartic Suffices ..... 2152
28.4.1 Bounding the difference between continuous and discrete ..... 2152
28.4.2 Bounding changes of $H$ when $w$ is in a small ball ..... 2154
28.4.3 Loss is decreasing while weights are not changing much ..... 2156
28.4.4 Convergence ..... 2161
28.4.5 Technical claims ..... 2166
28.5 Cubic Suffices ..... 2173
28.5.1 Technical claims ..... 2177
28.5.2 Main result ..... 2179
28.6 Quadratic Suffices ..... 2181
28.6.1 Main result ..... 2188
Chapter 29. Longest Common Subsequence I ..... 2190
29.1 Introduction ..... 2191
29.2 Preliminaries ..... 2195
29.3 Reducing to perfectly unbalanced case ..... 2197
29.4 Perfectly unbalanced strings ..... 2200
Chapter 30. Longest Common Subsequence II ..... 2212
30.1 Introduction ..... 2214
30.1.1 Our Results ..... 2216
30.1.2 Preliminaries ..... 2219
30.1.3 Techniques Overview ..... 2220
30.1.3.1 Summary of Previous ED Techniques ..... 2221
30.1.3.2 LCS ..... 2223
30.1.3.3 LIS ..... 2232
30.2 Organization of the Paper ..... 2237
30.3 LCS Step 1: Sparsification via Birthday Triangle Inequality ..... 2239
30.3.1 Sparsification for $\mathrm{O}_{\lambda^{2}}$ ..... 2239
30.3.2 Sparsification for $\mathrm{O}_{\lambda^{3}}$ ..... 2252
30.3.2.1 $\lambda^{3}$ Sparsification using Ics-cmp ..... 2253
30.3.2.2 Implementation of Ics-cmp ..... 2260
30.4 Longest Increasing Subsequence ..... 2263
30.4.1 Solution Domains ..... 2263
30.4.2 Constructing Approximately Optimal Pseudo-solutions ..... 2267
30.4.3 Evaluating the Pseudo-solutions ..... 2271
30.4.4 An $O\left(n^{\kappa}\right)$ Time Algorithm via Bootstrapping ..... 2277
30.5 Probability and Graph Tools ..... 2284
30.6 LCS Step 0: Window-compatible Solutions ..... 2286
30.6.1 Construction of the Windows ..... 2286
30.6.2 Optimality of the Windows ..... 2289
30.6.3 Dynamic Programming for Block-based LCS ..... 2292
30.7 LCS Step 2: Discovering the Sparse Graph ..... 2295
Chapter 31. Fourier Transform ..... 2308
31.1 One dimensional Fourier transform ..... 2308
31.2 High dimensional Fourier transform ..... 2309
31.3 Proof of Claim 31.2.2 ..... 2312
Chapter 32. Map-Reduce ..... 2313
32.1 Introduction ..... 2315
32.1.1 The MPC model ..... 2318
32.1.2 Our Results ..... 2319
32.1.3 Our Techniques ..... 2324
32.1.4 Roadmap ..... 2334
32.2 A Simplified Batch Algorithm for Connectivity ..... 2335
32.3 Notations ..... 2339
32.4 Graph Connectivity ..... 2340
32.4.1 Neighbor Increment Operation ..... 2340
32.4.2 Random Leader Selection ..... 2346
32.4.3 Tree Contraction Operation ..... 2349
32.4.4 Connectivity Algorithm ..... 2354
32.5 Spanning Forest ..... 2362
32.5.1 Local Shortest Path Tree ..... 2362
32.5.2 Multiple Local Shortest Path Trees ..... 2365
32.5.3 Path Generation and Root Changing ..... 2371
32.5.4 Spanning Forest Expansion ..... 2375
32.5.5 Spanning Forest Algorithm ..... 2377
32.6 Depth-First-Search Sequence for Tree and Applications ..... 2389
32.6.1 Lowest Common Ancestor and Multi-Paths Generation ..... 2389
32.6.2 Depth-First-Search Sequence for a Tree ..... 2394
32.6.2.1 Leaf Sampling ..... 2396
32.6.2.2 DFS Subsequence ..... 2398
32.6.2.3 DFS Sequence ..... 2404
32.6.3 Range Minimum Query ..... 2408
32.6.4 Applications of DFS Sequence ..... 2410
32.7 The MPC Model ..... 2411
32.7.1 Basic MPC Algorithms ..... 2413
32.7.2 Data Organization ..... 2417
32.7.3 Set Operations ..... 2420
32.7.4 Mapping Operations ..... 2424
32.7.5 Sequence Operations ..... 2427
32.7.6 Multiple Tasks ..... 2429
32.8 Implementations in MPC Model ..... 2432
32.8.1 Neighbor Increment Operation ..... 2432
32.8.2 Tree Contraction Operation ..... 2435
32.8.3 Graph Connectivity ..... 2437
32.8.4 Algorithms for Local Shortest Path Trees ..... 2440
32.8.5 Path Generation and Root Changing ..... 2444
32.8.6 Spanning Forest Algorithm ..... 2446
32.8.7 Lowest Common Ancestor and Multi-Paths Generation ..... 2451
32.8.8 Leaf Sampling ..... 2453
32.8.9 DFS Sequence ..... 2455
32.8.10Range Minimum Query ..... 2458
32.9 Minimum Spanning Forest ..... 2460
32.10Directed Reachability vs. Boolean Matrix Multiplication ..... 2466
32.11Discussion on a Previous Conjectured Fast Algorithm ..... 2469
32.12Alternative Approach for Leader Selection ..... 2471
32.12.1Min Parent Forest ..... 2471
32.12.2Leader Selection via Min Parent Forest ..... 2481
Appendices ..... 2504
Appendix A. Probability and Inequality Tools ..... 2505
A. 1 Scalar version probability tools ..... 2505
A.1.1 Reverse Chernoff bound ..... 2506
A. 2 Matrix version probability tools ..... 2510
A. 3 Inequalities ..... 2511
Appendix B. Coauthors Index ..... 2512
B. 1 Coauthors Index ..... 2512
Bibliography ..... 2517
Vita ..... 2689

## List of Tables

1.1 Let $\omega$ denote the exponent of the current matrix multiplication. "Iters" denotes the "number of iterations". "Cost/Iter" denotes the "cost per iteration". LP has $n$ variables, $d=n$ constraints, and can be encoded in $L$ input bits. We consider the case where $A$ is a dense full rank matrix. We remark that all of results have difference when $\operatorname{rank}(A)$ is low $\operatorname{rank}, \operatorname{nnz}(A)$ is sparse and $n \not \approx d$. However, when $d=n, \operatorname{rank}(A)=n$ and $\operatorname{nnz}(A)=n^{2}$, the running time remains to be $O\left(n^{2.5}\right)$ for three decades. . . . . . . . . . . . . . . . . . . . . .
1.2 Summary of results, $d$ is the dimension, $\kappa$ is the condition number of $\nabla^{2} f$. "Iters" denotes "the number of iterations / parallel depth". "Gra/Iter" denotes the "the number of gradients per iteration". We use the parallel depth of the algorithm as the number of iterations. We suppress polylogarithmic terms and dependence on the error parameter. Ball walk/hit-and-run apply to general logconcave distributions, the rest assume strongly logconcave with Lipschitz gradient and possibly more. "B./H." denotes "Ball Walk/Hit-and-run". "DL" denotes "Damped Langevin". In all previous work, for simplicity, we report the most favorable bounds by making various assumptions such as $\kappa \ll d$. . .
1.3 (A list of $\ell_{2} / \ell_{2}$-sparse recovery results). We ignore the " $O$ " for simplicity. LP denotes the time of solving Linear Programs [CLS19], and the state-of-the-art algorithm takes $n^{\omega}$ time where $\omega$ is the exponent of matrix multiplication. The results in [Don06, CRT06b, NT09a] do not explicitly state the $\ell_{2} / \ell_{2}$ guarantee, but their approach obtains it by an application of the JohnsonLindenstrauss Lemma; they also cannot facilitate $\epsilon<1$, obtaining thus only a 2-approximation. The $c$ in previous work is a sufficiently large constant, not explicitly stated, which is defined by probabilistically picking an errorcorrecting code of short length and iterating over all codewords. We estimate $c \geq 4$. We note that our runtime is (almost) achieved.
$1.4 n=p^{d}$. E. $k$ denotes the "Exactly $k$-sparse" setting, which is a noiseless setting. "Gua" denotes "Guarantee". We ignore the $O$ for simplicity. The $\ell_{\infty} / \ell_{2}$ is the strongest possible guarantee, with $\ell_{2} / \ell_{2}$ coming second, $\ell_{2} / \ell_{1}$ third and exactly $k$-sparse being the less strong. We note that [CT06, RV08, CGV13, Bou14, HR16] obtain a uniform guarantee, i.e. with $1-1 / \operatorname{poly}(n)$ they allow reconstruction of all vectors; $\ell_{\infty} / \ell_{2}$ and $\ell_{2} / \ell_{2}$ are impossible in the uniform case [CDD09]. We also note that [RV08, CGV13, Bou14, HR16] give improved analysis of the Restricted Isometry property; the algorithm is suggested and analyzed (modulo the RIP property) in [BD08]. The work in [HIKP12a] does not explicitly state the extension to the $d$-dimensional case, but can easily be inferred from the arguments. [HIKP12a, IK14, Kap16, KVZ19] work when the universe size in each dimension are powers of 2 . We also assume that the signal-to-noise ratio is bounded by a polynomial of $n$, which is a standard assumption in the sparse Fourier transform literature [HIKP12a, IK14, Kap16, Kap17, LN19].
2.1 The bound of each quantity under Assumption 2.4.1. For intuition, think $\epsilon \sim \epsilon_{m p} \sim 1 / 10$ and $k \sim \sqrt{n}$. . . . . . . . . . . . . . . . . . . . . . . . . . . .23
3.1 Bounding the changes of different variables ..... 100
3.2 Summary of data structure CentralPathMaintenance ..... 119
3.3 Summary of parameters. "Ntn" denotes Notation. ..... 153
4.1 Summary of results, $d$ is the dimension, $\kappa$ is the condition number of $\nabla^{2} f$."Iters" denotes "the number of iterations / parallel depth". "Gra/Iter" denotesthe "the number of gradients per iteration". We use the parallel depth of thealgorithm as the number of iterations. We suppress polylogarithmic terms anddependence on the error parameter. Ball walk/hit-and-run apply to generallogconcave distributions, the rest assume strongly logconcave with Lipschitzgradient and possibly more. "B. /H." denotes "Ball Walk/Hit-and-run". "DL"denotes "Damped Langevin". In all previous work, for simplicity, we reportthe most favorable bounds by making various assumptions such as $\kappa \ll d$. . .174
4.2 Summary of parameters of Theorem 4.2.4 ..... 224
10.1 (A list of $\ell_{2} / \ell_{2}$-sparse recovery results). We ignore the " $O$ " for simplicity. LP denotes the time of solving Linear Programs [CLS19], and the state-of-the-art algorithm takes $n^{\omega}$ time where $\omega$ is the exponent of matrix multiplication. The results in [Don06, CRT06b, NT09a] do not explicitly state the $\ell_{2} / \ell_{2}$ guarantee, but their approach obtains it by an application of the JohnsonLindenstrauss Lemma; they also cannot facilitate $\epsilon<1$, obtaining thus only a 2-approximation. The $c$ in previous work is a sufficiently large constant, not explicitly stated, which is defined by probabilistically picking an errorcorrecting code of short length and iterating over all codewords. We estimate $c \geq 4$. We note that our runtime is (almost) achieved
10.3 Summary of constants in Section 10.3, the column "Parameter" indicates which parameter is depending on that constant. Note that set $L$ is the input of the algorithm in Section 10.3 and the output of the algorithm in Section 10.2. . .
$13.1 n=p^{d}$. We ignore the $O$ for simplicity. The $\ell_{\infty} / \ell_{2}$ is the strongest possible guarantee, with $\ell_{2} / \ell_{2}$ coming second, $\ell_{2} / \ell_{1}$ third and exactly $k$-sparse being the less strong. We note that [CT06, RV08, CGV13, Bou14, HR16] obtain a uniform guarantee, i.e. with $1-1 / \operatorname{poly}(n)$ they allow reconstruction of all vectors; $\ell_{\infty} / \ell_{2}$ and $\ell_{2} / \ell_{2}$ are impossible in the uniform case [CDD09]. We also note that [RV08, CGV13, Bou14, HR16] give improved analysis of the Restricted Isometry property; the algorithm is suggested and analyzed (modulo the RIP property) in [BD08]. The work in [HIKP12a] does not explicitly state the extension to the $d$-dimensional case, but can easily be inferred from the arguments. [HIKP12a, IK14, Kap16, KVZ19] work when the universe size in each dimension are powers of 2 . We also assume that the signal-to-noise ratio is bounded by a polynomial of $n$, which is a standard assumption in the sparse Fourier transform literature [HIKP12a, IK14, Kap16, Kap17, LN19].
13.2 Summary of important constants. ..... 946
13.3 Summary of Lemmas. ..... 946
14.1 Results on set query. We ignore the " $O$ " for simplicity. "Space" means the space to store matrix, which doesn't include number of measurements. We note that the decoding time of the set query algorithms do not depend on $\epsilon$. ..... 965
14.2 Results on Fourier set query. We ignore the " $O$ " for simplicity. We assume signal is bounded by poly $(n)$. In Fourier set query, there is no update time concept. ..... 965
15.1 Set Query ..... 1010
20.1 Example functions satisfying both structural properties. ..... 1431
22.1 Synthetic tensor decomposition using the robust tensor power method. We use an order-3 normalized dense tensor with dimension $n=1200$ with $\sigma=0.01$ noise added. We run sketching-based and sampling-based methods to find the first eigenvalue and eigenvector by setting $L=50, T=30$ and varying $B$ and $b$. ..... 1817
22.2 Tensor decomposition in LDA on the wiki dataset. The tensor is generated by spectral LDA with dimension $200 \times 200 \times 200$. It is symmetric but not normalized. We fix $L=50, T=30$ and vary $B$ and $b$. ..... 1818
22.3 Sketching based robust power method: $n=1200$, inverse decay, $\|\mathbf{T}\|_{F}^{2}=1.01$ ..... 1880
22.4 Importance sampling based robust power method, without prescanning: $n=$ 1200, inverse decay, $\|\mathbf{T}\|_{F}^{2}=1.01$ ..... 1880
22.5 Importance sampling based robust power method with prescanning: $n=$ 1200, inverse decay, $\|\mathbf{T}\|_{F}^{2}=1.01$ ..... 1881
22.6 Sketching based robust power method: $n=1200$, inverse-square decay, $\|\mathbf{T}\|_{F}^{2}=1.01$ ..... 1883
22.7 Importance sampling based robust power method, without prescanning: $n=$ 1200, inverse-square decay, $\|\mathrm{T}\|_{F}^{2}=1.01$ ..... 1883
22.8 Importance sampling based robust power method with prescanning: $n=$ 1200, inverse-square decay, $\|\mathrm{T}\|_{F}^{2}=1.01$ ..... 1884
22.9 Sketching based robust power method: $n=\mathbf{1 2 0 0}$, linear decay, $\|\mathbf{T}\|_{F}^{2}=1.01$ ..... 1885
22.10Importance sampling based robust power method, without prescanning: $n=$ 1200, linear decay, $\|\mathbf{T}\|_{F}^{2}=1.01$ ..... 1885
22.11Importance sampling based robust power method with prescanning: $n=$ 1200, linear decay, $\|\mathrm{T}\|_{F}^{2}=1.01$ ..... 1886
22.12List of six real-life datasets ..... 1887
22.13Sketching based robust power method: dataset Enron, $\|$ T $\|_{F}^{2}=4.96 \mathrm{e}+07$ ..... 1888
22.14Importance sampling based robust power method, without prescanning: dataset Enron, $\|\mathrm{T}\|_{F}^{2}=4.96 \mathrm{e}+07$ ..... 1888
22.15Importance sampling based robust power method with prescanning: dataset Enron, $\|\mathbf{T}\|_{F}^{2}=4.96 \mathrm{e}+07$ ..... 1888
22.16Sketching based robust power method: dataset AP, $\|\mathbf{T}\|_{F}^{2}=6.905 \mathrm{e}+06$ ..... 1889
22.17Importance sampling based robust power method, without prescanning: dataset $\mathbf{A P},\|\mathbf{T}\|_{F}^{2}=6.905 \mathrm{e}+06$ ..... 1889
22.18Importance sampling based robust power method with prescanning: dataset $\mathbf{A P},\|\mathbf{T}\|_{F}^{2}=6.905 \mathrm{e}+06$ ..... 1889
22.19Sketching based robust power method: dataset NSF, $\|\mathbf{T}\|_{F}^{2}=4.765 \mathrm{e}+06$ ..... 1890
22.20Importance sampling based robust power method, without prescanning: dataset NSF, $\|\mathbf{T}\|_{F}^{2}=4.765 \mathrm{e}+06$ ..... 1890
22.21Importance sampling based robust power method with prescanning: dataset NSF, $\|\mathbf{T}\|_{F}^{2}=4.765 \mathrm{e}+06$ ..... 1890
22.22Sketching based robust power method: dataset NIPS, $\|\mathbf{T}\|_{F}^{2}=6.058 \mathrm{e}+09$ ..... 1891
22.23Importance sampling based robust power method, without prescanning: dataset NIPS, $\|\mathbf{T}\|_{F}^{2}=6.058 \mathrm{e}+09$ ..... 1891
22.24Importance sampling based robust power method with prescanning: dataset NIPS, $\|\mathbf{T}\|_{F}^{2}=6.058 \mathrm{e}+09$ ..... 1891
22.25 Sketching based robust power method: dataset KOS, $\|\mathbf{T}\|_{F}^{2}=5.016 \mathrm{e}+04$ ..... 1892
22.26Importance sampling based robust power method, without prescanning: dataset KOS, $\|\mathbf{T}\|_{F}^{2}=5.016 \mathrm{e}+04$ ..... 1892
22.27Importance sampling based robust power method with prescanning: dataset KOS, $\|\mathbf{T}\|_{F}^{2}=5.016 \mathrm{e}+04$ ..... 1892
25.1 $M$-estimators ..... 1967
25.2 Comparison between input-sparsity time linear regression algorithms ..... 1971
26.1 Results for $\ell_{2}$ and $\ell_{1}$-regression with respect to different sketch sizes $m$. ..... 2040
28.1 Summary of Convergence Result ..... 2144
28.2 Table of Parameters for the $m=\widetilde{\Omega}\left(n^{4}\right)$ result in Section 28.4. Nt. stands for notations. ..... 2156
28.3 Table of Parameters for the $m=\widetilde{\Omega}\left(n^{3}\right)$ result in Section 28.5. Nt. stands for notations. ..... 2175
28.4 Table of Parameters for the $m=\widetilde{\Omega}\left(n^{2}\right)$ result in Section 28.6. Nt. stands for notations. ..... 2183
29.1 Fill all the six cases in Eq. (29.5) into the whole space. Note that $1+2+3$means the combination of 1,2 and 3 covers it. 5, 6 means any one of themcovers it.2201

## List of Figures

1.1 Linear Programming is Chasing Matrix Multiplication ..... 11
1.2 Sampling Model. $x^{*}$ is a perfect signal. We are only allowed to samples from $x$ which is a signal that has noise. ..... 41
1.3 Recovered signal. Our goal is to output signal $x^{\prime}$ such that $x^{\prime}$ is close to $x^{*}$. ..... 42
1.4 Thought Experiments ..... 43
1.5 Moitra's lower bound [Moi15]. $\forall \epsilon>0$, if sample duration $T<\frac{(1-\epsilon)}{\eta}$, then $\exists$ two $k$-Fourier-sparse signals: $x(t)$ and $x^{\prime}(t)$, each has $\eta$ frequency gap. To tell them apart, need noise $\leq 2^{-\Omega(\epsilon k)}$. This example only implies that frequency recovery requires frequency gap, but what if we just want to interpolate signal? ..... 47
2.1 ClassicalStep happens with $n^{-2}$ probability ..... 18
$2.2 \psi(x), \psi(x)^{\prime}$ and $\psi(x)^{\prime \prime}$. For $\epsilon_{m p} \in(0,1)$ ..... 65
5.1 Landscapes of the CIFAR10 image-classification training objective $F(W)$ nearthe SGD training trajectory. The blue vertical stick marks the current point$W=W_{t}$ at the current iteration $t$. The $x$ and $y$ axes represent the gradientdirection $\nabla F\left(W_{t}\right)$ and the most negatively curved direction of the Hessianafter smoothing (approximately found by Oja's method [AL17, AL18]). The$z$ axis represents the objective value. Observation. As far as minimizingobjective is concerned, the (negative) gradient direction sufficiently decreasesthe training objective, and it is not needed to use second-order method to findnegative curvature. This is consistent with our findings Theorem 5.11.1 and5.12.1. Remark 1. Gradient norm does not tend to zero because cross-entropyloss is not strongly convex (see Section 5.6). Remark 2. The task is CIFAR10(for CIFAR100 or CIFAR10 with noisy label, see Figure 5.2 through 5.7 inappendix). Remark 3. Architecture is VGG19 (for Resnet-32 or ResNet-110,see Figure 5.2 through 5.7 in the later sections). Remark 4. The six plotscorrespond to epochs $5,40,90,120,130$ and 160 . We start with learning rate0.1 , and decrease it to 0.01 at epoch 81 , and to 0.001 at epoch 122 . SGD withmomentum 0.9 is used. The training code is unchanged from [Yan18] and weonly write new code for plotting such landscapes.265
5.2 ResNet-32 architecture [Yan18] landscape on CIFAR10 vs CIFAR100. ..... 344
5.3 ResNet-110 architecture [Yan18] landscape on CIFAR10 vs CIFAR100. ..... 345
5.4 VGG19 architecture (with BN) [Yan18] landscape on CIFAR10 vs CIFAR100. ..... 346

> 5.5 ResNet-32 architecture [Yan18] landscape on CIFAR10 vs CIFAR10 (20\% noise), means we have randomly perturbed $20 \%$ of the true labels in the training set.
> 5.6 ResNet-110 architecture [Yan18] landscape on CIFAR10 vs CIFAR10 (20\% noise), means we have randomly perturbed $20 \%$ of the true labels in the training set.
> 5.7 VGG19 architecture (with BN) [Yan18] landscape on CIFAR10 vs CIFAR10 ( $20 \%$ noise), means we have randomly perturbed $20 \%$ of the true labels in the training set.
> 7.1 The function $h(z)=-\frac{1+z}{1-z}+\sqrt{\left(\frac{1+z}{1-z}\right)^{2}+1}$ maps the unit disk $\{z \in \mathbb{C}:|z| \leq 1\}$ to the half disk $\{z \in \mathbb{C}:|z| \leq 1$ and $\operatorname{Re}(z) \geq 0\}$.
> 10.1 This is a visualization of part of the proof in Claim 10.4.7. We consider an example where there are $l=10$ blocks, $B_{1}=1, B_{2}=1, B_{1}=1, B_{3}=0$, $B_{4}=0, B_{5}=1, B_{6}=1, B_{7}=0, B_{8}=0, B_{9}=1$ and $B_{10}=0$. Recall the two important conditions in the proof of Claim 10.4.7, the first one is $B_{1}=1$ and the second one is, for all $j \in[l], \sum_{j^{\prime}=2}^{j} B_{j^{\prime}}>(j-1) / 2$. The number on the green arrow is $\sum_{j^{\prime}=2}^{j} B_{j^{\prime}}$. It is to see that the example we provided here is satisfying those two conditions. Recall the definition of set $S_{1}$ and $S_{0}$. Here $S_{1}=\{2,3,5,6,9\}$ and $S_{0}=\{4,7,8,10\}$. Then $S_{1}^{\prime}=\{2,3,5,6\}$. The mapping $\pi$ satisfies that $\pi(4)=2, \pi(7)=3, \pi(8)=5$ and $\pi(10)=6$.
> 11.1 Filter $\widehat{G^{\prime}}(f)$
> 11.2 For an arbitrary frequency interval $\left[l_{j}-\frac{\Delta l}{2}, l_{j}+\frac{\Delta l}{2}\right]$, we scale it by $2 \pi \sigma \beta$ to get a longer interval $\left[2 \pi \sigma \beta\left(l_{j}+\frac{\Delta l}{2}\right), 2 \pi \sigma \beta\left(l_{j}-\frac{\Delta l}{2}\right)\right]$. Then, we wrap the longer interval on a circle $[0,2 \pi)$. The number of folds after wrapping is $\lceil\sigma \beta \Delta l\rceil$. For any random sample, the observation $c_{j}$ is close to the true answer within $1 / \rho$ with some "good" probability.

$$
\begin{aligned}
& \text { 11.3 The number of "blue" regions is equal to the number of folds } m \text {. The number } \\
& \text { of total regions is } t \text {. For each observed } c_{j} \text {, instead of checking all the } t \text { regions, } \\
& \text { we only assign vote to the these "blue" regions. Since only these } m \text { "blue" } \\
& \text { regions can be the candidate region that contains frequency } f \text {. . . . . . } 727
\end{aligned}
$$

11.4 We demonstrate the algorithm for merging various stages $(R=O(\log k))$ on
2-dimensional data. Note that the true data should be 3-dimensional, since
for each tone $\left(v_{i}, f_{i}\right), v_{i} \in \mathbb{C}$ and $f_{i} \in \mathbb{R}$. The $x$-axis represents the frequency
and the $y$-axis represents the real part of magnitude.

11.5 $F(\theta)$ is a sinc function and has derivate function $F^{\prime}(\theta)$. ..... 745
$11.6 \int_{-\infty}^{+\infty} \min \left(T, \frac{1}{\left|\theta-f_{i}\right|}\right) \cdot \min \left(T, \frac{1}{\left|\theta-f_{j}\right|}\right) \mathrm{d} \theta$ ..... 747
12.1 A picture of a $\mathrm{Comb}_{s}$, rect $_{s}, \operatorname{sinc}_{s}$, Gaussian $_{\mu, \sigma}$ ..... 782
12.2 The filter function $(H(t), \widehat{H}(f))$ with a $k$-Fourier-sparse signal. The property I, II and III are presented in the bottom one, the property IV is presented in the top one. ..... 805
$12.3 G$ and $\widehat{G}$. [PS15] ..... 806
12.4 The Property of $\operatorname{sinc}(t)$. ..... 884
12.5 The light red area represents $\int_{\left(1 / 2-2 / s_{1}\right)}^{\left(1 / 2-2 / s_{1}\right)-t} s_{0} \cdot \operatorname{sinc}\left(s_{1}(\tau)\right)^{\ell} \mathrm{d} \tau$ and the light green area represents $\int_{1 / 2-2 / s_{1}}^{1 / 2-2 / s_{1}+t} s_{0} \cdot \operatorname{sinc}\left(s_{1}(\tau)\right)^{\ell} \mathrm{d} \tau$. ..... 886
$12.6 \widehat{H \cdot x^{*}}(f)$ and $H \cdot x^{*}(t)$. ..... 892
12.7 Property VI of filter function $H(t)$, first two figures of four figures. The light green area represents RHS(without scalar) of Property VI of filter $H$. ..... 894
12.8 Property VI of filter function $H(t)$, last two figures of four figures. The lightred area represents LHS of Property VI of filter $H$, the light yellow arearepresents the difference. Property VI says the light yellow area is only asmall constant fraction of the light green area.895
12.9 Parameters for $s_{1}, s_{3}$ and $\ell$. ..... 901$12.10 \widehat{G}_{\sigma, b}^{(j)}(f)$ where the top one is $j=0$ and the bottom one is $j=1, \mathcal{A}_{i, j}=$$\left[\frac{1}{\sigma}\left(2 \pi\left(i+\frac{j}{B}\right)-\frac{2 \pi}{2 B}\right), \frac{1}{\sigma}\left(2 \pi\left(i+\frac{j}{B}\right)+\frac{2 \pi}{2 B}\right)\right], \mathcal{B}_{i, j}=\left[\frac{1}{\sigma}\left(2 \pi\left(i+\frac{j}{B}\right)-\frac{2 \pi(1-\alpha)}{2 B}\right), \frac{1}{\sigma}(2 \pi(i+\right.$$\left.\left.\left.\frac{j}{B}\right)+\frac{2 \pi(1-\alpha)}{2 B}\right)\right]$906
13.1 Illustration of box $\mathcal{B}_{\infty}(c, r)$ and grid $\mathcal{G}_{r_{g}}$. Box $\mathcal{B}_{\infty}(c, r)$ refers to all the points in the square centered at $c$ with side length $2 r$. Grid $\mathcal{G}_{r_{g}}$ refers to all the solid round points, and the distance between origin $O$ and $A_{0}$ is $r_{g}$. Note that the dashed lines are decision boundaries of the projection $\Pi_{r_{g}}$, and all the points inside a minimum cell separated by the dashed lines are mapped (by $\Pi_{r_{g}}$ ) to the same grid point in $\mathcal{G}_{r_{g}}$ (which is the center of the cell). We have $\Pi_{r_{g}}(c)=A_{1}$ and $\Pi_{r_{g}}\left(\mathcal{B}_{\infty}(c, r)\right)=\left\{A_{1}, A_{2}, A_{3}, A_{4}\right\}$.
13.2 Illustration of the behavior of Line 16 to Line 20 in Algorithm 13.3. For any $f \in[p]^{d}$, we draw box $\mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}+z_{f}, 2^{1-h} \nu_{\ell}\right)$ after $h$ iterations of the for loop between Line 17 and Line 19 in Algorithm 13.3, where $h \in\{0,1, \ldots, H\}$. Conditioned on LinfinityReduce is correct, for every $h \in\{0,1, \ldots, H\}$, after $h$-th iteration we have $\widehat{x}_{f} \in \mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}+z_{f}, 2^{1-h} \nu_{\ell}\right)$. When $h=0$, i.e. before the loop between Line 17 and Line 19 starts, we know that $\widehat{x}_{f} \in$ $\mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}, 2 \nu_{\ell}\right)$ as depicted by (a). After each iteration in $h$, the radius of the box shrinks by half (and its center might change). Finally after $H$ iterations, as depicted by (c), we obtain $z^{(\ell-1)}$ such that $\widehat{x}_{f} \in \mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}, 2^{1-H} \nu_{\ell}\right)$.
13.3 Illustration of the iteration between Line 25 and Line 28 in Algorithm 13.3. The round solid points represent grid points in $\mathcal{G}_{\beta \nu}$, and the dashed lines represent decision boundaries of $\Pi_{\beta \nu_{\ell}}$. In this example we have $\mid \operatorname{supp}\left(y^{(\ell-1)}+\right.$ $\left.z^{(\ell)}\right) \mid=3$, and the dotted squares represent boxes $\mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}, 2^{1-H} \nu_{\ell}\right)$ for $f \in \operatorname{supp}\left(y^{(\ell-1)}+z^{(\ell)}\right)$. The algorithm repeatedly samples a random shift $s \sim \mathcal{B}_{\infty}\left(0, \alpha \nu_{\ell}\right)$, until all the shifted boxes $\left\{\mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}, 2^{1-H} \nu_{\ell}\right)+\right.$ $s\}_{f \in \operatorname{supp}\left(y^{(\ell-1)}+z^{(\ell)}\right)}$ do not intersect with the dashed lines (i.e. decision boundaries of $\Pi_{\beta \nu_{\ell}}$ ). In the figure, we color a shifted box in green if it does not intersect with dashed lines, and color in red otherwise. After a series of failed attempts from (a) to (b), we finally have a successful attempt in (c).
13.4 Illustration of good and bad shifts in Definition 13.2.6. In (a), the small square represents box $\mathcal{B}_{\infty}\left(c, r_{b}\right)$, and the dashed lines represent the decision boundary of $\Pi_{r_{g}}$. The arrows in (b) and (c) represent two different shifts, where the shift in (b) is an example of good shift, since the shifted box does not intersect with the decision boundaries of $\Pi_{r_{g}}$, while the shift in (c) is an example of bad shift, since the shifted box intersects with the decision boundaries of $\Pi_{r_{g}}$.
13.5 Illustration of Lemma 13.2.6. In (a) the smallest square represents box $\mathcal{B}_{\infty}\left(c, r_{b}\right)$, the medium-sized square represents $\mathcal{B}_{\infty}\left(c, r_{s}\right)$, and the dashed lines represent decision boundaries of $\Pi_{r_{g}}$. Note that for $s \sim \mathcal{B}_{\infty}\left(0, r_{s}\right)$, the center of the shifted box $s+\mathcal{B}_{\infty}\left(c, r_{b}\right)$ is $s+c \sim \mathcal{B}_{\infty}\left(c, r_{s}\right)$. Shift $s$ is good (recall in Definition 13.2.6) for box $\mathcal{B}_{\infty}\left(c, r_{b}\right)$ and grid $\mathcal{G}_{r_{g}}$ if and only if the distance between $s+c$ and decision boundaries of $\Pi_{r_{g}}$ is greater than $r_{b}$. In (b), we draw in red the set of points which are within distance at most $r_{b}$ to the decision boundaries of $\Pi_{r_{g}}$. Then in (b) the red part inside $\mathcal{B}_{\infty}\left(c, r_{s}\right)$ corresponds to bad shifts (plus $c$ ), and the green part corresponds to good shifts (plus c). Intuitively, the fraction of the green part is at least $\left(1-r_{b} / r_{s}\right)^{2}$ because the vertical red strips can cover a width of at most $2 r_{b}$ on the $x$-axis of $\mathcal{B}_{\infty}\left(c, r_{s}\right)$ (whose side length is $2 r_{s}$ ), and the horizontal red strips can cover a width of at most $2 r_{b}$ on the $y$-axis.
14.1 Filter $\widehat{G}^{\prime}$ ..... 978
15.1 (a) Huber function (top) (b) Tukey function (bottom) ..... 994
15.2 $\left\|x_{S \cap \bar{T}}\right\|_{H}-\left\|x_{T \cap \bar{S}}\right\|_{H}$ is the displacement error ..... 1000
$15.3\left\|x_{S \cap \bar{T}}\right\|_{H}-\left\|x_{T \cap \bar{S}}\right\|_{H}$ is the displacement error ..... 1021
17.1 The $x$-axis is $n$ where $A \in \mathbb{R}^{n \times n}$, and the $y$-axis is $\left\|A^{\prime}-A\right\|_{1}$ where $\operatorname{rank}\left(A^{\prime}\right)=$ $k$. This figure shows the performance of both our algorithm and [DZHZ06] on input matrix $A$ defined as Equation (17.40).1287
17.2 The $x$-axis is $n$ where $A \in \mathbb{R}^{n \times n}$, and the $y$-axis is $\left\|A^{\prime}-A\right\|_{1}$ where $\operatorname{rank}\left(A^{\prime}\right)=$ $k$. This figure shows the performance of both our algorithm and [BDB13] on input matrix $A$ defined as Equation (17.41).
17.3 The $x$-axis is $n$ where $A \in \mathbb{R}^{n \times n}$, and the $y$-axis is $\left\|A^{\prime}-A\right\|_{1}$ where $\operatorname{rank}\left(A^{\prime}\right)=$ $k$. This figure shows the performance of both our algorithm and [Kwa08] on input matrix $A$ defined as Equation (17.42). Algorithm [Kwa08] has two ways of initialization, which have similar performance on matrix $A$.
17.4 The $x$-axis is $n$ where $A \in \mathbb{R}^{n \times n}$, and the $y$-axis is $\left\|A^{\prime}-A\right\|_{1}$ where $\operatorname{rank}\left(A^{\prime}\right)=$ $k$. Algorithm [KK05] has two ways of initialization. The left figure shows the performance of both our algorithm and [KK05] with random vector initialization on input matrix $A$ defined as Equation (17.43). The right figure shows the performance of both our algorithm and [KK05] with top singular vector initialization on input matrix $A$ defined as Equation (17.45).
17.5 Let $A$ be $(2 n+2) \times(2 n+2)$ input matirx. (a) shows the performance of all the algorithms when the matrix dimension is growing. The $x$-axis is $n$, and the $y$-axis is $\left\|A^{\prime}-A\right\|_{1}$ where $A^{\prime}$ is the rank- 3 solution output by all the heuristic algorithms and also ours. The $\ell_{1}$ residual cost of all the other algorithms is growing much faster than ours, which is consistent with our theoretical results. (b) shows the running time (in seconds) of all the algorithms when the matrix dimension $n$ is growing. The $x$-axis is $n$ and the $y$-axis is time (seconds). The running time of some of the algorithms is longer than 3 seconds. For most of the algorithms (including ours), the running time is always less than 3 seconds. 1298
20.1 The blue curve is the Huber function which combines an $\ell_{2}$-like measure for small $x$ with an $\ell_{1}$-like measure for large $x$. The red curve is the "reverse" Huber function which combines an $\ell_{1}$-like measure for small $x$ with an $\ell_{2}$-like measure for large $x$.
21.1 A 3rd order tensor with size $8 \times 8 \times 8$.
21.2 Flattening. We flatten a third order $4 \times 4 \times 4$ tensor along the 1 st dimension to obtain a $4 \times 16$ matrix. The red blocks correspond to a column in the original third order tensor, the blue blocks correspond to a row in the original third order tensor, and the green blocks correspond to a tube in the original third order tensor.
21.3 A 3rd order tensor contains $n^{2}$ columns, $n^{2}$ rows, and $n^{2}$ tubes.
21.4 A third order tensor has three types of faces: the column-row faces, the column-tube faces, and the row-tube faces ..... 1500
21.5 Column subset selection, row subset selection and tube subset selection. ..... 1503
21.6 An example tensor CURT decomposition. ..... 1506
21.7 Let $W$ denote a tensor that has columns(red), rows(green) and tubes(blue). For each $i \in[3]$, let $W_{i}$ denote the matrix obtained by flattening tensor $W$ along the $i$-th dimension. ..... 1659
21.8 Each face $W_{*, *, i}$ is a column-row face. $W_{*, *, 1}$ is the bottom column-row face. $r=3$. The blue blocks represent column-tube faces, the red blocks represent column-tube faces. ..... 1661
21.9 Each face $W_{*, *, i}$ is a column-row face. $W_{*, *, 1}$ is the bottom column-row face. $r=3$. The blue blocks represent $\left|C_{3}\right|$ column-tube faces. The green blocks represet $\left|R_{3}\right|$ row-tube faces. In each column-row face, the intersection be- tween blue faces and green faces is a size $\left|R_{3}\right| \times\left|C_{3}\right|$ block, and all the entries in this block are the same. ..... 1666
21.10Cover number. For a 3SAT instance with $n$ variables and $m$ clauses, we can draw a bipartite graph which has $n$ nodes on the left and $m$ nodes on the right. Each node (blue) on the left corresponds to a variable $x_{i}$, each node (green) on the right corresponds to a clause $C_{j}$. If either $x_{i}$ or $\bar{x}_{i}$ belongs to clause $C_{j}$, then we draw a line between these two nodes. Consider an input string $y \in\{0,1\}^{7}$. There exists some unsatisfied clauses with respect to this input string $y$. For for example, let $C_{1}, C_{2}$ and $C_{3}$ denote those unsatisfied clauses. We want to pick a smallest set of nodes on the left partition of the graph to guarantee that for each unsatisfied clause in the right partition, there exists a node on the left to cover it. The cover number is defined to be the smallest such number over all possible input strings. ..... 1682
21.11There are $3 n+m$ column-row faces, $V_{i}, \forall i \in[n], S_{i}, \forall i \in[n], M_{i}, \forall i \in[n]$, $C_{l}, \forall l \in[m]$. In face $C_{l}$, each $u_{l, j}$ is either $x_{i}$ or $\bar{x}_{i}$ where $x_{i}=e_{2 i-1}$ and $\bar{x}_{i}=e_{2 i-1}+e_{2 i}$ ..... 1688
21.12Two possibilities for $V_{i}^{(1)}, \forall i \in[n], V^{(2)}, \forall i \in[n], M_{i}^{(1)}, \forall i \in[n]$. ..... 1690
$21.13 \widetilde{V}_{i}, \widetilde{S}_{i}, \widetilde{M}_{i}, \widetilde{C}_{l}$. ..... 169321.14There are $n+p$ matrices $A_{i} \in \mathbb{R}^{2 \times(2 n+p)}, \forall i \in[n+p]$ and $2 n+p$ matrices$B_{i} \in \mathbb{R}^{2 \times(n+p)}, \forall i \in[2 n+p]$. Tensor $A$ and tensor $B$ represet the same tensor,and for each $i \in[n+p], j \in[2], l \in[2 n+p],\left(A_{i}\right)_{j, l}=\left(B_{l}\right)_{j, i}$.1694
21.15For any $i \in[n], \beta_{i, 1} \in \mathbb{R}$, for any $l \in[m], \gamma_{l, 1}, \gamma_{l, 2} \in \mathbb{R}$, for any $l \in[m]$, if the first literal of clause $l$ is $x_{j}$, then row vector $u_{l, 1}=e_{2 i-1} \in \mathbb{R}^{2 n}$; if the first literal of clause $l$ is $\bar{x}_{j}$, then row vector $u_{l, 1}=e_{2 i-1}+e_{2 i} \in \mathbb{R}^{2 n}$. .
21.16For any $i \in[n], \beta_{i, 1} \in \mathbb{R}$. For any $i \in[q], \beta_{i, 2} \in \mathbb{R}$. For any $l \in[m]$, $\gamma_{l, 1}, \gamma_{l, 2} \in \mathbb{R}$. For any $l \in[m]$, if the first literal of clause $l$ is $x_{j}$, then row vector $u_{l, 1}=e_{2 i-1} \in \mathbb{R}^{2 n}$; if the first literal of clause $l$ is $\bar{x}_{j}$, then row vector $u_{l, 1}=e_{2 i-1}+e_{2 i} \in \mathbb{R}^{2 n}$.
22.1 Sketching v.s. importance sampling. Running time with growing dimension .

24.1 Our construction of $A$ and $b$ for the proof that Count-Sketch does not obey
the $\ell_{\infty}$ guarantee. $\alpha<d$. .
24.2 Count Sketch matrix $S \in \mathbb{R}^{m \times n}$. Event I, for any $k^{\prime} \in[d]$ and $k^{\prime} \neq k$, $\operatorname{supp}\left(S_{k}\right) \cap \operatorname{supp}\left(S_{k^{\prime}}\right)=\emptyset$. Event II, there exists a unique $k^{\prime} \in\{d+1, d+$ $2, \cdots, d+\alpha\}$ such that $S_{k}$ and $S_{k^{\prime}}$ intersect at exactly one location(row index).

```1951
```

25.1 Algorithm for Orlicz norm regression
27.1 The grid structure over the point set. From top to bottom, three levels of grids are
shown. Each cell splits into $2^{d}$ cells in the next level. . . . . . . . . . . . . . . . 2091
27.2 Random shift of grid brings down the number of heavy cells. In the left panel, we have a bad alignment of points and grids such that many cells contain lots of points. In the right panel, after the random shift, only two cells contain many points.
30.1 Birthday paradox for triangle inequality: let $w_{1}, w_{2}, w_{3}$ be three windows of length $d=8$ and assume $\lambda=1 / 2$. The LCS between $w_{1}$ and $w_{2}$ is $\lambda d=4$ and the LCS between $w_{2}$ and $w_{3}$ is $\lambda d=4$. Finally due to birthday paradox, we expect that the LCS between $w_{1}$ and $w_{3}$ is $\lambda^{2} d=2$.
30.2 Let $w_{i}, w_{a}, w_{b}$ and $w_{j}$ denote four windows and each of them has length $d=8$. This figure shows how the intersection of the edges of three windows are taken in order to construct a solution for the LCS of $w_{i}$ and $w_{j}$. If the size of the intersection is large, then such a tuple is called constructive. The solid lines represent LCS between two strings, and the dashed line represents the intersection of the three LCSs.

2229
30.3 Red rectangles show the elements of $\mathrm{sa}_{i}$ that contribute to lis $(A)$ and gray circles show the elements of sa that are sampled via our algorithm.
30.4 The graph on the left is an example of the string-based graph and the graph on the right is an example of the character-based graph.
30.5 computing a set of common subsequences between $w_{a}$ and $w_{i}$ for each $i$ such that every character of $w_{a}$ appears at most once among the sequences for a fixed $i$
30.6 The flowchart of the $O_{\lambda}\left(n^{\epsilon}\right)$ time algorithm is shown.
32.1 Each tree with green edges on the top-left is a rooted tree of each contracted component. For example, there are five components $\{1,2,3\},\{4,5,6,7\},\{8,9$, $10,11,12\},\{13,14,15\},\{16,17\}$. The dashed edges in the bottom-left figure is a root spanning tree of five components. The red edges in the top-right figure correspond to the dashed edges in the bottom-left figure before contraction. In bottom-right figure, by changing (see blue edges) the root of each contracted tree, we get a rooted spanning tree in the original graph
32.2 Given a tree that has 42 vertices (top-left), we label all the vertices from 1 to 42. Firstly, we sample some leaves (red vertices, i.e. $\{5,13,24,30,32,34,36,37$, $40,42\}$ ) in the tree (top-right tree). Then we find a DFS sequence of the tree (the tree formed by all the blue and red vertices in the bottom-left tree) which only contains all the sampled leaves and their ancestors. Finally, we recursively find the DFS sequences of remaining subtrees(bottom-right).
32.3 A hard example for [RMCS13]. For each $i \in\{2,3, \cdots, n / D-1\}$ and $j \in$ $\{1,2, \cdots, D-1\}$, node $(i-1) \cdot D+j$ has degree 4 . For node $D$ and $n$, they have degree 2. Node 0 has degree $D$. All the other nodes have degree 3. . . .
32.4 An example where \#roots $\approx \sum_{i=1}^{20} 1 /\left(d\left(v_{i}\right)+1\right)$. For each node, it has two numbers, the first number is the ID, and the second number is weight. $\sum_{i=1}^{20} 1 /\left(d\left(v_{i}\right)+1\right)=1 / 4+1 / 3+1 / 3+1 / 6+1 / 5+1 / 5+1 / 5+1 / 5+1 / 5+1 / 6$ $+1 / 4+1 / 6+1 / 8+1 / 7+1 / 6+1 / 9+1 / 8+1 / 7+1 / 6+1 / 4 \approx 3.89$ and \#roots= 4 .

## List of Algorithms

2.1 Stochastic Step ..... 19
2.2 Our Main Algorithm ..... 20
2.3 Projection Maintenance Data Structure - Initial, Query ..... 47
2.4 Projection Maintenance Data Structure - Update ..... 48
3.1 Central Path Maintenance Data Structure - Initial, Query, Move ..... 122
3.2 Central Path Maintenance Data Structure - Update and PartialUpdate ..... 123
3.3 Central Path Maintenance Data Structure - Full Update ..... 124
3.4 Central Path Maintenance Data Structure - Multiply and Move ..... 125
3.5 Robust Central Path ..... 156
3.6 Our main algorithm (More detailed version of RobustIPM in Section 3.4) ..... 157
4.1 Hamiltonian Monte Carlo Algorithm ..... 179
4.2 Collocation Method ..... 187
10.1 Interval-forest Sparse Recovery ..... 642
10.2 The Prune Procedure ..... 653
$10.3 \ell_{p}$-tail Estimation Algorithm ..... 668
10.4 Stronger $\ell_{2} / \ell_{2}$ Algorithm ..... 674
11.1 Continuous Fourier Sparse Recovery - Part 1 ..... 753
11.2 Continuous Fourier Sparse Recovery - Part 2 ..... 754
11.3 Continuous Fourier Sparse Recovery - Part 3 ..... 755
11.4 Continuous Fourier Sparse Recovery - Part 4 ..... 756
12.1 Linear Regression Algorithms ..... 881
12.2 Multipoint Evaluation of a Polynomial ..... 882
12.3 Get Empirical One Energy and Get Legal One Sample ..... 910
12.4 Locate One Signal, Locate One Inner, and Frequency Recovery One Cluster . ..... 911
12.5 Main Algorithm for One-cluster Recovery ..... 912
12.6 Locate $k$ Signal, Locate $k$ Inner, Hash To Bins ..... 913
12.7 Get Empirical $k$ Energy, Get Legal $k$ Sample, and Onestage ..... 914
12.8 Main Algorithm for $k$-cluster Recovery ..... 915
13.1 Fourier Sparse Recovery by Projection, $O(k \log n)$ Measurements When $k=$ $O(\log n)$ ..... 930
13.2 Fourier Sparse Recovery by Random Shift and Projection (Informal Version) ..... 934
13.3 Fourier Sparse Recovery by Random Shift and Projection ..... 944
13.4 Procedure for Reducing $\ell_{\infty}$-norm of the Residual Signal ..... 950
14.1 Iterative Set Query ..... 968
14.2 Fourier Set Query Algorithm ..... 986
15.1 Set Query Algorithm ..... 997
15.2 Sparse Recovery Algorithm ..... 1000
15.3 Fourier set query algorithm ..... 1020
16.1 Bicriteria Approximation Algorithm - Section 16.10 ..... 1105
16.2 Weighted Nonnegative Matrix Factorization ..... 1106
16.3 Weighted Nonnegative Matrix Factorization ..... 1107
17.1 Main Meta-Algorithm ..... 1128
17.2 Input Sparsity Time Algorithm ..... 1145
$17.3 \operatorname{poly}(k) \log d$-approximation Algorithm ..... 1150
17.4 $\widetilde{O}(k)$-approximation Algorithm ..... 1153
17.5 Bicriteria $O(1)$-approximation Algorithm ..... 1159
17.6 CUR Decomposition Algorithm ..... 1164
17.7 poly $(r, k)$-approximation Algorithm for Rank- $r$ matrix $B$ ..... 1172
17.8 Turnstile Streaming Algorithm ..... 1327
17.9 Row Update Streaming Algorithm ..... 1328
17.10Arbitrary Partition Distributed Protocol ..... 1329
17.11Row Partition Distributed Protocol ..... 1330
18.1 A Meta Algorithm ..... 1337
$19.1 \ell_{1}$-Low Rank Approximation with Input Assumption ..... 1378
20.1 Low Rank Approximation Algorithm for General Functions ..... 1437
$20.2 \ell_{0}$ regression [APY09] ..... 1457
21.1 Main Meta-Algorithm ..... 1489
21.2 Frobenius Norm Low-rank Approximation ..... 1522
21.3 Reducing the Size of the Objective Function from $\operatorname{poly}(n)$ to $\operatorname{poly}(k)$ ..... 1529
21.4 Frobenius Norm Tensor Multiple Regression ..... 1533
21.5 Frobenius Norm Bicriteria Low Rank Approximation Algorithm, rank- $O\left(k^{3} / \epsilon^{3}\right) 1538$
21.6 Frobenius Norm Low Rank Approximation Algorithm, rank- $O\left(k^{2} / \epsilon^{2}\right)$ ..... 1540
21.7 Generalized Matrix Row Subset Selection: Constructing $R$ with $r=O(k+k / \epsilon)$
Rows and a rank-k $U \in \mathbb{R}^{k \times r}$ ..... 1550
21.8 Frobenius Norm Tensor Column, Row and Tube Subset Selection, Polynomial Time ..... 1556
21.9 Frobenius Norm Tensor Column, Row and Tube Subset Selection, Input Spar- sity Time ..... 1560
21.10Fast Tensor Leverage Score Sampling ..... 1570
21.11Frobenius Norm CURT Decomposition Algorithm, Input Sparsity Time and Nearly Optimal Number of Samples ..... 1575
21.12Frobenius Norm CURT Decomposition Algorithm, Optimal Sample Complexity ..... 1580
21.13Frobenius Norm Tensor Column-row, Row-tube and Tube-column Face Subset Selection ..... 1581
21.14Frobenius Norm (Face-based) CURT Decomposition Algorithm, Optimal Sam- ple Complexity ..... 1586
21.15Fast Tensor Leverage Score Sampling, for General $q$-th Order ..... 1594
21.16General $q$-th Order Iterative Existential Proof ..... 1598
21.17General $q$-th Order Input Sparsity Reduction ..... 1600
21.18General $q$-th Order Bicriteria Algorithm ..... 1601
21.19General $q$-th Order CURT Decomposition ..... 1603
21.20Optimal Matrix CUR Decomposition Algorithm ..... 1605
21.21Reducing the Size of the Objective Function to poly $(k)$ ..... 1627
21.22生-Low Rank Approximation, Bicriteria Algorithm, rank- $\widetilde{O}\left(k^{3}\right)$, Nearly Input Sparsity Time ..... 1633
21.23 $\ell_{1}$-Low Rank Approximation, Bicriteria Algorithm, rank-poly $(k)$, Input Spar- sity Time ..... 1635
21.24生-Low Rank Approximation, Bicriteria Algorithm, rank- $\widetilde{O}\left(k^{2}\right)$, Nearly Input Sparsity Time ..... 1636
$21.25 \ell_{1}$-Low Rank Approximation, Bicriteria Algorithm, rank-poly $(k)$, Input Spar- sity Time ..... 1638
21.26 $\ell_{1}$-Low Rank Approximation, Bicriteria Algorithm, rank- $\widetilde{O}\left(k^{2}\right)$, Input Spar- sity Time ..... 1640
21.27 $\ell_{1}$-Low Rank Approximation, Input sparsity Time Algorithm ..... 1641
21.28建-Low Rank Approximation Algorithm, $\widetilde{O}\left(k^{3 / 2}\right)$-approximation ..... 1642
$21.29 \ell_{1}$-CURT Decomposition Algorithm ..... 1644
$21.30 \ell_{1}$-CURT decomposition algorithm ..... 1648
21.31Weighted Tensor Low-rank Approximation Algorithm when the Weighted Tensor has $r$ Distinct Faces in Each of the Three Dimensions. ..... 1652
21.32Weighted Tensor Low-rank Approximation Algorithm when the Weighted Tensor has $r$ Distinct Faces in Each of the Two Dimensions. ..... 1662
21.33Frobenius Norm Low (Tucker) Rank Approximation ..... 1707
21.34Frobenius Norm Low (Train) rank Approximation ..... 1712
$21.35 \ell_{p}$-Low Rank Approximation, Bicriteria Algorithm, rank- $\widetilde{O}\left(k^{2}\right)$, Input Spar- sity Time ..... 1732
$21.36 \ell_{1}$-Frobenius $\left(\ell_{1}-\ell_{2}-\ell_{2}\right)$ Low-rank Approximation Algorithm, poly $(k)$ approxi- mation ..... 1752
$21.37 \ell_{1}$-Frobenius $\left(\ell_{1}-\ell_{2}-\ell_{2}\right)$ Low-rank Approximation Algorithm, poly $(k, \log n)$ ap- proximation ..... 1759
$21.38 \ell_{1}-\ell_{1}-\ell_{2}$-Low Rank Approximation algorithm, input sparsity time ..... 1776
21.39 $\ell_{1}-\ell_{1}-\ell_{2}$-Low Rank Approximation Algorithm, $\widetilde{O}\left(k^{2 / 3}\right)$ ..... 1777
21.40 $\ell_{1}-\ell_{1}-\ell_{2}$-Low Rank Approximation Algorithm, Bicriteria Algorithm ..... 1778
21.41Turnstile Frobenius Norm Low Rank Approximation Algorithm ..... 1782
21.42Distributed Frobenius Norm Low Rank Approximation Protocol ..... 1794
21.43Distributed Frobenius Norm Low Rank Approximation Protocol ..... 1795
22.1 Generate importance sampling indices ..... 1820
22.3 Subroutine for approximate tensor contraction $A(I, v, w)$ ..... 1877
22.4 Subroutine for approximate tensor contraction $A(u, v, w)$ ..... 1877
22.5 Our main algorithm ..... 1878
22.2 Sampling Procedure for Rank-2 ..... 1882
26.1 Our $\ell_{2}$ Kronecker Product Regression Algorithm ..... 2032
26.2 Our $\ell_{p}$ Kronecker Product Regression Algorithm, $1 \leq p<2$. ..... 2033
26.3 Our All-Pairs Regression Algorithm ..... 2037
26.4 Algorithm to $\ell_{p}$ sample $\Theta\left(r_{2}\right)$ entires of $\rho=\left(A_{1} \otimes \cdots \otimes A_{q}\right) x^{\prime}-b$ ..... 2085
26.5 Algorithm to $\ell_{p}$ sample $\Theta(r)$ rows of $M=(F \otimes \mathbf{1}-\mathbf{1} \otimes F) R^{-1}$ ..... 2086
27.1 Sensitivity Estimation ..... 2102
27.2 Sensitivity Sampling Based Coreset Construction ..... 2103
27.3 Point-cell storing procedure ..... 2116
27.4 Estimating the number of points in each cell and in each level ..... 2119
27.5 Sensitivity sampling over a dynamic stream ..... 2125
27.6 Coreset construction over a dynamic stream ..... 2133
29.1 Approximate LCS algorithm ..... 2211
30.1 Sparsification for $\mathrm{O}_{\lambda^{2}}$ ..... 2241
30.2 Sparsification for $\mathrm{O}_{\lambda^{3}}$ ..... 2254
30.3 Ics-cmp data structure ..... 2299
30.4 Constructing the candidate domains ..... 2300
30.5 Constructing the pseudo solutions ..... 2300
30.6 Evaluate the pseudo solutions ..... 2301
30.7 Recursive Estimate LIS with Oracle ..... 2302
30.8 Recursive Estimate LIS ..... 2303
30.9 An algorithm that is able to generate multiple layers ..... 2304
30.10Dynamic programming algorithm for block-based LCS problem ..... 2305
30.11 ..... 2306
30.12Constructing $M$ based on $\widehat{M}$ ..... 2307
32.1 Neighbor Increment Operation ..... 2342
32.2 Tree Contraction Operation ..... 2351
32.3 Graph Connectivity ..... 2484
32.4 Local Complete Shortest Path Tree Expansion ..... 2485
32.5 Doubling Algorithm for Local Complete Shortest Path Trees ..... 2486
32.6 Large Local Shortest Path Trees ..... 2487
32.7 Depth and Ancestors of Every Vertex ..... 2488
32.8 Path in a Tree ..... 2489
32.9 Root Changing ..... 2489
32.10Spanning Forest Expansion ..... 2491
32.11Undirected Graph Spanning Forest ..... 2492
32.12Rooted Spanning Forest ..... 2493
32.13Lowest Common Ancestor ..... 2494
32.14Multiple Paths ..... 2495
32.15Leaf Sampling ..... 2496
32.16DFS subsequence ..... 2497
32.17DFS sequence ..... 2499
32.18A Sparser Table for RMQ ..... 2500
32.19A Sparse Table for RMQ ..... 2501
32.20Leader Selection via Min Parent Forest ..... 2502

## Chapter 1

## Introduction

### 1.1 Introduction

Many important questions in computer science can be formulated as problems on matrices. Having a better understanding of the matrix is crucial for devising fast algorithms for many fundamental problems, including a majority of the 10 most important algorithms of the 20th century [Cip00]. Examples of this include linear programming, Metropolis algorithm for Monte Carlo, matrix/tensor factorization, low-rank approximation, Fourier transforms, etc. In this thesis, we make progress on a number of long-standing open problems related to this areas. Along the way, we propose a number of new techniques which be could be useful in the future. We first summarize several core ideas, new insights, and new techniques related to the matrix. These techniques will eventually lead to fast algorithms. Later on in the introduction, we go to the details of each category.

- Sketching. A powerful technique we will use in many places throughout this thesis is sketching. Given a matrix $A \in \mathbb{R}^{n \times d}$, our goal is to produce a concise representation of $A$ to summarize informationa about $A$ without using $A$ directly. This is typically done by choosing a "small" sketching matrix $S$, and replacing $A$ with $S A$. We say $S A$ is a sketch of $A$. There are many previously studied choices for $S$, depending on the specific problem instance, including random Gaussian matrices, subsampled
randomized Hadamard matrices, and Count-Sketch matrices. By using $S A$, which is typically much smaller than $A$, we are able to speed up a number of algorithms.

This idea is often called "sketch and solve" in numerical linear algebra. Suppose that, for a given problem, we can produce a solver, which we wish to speed up. Oftentimes, we can apply a sketching matrix to reduce the dimensionality/freedom/size of the original problem, and then run the original solver on this smaller problem, and obtain a faster algorithm. Some classical examples are $\ell_{2}$ norm linear regression and Frobenius norm low-rank approximation. A number of results in this thesis use much more complicated versions of these sketching techniques, such as "guess and sketch" and "sketch and iterate," to achieve similar speed-ups in other problem settings. Many results in this thesis modify the classical iterating algorithms more decently with a better understanding of hashing-trick, sampling, and concentration techniques.

- Sampling. Given a matrix $A \in \mathbb{R}^{n \times d}$, our goal is to obtain a concise representation of $A$ by sampling some rows of $A$ and reweighting them. Let $D \in \mathbb{R}^{n \times n}$ denote a sparse diagonal matrix. Then $D A$ can be viewed as a sampled version of $A$. The goal of this is similar to sketching, as the goal is to produce a smaller representation of $A$, so that our algorithms run faster. In our context, the main difference between sketching and sampling is that in sketching, each row of $S A$ is a linear combination of rows in $A$, and whereas when we sample, each row of $D A$ is a row of $A$ with reasonable rescaling factor. There are several well-known examples, e.g. importance sampling, leverage scores sampling, and Lewis weights sampling.
- By sampling (in conjunction with some other ideas), we are able improve the
runtime for solving some convex problems, e.g. linear programs [CLS19]. We propose an algorithm that solves linear programs in the current matrix multiplication time. Our result breaks a thirty-year-old barrier.
- Sampling can also be used to solve some non-convex problems. We give a nearly optimal matrix CUR decomposition [SWZ17, SWZ19b] algorithm. We also extend the definition of matrix CUR problem to tensor CURT problem [SWZ19b] and provide a nearly optimal algorithm.
- Sampling can be used to speed up the iterating algorithm for orthogonal tensor decomposition [SWZ16]. This leads to a sublinear time finally.
- Matrix Maintainance. Given a matrix $A$ and a sequence of $T$ matrices $W^{1}, W^{2}, \cdots$, $W^{T}$, our goal is to approximately maintain $f\left(A, W^{t}\right)$ over all the iterations. Formally speaking, whenever we receives $W^{t}$, we want to generate $g\left(A, W^{t}\right)$ such that $g\left(A, W^{t}\right) \approx f\left(A, W^{t}\right)$. The objective is to do so faster than simply calculating $f\left(A, W^{t}\right)$ independently for each $t=1, \ldots, T$, by leveraging some structure in the sequence of update matrices.
- In [LSZ19], we propose a technique for doing this, which we call "iterate and sketch", when $f$ takes the form of a projection. Specifically, we consider functions $f$ of the form

$$
f(A, W)=\sqrt{W} A^{\top}\left(A W A^{\top}\right)^{-1} A \sqrt{W}
$$

When using the central path method to solve linear programs, in each iteration we need to multiply $f\left(A, W^{t}\right)$ by some vector. However, since the updates $W^{t}$ are
very structured, we don't need to compute this from scratch. By instead cleverly maintaining the sequence of updates, we are able to decrease the running time of the solver. Note that this idea is very different from "sketch and solve" described in the early paragraph. The main difference between "sketch and solve" and "iterate and sketch" is, the former just uses the original solver as a black-box, whereas the latter changes the original solver and uses a different sketching matrix over each iteration of the solver. There are other examples in the thesis can be thought of as "iterate and sketch", e.g. [CLS19, SWZ16].

- Matrix Concentration. The Chernoff bound for the concentration of scalar random variables is a fundamental and ubiquitous tool in the analysis of randomized algorithms. One common form of the bound is the following: given a list of independent random variables $x_{1}, \cdots, x_{k} \in[0,1]$ with mean $\mu$, then

$$
\operatorname{Pr}\left[\left|\frac{1}{k} \sum_{i=1}^{k} x_{i}-\mu\right|>\epsilon\right] \leq 2 \exp \left(-\Omega\left(k \epsilon^{2}\right)\right) .
$$

There is a natural generalization of this to matrices. Namely, if $X_{1}, \cdots, X_{k}$ are independent $n \times n$ complex Hermitian random matrices with $\left\|X_{i}\right\| \leq 1$ and mean $\mu$, then the following is true:

$$
\operatorname{Pr}\left[\left\|\frac{1}{k} \sum_{i=1}^{k} X_{i}-\mu\right\|>\epsilon\right] \leq 2 n \exp \left(-\Omega\left(k \epsilon^{2}\right)\right)
$$

Over the past decade, the matrix generalization of the Chernoff bound has also found widespread application, but this generalization is fairly restrictive. In particular, this bound requires strong independence of the individual $X_{i}$. In contrast, in the scalar setting, it is known that often times one can get away with some weak dependences
between the summands. For a number of decades, it has been open whether one can remove some of these restrictions in the matrix setting as well.

- In the scalar setting, it is well-known that if the random variables follow the dependence structure of a random walk on a graph, then the random variables still obey a Chernoff-style bound. It was conjectured by Wigderson and Xiao that a similar result holds in the matrix setting. In [GLSS18], we resolve this conjecture in the affirmative. Formally speaking, let $v_{1}, \cdots, v_{k}$ denote a random walk on expander graph $G$, let $f$ denote a mapping $f: V \rightarrow \mathbb{R}^{n \times n}$ with $\|f(v)\| \leq 1$, $\forall v \in V$, then $\sum_{i=1}^{k} f\left(v_{i}\right)$ follows matrix Chernoff bound.
- We show a matrix Chernoff bound for random spanning trees [KS18]. Such a result was conjectured by Batson, Spielman, Srivastava and Teng [BSST13]. Formally speaking, let $A_{1}, \cdots, A_{k} \in \mathbb{R}^{n \times n}$ denote a list of matrices where each of them corresponds to a random spanning tree. We show that $k=O\left(\log ^{2} n\right)$ suffices to imply that $\sum_{i=1}^{k} A_{i}$ gives spectral sparsifier for graph $G$.
- Sketch and Guess. The "sketch and guess" is the following, surprisingly powerful, paradigm: to solve matrix (or tensor) factorization problem, first demonstrate that the optimal solution may be represented as a concise sketch with fewer degrees of freedom, and then simply "guess" this sketch by exhaustively searching over all possible sketches. Since we have reduced the degrees of freedom, this exhaustive search can be done relatively efficiently. For instance, consider the problem of low-rank factorization. Given a matrix $A \in \mathbb{R}^{n \times n}$, our goal is to find two matrices $U \in \mathbb{R}^{n \times k}$ and $V \in \mathbb{R}^{k \times n}$ so that $\|U V-A\|_{F}$ is minimized. The idea is to write a polynomial system with a small
number of variables and constraints, abd low degree, and then run a polynomial system solver on this small polynomial system. For example, in order to solve min $\|U V-A\|_{F}^{2}$, we can create $2 n k$ variable (one for every entry of $U$ and $V$ ), and minimize the resulting degree-4 polynomial system. The polynomial system solver has exponential dependence on the number of variables, and thus has runtime $2^{O(n k)}$. The trick to reduce the number of variables is choosing two sketching matrices $S$ (a fat one, it has size $O(k) \times n$ ) and $R$ (a tall one, it has size $n \times O(k)$ ), then focusing on $\min _{X, Y}\|A S X Y R A-A\|_{F}^{2}$. To solve this problem, we can create $O\left(k^{2}\right)$ variables to minimize a degree-4 polynomial system, resulting in a runtime that, while exponential in $k^{2}$, is polynomial in $n$. Of course, for this specific problem, this technique is overkill, as we can solve this problem in time which is polynomial in both $n$ and $k$ via SVD. However, in situations where the problem we are trying to solve is NP-hard, this technique (and extensions thereof) can provide a nearly optimal approximation algorithms.
- We first propose the "guess a sketch" idea in [RSW16]. The combination of guessing and sketching can be used to get FPT-type algorithms for several matrix/tensor factorization problems [SWZ17, SWZ19b].
- For the entry-wise L1 norm, we propose "sketch and relax" idea in [SWZ17]. We want to solve min $\|U V-A\|_{1}$. To write a polynomial, we need to create variables for $U, V$ and also $n^{2}$ sign variables. To reduce the number of variables, we use L1 embedding matrix to reduce the dimension from $n$ to $r=\operatorname{poly}(k)$ and then relax the problem to Frobenius by losing some factors in the approximation ratio.

Eventually, we just need to solve

$$
\min _{X, Y}\left\|T_{1} A S X Y R A T_{2}-T_{1} A T_{2}\right\|_{1}
$$

Since the number of rows of $T_{1}$ and $T_{2}^{\top}$ is small, thus, we only need to create small number of sign variables, and as a result we obtain a faster runtime.

- For the tensor case, we propose "sketch and rotate" idea in [SWZ19b]. We want to solve min $\|U \otimes V \otimes W-A\|_{F}^{2}$. Slightly different from matrix case, we can create $3 n k$ variables to minimize a degree-6 polynomial system. Applying the sketching trick again, we can write

$$
\min _{X, Y, Z}\left\|A_{1} S X \otimes A_{2} R Y \otimes A_{3} T Z-A\right\|_{F}^{2}
$$

where $A_{1}, A_{2}, A_{3} \in \mathbb{R}^{n \times n^{2}}$ are the three possible flattening of tensor $A$ and $S, R, T$ are random sketching matrices.

- Iterating and Concentration. Deep neural networks (DNNs) have demonstrated dominating performance in many fields; since AlexNet, networks used in practice are going wider and deeper. We want to give some explanation about why stochastic gradient descent (SGD) can find global minima on the training objective of DNNs and Recurrent Neural Networks (RNNs). In previous paragraphs, we have discussed several techniques for improving convex, non-convex, and NP-hard problems. Unfortunately, the theory of deep learning cannot benefit from those "sketch and solve", "sketch and iterate", and "sketch and guess" techniques.
- One view of the over-parameterization theory of DNNs [AZLS19] is, $\|\phi(W x)\|_{2}$ has a good concentration when $x$ is some fixed unit vector, $W$ is random Gaussian matrix, and $\phi(t)=\max \{t, 0\}$ is the ReLU activation function. Technically
speaking, if we normalize $W$ properly, $\|\phi(W x)\| \approx(1+\epsilon)\|x\|_{2}$ with very high probability. In each of iteration of SGD, we update $W \leftarrow W+\Delta W$. There, we need to argue such concentration-type result for different $W$ s in each iteration. We call it "iterate and concentrate". A variation of this idea appears in [LSZ19] (not involved ReLU), where we use subsampled randomized Hadamard matrix to show concentration over each iteration.
- RNNs are multi-layer networks widely used in natural language processing. They are harder to analyze than feedforward neural networks, because the same recurrent unit is repeatedly applied across the entire time horizon. By doing a more careful randomness decomposition (in conjunction with other ideas, e.g., concentrations, Gram-Schmidt process), we can show the provable guarantees of RNNs [AZLS18].
- Hashing. For a given vector $x \in \mathbb{C}^{n}$ and a parameter $k$, we consider the problem where we observe $y=\Phi x$, and our goal is to figure out a $O(k)$-sparse approximation to $x$. There are two situations in which this problem is often considered. In one setting, we can design the matrix $\Phi$. In this case, we want $\Phi$ to implement some hashing-based function. The major things to minimize are, the number of rows of $\Phi$, the column sparsity of $\Phi$ (encoding time), and decoding time [GI10]. Alternatively, in some other settings, we cannot design the matrix $\Phi$. Typically, in this setting, $\Phi$ is a discrete Fourier matrix, and the problem is referred to as the sparse Fourier transform problem. In this case, we want to smartly take some samples from $\Phi x$, and those samples are able to do a hashing job. The major things to minimize are sample complexity and decoding time.

The sparse Fourier transform also can be defined in the continuous setting. This problem can be thought of as the mathematical abstraction of supperresolution problems [Moi15], the development of which was awarded the 2014 Nobel Prize in Chemistry. Intuitively, in the discrete setting we are working with a matrix $\Phi$ of finite size, but in the continuous setting, we need to understand the "infinite matrix" $\Phi$, which is more formally a linear operator over functions. However, by generalizing intuitions from the finite setting, many of the ideas can be transferred over to this infinite setting. In the continuous setting, there is one more important parameter to consider, in addition to the ones mentioned above. Namely, we should try to minimize the duration of the signal our algorithm needs to observe.

- The setting where we can design $\Phi$ is often called the compressed sensing problem [Don06, CRT06b]. We give a new construction of $\Phi$ and the corresponding algorithm that improves the classical L2/L2 compressed sensing task [NS19].
- We give a new iterating algorithm for the $d$-dimensional discrete sparse Fourier transform with nearly optimal sample complexity [NSW19a].
- In the continuous setting, in [PS15] we propose an algorithm for the robust sparse Fourier transform under the assumption of a frequency gap. The sample complexity is linear in $k$ and logarithmic in the signal-to-noise ratio and the frequency resolution. Previous results with similar sample complexities could not tolerate an infinitesimal amount of i.i.d. Gaussian noise, and even algorithms with higher sample complexities increased the noise by a polynomial factor. We also give new results for how precisely the individual frequencies of the signal can be recovered.
- The previously mentioned work of [PS15] relies on recovering the frequencies, and as a result, requires a dependence on the frequency gap. However, if all we care about is reconstructing a $k$-sparse signal which is close to the original, but which potentially uses different frequencies, such a gap is not necessary. In [CKPS16], we demonstrate an algorithm for the robust sparse Fourier transform that does not require any dependence on the frequency gap. In some ways, both [PS15] and [CKPS16] can be thought of generalizations of the aforementioned hashing techniques to the continuous setting.


### 1.2 Optimization

Linear programming is one of the key problems in computer science. In both theory and practice, many problems can be reformulated as linear programs to take advantage of fast algorithms. For an arbitrary linear program $\min _{A x=b, x \geq 0} c^{\top} x$ with $n$ variables and $d$ constraints ${ }^{1}$, the fastest algorithm takes $O^{*}\left(\sqrt{d} \cdot \mathrm{nnz}(A)+d^{2.5}\right)^{2}$ where $\mathrm{nnz}(A)$ is the number of non-zeros in $A$ [LS14, LS15].

For the generic case $d=\Omega(n)$ we focus in this thesis, the current fastest runtime is dominated by $O^{*}\left(n^{2.5}\right)$. This runtime has not been improved since the result by Vaidya on 1989 [Vai87, Vai89b]. The $n^{2.5}$ bound originated from two factors: the cost per iteration $n^{2}$ and the number of iterations $\sqrt{n}$. The $n^{2}$ cost per iteration looks optimal because this is the cost to compute $A x$ for a dense $A$. Therefore, many efforts [Kar84, Ren88, NN89, Vai89a,

[^0]

Figure 1.1: Linear Programming is Chasing Matrix Multiplication

LS14] have been focused on decreasing the number of iterations while maintaining the cost per iteration. As for many important linear programs (and convex programs), the number of iterations has been decreased, including maximum flow [Mad13, Mad16], minimum cost flow [CMSV17], geometric median [CLM $\left.{ }^{+} 16\right]$, matrix scaling and balancing [CMTV17], and $\ell_{p}$ regression [BCLL18]. Unfortunately, beating $\sqrt{n}$ iterations (or $\sqrt{d}$ when $d \ll n$ ) for the general case remains one of the biggest open problems in optimization.

Avoiding this open problem, this thesis develops a stochastic central path method that has a runtime of $O^{*}\left(n^{\omega}+n^{2.5-\alpha / 2}+n^{2+1 / 6}\right)$, where $\omega$ is the exponent of matrix multiplication
and $\alpha$ is the dual exponent of matrix multiplication ${ }^{3}$. For the current value of $\omega \sim 2.38$ and $\alpha \sim 0.31$, the runtime is simply $O^{*}\left(n^{\omega}\right)$. This achieves the natural barrier for solving linear programs because linear system is a special case of linear program and that the currently fastest way to solve general linear systems involves matrix multiplication. Despite the exact approach used in [CW87, Wil12, DS13, LG14] cannot give a bound on $\omega$ better than 2.3078 [AFLG15] and all known approaches cannot achieve the bound $\omega=2$ [AW18b], it is still possible that $\omega=2.01$ using all known approaches. Therefore, we believe improving the additive $2+1 / 6$ term remains an interesting open problem.

Our method is a stochastic version of the short step central path method. This short step method takes $O^{*}(\sqrt{n})$ steps and each step decreases $x_{i} s_{i}$ by a $1-1 / \sqrt{n}$ factor for all $i$ where $s$ is the dual variable [Ren88] (See the definition of $s$ in (2.1)). This results in $O^{*}(\sqrt{n}) \times n=O^{*}\left(n^{1.5}\right)$ coordinate updates. Our method takes the same number of step but only updates $\widetilde{O}(\sqrt{n})$ coordinates each step. Therefore, we only update $O^{*}(n)$ coordinates in total, which is nearly optimal.

Our framework is efficient enough to take a much smaller step while maintaining the same running time. For the current value of $\omega \sim 2.38$, we show how to obtain the same runtime of $O^{*}\left(n^{\omega}\right)$ by taking $O^{*}(n)$ steps and $\widetilde{O}(1)$ coordinates update per steps. This is because the complexity of each step decreases proportionally when the step size decreases. Beyond the cost per iteration, we remark that our algorithm is one of the very few central path algorithms [PRT02, Mad13, Mad16] that does not maintain $x_{i} s_{i}$ close to some ideal vector in $\ell_{2}$ norm. We are hopeful that our stochastic method and our proof will be useful

[^1]| Year | Author | Refs | \#Iters | Cost/Iter | Total |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 1984 | Karmarkar | $[$ Kar84] | $O(n L)$ | $\widetilde{O}\left(n^{2.2}\right)$ | $\widetilde{O}\left(n^{3.2} L\right)$ |
| 1986 | Renegar | $[$ Ren88] | $O(\sqrt{n} L)$ |  | $\widetilde{O}\left(n^{2.7} L\right)$ |
| 1989 | Vaidya | $[$ Vai89b] | $O(\sqrt{n} L)$ | $O\left(n^{2}\right)$ | $O\left(n^{2.5} L\right)$ |
| 1994 | Nesterov and Nemirovskii | [NN94] | $O(\sqrt{n} L)$ | $O\left(n^{2}\right)$ | $O\left(n^{2.5} L\right)$ |
| 2014 | Lee, Sidford | $[$ LS14] | $\widetilde{O}(\sqrt{n} L)$ | $\widetilde{O}\left(n^{2}\right)$ | $O\left(n^{2.5} L\right)$ |
| 2015 | Lee, Sidford | [LS15] | $\widetilde{O}(\sqrt{n} L)$ | $\widetilde{O}\left(n^{2}\right)$ | $O\left(n^{2.5} L\right)$ |
| 2019 | Cohen, Lee, Song | [CLS19](thesis) | $\widetilde{O}(\sqrt{n} L)$ | $\widetilde{O}\left(n^{\omega-1 / 2}\right)$ | $O\left(n^{\omega} L\right)$ |

Table 1.1: Let $\omega$ denote the exponent of the current matrix multiplication. "Iters" denotes the "number of iterations". "Cost/Iter" denotes the "cost per iteration". LP has $n$ variables, $d=n$ constraints, and can be encoded in $L$ input bits. We consider the case where $A$ is a dense full rank matrix. We remark that all of results have difference when $\operatorname{rank}(A)$ is low $\operatorname{rank}, \operatorname{nnz}(A)$ is sparse and $n \not \approx d$. However, when $d=n, \operatorname{rank}(A)=n$ and $\operatorname{nnz}(A)=n^{2}$, the running time remains to be $O\left(n^{2.5}\right)$ for three decades.
for future research on interior point methods. In particular, it would be interesting to see how this can be combined with techniques in [Cla95, LS14] to get a faster algorithm for linear programs with $d \ll n$.

Besides the applications to linear programs, some of our techniques are probably useful for studying other important problems in convex optimization. In particular, our framework should be naturally extendable to a larger class of convex programs.

Empirical Risk Minimization Empirical Risk Minimization (ERM) problem is a fundamental question in statistical machine learning. There are a huge number of papers that have considered this topic [Nes83, Vap92, PJ92, Nes04, BBM05, BB08, NJLS09, MB11, FGRW12, LRSB12, JZ13, Vap13, SSZ13, DB14, DBLJ14, FGKS15, DB16, SLC+17, ZYJ17, ZX17, ZWX ${ }^{+} 17$, GSS17, MS17, NS17, AKK+17, Csi18, JLGJ18] as almost all convex optimization
machine learning can be phrased in the ERM framework [SSBD14, Vap92]. While the statistical convergence properties and generalization bounds for ERM are well-understood, a general runtime bound for general ERM is not known although fast runtime bounds do exist for specific instances [AKPS19].

Examples of applications of ERM include linear regression, LASSO [Tib96], elastic net [ZH05], logistic regression [Cox58, HJLS13], support vector machines [CV95], $\ell_{p}$ regression [Cla05, $\mathrm{DDH}^{+} 09$, BCLL18, AKPS19], quantile regression [Koe00, KH01, Koe05], AdaBoost [FS97], kernel regression [Nad64, Wat64], and mean-field variational inference [XJR02].

The classical Empirical Risk Minimization problem is defined as

$$
\min _{x} \sum_{i=1}^{m} f_{i}\left(a_{i}^{\top} x+b_{i}\right)
$$

where $f_{i}: \mathbb{R} \rightarrow \mathbb{R}$ is a convex function, $a_{i} \in \mathbb{R}^{d}$, and $b_{i} \in \mathbb{R}, \forall i \in[m]$. Note that this formulation also captures most standard forms of regularization as well.

Letting $y_{i}=a_{i}^{\top} x+b_{i}$, and $z_{i}=f_{i}\left(a_{i}^{\top} x+b_{i}\right)$ allows us to rewrite the original problem in the following sense,

$$
\begin{array}{ll}
\min _{x, y, z} & \sum_{i=1}^{m} z_{i}  \tag{1.1}\\
\text { s.t. } & A x+b=y \\
& \left(y_{i}, z_{i}\right) \in K_{i}=\left\{\left(y_{i}, z_{i}\right): f_{i}\left(y_{i}\right) \leq z_{i}\right\}, \forall i \in[m]
\end{array}
$$

We can consider a more general version where dimension of $K_{i}$ can be arbitrary, e.g. $n_{i}$. Therefore, we come to study the general $n$-variable form

$$
\min _{x \in \prod_{i=1}^{m} K_{i}, A x=b} c^{\top} x
$$

where $\sum_{i=1}^{m} n_{i}=n$. We give a result for solving the general model.
First-order algorithms for ERM are well-studied and a long series of accelerated stochastic gradient descent algorithms have been developed and optimized [Nes98, JZ13, XZ14, SSZ14, FGKS15, LMH15, MLF15, AY16, RHS+16, SS16, AH16, SLRB17, MS17, LMH17, LJCJ17, All17b, All17a, All18b, All18a]. However, these rates depend polynomially on the Lipschitz constant of $\nabla f_{i}$ and in order to achieve a $\log (1 / \epsilon)$ dependence, the runtime will also have to depend on the strong convexity of the $\sum_{i} f_{i}$. In this thesis, we want to focus on algorithms that depend logarithmically on diameter/smoothness/strong convexity constants, as well as the error parameter $\epsilon$. Note that gradient descent and a direct application of Newton's method do not belong to these class of algorithms, but for example, interior point method and ellipsoid method does.

Therefore, in order to achieve high-accuracy solutions for non-smooth and non strongly convex case, most convex optimization problems will rely on second-order methods, often under the general interior point method (IPM) or some sort of iterative refinement framework. So, we note that our algorithm is thus optimal in this general setting since second-order methods require at least $n^{\omega}$ runtime for general matrix inversion.

Our algorithm applies the interior point method framework to solve ERM. The most general interior point methods require $O(\sqrt{n})$-iterations of linear system solves [Nes98], requiring a naive runtime bound of $O\left(n^{\omega+1 / 2}\right)$. Using the inverse maintenance technique [Vai89b, CLS19], one can improve the running time for LP to $O\left(n^{\omega}\right)$. This essentially implies that almost all convex optimization problems can be solved, up to subpolynomial factors, as fast as linear regression or matrix inversion!

The specific case of $\ell_{2}$ regression can be solved in $O\left(n^{\omega}\right)$ time since the solution is explicitly given by solving a linear system. In the more general case of $\ell_{p}$ regression, [BCLL18] proposed a $\widetilde{O}_{p}\left(n^{|1 / 2-1 / p|}\right)$-iteration iterative solver with a naive $O\left(n^{\omega}\right)$ system solve at each step. Recently, [AKPS19] improved the runtime to $\widetilde{O}_{p}\left(n^{\max (\omega, 7 / 3)}\right)$, which is current matrix multiplication time as $\omega>7 / 3$. However, both these results depend exponentially on $p$ and fail to be impressive for large $p$. Otherwise, we are unaware of other ERM formulations that have have general runtime bounds for obtaining high-accuracy solutions.

Recently several works [AW18a, AW18b, Alm19] try to show the limitation of current known techniques for improving matrix multiplication time. Alman and Vassilevska Williams [AW18b] proved limitations of using the Galactic method applied to many tensors of interest (including Coppersmith-Winograd tensors [CW87]). More recently, Alman [Alm19] proved that by applying the Universal method on those tensors, we cannot hope to achieve any running time better than $n^{2.168}$ which is already above our $n^{2+1 / 6}$.

Algorithmic Theory of ODEs The complexity of sampling a high-dimensional density of the form $e^{-f(x)}$ where $f$ is a convex function is a fundamental problem with many applications [LS90, LS92, LS93, LV06a, LV06b, Dal17, DK17, DRD18, DCWY18]. The focus of this part is to give very fast, i.e., nearly linear time algorithms, for a large subclass of such densities. A motivating and important case is the loss function for logistic regression, widely used in machine learning applications [Ber44, Paa00, NJ02, HJLS13, Bou14]:

$$
\sum_{i=1}^{n} \phi_{i}\left(a_{i}^{\top} x\right)
$$

where $\phi_{i}$ are convex functions; a popular choice is $\phi(t)=\log \left(1+e^{-t}\right)$. Sampling according to $e^{-f}$ for this choice of $f$ corresponds to sampling models according to their KL-divergence,
a natural and effective choice for classification problems [HJLS13].
A general approach to sampling is by an ergodic Markov chain whose stationary distribution is designed to have the desired density. Traditionally, this is done via a Metropolis filter, which accepts a proposed (random) next step $y$ from the current point $x$ with probability $\min \left\{1, \frac{f(y)}{f(x)}\right\}$. While very general, one downside of this approach is the possibility of high rejection probabilities, which typically force local steps to be very small. Nevertheless, for arbitrary logconcave functions (including nonsmooth ones), this approach has the current best guarantees [LV06b].

Another family of algorithms is derived from an underlying continuous stochastic process with the desired stationary density. A classic example of such a continuous process is Brownian motion. To sample a convex body for example, one could use Brownian motion with a boundary reflection condition. This is written as the stochastic equation:

$$
\mathrm{d} X_{t}=\mathrm{d} W_{t}
$$

with reflection at the boundary of the domain, and $d W_{t}$ being infinitesimal Brownian motion. To sample from the density proportional to $e^{-f(x)}$, one can use the stochastic differential equation,

$$
\mathrm{d} X_{t}=-\nabla f\left(X_{t}\right) \mathrm{d} t+\sqrt{2} \mathrm{~d} W_{t} .
$$

By the classical Fokker-Planck equation, under mild assumptions on $f$, the stationary density of this process is proportional to $e^{-f(x)}$.

How can we turn these continuous processes into algorithms? One approach is to take small rather than infinitesimal steps, and this leads to the Langevin dynamics, of which there are multiple flavors [Dal17, DK17, ZLC17, RRT17, DRD18, CCBJ18, $\left.\mathrm{CCAY}^{+} 18, \mathrm{CFM}^{+} 18\right]$.

| Year | Refs. | Method | Iters. | Gra/Iter | Total time |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 2006 | $[$ LV06b $]$ | B./H. $^{*}$ | $d^{3}, d^{4}$ | 1 | $d^{5}, d^{6}$ |
| 2017 | $[$ Dal17] | LMC $^{*}$ | $\kappa^{2} d, \kappa^{3} d^{3}$ | 1 | $\kappa^{2} d^{2}, \kappa^{3} d^{4}$ |
| 2017 | $[$ Dal17] | LMCO $^{*}$ | $\kappa^{2} d, \kappa^{2} d^{2.5}$ | 1 | $\kappa^{2} d^{4}, \kappa^{2} d^{5.5}$ |
| 2018 | $[$ CCBJ18] | DL | $\kappa^{2} d^{0.5}$ | 1 | $\kappa^{2} d^{1.5}$ |
| 2018 | $[$ DCWY18] | MALA $^{*}$ | $\kappa d, \kappa d^{2}$ | 1 | $\kappa d^{2}, \kappa d^{3}$ |
| 2017 | $[$ MS17] | HMC $^{*, \dagger}$ | $\kappa^{6.5} d^{0.5}$ | 1 | $\kappa^{6.5} d^{1.5}$ |
| 2018 | $[$ MV18] | HMC $^{2.7}$ | $\kappa^{2.75} d^{0.25}, \kappa^{3.5} d^{0.25}$ | 1 | $\kappa^{2.75} d^{1.25}, \kappa^{3.5} d^{1.25}$ |
| 2018 | [LSV18] <br> (thesis) | HMC <br>  <br> HMC | $\kappa^{1.5}$ <br> $\kappa^{1.5}$ | $\kappa^{1.5} d$ <br> $\kappa^{0.25} d^{0.5}$ | $\kappa^{1.75} d^{1.5}$ |

* have different bounds for warm start and general (cold) start. We stated the runtime for cold start in green color.
${ }^{\dagger}$ make smoothness and incoherence assumptions motivated by and applicable to Bayesian logistic regression.

Table 1.2: Summary of results, $d$ is the dimension, $\kappa$ is the condition number of $\nabla^{2} f$. "Iters" denotes "the number of iterations / parallel depth". "Gra/Iter" denotes the "the number of gradients per iteration". We use the parallel depth of the algorithm as the number of iterations. We suppress polylogarithmic terms and dependence on the error parameter. Ball walk/hit-and-run apply to general logconcave distributions, the rest assume strongly logconcave with Lipschitz gradient and possibly more. "B./H." denotes "Ball Walk/Hit-andrun". "DL" denotes "Damped Langevin". In all previous work, for simplicity, we report the most favorable bounds by making various assumptions such as $\kappa \ll d$.

Starting with Dalalyan [Dal17], it has been established that these dynamics converge in polynomial (in dimension) time for strongly logconcave functions, with the underdamped version converging in $O(\sqrt{d})$ iterations (and polynomial dependences on appropriate condition numbers) [CCBJ18]. The dependence on dimension seems unavoidable in the discretized algorithm, even though the continuous process has no such dependence.

Hamiltonian Monte Carlo. HMC is a random process that maintains a position $x$ and velocity pair $v$. To sample according to $e^{-f}$, we define a Hamiltonian $H(x, v)=f(x)+\frac{1}{2}\|v\|^{2}$.

At each step $v$ is chosen randomly from $N(0, I)$ and $x$ is updated using the following Ordinary Differential Equation (ODE) for a some fixed time interval.

$$
\begin{equation*}
\frac{\mathrm{d} x(t)}{\mathrm{d} t}=v(t), \quad \frac{\mathrm{d} v(t)}{\mathrm{d} t}=-\nabla f(x(t)) \tag{1.2}
\end{equation*}
$$

This process has the particularly nice property that it conserves the value of $H$, and as a result there is no need to apply a Metropolis filter. HMC has been studied in many works [MS17, MV18, LV18]. Mangoubi and Smith [MS17] gave the following guarantee for strongly logconcave densities.

Note that Eq. (1.2) is a special case of the following first order ODE

$$
\frac{\mathrm{d}}{\mathrm{~d} t} x(t)=F(x(t), t), x(0)=v
$$

where $F: \mathbb{R}^{d+1} \rightarrow \mathbb{R}^{d}, x(t) \in \mathbb{R}^{d}$ and $v \in \mathbb{R}^{d}$.

Moreover, we can write the $k$-th order ODE in the following way

$$
\begin{aligned}
\frac{\mathrm{d}^{k}}{\mathrm{~d} t^{k}} x(t) & =F\left(\frac{\mathrm{~d}^{k-1}}{\mathrm{~d} t^{k-1}} x(t), \cdots, x(t), t\right) \\
\frac{\mathrm{d}^{i}}{\mathrm{~d} t^{i}} x(0) & =v_{i}, \forall i \in\{k-1, \cdots, 1,0\} .
\end{aligned}
$$

where $F: \mathbb{R}^{k d+1} \rightarrow \mathbb{R}^{d}, x(t) \in \mathbb{R}^{d}$, and $v_{0}, v_{1}, \cdots, v_{k-1} \in \mathbb{R}^{d}$.
To solve HMC such special case ODE efficiently, it requires a better understanding of the ODE system. We analyze the collocation method for solving ODEs [Ise09]. This method is classical in numerical analysis.

We present the multivariate high-oder ODE guarantee. This generalizes and improves on the guarantee from [LV17]. Note that our result is a general result about solving ODE
efficiently, independent of the application to sampling. The only assumptions needed are that the ODE function is Lipschitz and that the solution is close to the span of small number of basis of functions. These natural assumptions suffice to get around the worst-case complexity lower bounds for solving such general ODEs [KF82, Ko83, Ko10, Kaw10, KC12].

Deep Neural Networks Neural networks have demonstrated a great success in numerous machine-learning tasks [KSH12, GMH13, $\left.\mathrm{LHP}^{+} 15, \mathrm{AAA}^{+} 16, \mathrm{HZRS}^{2} 16, \mathrm{SHM}^{+} 16, \mathrm{SSS}^{+} 17\right]$. One of the empirical findings is that neural networks, trained by first-order methods from random initialization, have a remarkable ability to fit training data $\left[\mathrm{ZBH}^{+} 17\right]$.

From an expressibility perspective, this may not be surprising since modern neural networks are often over-parameterized: they have much more parameters than the number of training samples. There certainly exist parameter choices with zero training error as long as data is non-degenerate.

Yet, from an optimization perspective, the fact that randomly-initialized first-order methods can find global minima on the training data is quite non-trivial: neural networks are often equipped with the ReLU activation, making the training objective not only nonconvex, but even non-smooth. Even the general convergence for finding approximate critical points of a non-convex, non-smooth function is not fully-understood [BLO05] and appears to be a challenging question on its own. This is in direct contrast to practice, in which ReLU networks trained by stochastic gradient descent (SGD) from random initialization almost never suffer from non-smoothness or non-convexity, and can avoid local minima for a variety of network architectures (see [GVS15]). A theoretical justification was missing to explain this phenomenon.

There are quite a few papers trying to understand the success of neural networks from optimization perspective. Many of them focus on the case when the inputs are random Gaussian, and work only for two-layer neural networks [BG17, Son17, Tia17, LY17, DLT ${ }^{+}$18, GLM17, PRSZ18, ZSJ ${ }^{+}$17, ZSD17]. [LL18] show that for a two-layer network with ReLU activation, SGD finds nearly-global optimal (say, 99\% classification accuracy) solutions on the training data, as long as the network is over-parameterized, meaning that the number of neurons is polynomially large comparing to the input size. Moreover, if the data is sufficiently structured (say, coming from mixtures of separable distributions), this accuracy extends also to test data. As a separate note, over-parameterization is suggested as the possible key to avoid bad local minima by [SS18] even for two-layer networks.

There are also results that go beyond two-layer networks with limitations. Some consider deep linear neural networks without any activation functions [HM17, ACGH18, BHL18, Kaw16]. [Dal17] studies multi-layer neural networks but essentially only with respect to the convex task of training the last layer. ${ }^{4}$ [SC16] show that under over-parameterization and under random input perturbation, there is bad local minima for multi-layer neural networks. [JGH18] derive global convergence using neural tangent kernel for infinite-width neural networks.

In this thesis, we study the following fundamental questions

Can DNN be trained close to zero training error efficiently under mild assumptions?
If so, can the running time depend only polynomially in the network depth and input size?

[^2]Motivation In 2012, AlexNet was born with 5 convolutional layers [KSH12]. The later VGG network uses 19 layers [SSZ14], and GoogleNet uses 22 layers [SLJ ${ }^{+}$15]. In practice, we cannot go deeper by naively stacking layers together, due to the so-called vanishing/exploding gradient problem. To deal with this issue, networks with residual links (ResNet) were proposed with the capability of handling at least 152 layers [HZRS16]. Compared with practical networks that go much deeper, existing theory has been mostly around two-layer (thus one-hidden-layer) neural networks, even just for the training process alone. Thus,

Can we theoretically justify how the training process has worked for multi-layer neural networks?

In this thesis, we extend the over-parameterization theory to multi-layer neural networks.

Recurrent Neural Networks Among different architectures of neural networks, one of the least theoretically-understood structure is the recurrent one [Elm90]. A recurrent neural network recurrently applies the same network unit to a sequence of input tokens, such as a sequence of words in a language sentence. RNN is particularly useful when there are longterm, non-linear interactions between input tokens in the same sequence. These networks are widely used in practice for natural language processing, language generation, machine translation, speech recognition, video and music processing, and many other tasks [MKB ${ }^{+} 10$, MKB $^{+}$11, SSN12, KB13, SSB14, SVL14, CVMBB14, CGCB14]. On the theory side, while there are some attempts to show that an RNN is more expressive than a feedforward neural network [KNO18], when and how an RNN can be efficiently learned has little theoretical explanation.

In practice, RNN is usually trained by simple local-search algorithms such as SGD. However, unlike shallow networks, the training process of RNN often runs into the trouble of vanishing or exploding gradient [SMH11]. That is, the value of the gradient becomes exponentially small or large in the time horizon, even when the training objective is still constant. ${ }^{5}$

In practice, one of the popular ways to resolve this is by the long short term memory (LSTM) structure [HS97]. However, one can also use rectified linear units (ReLUs) as activation functions to avoid vanishing or exploding gradient [ $\left.\mathrm{SBS}^{+} 17\right]$. In fact, one of the earliest adoptions of ReLUs was on applications of RNNs for this purpose twenty years ago [Hah98, SA96].

Since the RNN structure was proposed, a large number of variations have been designed over past decades. For a more detailed survey, we refer the readers to [ $\left.\mathrm{SBS}^{+} 17\right]$.

In this thesis, we study the following general question

- Can ReLU provably stabilize the training process and avoid vanishing/exploding gradient?
- Can RNN be trained close to zero training error efficiently under mild assumptions?

Remark 1.2.1. When there is no activation function, RNN is known as linear dynamical system. Hardt, Ma and Recht [HMR18] first proved the convergence of finding global minima for

[^3]such linear dynamical systems. Followups in this line of research include [HSZ17, HLS ${ }^{+}$18].

This part is based on the following papers:

- Michael B. Cohen, Yin Tat Lee, Zhao Song

Solving Linear Programs in the Current Matrix Multiplication Time.
STOC 2019 [CLS19]

- Yin Tat Lee, Zhao Song, Qiuyi Zhang

Solving Empirical Risk Minimization in the Current Matrix Multiplication.
COLT 2019 [LSZ19]

- Yin Tat Lee, Zhao Song, Santosh S. Vempala Algorithmic Theory of ODEs and Sampling from Well-conditioned Logconcave Densities.

Manuscript 2018 [LSV18]

- Zeyuan Allen-Zhu, Yuanzhi Li, Zhao Song

A Convergence Theory for Deep Learning via Over-Parameterization.
ICML 2019 [AZLS19]

- Zeyuan Allen-Zhu, Yuanzhi Li, Zhao Song

On the convergence rate of training recurrent neural networks.
Manuscript 2018 [AZLS18]

### 1.3 Concentration

The study of concentration of sums of random variables dates back to Central Limit Theorems, and hence de Moivre and Laplace [Tij], while modern concentration bounds for sums of random variables were perhaps first established by Bernstein [Ber24], and a popular variant now known as Chernoff bounds was introduced by Rubin and published by Chernoff [Che52].

Concentration of measure for matrix-valued random variables is the phenomenon that many matrix valued distributions are to close their mean with high probability, closeness usually being measured by spectral norm. Modern quantitative bounds of the form often used in theoretical computer science were derived by Rudelson [Rud99], while Ahlswede and Winter [AW02] established a useful matrix-version of the Laplace transform that plays a central role in scalar concentration results such as those of Bernstein. [AW02] combined this with the Golden-Thompson trace inequality to prove matrix concentration results. Tropp refined this approach, and by replacing the use of Golden-Thompson with deep a theorem on concavity of certain trace functions due to Lieb, Tropp was able to recover strong versions of a wide range of scalar concentration results, including matrix Chernoff bounds, Azuma and Freedman's inequalities for matrix martingales [Tro12].

Matrix concentration results have had an enormous range of applications in computer science, and are ubiquitous throughout spectral graph theory [ST04, SS11, CKP ${ }^{+}$17], sketching [Coh16a], approximation algorithms [HSSS16], numerical linear algebra, quantum information theory, and deep learning [ZSJ ${ }^{+}$17, ZSD17, SY19]. Most applications are based on results for independent random matrices, but more flexible bounds, such as Tropp's Matrix Freedman Inquality [Tro11a], have been used to greatly simplify algorithms, e.g.
for solving Laplacian linear equations [KS16] and for semi-streaming graph sparsification [AG09, KPPS17]. Matrix concentration results are also closely related to other popular tools sampling tools, such as Karger's techniques for generating sparse graphs that approximately preserve the cuts of denser graphs [BK96].

We formally state the natural generalization of the Chernoff bound appeared in the works of Rudelson [Rud99], Ahlswede-Winter [AW02], and Tropp [Tro12]. They showed that a similar concentration phenomenon is true for matrix-valued random variables. In particular, if $X_{1}, \ldots, X_{k}$ are independent $d \times d$ complex Hermitian random matrices with $\left\|X_{i}\right\| \leq 1$, then the following is true:

$$
\begin{equation*}
\operatorname{Pr}\left[\left\|\frac{1}{k} \sum_{i=1}^{k} X_{i}-\mathbb{E}[X]\right\|>\epsilon\right] \leq 2 d \cdot \exp \left(-\Omega\left(k \epsilon^{2}\right)\right) \tag{1.3}
\end{equation*}
$$

The only difference between this and the usual Chernoff bound is the factor of $d$ in front of the deviation probability; to see that it is necessary, notice that the diagonal case simply corresponds to a direct sum of $d$ arbitrarily correlated instances of the scalar Chernoff bound, so by the union bound the probability should be $d$ times as large in the worst case. This so called "Matrix Chernoff Bound".

Expander Walk An important generalization of this bound was achieved by Gillman [Gil98] (with refinements later by [Lez98, Kah97, LP04, WX05, Hea08, Wag08, CLLM12, RR17]), who significantly relaxed the independence assumption to Markov dependence. In particular, suppose $G$ is a regular graph with vertex set $V=[n], X: V \rightarrow \mathbb{C}$ is a bounded function, and $v_{1}, \ldots, v_{k}$ is a stationary random walk ${ }^{6}$ of length $k$ on $G$. Then, even though

[^4]the random variables $X\left(v_{i}\right)$ are in not independent (except when $G$ is the complete graph with self loops), it is shown that:
\[

$$
\begin{equation*}
\operatorname{Pr}\left[\left|\frac{1}{k} \sum_{i=1}^{k} X\left(v_{i}\right)-\mathbb{E}[X]\right|>\epsilon\right] \leq 2 \cdot \exp \left(-\Omega\left((1-\lambda) k \epsilon^{2}\right)\right), \tag{1.4}
\end{equation*}
$$

\]

where $1-\lambda$ is the spectral gap of the transition matrix of the random walk. The gain here is that sampling a stationary random walk of length $k$ on a constant degree graph with constant spectral gap requires $\log (n)+O(k)$ random bits, which is much less than the $k \log (n)$ bits required to produce $k$ independent samples. Since such graphs can be explicitly constructed, this leads to a generic "derandomization" of the Chernoff bound, which has had several important applications (see [WX05] for a detailed discussion). In particular, it leads to the following randomness efficient sampler for scalar-valued functions ([Gil98]) using known strongly explicit constructions of expander graphs [RVW00, LPS88]:

Theorem 1.3.1 ([Gil98]). For any $\epsilon>0$ and $k \geq 1$, there is a poly $(r)$-time computable sampler $\sigma:\{0,1\}^{r} \rightarrow[n]^{k}$, where $r=\log (n)+O(k)$ s.t. for all functions $f:[n] \rightarrow[-1,1]$ satisfying $\mathbb{E} f=0$, we have that

$$
\operatorname{Pr}_{w \in \in_{R 0,1\}^{r}}}\left[\left|\frac{1}{k} \sum_{i=1}^{k} f\left(\sigma(w)_{i}\right)\right| \geq \epsilon\right] \leq 2 \exp \left(-\Omega\left(\epsilon^{2} k\right)\right)
$$

In many applications of interest $k$ is about $\log (n)$, and going from $O\left(\log ^{2}(n)\right)$ to $O(\log (n))$ random bits leads to a complete derandomization by cycling over all seeds $w \in$ $\{0,1\}^{r}$.

It is natural to wonder whether there is a common generalization of (1.4) and (1.3), i.e., a "Matrix Expander Chernoff Bound". Such a result was conjectured by Wigderson and

Xiao in [WX06]. In this thesis, we prove the Wigderson and Xiao conjecture. In order to prove that conjecture, we also propose a new Golden-Thompson inequality. The original Golden-Thompson inequality only holds for two matrices [Gol65, Tho65]. Our new GoldenThompson inequality holds for multiple matrices.

## Strongly Rayleigh and Matrix Chernoff for Random Spanning Trees Negative

 dependence of random variables is an appealing property that intuition suggests should help with concentration of measure. Notions of negative dependence can be formalized in many ways. Roughly speaking, these notions characterize distributions where where some event occurring ensures that other events of interest become less likely. A simple example is the distribution of a sequence of coin flips, conditioned on the total number of heads in the outcome. In this distribution, conditioning on some coin coming out heads makes all other coins less likely to come out heads. Unfortunately, negative dependence phenomena are not as robust as positive association which can be established from local conditions using the powerful FKG theorem [FKG71].Strongly Rayleigh distributions were introduced recently by Borcea, Brändén, and Liggett [BBL09] as a class of negatively dependent distributions of binary-valued random variables with many useful properties. Strongly Rayleigh distributions satisfy useful negative dependence properties, and retain these properties under natural conditioning operations. Strongly Rayleigh distributions also satisfy a powerful stability property under conditioning known as Stochastic Covering [PP14], which is useful for analyzing them through martingale techniques. A measure on $\{0,1\}^{n}$ is said to be Strongly Rayleigh if its generating polynomial is real stable [BBL09]. There are many interesting examples of Strongly Rayleigh distribu-
tions [PP14]: The example mentioned earlier of heads of independent coin flips conditional on the total number of heads in the outcome; symmetric exclusion processes; determinental point processes and determinental measures on a boolean lattice. An example of particular interest to us is the edges of uniform or weighted random spanning trees, which form a Strongly Rayleigh distribution.

We prove a Matrix Chernoff bound for the case of $k$-homogeneous Strongly Rayleigh distributions. Our bound is slightly weaker than the bound for independent variables. We give lower bounds that show our bounds are close to tight in some regimes, but importantly, our lower bounds do not establish separation from the behaviour of indepedent random matrices, leaving open the question of whether the true bound should match the independent case in all regimes - which seems plausible.

We use our bound to show new concentration results related to random spanning trees of graphs. Random spanning trees are one among the most well-studied probabilistic objects in graph theory, going back to the work of Kirchoff [Kir47] in 1847, who gave formula relating the number of spanning trees in a graph to the determinant of the Laplacian of the same graph.

Algorithms for sampling of random spanning trees have been studied extensively, [Gue83, Bro89, Ald90, Kul90, Wil96, CMN96, KM09, MST15, HX16, DKP+17, DPPR17, Sch18], and a random spanning tree can now be sampled in almost linear time [Sch18]. There is also a long line [BK96, ST11, SS11, BSS12, Zou12, LS17, LS18] of research about generating graph spectral sparsifier as fast as possible. However, the connection between those two objects is open for decades. We answer the question positively by providing the following result

Theorem 1.3.2. Given as input a weighted graph $G$ with $n$ vertices and a parameter $\epsilon>0$, let $T_{1}, T_{2}, \cdots, T_{t}$ denote $t$ independent inverse leverage score weighted random spanning trees, if we choose $t=O\left(\epsilon^{-2} \log ^{2} n\right)$ then with probability $1-1 / \operatorname{poly}(n)$,

$$
(1-\epsilon) L_{G} \preceq \frac{1}{t} \sum_{i=1}^{t} L_{T_{i}} \preceq(1+\epsilon) L_{G} .
$$

In addition, we also show $t=\Omega\left(\epsilon^{-2} \log n\right)$ is necessary.

Discrepancy Discrepancy theory is an area of combinatorics that studies how well continuous objects can be approximated by discrete ones. It lies at the heart of numerous problems in mathematics and computer science [Cha00]. Although closely tied to probability theory, direct randomized approaches rarely yield the best bounds. In a classical formulation in discrepancy theory, we have $n$ sets on $n$ elements, and would like to two-color the elements so that each set has roughly the same number of elements of each color. Using a simple random coloring, it is an easy consequence of Chernoff's bound that there exists a coloring such that the discrepancy in all $n$ sets is $O(\sqrt{n \log n})$ [AS16]. However, in a celebrated result, Spencer showed that in fact there is a coloring with discrepancy at most $6 \sqrt{n}$ [Spe85].

Recently, there has been significant success in generalizing Chernoff [Che52], Hoeffding, Bernstein, and Bennett-type concentration bounds for scalar random variables to matrix-valued random variables [Rud99, AW02, Tro12]. Consider the following matrix concentration bound, which is a direct consequence of a matrix Hoeffding bound.

Theorem 1.3.3 ([Tro12]). Let $\xi_{i} \in\{ \pm 1\}$ be independent, symmetric random signs and $A_{1}, \ldots, A_{n} \in \mathbb{C}^{m \times m}$ be positive semi-definite matrices. Suppose $\max _{i \in[n]}\left\|A_{i}\right\| \leq \epsilon$ and
$\left\|\sum_{i=1}^{n} A_{i}\right\| \leq 1$. Then,

$$
\operatorname{Pr}\left[\left\|\sum_{i=1}^{n} \xi_{i} A_{i}\right\| \geq t \sqrt{\epsilon}\right] \leq 2 m \exp \left(-t^{2} / 2\right)
$$

A consequence of this theorem is that with high probability

$$
\begin{equation*}
\left\|\sum_{i=1}^{n} \xi_{i} A_{i}\right\|=O(\sqrt{\log m}) \sqrt{\epsilon} \tag{1.5}
\end{equation*}
$$

The Kadison-Singer theorem of [MSS15b] is essentially equivalent to the following statement (which can readily be derived from the bipartition statement in [MSS15b]).

Theorem 1.3.4 (Kadison-Singer [MSS15b]). Let $u_{1}, \ldots, u_{n} \in \mathbb{C}^{m}$ and suppose $\max _{i \in[n]}\left\|u_{i} u_{i}^{*}\right\| \leq$ $\epsilon$ and $\sum_{i=1}^{n} u_{i} u_{i}^{*}=I$. Then, there exists signs $\xi_{i} \in\{ \pm 1\}$ s.t.

$$
\left\|\sum_{i=1}^{n} \epsilon_{i} u_{i} u_{i}^{*}\right\| \leq O(\sqrt{\epsilon})
$$

Thus, for rank 1 matrices the theorem improves on the norm bound in Equation (1.5), by a factor $\sqrt{\log m}$, in a manner analogous to the improvement of Spencer's theorem over the bound based on the scalar Chernoff bound.

For random signings of matrices one can establish bounds in some cases that are much stronger than Theorem 1.3.3.

Theorem 1.3.5 ([Tro12]). Let $\xi_{i} \in\{ \pm 1\}$ be independent random signs, and let $A_{1}, \ldots, A_{n} \in$ $\mathbb{C}^{m \times m}$ be Hermitian matrices. Let $\sigma^{2}=\left\|\sum_{i=1}^{n} \mathbb{V}\left[\xi_{i}\right] A_{i}^{2}\right\|$. Then,

$$
\operatorname{Pr}\left[\left\|\sum_{i=1}^{n} \mathbb{E}\left[\xi_{i}\right] A_{i}-\sum_{i=1}^{n} \xi_{i} A_{i}\right\| \geq t \cdot \sigma\right] \leq 2 m \exp \left(-t^{2} / 2\right)
$$

From this theorem we deduce that with high probability

$$
\begin{equation*}
\left\|\sum_{i=1}^{n} \mathbb{E}\left[\xi_{i}\right] A_{i}-\sum_{i=1}^{n} \xi_{i} A_{i}\right\|=O(\sqrt{\log m}) \sigma . \tag{1.6}
\end{equation*}
$$

Of course, this implies that there exists a choice of signs $\epsilon_{1}, \ldots, \epsilon_{n} \in\{ \pm 1\}$ such that

$$
\left\|\sum_{i=1}^{n} \mathbb{E}\left[\xi_{i}\right] A_{i}-\sum_{i=1}^{n} \epsilon_{i} A_{i}\right\|=O(\sqrt{\log m}) \sigma
$$

For rank-1 matrices, we want to show there exists a choice of signs with a stronger guarantee. Formally speaking, we can prove the following result.

Theorem 1.3.6. Consider any independent scalar random variables $\xi_{1}, \ldots, \xi_{n}$ with finite support. Let $u_{1}, \ldots, u_{n} \in \mathbb{C}^{m}$ and

$$
\sigma^{2}=\left\|\sum_{i=1}^{n} \mathbb{V}\left[\xi_{i}\right]\left(u_{i} u_{i}^{*}\right)^{2}\right\|
$$

Then there exists a choice of outcomes $\epsilon_{1}, \ldots, \epsilon_{n}$ in the support of $\xi_{1}, \ldots, \xi_{n}$

$$
\left\|\sum_{i=1}^{n} \mathbb{E}\left[\xi_{i}\right] u_{i} u_{i}^{*}-\sum_{i=1}^{n} \epsilon_{i} u_{i} u_{i}^{*}\right\| \leq O(\sigma) .
$$

This part is based on the following papers:

- Ankit Garg, Yin Tat Lee, Zhao Song, Nikhil Srivastava

A Matrix Expander Chernoff Bound.
STOC 2018 [GLSS18]

- Rasmus Kyng, Zhao Song

A Matrix Chernoff Bound for Strongly Rayleigh Distributions and Spectral Sparsifiers from a few Random Spanning Trees.

FOCS 2018 [KS18]

- Rasmus Kyng, Kyle Luh, Zhao Song

Four Deviations Suffice for Rank 1 Matrices.
Manuscript 2019 [KLS19]

### 1.4 Compressed Sensing and Sparse Fourier Transform

Compressed Sensing, or sparse recovery, is a powerful mathematical framework the goal of which is to reconstruct an approximately $k$-sparse vector $x \in \mathbb{R}^{n}$ from linear measurements $y=\Phi x$, where $\Phi \in \mathbb{R}^{m \times n}$. The most important goal is to reduce the number of measurements $m$ needed to approximate the vector $x$, avoiding the linear dependence on $n$. In discrete signal processing, where this framework was initiated [CRT06b, Don06], the core principle that the sparsity of a signal can be exploited to recover it using much fewer samples than the Shannon-Nyquist Theorem. We refer to the matrix $\Phi$ as the sketching or sensing matrix, and $y=\Phi x$ as the sketch of vector $x$.

Sparse recovery is the primary task of interest in a number of applications, such as image processing $\left[\mathrm{TLW}^{+} 06, \mathrm{LDP} 07 \mathrm{a}, \mathrm{DDT}^{+} 08\right]$, design pooling schemes for biological tests [ECG $\left.{ }^{+} 09, \mathrm{DWG}^{+} 13\right]$, pattern matching [CEPR07], combinatorial group testing [SAZ09, ESAZ09, $\left.\mathrm{KBG}^{+} 10\right]$, localizing sources in sensor networks [ZBSG05, ZPB06], as well as neuroscience [GS12]. Furthermore, not surprisingly, tracking heavy hitters in data streams, also known as frequent items, can be captured by the sparse recovery framework [Mut05a, CH09, KSZC03, Ind07]. In practice, streaming algorithms for detecting heavy hitters have been used to find popular destination addresses and heavy bandwidth users by AT\&T [CJK $\left.{ }^{+} 04\right]$ or answer "iceberg queries" in databases [FSGM ${ }^{+} 99$ ].

Sparse recovery attracts researchers from different communities, from both theoret-
ical and practical perspective. During the last ten years, hundreds of papers have been published by theoretical computer scientists, applied mathematicians and electrical engineers that specialize in compressed sensing. While numerous algorithms using space linear in the universe size $n$ are known, [Don06, CRT06b, IR08a, NT09a, BIR08a, BD09a, SV16] to name a few, our goal is to obtain algorithms that are sublinear, something that is crucial in many applications.

The desirable quantities we want to optimize may vary depending on the application. For example, in network management, $x_{i}$ could denote the total number of packets with destination $i$ passing through a network router. In such an application, storing the sketching matrix explicitly is typically not a tenable solution, since this would lead to an enormous space consumption; the number of possible IP addresses is $2^{32}$. Moreover, both the query and the update time should be very fast, in order to avoid congestion on the network. Incremental updates to $x$ come rapidly, and the changes to the sketch should also be implemented very fast; we note that in this case, even poly-logarithmic factors might be prohibitive. Interested readers can refer to [KSZC03, EV03] for more information about streaming algorithms for network management applications.
"The goal of that research is to obtain encoding and recovery schemes with good compression rate (i.e., short sketch lengths) as well as good algorithmic properties (i.e., low encoding, update and recovery times)." - Anna Gilbert and Piotr Indyk [GI10]

In this thesis, we consider the extensively studied problem of L2/L2 compressed sensing. The main contribution of our work is an improvement over [GLPS10] with faster decoding time and significantly smaller column sparsity, answering two open questions of the aforementioned work.

| Year | Reference | Measurements | Decoding Time | Encoding Time |
| :--- | :--- | :--- | :--- | :--- |
| 2006 | [Don06, CRT06b] | $k \log (n / k)$ | LP | $k \log (n / k)$ |
| 2006 | [CCF02, CM06] | $\epsilon^{-2} k \log n$ | $\epsilon^{-1} n \log n$ | $\log n$ |
| 2008 | [NT09a] | $k \log (n / k)$ | $n k \log (n / k)$ | $\log (n / k)$ |
| 2004 | [CM04] | $\epsilon^{-2} k \log ^{2} n$ | $\epsilon^{-1} k \log ^{c} n$ | $\log ^{2} n$ |
| 2006 | [CCF02, CM06] | $\epsilon^{-2} k \log ^{c} n$ | $\epsilon^{-1} k \log ^{2} n$ | $\log ^{c} n$ |
| 2010 | [GLPS10] | $\epsilon^{-1} k \log (n / k)$ | $\epsilon^{-1} k \log ^{c} n$ | $\log (n / k) \cdot \log ^{2} k$ |
| 2019 | [NS19] (thesis) | $\epsilon^{-1} k \log (n / k)$ | $\epsilon^{-1} k \log ^{2}(n / k)$ | $\log (n / k)$ |

Table 1.3: (A list of $\ell_{2} / \ell_{2}$-sparse recovery results). We ignore the " $O$ " for simplicity. LP denotes the time of solving Linear Programs [CLS19], and the state-of-the-art algorithm takes $n^{\omega}$ time where $\omega$ is the exponent of matrix multiplication. The results in [Don06, CRT06b, NT09a] do not explicitly state the $\ell_{2} / \ell_{2}$ guarantee, but their approach obtains it by an application of the Johnson-Lindenstrauss Lemma; they also cannot facilitate $\epsilon<1$, obtaining thus only a 2 -approximation. The $c$ in previous work is a sufficiently large constant, not explicitly stated, which is defined by probabilistically picking an error-correcting code of short length and iterating over all codewords. We estimate $c \geq 4$. We note that our runtime is (almost) achieved.

Previous work on sublinear-time compressed sensing employed an iterative procedure, recovering the heavy coordinates in phases. We completely depart from that framework, and give the first sublinear-time L2/L2 scheme which achieves the optimal number of measurements without iterating; this new approach is the key step to our progress. Towards that, we satisfy the L2/L2 guarantee by exploiting the heaviness of coordinates in a way that was not exploited in previous work. Via our techniques we obtain improved results for various sparse recovery tasks, and indicate possible further applications to problems in the field, to which the aforementioned iterative procedure creates significant obstructions.

High-dimensional Discrete Fourier Transform Probably the most important subtopic of compressed sensing/sparse recovery is the sparse Fourier transform, where one desires to
reconstruct a $k$-sparse vector from Fourier measurements. In other words, measurements are not allowed to be generic, but have to belong to the so-called Fourier ensemble. In Optics imaging [Goo05, Voe11] and Magnetic resonance imaging (MRI) [ASSN08], the physics [Rey89] of the underlying device restricts us to the Fourier ensemble, where the sparse Fourier problem becomes highly relevant. In fact, one of the initial motivations of Candes, Romberg and Tao came out due to the aforementioned applications. The number of samples plays a crucial role: they determine the amount of radiation a patient receives in CT scans, and taking fewer samples can reduce the amount of time the patient needs to stay in the machine. The framework has found its way in practical life-changing applications. Software includes the Compressed Sensing GRAB-VIBE, CS SPACE, CS SEMAC and CS TOF by Siemens [Sie], as well as Compressed Sense by Phillips [Phi]. Its incorporation in the MRI technology allows faster acquisition rates, depiction of dynamic processes or moving organs, as well as acceleration of MRI scanning up to a factor of 40 . On the webpage of SIEMENS Healthineers, for example, one can see the following, as well as numerous similar statements.

This allows bringing the advantages of Compressed Sensing GRASP-VIBE to daily clinical routine.

- Perform push-button, free-breathing liver dynamics.
- Overcome timing challenges in dynamic imaging and respiratory artifacts.
- Expand the patient population eligible for abdominal MRI.

The Fourier transform is in fact ubiquitous: image processing, audio processing, telecommunications, seismology, polynomial multiplication, SUBSET Sum and other text-

| Year | Reference | Samples | Time | Filter | RIP | Gua. |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 2005 | [GMS05] | $k \log ^{O(d)} n$ | $k \log ^{O(d)} n$ | Yes | No | $\ell_{2} / \ell_{2}$ |
| 2006 | $[$ CT06] | $k \log ^{6} n$ | poly $(n)$ | No | Yes | $\ell_{2} / \ell_{1}$ |
| 2008 | $[$ RV08] | $k \log ^{2} k \log (k \log n) \log n$ | $\widetilde{O}(n)$ | No | Yes | $\ell_{2} / \ell_{1}$ |
| 2012 | [HIKP12a] | $k \log ^{d} n \log (n / k)$ | $k \log ^{d} n \log (n / k)$ | Yes | No | $\ell_{2} / \ell_{2}$ |
| 2013 | [CGV13] | $k \log ^{3} k \log n$ | $\widetilde{O}(n)$ | No | Yes | $\ell_{2} / \ell_{1}$ |
| 2014 | [IK14] | $2^{d \log d} k \log n$ | $\widetilde{O}(n)$ | Yes | No | $\ell_{\infty} / \ell_{2}$ |
| 2014 | [Bou14] | $k \log k \log { }^{2} n$ | $\widetilde{O}(n)$ | No | Yes | $\ell_{2} / \ell_{1}$ |
| 2016 | [HR16] | $k \log ^{2} k \log n$ | $\widetilde{O}(n)$ | No | Yes | $\ell_{2} / \ell_{1}$ |
| 2016 | [Kap16] | $2^{d^{2} k \log n \log \log n}$ | $2^{d^{2}} k \log ^{d+3} n$ | Yes | No | $\ell_{2} / \ell_{2}$ |
| 2019 | [KVZ19] | $k^{3} \log ^{2} k \log { }^{2} n$ | $k^{3} \log ^{2} k \log ^{2} n$ | Yes | Yes | E. $k$ |
| 2019 | [NSW19a] <br> (thesis) | $k \log k \log n$ | $\widetilde{O}(n)$ | No | No | $\ell_{\infty} / \ell_{2}$ |

Table 1.4: $n=p^{d}$. E. $k$ denotes the "Exactly $k$-sparse" setting, which is a noiseless setting. "Gua" denotes "Guarantee". We ignore the $O$ for simplicity. The $\ell_{\infty} / \ell_{2}$ is the strongest possible guarantee, with $\ell_{2} / \ell_{2}$ coming second, $\ell_{2} / \ell_{1}$ third and exactly $k$-sparse being the less strong. We note that [CT06, RV08, CGV13, Bou14, HR16] obtain a uniform guarantee, i.e. with $1-1 / \operatorname{poly}(n)$ they allow reconstruction of all vectors; $\ell_{\infty} / \ell_{2}$ and $\ell_{2} / \ell_{2}$ are impossible in the uniform case [CDD09]. We also note that [RV08, CGV13, Bou14, HR16] give improved analysis of the Restricted Isometry property; the algorithm is suggested and analyzed (modulo the RIP property) in [BD08]. The work in [HIKP12a] does not explicitly state the extension to the $d$-dimensional case, but can easily be inferred from the arguments. [HIKP12a, IK14, Kap16, KVZ19] work when the universe size in each dimension are powers of 2 . We also assume that the signal-to-noise ratio is bounded by a polynomial of $n$, which is a standard assumption in the sparse Fourier transform literature [HIKP12a, IK14, Kap16, Kap17, LN19].
book algorithms are a few of the examples where the Fast Fourier Transform finds applications. The Fast Fourier Transform by Cooley and Tukey [CT65] runs in $O(n \log n)$ time, and has far-reaching applications in all of the aforementioned cases. It is thus expected that algorithms which exploit sparsity assumptions about the input, and can outperform FFT in applications are of high practical value. Generally, the two most important parameters one
would like to optimize are the sample complexity, i.e. the numbers needed to obtain from the time domain, as well as time needed to approximate the Fourier Transform.

Two different lines of research exist for the problem: the one focuses solely on sample complexity, while the other tries to achieve sublinear time while keeping the sample complexity as low as possible. The first line of research operates via the renowned Restricted Isometry Property (RIP), which proceeds by taking random samples and solving a linear/convex program, or an iterative thresholding procedure [CT06, DDTS06, TG07, BD08, DM08, RV08, BD09b, BD09a, NT09b, NV09, GK09, BD10, NV10, Fou11, Bou14, HR16]. The analysis of the algorithms is performed in the following way, in two steps. The first step ensures that, after sampling an appropriate number of points from the time domain, the inverse DFT matrix restricted on the rows indexed by those points acts as a near isometry on the space of $k$-sparse vectors. All of the state of the art results [CT06, RV08, Bou14, HR16] employ chaining arguments to make the analysis of this sampling procedure as tight as possible. The second part is how to exploit the aforementioned near-isometry property to find the best $k$-sparse approximation to the signal. There the approaches either follow an iterative procedure which gradually denoise the signal [BD08, NT09b, NV09], or perform $\ell_{1}$ minimization [CT06], a method that promotes sparsity of solutions.

The second line of research tries to implement arbitrary linear measurements via sampling Fourier coefficients [GL89, Man92, KM93, GGI ${ }^{+}$02, AGS03, GMS05, Iwe08, Iwe10, HIKP12a, HIKP12b, LWC13, Iwe13, PR14, IKP14, IK14, Kap16, Kap17, CI17, BZI17, MZIC17, LN19] and use sparse functions (in the time domain) which behave like bandpass filters in the frequency domain. The seminal work of Kapralov [Kap17] achieves $O(k \log n)$ samples and running time that is some log factors away from the sample complexity. This
would be the end of the story, apart from the fact that this algorithm does not scale well with dimension, since it has an exponential dependence on $d$. Indeed, in many applications, one is interested in higher dimensions, rather than the one-dimensional case. The main reason ${ }^{7}$ why this curse of dimensionality appears is due to the lack of dimension-independent ways to construct functions that approximate the $\ell_{\infty}$ ball and are sufficiently sparse in the time domain. A very nice work of Kapralov, Velingker and Zandieh [KVZ19] tries to remedy that by combining the standard execution of FFT with careful aliasing, but their algorithm works in a noiseless setting, and has a polynomial, rather than linear, dependence on $k$; the running time is polynomial in $k, \log n$ and the exponential dependence is avoided. It is an important and challenging question whether a robust and more efficient algorithm can be found.

We note that in many applications, such as MRI or computed tomography (CT), the main focus is the sample complexity; the algorithms that have found their way to industry are, to the best of our knowledge, not concerned with sublinear running time, but with the number of measurements, which determine the acquisition time, or in CT the radiation dose the patient receives.

In this work, we consider the extensively studied problem of computing a $k$-sparse approximation to the $d$-dimensional Fourier transform of a length $n$ signal. Our algorithm uses $O(k \log k \log n)$ samples, is dimension-free, operates for any universe size, and achieves the strongest $\ell_{\infty} / \ell_{2}$ guarantee, while running in time comparable to the Fast Fourier Transform. All previous algorithms proceed either via the Restricted Isometry Property or via

[^5]| Year | Refs. | Duration | Sample/Time | Approx. Ratio |
| :--- | :--- | :--- | :--- | :--- |
| 2012 | $\left[\mathrm{BCG}^{+} 12\right]$ | $>\frac{1}{\eta} \cdot k$ | sublinear | NO |
| 2015 | $[$ Moi15 $]$ | $>\frac{1}{\eta}$ | linear | $\operatorname{poly}(k)$ |
| 2015 | $[$ PS15] (thesis) | $>\frac{1}{\eta} \cdot \log (k)$ | sublinear | $O(\log k)$ |
| 2015 | $[$ PS15] (thesis) | $>\frac{1}{\eta} \cdot \log ^{2}(k)$ | sublinear | $O(1)$ |
| 2016 | $[$ CKPS16] (thesis) | $>0$ | sublinear | $O(1)$ |

Table 1.5: Reconstructing the signal in the continuous setting. Let $\eta$ denote the frequency gap. The first three results requires a frequency gap, and the last result doesn't require a frequency gap. We reconstruct an $x^{\prime}(t)$ such that $\frac{1}{T} \int_{0}^{T}\left|x^{\prime}(t)-x(t)\right|^{2} \mathrm{~d} t \leq \alpha \cdot \frac{1}{T} \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t$, where $\alpha$ is the Approx. Ratio.
filter functions. Our approach totally departs from the aforementioned techniques, and we believe it is a fresh look to the sparse Fourier transform problem.

Continuous Fourier Transform In many situations, much of the reason for using Fourier transforms is because the transformed signal is sparse -i.e., the energy is concentrated in a small set of $k$ locations. In such situations, one could hope for a dependency that depends nearly linearly on $k$ rather than $n$. Moreover, one may be able to find these frequencies while only sampling the signal for some period of time. This idea has led to a number of results on sparse Fourier transforms, including [GGI ${ }^{+}$02, GMS05, HIKP12a, IK14], that can achieve $O(k \log (n / k) \log n)$ running time and $O(k \log (n / k))$ sample complexity (although not quite both at the same time) in a robust setting.

These works apply to the discrete Fourier transform, but lots of signals including audio or radio originally come from a continuous domain. The standard way to convert a continuous Fourier transform into a discrete one is to apply a window function then subsample. Unfortunately, doing so "smears out" the frequencies, blowing up the sparsity. Thus,


Figure 1.2: Sampling Model. $x^{*}$ is a perfect signal. We are only allowed to samples from $x$ which is a signal that has noise.
one can hope for significant efficiency gains by directly solving the sparse Fourier transform problem in the continuous setting. This has led researchers to adapt techniques from the discrete setting to the continuous both in theory $\left[\mathrm{BCG}^{+} 12\right.$, TBSR13, CF14, DB13] and in practice $\left[\mathrm{SAH}^{+} 13\right]$. However, these results are not robust to noise: if the signal is sampled with a tiny amount of Gaussian noise or decays very slightly over time, no method has been known for computing a sparse Fourier transform in the continuous setting with sample complexity linear in $k$ and logarithmic in other factors. That is what we present in this paper.

Formally, a vector $x^{*}(t)$ has a $k$-sparse Fourier transform if it can be written as

$$
x^{*}(t)=\sum_{i=1}^{k} v_{i} e^{2 \pi \mathrm{i} \mathrm{f}_{i} t}
$$

for some tones $\left\{\left(v_{i}, f_{i}\right)\right\}$. We consider the problem where we can sample some signal

$$
x(t)=x^{*}(t)+g(t)
$$

at any $t$ we choose in some interval $[0, T]$, where $x^{*}(t)$ has a $k$-sparse Fourier transform and


Figure 1.3: Recovered signal. Our goal is to output signal $x^{\prime}$ such that $x^{\prime}$ is close to $x^{*}$.
$g(t)$ is arbitrary noise. As long as $g$ is "small enough," one would like to recover a good approximation to $x$ (or to $x^{*}$, or to $\left\{\left(v_{i}, f_{i}\right)\right\}$ ) using relatively few samples $t \in[0, T]$ and fast running time. Our algorithm achieves several results of this form, but a simple one is an $\ell_{2} / \ell_{2}$ guarantee: we reconstruct an $x^{\prime}(t)$ with $k$-sparse Fourier transform such that

$$
\frac{1}{T} \int_{0}^{T}\left|x^{\prime}(t)-x(t)\right|^{2} \mathrm{~d} t \lesssim \frac{1}{T} \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t
$$

using a number of samples that is $k$ times logarithmic factors ${ }^{8}$. To the best of our knowledge, this is the first algorithm achieving such a constant factor approximation with a sample complexity sublinear in $T$ and the signal-to-noise ratio.

Our algorithm also gives fairly precise estimates of the individual tones $\left(v_{i}, f_{i}\right)$ of the signal $x^{*}$. To demonstrate what factors are important, it is helpful to think about a concrete setting. Let us consider sound from a simplified model of a piano.

Thought experiment: piano tuning In a simplified model of a piano, we have keys corresponding to frequencies over some range $[-F, F]$. The noise $g(t)$ comes from ambient

[^6]

Figure 1.4: Thought Experiments
noise and the signals not being pure tones (because, for example, the notes might decay slowly over time). For concrete numbers, a modern piano has 88 keys spaced from about 27.5 Hz to $F=4200 \mathrm{~Hz}$. The space between keys ranges from a few Hz to a few hundred Hz , but most chords will have an $\eta=30 \mathrm{~Hz}$ or more gap between the frequencies being played. One typically would like to tune the keys to within about $\pm \nu=1 \mathrm{~Hz}$. And piano music typically has $k$ around 5 .

Now, suppose you would like to build a piano tuner that can listen to a chord and tell you what notes are played and how they are tuned. For such a system, how long must we wait for the tuner to identify the frequencies? How many samples must the tuner take? And how robust is it to the noise?

If you have a constant signal-to-noise ratio, you need to sample for a time $T$ of at
least order $1 / \nu=1$ second in order to get 1 Hz precision-frequencies within 1 Hz of each other will behave very similarly over small fractions of a second, which noise can make indistinguishable. You also need at least $\Omega\left(k \log \left(\frac{F}{k \nu}\right)\right) \approx 50$ samples, because the support of the signal contains that many bits of information and you only get a constant number per measurement (at constant SNR). At higher signal-to-noise ratios $\rho$, these results extend to $\Omega\left(\frac{1}{\nu \rho}\right)$ duration and $\Omega\left(k \log _{\rho}\left(\frac{F}{k \nu}\right)\right)$ samples. But as the signal-to-noise ratio gets very high, there is another constraint on the duration: for $T<\frac{1}{\eta} \approx 33$ milliseconds the different frequencies start becoming hard to distinguish, which causes the robustness to degrade exponentially in $k$ [Moi15] (though the lower bound there only directly applies to a somewhat restricted version of our setting).

This suggests the form of a result: with a duration $T>\frac{1}{\eta}$, one can hope to recover the frequencies to within $\frac{1}{\rho T}$ using $O\left(k \log _{\rho}\left(\frac{F T}{k}\right)\right)$ samples. We give an algorithm that is within logarithmic factors of this ideal: with a duration $T>\frac{O(\log (k / \delta))}{\eta}$, we recover the frequencies to within $O\left(\frac{1}{\rho T}\right)$ using $O\left(k \log _{\rho}(F T) \cdot \log (k / \delta) \cdot \log k\right)$ samples, where $\rho$ and $1 / \delta$ are (roughly speaking) the minimum and maximum signal-to-noise ratios that you can tolerate, respectively.

Instead of trying to tune the piano by recovering the frequencies precisely, one may simply wish to record the sound for future playback with relatively few samples. Our algorithm works for this as well: the combination $x^{\prime}(t)$ of our recovered frequencies satisfies

$$
\frac{1}{T} \int_{0}^{T}\left|x^{\prime}(t)-x(t)\right|^{2} \mathrm{~d} t \lesssim \frac{1}{T} \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t
$$

Continuous Fourier Transform without Frequency Gap In an interpolation problem, one can observe $x(t)=x^{*}(t)+g(t)$, where $x^{*}(t)$ is a structured signal and $g(t)$ denotes noise,
at points $t_{i}$ of one's choice in some interval $[0, T]$. The goal is to recover an estimate $\widetilde{x}$ of $x^{*}$ (or of $x$ ). Because we can sample over a particular interval, we would like our approximation to be good on that interval, so for any function $y(t)$ we define

$$
\|y\|_{T}^{2}=\frac{1}{T} \int_{0}^{T}|y(t)|^{2} \mathrm{~d} t
$$

to be the $\ell_{2}$ error on the sample interval. For some parameters $C$ and $\delta$, we would then like to get

$$
\begin{equation*}
\left\|\widetilde{x}-x^{*}\right\|_{T} \leq C\|g\|_{T}+\delta\left\|x^{*}\right\|_{T} \tag{1.7}
\end{equation*}
$$

while minimizing the number of samples and running time. Typically, we would like $C$ to be $O(1)$ and to have $\delta$ be very small (either zero, or exponentially small). Note that, if we do not care about changing $C$ by $O(1)$, then by the triangle inequality it doesn't matter whether we want to estimate $x^{*}$ or $x$ (i.e. we could replace the LHS of (1.7) by $\|\widetilde{x}-x\|_{T}$ ).

Of course, to solve an interpolation problem one also needs $x^{*}$ to have structure. One common form of structure is that $x^{*}$ have a sparse Fourier representation. We say that a function $x^{*}$ is $k$-Fourier-sparse if it can be expressed as a sum of $k$ complex exponentials:

$$
x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathbf{i} f_{j} t}
$$

for some $v_{j} \in \mathbb{C}$ and $f_{j} \in[-F, F]$, where $F$ is the "bandlimit". Given $F, T$, and $k$, how many samples must we take for the interpolation (1.7)?

If we ignore sparsity and just use the bandlimit, then Nyquist sampling and ShannonWhittaker interpolation uses $F T+1 / \delta$ samples to achieve (1.7). Alternatively, in the absence of noise, $x^{*}$ can be found from $O(k)$ samples by a variety of methods, including Prony's
method from 1795 or Reed-Solomon syndrome decoding [Mas69], but these methods are not robust to noise.

If the signal is periodic with period $T$-i.e., the frequencies are multiples of $1 / T$-then we can use sparse discrete Fourier transform methods, which take $O\left(k \log ^{c}(F T / \delta)\right)$ time and samples (e.g. [GGI ${ }^{+} 02$, HIKP12a, IKP14]). If the frequencies are not multiples of $1 / T$ (are "off the grid"), then the discrete approximation is only $k / \delta$ sparse, making the interpolation less efficient; and even this requires that the frequencies be well separated.

A variety of algorithms have been designed to recover off-grid frequencies directly, but they require the minimum gap among the frequencies to be above some threshold. With frequency gap at least $1 / T$, we can achieve a $k^{c}$ approximation factor using $O(F T)$ samples [Moi15], and with gap above $O\left(\log ^{2} k\right) / T$ we can get a constant approximation using $O\left(k \log ^{c}(F T / \delta)\right)$ samples and time [PS15].

Having a dependence on the frequency gap is natural. If two frequencies are very close together-significantly below $1 / T$-then the corresponding complex exponentials will be close on $[0, T]$, and hard to distinguish in the presence of noise. In fact, from a lower bound in [Moi15], below $1 / T$ frequency gap one cannot recover the frequencies in the presence of noise as small as $2^{-\Omega(k)}$. The lower bound proceeds by constructing two signals using significantly different frequencies that are exponentially close over $[0, T]$.

But if two signals are so close, do we need to distinguish them? Such a lower bound doesn't apply to the interpolation problem, it just says that you can't solve it by finding the frequencies. Our question becomes: can we benefit from Fourier sparsity in a regime where we can't recover the individual frequencies?


Figure 1.5: Moitra's lower bound [Moi15]. $\forall \epsilon>0$, if sample duration $T<\frac{(1-\epsilon)}{\eta}$, then $\exists$ two $k$-Fourier-sparse signals: $x(t)$ and $x^{\prime}(t)$, each has $\eta$ frequency gap. To tell them apart, need noise $\leq 2^{-\Omega(\epsilon k)}$. This example only implies that frequency recovery requires frequency gap, but what if we just want to interpolate signal?

We answer in the affirmative, giving an algorithm for the interpolation using $O(\operatorname{poly}(k \log (F T / \delta))$ samples.

This part is based on the following papers:

- Eric Price, Zhao Song

A Robust Sparse Fourier Transform in the Continuous Setting.
FOCS 2015 [PS15]

- Xue Chen, Eric Price, Daniel Kane, Zhao Song

Fourier-sparse Interpolation Without a Frequency Gap.
FOCS 2016 [CKPS16]

- Vasileios Nakos, Zhao Song, Zhengyu Wang

The Power of (careful) Iterative Loop Analysis in Compressed Sensing.
Manuscript 2018 [NSW18]

- Vasileios Nakos, Zhao Song

Stronger L2/L2 Compressed Sensing; Without Iterating.
STOC 2019 [NS19]

- Vasileios Nakos, Zhao Song, Zhengyu Wang
(Nearly) Sample-Optimal Sparse Fourier Transform in Any Dimension; RIPless and Filterless.

FOCS 2019 [NSW19a]

### 1.5 Factorization

Low rank approximation is arguably one of the most well-studied problems in randomized numerical linear algebra, with diverse applications to clustering [ $\mathrm{DFK}^{+} 04$, FSS 13 , LBKW14, $\mathrm{CEM}^{+}$15], data mining [ $\mathrm{AFK}^{+} 01$ ], distance matrix completion [Cha12], information retrieval [PRTV00], learning mixtures of distributions [AM05, KSV08], recommendation systems [DKR02], and web search [AFKM01, Kle99]. In practice one often has a low rank matrix which has been corrupted with noise of bounded norm, and low rank approximation allows one to approximately recover the original matrix. Low rank approximation may also help explain a dataset, revealing low dimensional structure in high dimensional data. Given a low rank approximation, one can store a matrix and compute a matrix-vector product much more efficiently by storing the corresponding factorization. It can also be used as a preprocessing step in applications, that is, by first projecting data onto a lower-dimensional
subspace one preserves important properties of the input, but can now run subsequent algorithms in the lower-dimensional space. For example, it has been proposed to reduce the data dimension in Non-Negative Matrix Factorization [LS00] (more on this below), and Latent Dirichlet Allocation (LDA) [BNJ01].

The basic low rank approximation problem is: given an $n \times n$ matrix $A$, find a matrix $\widehat{A}$ of rank at most $k$ for which $\|A-\widehat{A}\|_{F}$ is minimized, where for a matrix $B$, $\|B\|_{F}=\left(\sum_{i, j} B_{i, j}^{2}\right)^{1 / 2}$ is its Frobenius norm. This formulation intuitively corresponds to the matrix $\widehat{A}$ capturing as much of the variance of $A$ as possible. It is well-known that the optimal solution is given by $A_{k}$, which if $U \Sigma V^{\top}$ is the singular value decomposition (SVD) of $A$, where $U$ and $V$ are orthogonal matrices and $\Sigma$ is a non-negative diagonal matrix with $\Sigma_{1,1} \geq \Sigma_{2,2} \geq \cdots \geq \Sigma_{n, n}$, then $A_{k}=U \Sigma_{k} V^{\top}$, where $\Sigma_{k}$ agrees with $\Sigma$ on its first $k$ diagonal entries and is 0 otherwise. Although the SVD is computable in polynomial time, it is often acceptable to output a matrix $\widehat{A}$ for which $\|A-\widehat{A}\|_{F} \leq(1+\epsilon)\left\|A-A_{k}\right\|_{F}$ with high probability. In the latter case, much more efficient algorithms are known, and it is possible to compute such an $\widehat{A}$ in $n n z(A)+n \cdot \operatorname{poly}(k / \epsilon)$ time, where $n n z(A)$ denotes the number of non-zero entries of $A$ [CW13, MM13, NN13a]. We note that for typical applications $k$ and $1 / \epsilon$ are assumed to be much smaller than $n$, e.g., in [Har14] they are treated as absolute constants.

Weighted Despite the large body of work on low rank approximation, the weighted case is not well understood. In this case one is given an $n \times n$ matrix $A$ and an $n \times n$ matrix $W$
with $W_{i, j} \geq 0$, and one seeks to solve:

$$
\min _{\text {rank-k matrices } \widehat{A}}\|W \circ(A-\widehat{A})\|_{F}^{2}=\min _{\text {rank-k matrices }} \sum_{i, j} W_{i, j}^{2}\left(A_{i, j}-\widehat{A}_{i, j}\right)^{2} .
$$

The classical low rank approximation is a special case in which $W_{i, j}=1$ for all $i$ and $j$. However, in general there may not be a good reason to weight all elements of the approximation error $A-\widehat{A}$ equally, especially if one is given prior knowledge about the distribution of the errors. For example, suppose the columns of $A$ each come from a low-dimensional subspace but one of the columns is then shifted by a fixed large vector so that its mean is different. One may first want to recenter the data by subtracting off the mean from each of the columns. While this is possible without weighted low rank approximation, suppose instead that each of the columns of $A$ comes from a perturbation of columns in a low dimensional subspace but one of the columns has a much larger variance. Then if all weights were equal, it would be enough for $\widehat{A}$ to fit this one single large variance column, which fails to capture the entire low-dimensional subspace. One way of fixing this is to reweight each entry of $A$ by the inverse of its variance. This is a common technique used in gene expression analysis, where the error model for microarray measurements provides entry-specific noise estimates, or when entries of $A$ represent aggregates of many samples such as in word co-occurrence matrices and non-uniform weights are needed to appropriately capture any differences in the sample sizes; see [SJ03] for a discussion, and also the Wikipedia entry on weighted low rank approximation for a brief introduction.

While the extension of low rank approximation to the weighted case goes back to work of Young in 1940 [You40], its complexity is not well-understood, partly because the weighted case does not admit a solution via the SVD and may have many local minima [SJ03]. Early
work by Shpak [Shp90] looked at gradient-based approaches while Lu et al. [LPW97, LA03] looked at alternating minimization methods. These were significantly sped up in practice by the work of Srebro and Jaakkola [SJ03], with success in various applications such as color image restoration [MES08], though there are no provable time bounds and in the worst case the running times could be exponential or worse. In fact, weighted low rank approximation is known to be NP-hard to approximate up to a $(1 \pm 1 / \operatorname{poly}(n))$ factor [GG11]. We note that this also follows from the fact that matrix completion, arguably one of the most important special cases of weighted low rank approximation in which case all weights are 0 or 1 , which we discuss more below, is also known to be NP-hard [Pee96, HMRW14]. Typically, though, assumptions such as incoherence and randomly sampled entries allow one to circumvent this hardness [CR09, LLR16]. There is some debate as to whether these assumptions are valid, for instance in [SW15] an argument is made why randomly missing entries may not hold for real-world datasets.

Many natural questions are left open from previous work. In particular one question as we see it is the following:

For which weight matrices $W$ is the problem tractable? More generally, is it possible to identify a natural parameter of $W$ and to obtain parameterized complexity bounds in terms of that parameter?

Entry-wise L1 norm Two well-studied problems in numerical linear algebra are regression and low rank approximation. In regression, one is given an $n \times d$ matrix $A$, and an $n \times 1$ vector $b$, and one seeks an $x \in \mathbb{R}^{d}$ which minimizes $\|A x-b\|$ under some norm. For example, for least squares regression one minimizes $\|A x-b\|_{2}$. In low rank approximation, one is
given an $n \times d$ matrix $A$, and one seeks a rank- $k$ matrix $\widehat{A}$ which minimizes $\|A-\widehat{A}\|$ under some norm. For example, in Frobenius norm low rank approximation, one minimizes $\|A-\widehat{A}\|_{F}=\left(\sum_{i, j}\left(A_{i, j}-\widehat{A}_{i, j}\right)^{2}\right)^{1 / 2}$. Algorithms for regression are often used as subroutines for low rank approximation. Indeed, one of the main insights of [DMM06c, DMM06b, Sar06, DMM08, CW09] was to use results for generalized least squares regression for Frobenius norm low rank approximation. Algorithms for $\ell_{1}$-regression, in which one minimizes $\|A x-b\|_{1}=$ $\sum_{i}\left|(A x)_{i}-b_{i}\right|$, were also used [BD13, SW11] to fit a set of points to a hyperplane, which is a special case of entrywise $\ell_{1}$-low rank approximation, the more general problem being to find a rank- $k$ matrix $\widehat{A}$ minimizing $\sum_{i, j}\left|A_{i, j}-\widehat{A}_{i, j}\right|$.

Randomization and approximation were introduced to significantly speed up algorithms for these problems, resulting in algorithms achieving relative error approximation with high probability. Such algorithms are based on sketching and sampling techniques; we refer to [Woo14b] for a survey. For least squares regression, a sequence of work [Sar06, CW13, MM13, NN13a, LMP13, BDN15, Coh16a] shows how to achieve algorithms running in nnz $(A)+\operatorname{poly}(d)$ time. For Frobenius norm low rank approximation, using the advances for regression this resulted in $n n z(A)+(n+d) \operatorname{poly}(k)$ time algorithms. For $\ell_{1}$-regression, sketching and sampling-based methods [Cla05, SW11, CDMI ${ }^{+}$13, CW13, MM13, LMP13, WZ13, CW15b, CP15] led to an $n n z(A)+\operatorname{poly}(d)$ time algorithm.

Just like Frobenius norm low rank approximation is the analogue of least squares regression, entrywise $\ell_{1}$-low rank approximation is the analogue of $\ell_{1}$-regression. Despite this analogy, no non-trivial upper bounds with provable guarantees are known for $\ell_{1}$-low rank approximation. Unlike Frobenius norm low rank approximation, which can be solved exactly using the singular value decomposition, no such algorithm or closed-form solution is known
for $\ell_{1}$-low rank approximation. Moreover, the problem was recently shown to be NP-hard [GV15]. A major open question is whether there exist approximation algorithms, sketchingbased or otherwise, for $\ell_{1}$-low rank approximation. Indeed, the question of obtaining betters algorithms was posed in section 6 of [GV15], in [Exc13], and as the second part of open question 2 in [Woo14b], among other places. The earlier question of NP-hardness was posed in Section 1.4 of [KV09], for which the question of obtaining approximation algorithms is a natural followup. The goal of our work is to answer this question.

We now formally define the $\ell_{1}$-low rank approximation problem: we are given an $n \times d$ matrix $A$ and approximation factor $\alpha \geq 1$, and we would like, with large constant probability, to output a rank- $k$ matrix $\widehat{A}$ for which

$$
\|A-\widehat{A}\|_{1} \leq \alpha \cdot \min _{\text {rank-k matrices } A^{\prime}}\left\|A-A^{\prime}\right\|_{1},
$$

where for an $n \times d$ matrix $C$, we let $\|C\|_{1}=\sum_{i=1}^{n} \sum_{j=1}^{d}\left|C_{i, j}\right|$. This notion of low rank approximation has been proposed as a more robust alternative to Frobenius norm low rank approximation $\left[K K 03, ~ K K 05\right.$, KLC $^{+} 15$, Kwa08, ZLS ${ }^{+} 12$, BJ12, BD13, BDB13, MXZZ13, MKP13, MKP14, MKCP16, PK16], and is sometimes referred to as $\ell_{1}$-matrix factorization or robust PCA. $\ell_{1}$-low rank approximation gives improved results over Frobenius norm low rank approximation since outliers are less exaggerated, as one does not square their contribution in the objective. The outlier values are often erroneous values that are far away from the nominal data, appear only a few times in the data matrix, and would not appear again under normal system operation. These works also argue $\ell_{1}$-low rank approximation can better handle missing data, is appropriate in noise models for which the noise is not Gaussian, e.g., it produces the maximum likelihood estimator for Laplacian noise [Gao08, $\mathrm{KAC}^{+}$08, VT01], and can be used in image processing to prevent image occlusion [YZD12].

To see that $\ell_{1}$-low rank approximation and Frobenius norm low rank approximation can give very different results, consider the $n \times n$ matrix $A=\left[\begin{array}{cc}n & 0 \\ 0 & B\end{array}\right]$, where $B$ is any ( $n-$ 1) $\times(n-1)$ matrix with $\|B\|_{F}<n$. The best rank-1 approximation with Frobenius norm error is given by $\widehat{A}=n \cdot e_{1} e_{1}^{\top}$, where $e_{1}$ is the first standard unit vector. Here $\widehat{A}$ ignores all but the first row and column of $A$, which may be undesirable in the case that this row and column represent an outlier. Note $\|A-\widehat{A}\|_{1}=\|B\|_{1}$. If, for example, $B$ is the all 1 s matrix, then $\widehat{A}=[0,0 ; 0, B]$ is a rank- 1 approximation for which $\|A-\widehat{A}\|_{1}=n$, and therefore this solution is a much better solution to the $\ell_{1}$-low rank approximation problem than $n \cdot e_{1} e_{1}^{\top}$, for which $\left\|A-n \cdot e_{1} e_{1}^{\top}\right\|_{1}=(n-1)^{2}$.

Despite the advantages of $\ell_{1}$-low rank approximation, its main disadvantage is its computationally intractability. It is not rotationally invariant and most tools for Frobenius low rank approximation do not apply. To the best of our knowledge, all previous works only provide heuristics. Using that for an $n \times d$ matrix $C,\|C\|_{F} \leq\|C\|_{1} \leq \sqrt{n d}\|C\|_{F}$, a Frobenius norm low rank approximation gives a $\sqrt{n d}$ approximation for $\ell_{1}$-low rank approximation. A bit better is to use algorithms for low rank approximation with respect to the sum of distances, i.e., to find a rank- $k$ matrix $\widehat{A}$ minimizing $\|A-\widehat{A}\|_{1,2}$, where for an $n \times d$ matrix $C,\|C\|_{1,2}=\sum_{i=1}^{n}\left\|C_{i}\right\|_{2}$, where $C_{i}$ is the $i$-th row of $C$. A sequence of work [DV07, FMSW10, FL11, SV12a, CW15a] shows how to obtain an $O(1)$-approximation to this problem in $n n z(A)+(n+d) \operatorname{poly}(k)+\exp (k)$ time, and using that $\|C\|_{1,2} \leq\|C\|_{1} \leq \sqrt{d}\|C\|_{1,2}$ results in an $O(\sqrt{d})$-approximation.

Zero-one Law To understand the role of the Frobenius norm in the algorithms above, we recall a standard motivation for this error measure. Suppose one has $n$ data points in
a $k$-dimensional subspace of $\mathbb{R}^{d}$, where $k \ll d$. We can write these points as the rows of an $n \times d$ matrix $A^{*}$ which has rank $k$. The matrix $A^{*}$ is often called the ground truth matrix. In a number of settings, due to measurement noise or other kinds of noise, we only observe the matrix $A=A^{*}+\Delta$, where each entry of the noise matrix $\Delta \in \mathbb{R}^{n \times n}$ is an i.i.d. random variable from a certain mean-zero noise distribution $\mathcal{D}$. One method for approximately recovering $A^{*}$ from $A$ is maximum likelihood estimation. Here one tries to find a matrix $B$ maximizing the $\log$-likelihood: $\max _{\text {rank-k }} \sum_{i, j} \log p\left(A_{i, j}-B_{i, j}\right)$, where $p(\cdot)$ is the probability density function of the underlying noise distribution $\mathcal{D}$. For example, when the noise distribution is Gaussian with mean zero and variance $\sigma^{2}$, denoted by $N\left(0, \sigma^{2}\right)$, then the optimization problem is $\max _{\text {rank-k }} B \sum_{i, j}\left(\log \left(1 / \sqrt{2 \pi \sigma^{2}}\right)-\left(A_{i, j}-B_{i, j}\right)^{2} /\left(2 \sigma^{2}\right)\right)$, which is equivalent to solving the Frobenius norm loss low rank approximation problem defined above.

The Frobenius norm loss, while having nice statistical properties for Gaussian noise, is well-known to be sensitive to outliers. Applying the same maximum likelihood framework above to other kinds of noise distributions results in minimizing other kinds of loss functions. In general, if the density function of the underlying noise $\mathcal{D}$ is $p(z)=c \cdot e^{-g(z)}$, where $c$ is a normalization constant, then the maximum likelihood estimation problem for this noise distribution becomes the following generalized entry-wise loss low rank approximation problem: $\min _{\text {rank-k } B} \sum_{i, j} g\left(A_{i, j}-B_{i, j}\right)=\min _{\text {rank-k } B}\|A-B\|_{g}$, which is a central topic of recent work on generalized low-rank models $\left[\mathrm{UHZ}^{+} 16\right]$. For example, when the noise is Laplacian, the entrywise $\ell_{1}$ loss is the maximum likelihood estimation, which is also robust to sparse outliers. A natural setting is when the noise is a mixture of small Gaussian noise and sparse outliers; this noise distribution is referred to as the Huber density. In this case
the Huber loss function gives the maximum likelihood estimate $\left[\mathrm{UHZ}^{+} 16\right]$, where the Huber function [Hub64] is defined to be: $g(x)=x^{2} /(2 \tau)$ if $|x|<\tau / 2$, and $g(x)=|x|-\tau / 2$ if $|x| \geq \tau$. Another nice property of the Huber error measure is that it is differentiable everywhere, unlike the $\ell_{1}$-norm, yet still enjoys the robustness properties as one moves away from the origin, making it less sensitive to outliers than the $\ell_{2}$-norm. There are many other kinds of loss functions, known as $M$-estimators [Zha97], which are widely used as loss functions in robust statistics [HRRS11].

Although several specific cases have been studied, such as entry-wise $\ell_{p}$ loss [CLMW11, SWZ17, $\mathrm{CGK}^{+}$17a, BKW17, $\mathrm{BBB}^{+}$19a], weighted entry-wise $\ell_{2}$ loss [RSW16], and cascaded $\ell_{p}\left(\ell_{2}\right)$ loss [DVTV09, CW15a], the landscape of general entry-wise loss functions remains elusive. There are no results known for any loss function which is not scale-invariant, much less any kind of characterization of which loss functions admit efficient algorithms. This is despite the importance of these loss functions; we refer the reader to [UHZ $\left.{ }^{+} 16\right]$ for a survey of generalized low rank models. This motivates the main question in our work:

Question 1.5.1 (General Loss Functions). For a given approximation factor $\alpha>1$, which functions $g$ allow for efficient low-rank approximation algorithms? Formally, given an $n \times d$ matrix $A$, can we find a rank-k matrix $B$ for which $\|A-B\|_{g} \leq \alpha \min _{\text {rank-k } B^{\prime}}\left\|A-B^{\prime}\right\|_{g}$, where for a matrix $C,\|C\|_{g}=\sum_{i \in[n], j \in[d]} g\left(C_{i, j}\right)$ ? What if we also allow $B$ to have rank $\operatorname{poly}(k \log n)$ ?

For Question 1.5.1, one has $g(x)=|x|^{p}$ for $p$-norms, and note the Huber loss function also fits into this framework. Allowing $B$ to have slightly larger rank than $k$, namely, $\operatorname{poly}(k \log n)$, is often sufficient for applications as it still allows for the space savings and
computational gains outlined above. These are referred to as bicriteria approximations and are the focus of our work.

Tensor Tensors are often more useful than matrices for capturing higher order relations in data. Computing low rank factorizations of approximations of tensors is the primary task of interest in a number of applications, such as in psychology[Kro83], chemometrics [Paa00, SBG04], neuroscience [AAB ${ }^{+} 07, \mathrm{~KB} 09, \mathrm{CLK}^{+}$15], computational biology [CV15, SC15], natural language processing [CYYM14, LZBJ14, LZMB15, BNR ${ }^{+}$15], computer vision [VT02, WA03, SH05, HPS05, HD08, AFdLGTL09, PLY10, LFC ${ }^{+}$16, CLZ17], computer graphics [VT04, $\mathrm{WWS}^{+} 05$, Vas09], security [AÇKY05, ACY06, KB06], cryptography [FS99, Sch12, KYFD15, SHW ${ }^{+}$16] data mining [KS08, RST10, KABO10, Mør11], machine learning applications such as learning hidden Markov models, reinforcement learning, community detection, multi-armed bandit, ranking models, neural network, Gaussian mixture models and Latent Dirichlet allocation [MR05, AFH ${ }^{+}$12, HK13, ALB13, ABSV14, AGH ${ }^{+} 14$, AGHK14, BCV14, JO14a, GHK15, PBLJ15, JSA15, ALA16, AGMR16, ZSJ+ 17], programming languages [RTP16], signal processing [Wes94, DLDM98, Com09, CMDL+15], and other applications [YCS11, LMWY13, OS14, ZCZJ14, STLS14, YCS16, RNSS16].

Despite the success for matrices, the situation for order- $q$ tensors for $q>2$ is much less understood. There are a number of works based on alternating minimization [CC70, Har70, FMPS13, FT15, ZG01, BS15] gradient descent or Newton methods [ES09, ZG01], methods based on the Higher-order SVD (HOSVD) [LMV00a] which provably incur $\Omega(\sqrt{n})$ inapproximability for Frobenius norm error [LMV00b], the power method or orthogonal iteration method [LMV00b], additive error guarantees in terms of the flattened (unfolded)
tensor rather than the original tensor [MMD08], tensor trains [Ose11], the tree Tucker decomposition [OT09], or methods specialized to orthogonal tensors [KM11, $\mathrm{AGH}^{+} 14$, MHG15, WTSA15, WA16, SWZ16]. There are also a number of works on the problem of tensor completion, that is, recovering a low rank tensor from missing entries [WM01, AKDM10, TSHK11, LMWY13, MHWG14, JO14b, BM16]. There is also another line of work using the sum of squares (SOS) technique to study tensor problems [BKS15a, GM15, HSS15, HSSS16, MSS16, PS17, SS17], other recent work on tensor PCA [All12b, All12a, RM14, JMZ15, ADGM16, ZX17], and work applying smoothed analysis to tensor decomposition [BCMV14]. Several previous works also consider more robust norms than the Frobenius norm for tensors, e.g., the $R_{1}$ norm ( $\ell_{1}-\ell_{2}-\ell_{2}$ norm in our work) [HD08], $\ell_{1}$-PCA [PLY10], entry-wise $\ell_{1}$ regularization [GGH14], M-estimator loss [YFS16], weighted approximation [Paa97, TK11, LRHG13], tensor-CUR [OST08, MMD08, CC10, FMMN11, FT15], or robust tensor PCA [GQ14, LFC ${ }^{+}$16, CLZ17].

Some of the above works, such as ones based on the tensor power method or alternating minimization, require incoherence or orthogonality assumptions. Others, such as those based on the simultaneous SVD, require an assumption on the minimum singular value. See the monograph of Moitra [Moi14a] for further discussion. Unlike the situation for matrices, there is no work for tensors that is able to achieve the following natural relative error guarantee: given a $q$-th order tensor $A \in \mathbb{R}^{n^{\otimes q}}$ and an arbitrary accuracy parameter $\epsilon>0$, output a rank- $k$ tensor $B$ for which

$$
\begin{equation*}
\|A-B\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT} \tag{1.8}
\end{equation*}
$$

where $\mathrm{OPT}=\inf _{\text {rank-k } B^{\prime}}\left\|A-B^{\prime}\right\|_{F}^{2}$, and where recall the rank of a tensor $B$ is the minimal integer $k$ for which $B$ can be expressed as $\sum_{i=1}^{k} U_{1}^{i} \otimes U_{2}^{i} \otimes \cdots \otimes U_{q}^{i}$. General speaking, a $q$-th
order tensor can have rank at most $n^{q-1}$. A third order tensor, for example, has rank which is an integer in $\left\{0,1,2, \ldots, n^{2}\right\}$.

For notational simplicity, we will start by assuming third order tensors with all dimensions of equal size, but we extend all of our main theorems below to tensors of any constant order $q>3$ and dimensions of different sizes.

The first caveat regarding (1.8) for tensors is that an optimal rank- $k$ solution may not even exist! This is a well-known problem for tensors (see, e.g., [KHL89, Paa00, KDS08, Ste06, Ste08] and more details in section 4 of [DSL08]), for which for any rank- $k$ tensor $B$, there always exists another rank- $k$ tensor $B^{\prime}$ for which $\left\|A-B^{\prime}\right\|_{F}^{2}<\|A-B\|_{F}^{2}$. If OPT $=0$, then in this case for any rank- $k$ tensor $B$, necessarily $\|A-B\|_{F}^{2}>0$, and so (1.8) cannot be satisfied. This fact was known to algebraic geometers as early as the 19th century, which they refer to as the fact that the locus of $r$-th secant planes to a Segre variety may not define a (closed) algebraic variety [DSL08, Lan12]. It is also known as the phenomenon underlying the concept of border $\operatorname{rank}^{9}$ [Bin80, Bin86, BCS97, Knu98, Lan06]. In this case it is natural to allow the algorithm to output an arbitrarily small $\gamma>0$ amount of additive error. Note that unlike several additive error algorithms for matrices, the additive error here can in fact be an arbitrarily small positive function of $n$. If, however, OPT $>0$, then for any $\epsilon>0$, there exists a rank- $k$ tensor $B$ for which $\|A-B\|_{F}^{2} \leq(1+\epsilon)$ OPT, and in this case we should still require the algorithm to output a relative-error solution. If an optimal rank- $k$ solution $B$ exists, then as for matrices, it is natural to require the algorithm to output a relative-error solution.

[^7]Besides the above definitional issue, a central reason that (1.8) has not been achieved is that computing the rank of a third order tensor is well-known to be NP-hard [Hås90, HL13]. Thus, if one had such a polynomial time procedure for solving the problem above, one could determine the rank of $A$ by running the procedure on each $k \in\left\{0,1,2, \ldots, n^{2}\right\}$, and check for the first value of $k$ for which $\|A-B\|_{F}^{2}=0$, thus determining the rank of $A$. However, it is unclear if approximating the tensor rank is hard. This question will also be answered in this work.

This part is based on the following papers:

- Ilya Razenshteyn, Zhao Song, David Woodruff

Weighted low rank approximations with provable guarantees.
STOC 2016 [RSW16]

- Zhao Song, David Woodruff, Peilin Zhong

Low Rank Approximation with Entrywise $\ell_{1}$-Norm Error.
STOC 2017 [SWZ17]

- Zhao Song, David Woodruff, Peilin Zhong

Towards a Zero-One Law for Entrywise Low Rank Approximation.
Manuscript 2018 [SWZ18]

- Zhao Song, David Woodruff, Peilin Zhong

Relative Error Tensor Low Rank Approximation.
SODA 2019 [SWZ19b]

- Zhao Song, David Woodruff, Huan Zhang

Sublinear Time Orthogonal Tensor Decomposition.
NeurIPS 2016 [SWZ16]

### 1.6 More Algorithms

Regression Sketching has emerged as a powerful technique for speeding up problems in numerical linear algebra, such as regression. In the overconstrained regression problem, one is given an $n \times d$ matrix $A$, with $n \gg d$, as well as an $n \times 1$ vector $b$, and one wants to find a vector $\widehat{x}$ so as to minimize the residual error $\|A x-b\|_{2}$. Using the sketch and solve paradigm, one first computes $S \cdot A$ and $S \cdot b$ for a randomly chosen matrix $S$, then outputs $x^{\prime}=(S A)^{\dagger} S b$ so as to minimize $\left\|S A x^{\prime}-S b\right\|_{2}$.

The sketch-and-solve paradigm gives a bound on $\left\|x^{\prime}-x^{*}\right\|_{2}$ when $A$ is well-conditioned. One of result cares about $\left\|x^{\prime}-x^{*}\right\|_{\infty}$, we call it $\ell_{\infty}$ regression problem. We also studied regression problem under general norms (i.e. symmetric norm regression min $\|A x-b\|_{\text {symmetric norm }}$ ) and general structure (i.e. tensor regression $\left.\min _{x}\left\|\left(A_{1} \otimes A_{2}\right) x-b\right\|\right)$.

Clustering We consider the $k$-means clustering problem in the dynamic streaming setting, where points from a discrete Euclidean space $\{1,2, \ldots, \Delta\}^{d}$ can be dynamically inserted to or deleted from the dataset. For this problem, we provide a one-pass coreset construction algorithm using space $\widetilde{O}(k \cdot \operatorname{poly}(d, \log \Delta))$, where $k$ is the target number of centers. To our knowledge, this is the first dynamic geometric data stream algorithm for $k$-means using space polynomial in dimension and nearly optimal (linear) in $k$.

LCS Binary String Given a pair of $n$-character strings, the problems of computing their Longest Common Subsequence and Edit Distance have been extensively studied for decades. For exact algorithms, LCS and Edit Distance (with character insertions and deletions) are equivalent; the state of the art running time is (almost) quadratic in $n$, and this is tight under plausible fine-grained complexity assumptions. But for approximation algorithms the picture is different: there is a long line of works with improved approximation factors for Edit Distance, but for LCS (with binary strings) only a trivial 1/2-approximation was known. In this work we give a reduction from approximate LCS to approximate Edit Distance, yielding the first efficient $(1 / 2+\epsilon)$-approximation algorithm for LCS for some constant $\epsilon>0$.

LCS Longest common subsequence (LCS) is a classic and central problem in combinatorial optimization. While LCS admits a quadratic time solution, recent evidence suggests that solving the problem may be impossible in truly subquadratic time. A special case of LCS wherein each character appears at most once in every string is equivalent to the longest increasing subsequence problem (LIS) which can be solved in quasilinear time. In this work, we present novel algorithms for approximating LCS in truly subquadratic time and LIS in truly sublinear time. Our approximation factors depend on the ratio of the optimal solution size over the input size. We denote this ratio by $\lambda$ and obtain the following results for LCS and LIS without any prior knowledge of $\lambda$.

- A truly subquadratic time algorithm for LCS with approximation factor $O\left(\lambda^{3}\right)$.
- A truly sublinear time algorithm for LIS with approximation factor $O\left(\lambda^{3}\right)$.

Triangle inequality was recently used by Boroujeni et al. $\left[\mathrm{BEG}^{+} 18\right]$ and Chakraborty
et al. $\left[\mathrm{CDG}^{+} 18\right]$ to present new approximation algorithms for edit distance. Our techniques for LCS extend the notion of triangle inequality to non-metric settings.

Map-Reduce Many modern parallel systems, such as MapReduce, Hadoop and Spark, can be modeled well by the MPC model. The MPC model captures well coarse-grained computation on large data - data is distributed to processors, each of which has a sublinear (in the input data) amount of memory and we alternate between rounds of computation and rounds of communication, where each machine can communicate an amount of data as large as the size of its memory. This model is stronger than the classical PRAM model, and it is an intriguing question to design algorithms whose running time is smaller than in the PRAM model.

One fundamental graph problem is connectivity. On an undirected graph with $n$ nodes and $m$ edges, $O(\log n)$ round connectivity algorithms have been known for over 35 years. However, no algorithms with better complexity bounds were known. In this work, we give fully scalable, faster algorithms for the connectivity problem, by parameterizing the time complexity as a function of the diameter of the graph. Our main result is a $O\left(\log D \log \log _{m / n} n\right)$ time connectivity algorithm for diameter- $D$ graphs, using $\Theta(m)$ total memory. If our algorithm can use more memory, it can terminate in fewer rounds, and there is no lower bound on the memory per processor.

We extend our results to related graph problems such as spanning forest, finding a DFS sequence, exact/approximate minimum spanning forest, and bottleneck spanning forest. We also show that achieving similar bounds for reachability in directed graphs would imply faster boolean matrix multiplication algorithms.

We introduce several new algorithmic ideas. We describe a general technique called double exponential speed problem size reduction which roughly means that if we can use total memory $N$ to reduce a problem from size $n$ to $n / k$, for $k=(N / n)^{\Theta(1)}$ in one phase, then we can solve the problem in $O\left(\log \log _{N / n} n\right)$ phases. In order to achieve this fast reduction for graph connectivity, we use a multistep algorithm. One key step is a carefully constructed truncated broadcasting scheme where each node broadcasts neighbor sets to its neighbors in a way that limits the size of the resulting neighbor sets. Another key step is random leader contraction, where we choose a smaller set of leaders than many previous works do.

## Part I

## Optimization

## Chapter 2

## Linear Programs

This chapter shows how to solve linear programs of the form $\min _{A x=b, x \geq 0} c^{\top} x$ with $n$ variables in time

$$
O^{*}\left(\left(n^{\omega}+n^{2.5-\alpha / 2}+n^{2+1 / 6}\right) \log (n / \delta)\right)
$$

where $\omega$ is the exponent of matrix multiplication, $\alpha$ is the dual exponent of matrix multiplication, and $\delta$ is the relative accuracy. For the current value of $\omega \sim 2.37$ and $\alpha \sim 0.31$, our algorithm takes $O^{*}\left(n^{\omega} \log (n / \delta)\right)$ time. When $\omega=2$, our algorithm takes $O^{*}\left(n^{2+1 / 6} \log (n / \delta)\right)$ time.

Our algorithm utilizes several new concepts that we believe may be of independent interest:

- We define a stochastic central path method.
- We show how to maintain a projection matrix $\sqrt{W} A^{\top}\left(A W A^{\top}\right)^{-1} A \sqrt{W}$ in sub-quadratic time under $\ell_{2}$ multiplicative changes in the diagonal matrix $W$.


### 2.1 Introduction

Linear programming is one of the key problems in computer science. In both theory and practice, many problems can be reformulated as linear programs to take advantage of fast algorithms. For an arbitrary linear program $\min _{A x=b, x \geq 0} c^{\top} x$ with $n$ variables and $d$ constraints ${ }^{1}$, the fastest algorithm takes $O^{*}\left(\sqrt{d} \cdot \mathrm{nnz}(A)+d^{2.5}\right)^{2}$ where $\mathrm{nnz}(A)$ is the number of non-zeros in $A$ [LS14, LS15].

For the generic case $d=\Omega(n)$ we focus in this paper, the current fastest runtime is dominated by $O^{*}\left(n^{2.5}\right)$. This runtime has not been improved since the result by Vaidya on 1989 [Vai87, Vai89b]. The $n^{2.5}$ bound originated from two factors: the cost per iteration $n^{2}$ and the number of iterations $\sqrt{n}$. The $n^{2}$ cost per iteration looks optimal because this is the cost to compute $A x$ for a dense $A$. Therefore, many efforts [Kar84, Ren88, NN89, Vai89a, LS14] have been focused on decreasing the number of iterations while maintaining the cost per iteration. As for many important linear programs (and convex programs), the number of iterations has been decreased, including maximum flow [Mad13, Mad16], minimum cost flow [CMSV17], geometric median [CLM $\left.{ }^{+} 16\right]$, matrix scaling and balancing [CMTV17], and $\ell_{p}$ regression [BCLL18]. Unfortunately, beating $\sqrt{n}$ iterations (or $\sqrt{d}$ when $d \ll n$ ) for the general case remains one of the biggest open problems in optimization.

Avoiding this open problem, we develop a stochastic central path method that has a runtime of $O^{*}\left(n^{\omega}+n^{2.5-\alpha / 2}+n^{2+1 / 6}\right)$, where $\omega$ is the exponent of matrix multiplication

[^8]and $\alpha$ is the dual exponent of matrix multiplication ${ }^{3}$. For the current value of $\omega \sim 2.38$ and $\alpha \sim 0.31$, the runtime is simply $O^{*}\left(n^{\omega}\right)$. This achieves the natural barrier for solving linear programs because linear system is a special case of linear program and that the currently fastest way to solve general linear systems involves matrix multiplication. Despite the exact approach used in [CW87, Wil12, DS13, LG14] cannot give a bound on $\omega$ better than 2.3078 [AFLG15] and all known approaches cannot achieve the bound $\omega=2$ [AW18b], it is still possible that $\omega=2.01$ using all known approaches. Therefore, we believe improving the additive $2+1 / 6$ term remains an interesting open problem.

Our method is a stochastic version of the short step central path method. This short step method takes $O^{*}(\sqrt{n})$ steps and each step decreases $x_{i} s_{i}$ by a $1-1 / \sqrt{n}$ factor for all $i$ where $s$ is the dual variable [Ren88] (See the definition of $s$ in (2.1)). This results in $O^{*}(\sqrt{n}) \times n=O^{*}\left(n^{1.5}\right)$ coordinate updates. Our method takes the same number of step but only updates $\widetilde{O}(\sqrt{n})$ coordinates each step. Therefore, we only update $O^{*}(n)$ coordinates in total, which is nearly optimal.

Our framework is efficient enough to take a much smaller step while maintaining the same running time. For the current value of $\omega \sim 2.38$, we show how to obtain the same runtime of $O^{*}\left(n^{\omega}\right)$ by taking $O^{*}(n)$ steps and $\widetilde{O}(1)$ coordinates update per steps. This is because the complexity of each step decreases proportionally when the step size decreases. Beyond the cost per iteration, we remark that our algorithm is one of the very few central path algorithms [PRT02, Mad13, Mad16] that does not maintain $x_{i} s_{i}$ close to some ideal vector in $\ell_{2}$ norm. We are hopeful that our stochastic method and our proof will be useful

[^9]for future research on interior point methods. In particular, it would be interesting to see how this can be combined with techniques in [Cla95, LS14] to get a faster algorithm for linear programs with $d \ll n$.

Besides the applications to linear programs, some of our techniques are probably useful for studying other important problems in convex optimization. In particular, our framework should be naturally extendable to a larger class of convex programs.

### 2.1.1 Related Work

Interior point method has a long history, for more detailed surveys, we refer the readers to [Wri97, Ye97, Ren01, RTV05, Meg12, Ter13]. This paper in part inspired by the use of data-structure in Laplacian solvers [ST04, KMP10, KMP11, CKM ${ }^{+}$11, KOSZ13, $\mathrm{CKM}^{+} 14, \mathrm{KLP}^{+} 16, \mathrm{KS} 16, \mathrm{CKK}^{+} 18, \mathrm{KPSZ18]}$, in particular the cycle update in [KOSZ13].

### 2.2 Results and Techniques

Theorem 2.2.1 (Main result). Given a linear program $\min _{A x=b, x \geq 0} c^{\top} x$ with no redundant constraints. Assume that the polytope has diameter $R$ in $\ell_{1}$ norm, namely, for any $x \geq 0$ with $A x=b$, we have $\|x\|_{1} \leq R$.

Then, for any $0<\delta \leq 1$, $\operatorname{Main}(A, b, c, \delta)$ outputs $x \geq 0$ such that

$$
c^{\top} x \leq \min _{A x=b, x \geq 0} c^{\top} x+\delta \cdot\|c\|_{\infty} R \quad \text { and } \quad\|A x-b\|_{1} \leq \delta \cdot\left(R \sum_{i, j}\left|A_{i, j}\right|+\|b\|_{1}\right)
$$

in expected time

$$
\left(n^{\omega+o(1)}+n^{2.5-\alpha / 2+o(1)}+n^{2+1 / 6+o(1)}\right) \cdot \log \left(\frac{n}{\delta}\right)
$$

where $\omega$ is the exponent of matrix multiplication, $\alpha$ is the dual exponent of matrix multiplication.

For the current value of $\omega \sim 2.38$ and $\alpha \sim 0.31$, the expected time is simply $n^{\omega+o(1)} \log \left(\frac{n}{\delta}\right)$.

Remark 2.2.1. See [Ren88] and [LS13, Sec E, F] on the discussion on converting an approximation solution to an exact solution. For integral $A, b, c$, it suffices to pick $\delta=2^{-O(L)}$ to get an exact solution where $L=\log \left(1+d_{\max }+\|c\|_{\infty}+\|b\|_{\infty}\right)$ is the bit complexity and $d_{\text {max }}$ is the largest absolute value of the determinant of a square sub-matrix of $A$. For many combinatorial problems, $L=O\left(\log \left(n+\|b\|_{\infty}+\|c\|_{\infty}\right)\right)$.

In this paper, we assume all floating point calculations are done exactly for simplicity. In general, the algorithm can be carried out with $O(L)$ bits of accuracy. This is necessary because each coordinate in the solution could require as much as $\Omega(L)$ bits to represent. See [Ren88] for some discussions on the numerical stability of the interior point methods.

If $T(n)$ is the current cost of matrix multiplication and inversion with $T(n) \sim n^{2.38}$, our runtime is simply $O\left(T(n) \log n \log \left(\frac{n}{\delta}\right)\right)$. The $\log \left(\frac{n}{\delta}\right)$ comes from iteration count and the $\log n$ factor comes from the doubling trick $\left(\left|y_{\pi(1.5 r)}\right| \geq(1-1 / \log n)\left|y_{\pi(r)}\right|\right)$ in the projection maintenance section. We left the problem of obtaining $O\left(T(n) \log \left(\frac{n}{\delta}\right)\right)$ as an open problem.

Finally, we note that our runtime holds for any square and rectangular matrix multiplication algorithm as long as $\omega \leq 3-\alpha$ (See Lemma 2.6.3) For example, Strassen algorithm together with a simple rectangular multiplication algorithm gives a runtime of roughly $n^{2.807}$.

### 2.2.1 Central Path Method

Our algorithm relies on two new ingredients: stochastic central path and projection maintenance. The central path method consider the linear programs

$$
\min _{A x=b, x \geq 0} c^{\top} x \quad \text { (primal) and } \quad \max _{A^{\top} y \leq c} b^{\top} y \quad \text { (dual) }
$$

with $A \in \mathbb{R}^{d \times n}$. Any solution of the linear program satisfies the following optimality conditions:

$$
\begin{align*}
x_{i} s_{i} & =0 \text { for all } i,  \tag{2.1}\\
A x & =b, \\
A^{\top} y+s & =c, \\
x_{i}, s_{i} & \geq 0 \text { for all } i .
\end{align*}
$$

We call $(x, s, y)$ feasible if it satisfies the last three equations above. For any feasible $(x, s, y)$, the duality gap is $\sum_{i} x_{i} s_{i}$. The central path method find a solution of the linear program by following the central path which uniformly decrease the duality gap. The central path $\left(x_{t}, s_{t}, y_{t}\right) \in \mathbb{R}^{n+n+d}$ is a path parameterized by $t$ and defined by

$$
\begin{align*}
x_{t, i} s_{t, i} & =t \text { for all } i,  \tag{2.2}\\
A x_{t} & =b, \\
A^{\top} y_{t}+s_{t} & =c, \\
x_{t, i}, s_{t, i} & \geq 0 \text { for all } i .
\end{align*}
$$

It is known [YTM94] how to transform linear programs by adding $O(n)$ many variables and constraints so that:

- The optimal solution remains the same.
- The central path at $t=1$ is near $\left(1_{n}, 1_{n}, 0_{d}\right)$ where $1_{n}$ and $0_{d}$ are all 1 and all 0 vectors with lengths $n$ and $d$.
- It is easy to convert an approximate solution of the transformed program to the original one.

For completeness, a theoretical version of such result is included in Lemma 3.8.2. This result shows that it suffices to move gradually $\left(x_{1}, s_{1}, y_{1}\right)$ to $\left(x_{t}, s_{t}, y_{t}\right)$ for small enough $t$.

### 2.2.1.1 Short Step Central Path Method

The short step central path method maintains $x_{i} s_{i}=\mu_{i}$ for some vector $\mu$ such that

$$
\begin{equation*}
\sum_{i}\left(\mu_{i}-t\right)^{2}=O\left(t^{2}\right) \quad \text { for some scalar } t>0 \tag{2.3}
\end{equation*}
$$

Since the duality gap is $\sum_{i} \mu_{i}$, it suffices to find $x$ and $s$ satisfying the above equation with small enough $t$. There are many variants of central path methods. We will focus on the version that decreases $t$ and takes a step of $\mu$ at the same time. The purpose of moving $\mu$ is to maintain the invariant (2.3) and the purpose of decreasing $t$ is decrease the duality gap, which is roughly $n t$. One natural way to maintain the invariant (2.3) is to do a gradient descent step on the energy $\sum_{i}\left(\mu_{i}-t\right)^{2}$ defined in (2.3), namely, moving $\mu$ to $\mu-h(\mu-t)$ with step size $h^{4}$. Compared to other versions, we only take one step instead of multiple steps to move $\mu$ closer to the central path $t$ per update of $t$.

[^10]More generally, say we want to move from $\mu$ to $\mu+\delta_{\mu}$, we approximate the term $\left(x+\delta_{x}\right)_{i}\left(s+\delta_{s}\right)_{i}$ by $x_{i} s_{i}+x_{i} \delta_{s, i}+s_{i} \delta_{x, i}$ and obtain the following system:

$$
\begin{align*}
X \delta_{s}+S \delta_{x} & =\delta_{\mu} \\
A \delta_{x} & =0  \tag{2.4}\\
A^{\top} \delta_{y}+\delta_{s} & =0,
\end{align*}
$$

where $X=\operatorname{diag}(x)$ and $S=\operatorname{diag}(s)$. This equation is the linear approximation of the original goal (moving from $\mu$ to $\mu+\delta_{\mu}$ ), and that the step is explicitly given by the formula

$$
\begin{equation*}
\delta_{x}=\frac{X}{\sqrt{X S}}(I-P) \frac{1}{\sqrt{X S}} \delta_{\mu} \text { and } \delta_{s}=\frac{S}{\sqrt{X S}} P \frac{1}{\sqrt{X S}} \delta_{\mu} \tag{2.5}
\end{equation*}
$$

where $P=\sqrt{\frac{X}{S}} A^{\top}\left(A \frac{X}{S} A^{\top}\right)^{-1} A \sqrt{\frac{X}{S}}$ is an orthogonal projection and the formulas $\frac{X}{\sqrt{X S}}, \frac{X}{S}, \cdots$ are the diagonal matrices of the corresponding vectors.

In turns out that one can decrease $t$ by $1-\frac{1}{\sqrt{n}}$ multiplicative factor every iteration while maintain the invariant (2.3). This requires $\widetilde{O}(\sqrt{n})$ iterations to converge. Combining this with the inverse maintenance technique [Vai87], this gives a total runtime of $n^{2.5}$. More precisely, the algorithm maintains the invariant $\sum_{i}\left(\mu_{i}-t\right)^{2}=O\left(t^{2}\right)$ by making steps bring $\mu_{i}$ closer to $t$ while taking steps to decrease $\mu_{i}$ uniformly. The progress of the whole algorithm is measured by $t$ because the duality gap is bounded by $n t$.

### 2.2.1.2 Stochastic Central Path Method

This part discuss how to modify the short step central path to decrease the cost per iteration to roughly $n^{\omega-\frac{1}{2}}$. Since our goal is to implement a central path method in sub-quadratic time per iteration, we even do not have the budget to compute $A x$ every
iterations. Therefore, instead of maintaining $\left(A \frac{X}{S} A^{\top}\right)^{-1}$ shown in previous papers, we will study the problem of maintaining a projection matrix $P=\sqrt{\frac{X}{S}} A^{\top}\left(A \frac{X}{S} A^{\top}\right)^{-1} A \sqrt{\frac{X}{S}}$ due to the formula of $\delta_{x}$ and $\delta_{s}(2.5)$.

However, even if the projection matrix $P$ is given explicitly for free, it is difficult to multiply the dense projection matrix with a dense vector $\delta_{\mu}$ in time $o\left(n^{2}\right)$. To avoid moving along a dense $\delta_{\mu}$, we move along an $O(k)$ sparse direction $\widetilde{\delta}_{\mu}$ defined by

$$
\widetilde{\delta}_{\mu, i}= \begin{cases}\delta_{\mu, i} / p_{i}, & \text { with probability } p_{i} \stackrel{\text { def }}{=} k \cdot\left(\frac{\delta_{\mu, i}^{2}}{\sum_{l} \delta_{\mu, l}^{2}}+\frac{1}{n}\right)  \tag{2.6}\\ 0, & \text { else. }\end{cases}
$$

The sparse direction is defined so that we are moving in the same direction in expectation $\left(\mathbb{E}\left[\widetilde{\delta}_{\mu, i}\right]=\delta_{\mu, i}\right)$ and that the direction has as small variance as possible $\left(\mathbb{E}\left[\widetilde{\delta}_{\mu, i}^{2}\right] \leq \frac{\sum_{i} \delta_{\mu, i}^{2}}{k}\right)$. If the projection matrix is given explicitly, we can apply the projection matrix on $\widetilde{\delta}_{\mu}$ in time $O(n k)$. This paper picks $k \sim \sqrt{n}$ and the sum of the cost of projection vector multiplications in the whole algorithm is about $n k^{2}=n^{2}$.

During the whole algorithm, we maintain a projection matrix

$$
\bar{P}=\sqrt{\frac{\bar{X}}{\bar{S}}} A^{\top}\left(A \frac{\bar{X}}{\bar{S}} A^{\top}\right)^{-1} A \sqrt{\frac{\bar{X}}{\bar{S}}}
$$

for vectors $\bar{x}$ and $\bar{s}$ such that $\bar{x}_{i}$ and $\bar{s}_{i}$ are multiplicative approximations of $x_{i}$ and $s_{i}$ respectively for all $i$. Since we maintain the projection at a nearby point $(\bar{x}, \bar{s})$, our stochastic step $x \leftarrow x+\widetilde{\delta}_{x}, s \leftarrow s+\widetilde{\delta}_{s}$ and $y \leftarrow y+\widetilde{\delta}_{y}$ are defined by

$$
\begin{align*}
\bar{X} \widetilde{\delta}_{s}+\bar{S} \widetilde{\delta}_{x} & =\widetilde{\delta}_{\mu} \\
A \widetilde{\delta}_{x} & =0  \tag{2.7}\\
A^{\top} \widetilde{\delta}_{y}+\widetilde{\delta}_{s} & =0
\end{align*}
$$

which is different from (2.4) on both sides of the first equation. Note that this system use $\bar{X}$ and $\bar{S}$ because we have only maintained this projection matrix. The main goal of Section 2.4 is to show $\bar{X}=\Theta(X)$ and $\bar{S}=\Theta(S)$ is good enough for our interior point method. Similar to (2.5), Lemma 2.4.2 shows that

$$
\begin{equation*}
\widetilde{\delta}_{x}=\frac{\bar{X}}{\sqrt{\overline{X S}}}(I-\bar{P}) \frac{1}{\sqrt{\overline{X S}}} \widetilde{\delta}_{\mu} \text { and } \widetilde{\delta}_{s}=\frac{\bar{S}}{\sqrt{\overline{X S}}} \bar{P} \frac{1}{\sqrt{\overline{X S}}} \widetilde{\delta}_{\mu} . \tag{2.8}
\end{equation*}
$$

The previously fastest algorithm involves maintaining the matrix inverse $\left(A \frac{X}{S} A^{\top}\right)^{-1}$ using subspace embedding techniques [Sar06, CW13, NN13a] and leverage score sampling [SS11]. In this paper, we maintain the projection directly using lazy update.

The key departure from the central path we present is that we can only maintain

$$
0.9 t \leq \mu_{i}=x_{i} s_{i} \leq 1.1 t \quad \text { for some } t>0
$$

instead of $\mu$ close to $t$ in $\ell_{2}$ norm. We will further explain the proof in Section 2.4.1.

### 2.2.2 Projection Maintenance via Lazy Update

The projection matrix we maintain is of the form $\sqrt{W} A^{\top}\left(A W A^{\top}\right)^{-1} A \sqrt{W}$ where $W=\operatorname{diag}(x / s)$. For intuition, we only explain how to maintain the matrix

$$
M_{w} \stackrel{\text { def }}{=} A^{\top}\left(A W A^{\top}\right)^{-1} A
$$

for the short step central path step here. In this case, we have $\sum_{i}\left(\frac{w_{i}^{\mathrm{new}}-w_{i}}{w_{i}}\right)^{2}=O(1)$ for each step. Given this, there are mainly two extreme cases, $w$ changes uniformly on all coordinates and $w$ changes only on a few coordinates.

If the changes $\left(\frac{w_{i}^{\text {new }}-w_{i}}{w_{i}}\right)^{2}$ is uniformly across all the coordinates, then $w_{i}^{\text {new }}=(1 \pm$ $\left.\frac{1}{\sqrt{n}}\right) w_{i}$ for all $i$. Since it takes $\sqrt{n}$ steps to change all coordinates by a constant factor and we only need to maintain $M_{v}$ for some $v_{i}=\Theta\left(w_{i}\right)$ for all $i$, we can update the matrix every $\sqrt{n}$ steps. Hence, the average cost per iteration of maintaining the projection matrix is $n^{\omega-\frac{1}{2}}$, which is exactly what we desired.

For the other extreme case that the "adversary" puts all of his $\ell_{2}$ budget on few coordinates, only $\sqrt{n}$ coordinates are changed by a constant factor during all $\sqrt{n}$ iterations. In this case, instead of updating $M_{w}$ every step, we can compute $M_{w} h$ online by the Woodbury matrix identity.

Fact 2.2.2 ([Woo50]). The Woodbury matrix identity is

$$
(M+U C V)^{-1}=M^{-1}-M^{-1} U\left(C^{-1}+V M^{-1} U\right)^{-1} V M^{-1} .
$$

Let $S \subset[n]$ denote the set of coordinates that is changed by more than a constant factor and $r=|S|$. Using the identity above, we have that

$$
\begin{equation*}
M_{w^{\mathrm{new}}}=M_{w}-\left(M_{w}\right)_{S}\left(\Delta_{S, S}^{-1}+\left(M_{w}\right)_{S, S}\right)^{-1}\left(\left(M_{w}\right)_{S}\right)^{\top} \tag{2.9}
\end{equation*}
$$

where $\Delta=\operatorname{diag}\left(w^{\text {new }}-w\right),\left(M_{w}\right)_{S} \in \mathbb{R}^{n \times r}$ is the $r$ columns from $S$ of $M_{w}$ and $\left(M_{w}\right)_{S, S}, \Delta_{S, S} \in$ $\mathbb{R}^{r \times r}$ are the $r$ rows and columns from $S$ of $M_{w}$ and $\Delta$.

As long as there are only few coordinates violating $v_{i}=\Theta\left(w_{i}\right)$, (2.9) can be applied online efficiently. In another case, we can use (2.9) instead to update the matrix $M_{w}$ and the cost is dominated by multiplying a $n \times n$ matrix with a $n \times n^{r}$ matrix.

Theorem 2.2.3 (Rectangular matrix multiplication, [LGU18]). Let the dual exponent of matrix multiplication $\alpha$ be the supremum among all $a \geq 0$ such that it takes $n^{2+o(1)}$ time to multiply an $n \times n$ matrix by an $n \times n^{a}$ matrix.

Then, for any $n \geq r$, multiplying an $n \times r$ with an $r \times n$ matrix or $n \times n$ with $n \times r$ takes time

$$
n^{2+o(1)}+r^{\frac{\omega-2}{1-\alpha}} n^{2-\frac{\alpha(\omega-2)}{1-\alpha}+o(1)} .
$$

Furthermore, we have $\alpha>0.31389$.

See Lemma 2.6.4 for the origin of the formula. Since the cost of multiplying $n \times n$ matrix by a $n \times 1$ matrix is same as the cost for $n \times n$ with $n \times n^{0.31},(2.9)$ should be used to update at least $n^{0.31}$ coordinates. In the extreme case only few $w_{i}$ are changing, we only need to update the matrix $n^{\frac{1}{2}-0.31}$ times during the whole algorithm and each takes $n^{2}$ time, and hence the total cost is less than $n^{\omega}$ for the current value of $\omega \sim 2.37$.

In previous papers [Kar84, Vai89b, NN91, NN94, LS14, LS15], the matrix is updated in a fixed schedule independent of the input sequence $w$. This leads to sub-optimal bounds if used in this paper. We instead define a potential function to measure the distance between the approximate vector $v$ and the target vector $w$. When there are less than $n^{\alpha}$ coordinates
of $v$ that is far from $w$, we are lazy and do not update the matrix. We simply apply the Woodbury matrix identity online. When there are more than $n^{\alpha}$ coordinates, we update $v$ by a certain greedy step. As in the extreme cases, the worst case of our algorithm is that the "adversary" puts his $\ell_{2}$ budget across all coordinates uniformly and hence the worst case runtime is $n^{\omega-\frac{1}{2}}$ per iteration. We will further explain the potential in Section 2.5.1.

### 2.3 Notations

For notation convenience, we assume the number of variables $n \geq 10$ and there is no redundant constraints. In particular, this implies that the constraint matrix $A$ is full rank and $n \geq d$.

For a positive integer $n$, let $[n]$ denote the set $\{1,2, \cdots, n\}$.
For any function $f$, we define $\widetilde{O}(f)$ to be $f \cdot \log ^{O(1)}(f)$. In addition to $O(\cdot)$ notation, for two functions $f, g$, we use the shorthand $f \lesssim g$ (resp. $\gtrsim$ ) to indicate that $f \leq C g$ (resp. $\geq)$ for some absolute constant $C$.

We use $\sinh x$ to denote $\frac{e^{x}-e^{-x}}{2}$ and $\cosh x$ to denote $\frac{e^{x}+e^{-x}}{2}$.
For vectors $a, b \in \mathbb{R}^{n}$ and accuracy parameter $\epsilon \in(0,1)$, we use $a \approx_{\epsilon} b$ to denote that $(1-\epsilon) b_{i} \leq a_{i} \leq(1+\epsilon) b_{i}, \forall i \in[n]$. Similarly, for any scalar $t$, we use $a \approx_{\epsilon} t$ to denote that $(1-\epsilon) t \leq a_{i} \leq(1+\epsilon) t, \forall i \in[n]$.

For a vector $x \in \mathbb{R}^{n}$ and $s \in \mathbb{R}^{n}$, we use $x s$ to denote a length $n$ vector with the $i$-th coordinate $(x s)_{i}$ is $x_{i} \cdot s_{i}$. Similarly, we extend other scalar operations to vector coordinatewise.

Given vectors $x, s \in \mathbb{R}^{n}$, we use $X$ and $S$ to denote the diagonal matrix of those two vectors. We use $\frac{X}{S}$ to denote the diagonal matrix given $\left(\frac{X}{S}\right)_{i, i}=x_{i} / s_{i}$. Similarly, we extend other scalar operations to diagonal matrix diagonal-wise. Note that matrix $\sqrt{\frac{X}{S}} A^{\top}\left(A \frac{X}{S} A^{\top}\right)^{-1} A \sqrt{\frac{X}{S}}$ is an orthogonal projection matrix.


Figure 2.1: ClassicalStep happens with $n^{-2}$ probability

### 2.4 Stochastic Central Path Method

### 2.4.1 Proof Outline

The short step central path method is defined using the approximation $\left(x+\delta_{x}\right)_{i}(s+$ $\left.\delta_{s}\right)_{i} \sim x_{i} s_{i}+x_{i} \delta_{s, i}+s_{i} \delta_{x, i}$. This approximate is only accurate if $\left\|X^{-1} \delta_{x}\right\|_{\infty} \leq 1 / 2$ and $\left\|S^{-1} \delta_{s}\right\|_{\infty} \leq 1 / 2$. For the $\delta_{x}$ step, we have

$$
\begin{equation*}
X^{-1} \delta_{x}=\frac{1}{\sqrt{X S}}(I-P) \frac{1}{\sqrt{X S}} \delta_{\mu} \sim \frac{1}{t}(I-P) \delta_{\mu} \tag{2.10}
\end{equation*}
$$

where we used $x_{i} s_{i} \sim t$ for all $i$.
If we know that $\left\|\delta_{\mu}\right\|_{2} \leq t / 4$, then the $\ell_{\infty}$ norm can be bounded as follows:

$$
\left\|X^{-1} \delta_{x}\right\|_{\infty} \leq\left\|X^{-1} \delta_{x}\right\|_{2} \lesssim \frac{1}{t}\left\|(I-P) \delta_{\mu}\right\|_{2} \leq \frac{1}{t}\left\|\delta_{\mu}\right\|_{2} \leq 1 / 2
$$

where we used that $I-P$ is an orthogonal projection matrix. This is the reason why a standard choice of $\delta_{\mu, i}$ is $-c t / \sqrt{n}$ for all $i$ for some small constant $c$.

For the stochastic step, $\widetilde{\delta}_{\mu, i} \sim-\frac{t}{\sqrt{n}} \frac{n}{k}$ for roughly $k$ coordinates where the term $\frac{n}{k}$ is used to preserve the expectation of the step. Therefore, the $\ell_{2}$ norm of $\widetilde{\delta}_{\mu}$ is very large

```
Algorithm 2.1 Stochastic Step
    procedure StochasticStep \(\left(\mathrm{mp}, x, s, \delta_{\mu}, k, \epsilon\right) \quad \triangleright\) Lemma 2.4.2,2.4.3,2.4.7
        \(w \leftarrow \frac{x}{s}, \widetilde{v} \leftarrow \operatorname{mp.UPDATE}(w) \quad \triangleright\) Algorithm 2.3
        \(\bar{x} \leftarrow x \sqrt{\frac{\widetilde{v}}{w}}, \bar{s} \leftarrow s \sqrt{\frac{w}{\tilde{v}}} \quad \triangleright\) It guarantees that \(\frac{\bar{x}}{\bar{s}}=\widetilde{v}\) and \(\overline{x s}=x s\)
        repeat
            Generate \(\widetilde{\delta}_{\mu}\) such that \(\triangleright\) Compute a sparse direction
            \(\widetilde{\delta}_{\mu, i} \leftarrow \begin{cases}\delta_{\mu, i} / p_{i}, & \text { with prob. } p_{i}=\min \left(1, k \cdot\left(\left(\delta_{\mu, i}^{2} / \sum_{l=1}^{n} \delta_{\mu, l}^{2}\right)+1 / n\right)\right) ; \\ 0 & \text { else. }\end{cases}\)
                                    \(\triangleright\) Compute an approximate step
            \(\triangleright\) Find \(\left(\widetilde{\delta}_{x}, \widetilde{\delta}_{s}, \widetilde{\delta}_{y}\right)\) such that these three equations hold
\[
\begin{aligned}
\bar{X} \widetilde{\delta}_{s}+\bar{S} \widetilde{\delta}_{x} & =\widetilde{\delta}_{\mu}, \\
A \widetilde{\delta}_{x} & =0 \\
A^{\top} \widetilde{\delta}_{y}+\widetilde{\delta}_{s} & =0
\end{aligned}
\]
9: \(\quad p_{\mu} \leftarrow \operatorname{mp.QUERY}\left(\frac{1}{\sqrt{\overline{X S}}} \widetilde{\delta}_{\mu}\right) \quad \triangleright\) Algorithm 2.3
10: \(\quad \widetilde{\delta}_{s} \leftarrow \frac{\bar{S}}{\sqrt{\overline{X S}}} p_{\mu} \quad \triangleright\) According to (2.11)
11: \(\widetilde{\delta}_{x} \leftarrow \frac{1}{\bar{S}} \widetilde{\delta}_{\mu}-\frac{\bar{X}}{\sqrt{\overline{X S}}} p_{\mu} \quad \triangleright\) According to (2.12)
until \(\left\|\bar{s}^{-1} \widetilde{\delta}_{s}\right\|_{\infty} \leq \frac{1}{100 \log n}\) and \(\left\|\bar{x}^{-1} \widetilde{\delta}_{x}\right\|_{\infty} \leq \frac{1}{100 \log n}\)
return \(\left(x+\widetilde{\delta}_{x}, s+\widetilde{\delta}_{s}\right)\)
end procedure
```

$\left(\left\|\widetilde{\delta}_{\mu}\right\|_{2} \sim t \sqrt{\frac{n}{k}}\right)$. After the projection, we have $\left\|X^{-1} \delta_{x}\right\|_{2} \sim \frac{1}{t}\left\|(I-P) \delta_{\mu}\right\|_{2} \sim \sqrt{\frac{n}{k}}$. Hence, the bound of $\left\|X^{-1} \delta_{x}\right\|_{\infty}$ using $\left\|X^{-1} \delta_{x}\right\|_{2}$ is too weak. To improve the bound, we use Chernoff bounds to estimate $\left\|X^{-1} \delta_{x}\right\|_{\infty}$.

Beside the $\ell_{\infty}$ norm bound, the proof sketch in (2.10) also requires using $x_{i} s_{i} \sim t$ for all $i$. The short step central path proof maintains an invariant that $\sum_{i}\left(x_{i} s_{i}-t\right)^{2}=O\left(t^{2}\right)$. However, since our stochastic step has a stochastic noise with $\ell_{2}$ norm as large as $t \sqrt{\frac{n}{k}}$, one cannot hope to maintain $x_{i} s_{i}$ close to $t$ in $\ell_{2}$ norm. Instead, we follow an idea in

```
Algorithm 2.2 Our Main Algorithm
    procedure \(\operatorname{Main}(A, b, c, \delta) \quad \triangleright\) Theorem 2.2.1
        \(\epsilon \leftarrow \frac{1}{40000 \log n}, \epsilon_{m p} \leftarrow \frac{1}{40000}, k \leftarrow \frac{1000 \epsilon \sqrt{n} \log ^{2} n}{\epsilon_{m p}}\).
        \(\lambda \leftarrow 40 \log n, \delta \leftarrow \min \left(\frac{\delta}{2}, \frac{1}{\lambda}\right), a \leftarrow \min (\alpha, 2 / 3)\).
        Modify the linear program and obtain an initial \(x\) and \(s\) according to Lemma 3.8.2.
        MaintainProjection mp
        mp.Initialize \(\left(A, \frac{x}{s}, \epsilon_{m p}, a\right) \quad \triangleright\) Algorithm 2.3
        \(t \leftarrow 1 \quad \triangleright\) Initialize \(t\)
        while \(t>\delta^{2} /(2 n)\) do \(\quad \triangleright\) We stopped once the precision is good
            \(t^{\text {new }} \leftarrow\left(1-\frac{\epsilon}{3 \sqrt{n}}\right) t\)
            \(\mu \leftarrow x s\)
            \(\delta_{\mu} \leftarrow\left(\frac{t^{\text {new }}}{t}-1\right) x s-\frac{\epsilon}{2} \cdot t^{\text {new }} \cdot \frac{\nabla \Phi_{\lambda}(\mu / t-1)}{\left\|\nabla \Phi_{\lambda}(\mu / t-1)\right\|_{2}} \quad \triangleright \Phi_{\lambda}\) is defined in Lemma 2.4.11
            \(\left(x^{\text {new }}, s^{\text {new }}\right) \leftarrow\) StochasticStep \(\left(\mathrm{mp}, x, s, \delta_{\mu}, k, \epsilon\right) \quad \triangleright\) Algorithm 2.1
            if \(\Phi_{\lambda}\left(\mu^{\text {new }} / t^{\text {new }}-1\right)>n^{3}\) then \(\quad \triangleright\) When potential function is large
                \(\left(x^{\text {new }}, s^{\text {new }}\right) \leftarrow \operatorname{CLASSICALStEP}\left(x, s, t^{\text {new }}\right) \quad \triangleright \operatorname{Lemma} 2.6 .2\), [Vai89b]
                    mp.Initialize \(\left(A, \frac{x^{\text {new }}}{s^{\text {new }}}, \epsilon_{m p}, a\right) \quad \triangleright\) Restart the data structure
            end if
            \((x, s) \leftarrow\left(x^{\text {new }}, s^{\text {new }}\right), t \leftarrow t^{\text {new }}\)
        end while
        Return an approximate solution of the original linear program according to Lemma
    3.8.2.
20: end procedure
```

[LS14, LSW15] and maintain the following potential

$$
\sum_{i=1}^{n} \cosh \left(\lambda\left(\frac{x_{i} s_{i}}{t}-1\right)\right)=n^{O(1)}
$$

with $\lambda=\Theta(\log n)$. This potential is a variant of soft-max. Note that the potential bounded by $n^{O(1)}$ implies that $x_{i} s_{i}$ is a multiplicative approximation of $t$. To bound the potential, consider $r_{i}=\frac{x_{i} s_{i}}{t}$ and $\Phi(r)$ be the potential above. Then, we have that

$$
\mathbb{E}\left[\Phi\left(r^{\text {new }}\right)\right] \leq \Phi(r)+\left\langle\nabla \Phi(r), \mathbb{E}\left[r^{\text {new }}-r\right]\right\rangle+O(1) \mathbb{E}\left\|r^{\text {new }}-r\right\|_{\nabla^{2} \Phi(r)}^{2}
$$

The first order term can be bounded efficiently because $\mathbb{E}\left[r^{\text {new }}-r\right]$ is close to the short step central path step. The second term is a variance term which scales like $1 / k$ due to the $k$ independent coordinates. Therefore, the potential changed by $1 / k \sim 1 / \sqrt{n}$ factor each step. Hence, we can maintain it for roughly $\sqrt{n}$ steps.

To make sure the potential $\Phi$ is bounded during the whole algorithm, our step is the mixtures of two steps of the form $\delta_{\mu} \sim-\frac{t}{\sqrt{n}}-t \frac{\nabla \Phi}{\|\nabla \Phi\|_{2}}$. The first term is to decrease $t$ and the second term is to decrease $\Phi$.

Since the algorithm is randomized, there is a tiny probability that $\Phi$ is large. In that case, we switch to a short step central path method. See Figure 2.1, Algorithm 2.1, and Algorithm 3.6. The first part of the proof involves bounding every quantity listed in Table 2.1. In the second part, we are using these quantities to bound the expectation of $\Phi$.

To decouple the proof in both parts, we will make the following assumption in first part. It will be verified in the second part.

Assumption 2.4.1. Assume the following for the input of the procedure STOCHASTICSTEP (see Algorithm 2.1):

- $x s \approx_{0.1} t$ with $t>0$.
- mp. $\operatorname{UPDATE}(w)$ outputs $\widetilde{v}$ such that $w \approx_{\epsilon_{m p}} \widetilde{v}$ with $\epsilon_{m p} \leq 1 / 40000$.
- $\left\|\delta_{\mu}\right\|_{2} \leq \epsilon$ with $0<\epsilon<1 /(40000 \log n)$.
- $k \geq 1000 \epsilon \sqrt{n} \log ^{2} n / \epsilon_{m p}$.

| Quantity | Bound | Place |
| :--- | :--- | :--- |
| $\left\\|\mathbb{E}\left[s^{-1} \widetilde{\delta}_{s}\right]\right\\|_{2},\left\\|\mathbb{E}\left[x^{-1} \widetilde{\delta}_{s}\right]\right\\|_{2},\left\\|\mathbb{E}\left[\mu^{-1} \widetilde{\delta}_{\mu}\right]\right\\|_{2}$ | $O(\epsilon)$ | Part 1, Lemma 2.4.3 |
| $\left\\|\mathbb{E}\left[\mu^{-1}\left(\mu^{\text {new }}-\mu-\widetilde{\delta}_{\mu}\right)\right]\right\\|_{2}$ | $O\left(\epsilon_{m p} \cdot \epsilon\right)$ | Part 1, Lemma 2.4.7 |
| $\left\\|\mathbb{E}\left[\mu^{-1}\left(\mu^{\text {new }}-\mu\right)\right]\right\\|_{2}$ | $O(\epsilon)$ | Part 1, Lemma 2.4.7 |
| $\mathbb{V}\left[s_{i}^{-1} \widetilde{\delta}_{s, i}\right], \mathbb{V}\left[x_{i}^{-1} \widetilde{\delta}_{x, i}\right], \mathbb{V}\left[\mu_{i}^{-1} \widetilde{\delta}_{\mu, i}\right]$ | $O\left(\epsilon^{2} / k\right)$ | Part 2, Lemma 2.4.3 |
| $\mathbb{V}\left[\mu_{i}^{-1} \mu^{\text {new }}\right]$ | $O\left(\epsilon^{2} / k\right)$ | Part 2, Lemma 2.4.7 |
| $\left\\|s^{-1} \widetilde{\delta}_{s}\right\\|_{\infty},\left\\|x^{-1} \widetilde{\delta}_{x}\right\\|_{\infty},\left\\|\mu^{-1} \widetilde{\delta}_{\mu}\right\\|_{\infty}$ | $O(1 / \log n)$ | Part 3, Lemma 2.4.3 |
| $\left\\|\mu^{-1}\left(\mu^{\text {new }}-\mu\right)\right\\|_{\infty}$ | $O(1 / \log n)$ | Part 3, Lemma 2.4.7 |

Table 2.1: The bound of each quantity under Assumption 2.4.1. For intuition, think $\epsilon \sim$ $\epsilon_{m p} \sim 1 / 10$ and $k \sim \sqrt{n}$.

### 2.4.2 Bounding each quantity of stochastic step

First, we give an explicit formula for our step, which will be used in all subsequent calculations.

Lemma 2.4.2. The procedure $\operatorname{StochasticStep}\left(\mathrm{mp}, x, s, \delta_{\mu}, k, \epsilon\right)$ (see Algorithm 2.1) finds a solution $\widetilde{\delta}_{x}, \widetilde{\delta}_{s} \in \mathbb{R}^{n}$ to (2.7) by the formula

$$
\begin{align*}
& \widetilde{\delta}_{x}=\frac{\bar{X}}{\sqrt{X S}}(I-\bar{P}) \frac{1}{\sqrt{X S}} \widetilde{\delta}_{\mu}  \tag{2.11}\\
& \widetilde{\delta}_{s}=\frac{\bar{S}}{\sqrt{\overline{X S}}} \bar{P} \frac{1}{\sqrt{X S}} \widetilde{\delta}_{\mu} \tag{2.12}
\end{align*}
$$

with

$$
\begin{equation*}
\bar{P}=\sqrt{\frac{\bar{X}}{\bar{S}}} A^{\top}\left(A \frac{\bar{X}}{\bar{S}} A^{\top}\right)^{-1} A \sqrt{\frac{\bar{X}}{\bar{S}}} \tag{2.13}
\end{equation*}
$$

Proof. For the first equation of (2.7), we multiply $A \bar{S}^{-1}$ on both sides,

$$
A \bar{S}^{-1} \bar{X} \widetilde{\delta}_{s}+A \widetilde{\delta}_{x}=A \bar{S}^{-1} \widetilde{\delta}_{\mu}
$$

Since the second equation gives $A \widetilde{\delta}_{x}=0$, then we know that $A \bar{S}^{-1} \bar{X} \widetilde{\delta}_{s}=A \bar{S}^{-1} \widetilde{\delta}_{\mu}$. Multiplying $A \bar{S}^{-1} \bar{X}$ on both sides of the third equation of (2.7), we have

$$
-A \bar{S}^{-1} \bar{X} A^{\top} \widetilde{\delta}_{y}=A \bar{S}^{-1} \bar{X} \widetilde{\delta}_{s}=A \bar{S}^{-1} \widetilde{\delta}_{\mu}
$$

Thus,

$$
\begin{aligned}
& \widetilde{\delta}_{y}=-\left(A \bar{S}^{-1} \bar{X} A^{\top}\right)^{-1} A \bar{S}^{-1} \widetilde{\delta}_{\mu}, \\
& \widetilde{\delta}_{s}=A^{\top}\left(A \bar{S}^{-1} \bar{X} A^{\top}\right)^{-1} A \bar{S}^{-1} \widetilde{\delta}_{\mu}, \\
& \widetilde{\delta}_{x}=\bar{S}^{-1} \widetilde{\delta}_{\mu}-\bar{S}^{-1} \bar{X} A^{\top}\left(A \bar{S}^{-1} \bar{X} A^{\top}\right)^{-1} A \bar{S}^{-1} \widetilde{\delta}_{\mu} .
\end{aligned}
$$

Recall we define $\bar{P}$ as (2.13), then we have

$$
\widetilde{\delta}_{s}=\frac{\bar{S}}{\sqrt{\overline{X S}}} \cdot \sqrt{\frac{\bar{X}}{\bar{S}}} A^{\top}\left(A \frac{\bar{X}}{\bar{S}} A^{\top}\right)^{-1} \sqrt{\frac{\bar{X}}{\bar{S}}} \cdot \frac{1}{\sqrt{\overline{X S}}} \widetilde{\delta}_{\mu}=\frac{\bar{S}}{\sqrt{\overline{X S}}} \bar{P} \frac{1}{\sqrt{X S}} \widetilde{\delta}_{\mu}
$$

and

$$
\widetilde{\delta}_{x}=\bar{S}^{-1} \widetilde{\delta}_{\mu}-\frac{\bar{X}}{\sqrt{\overline{X S}}} \cdot \sqrt{\frac{\bar{X}}{\bar{S}}} A^{\top}\left(A \frac{\bar{X}}{\bar{S}} A^{\top}\right)^{-1} \sqrt{\frac{\bar{X}}{\bar{S}}} \cdot \frac{1}{\sqrt{\overline{X S}}} \widetilde{\delta}_{\mu}=\frac{\bar{X}}{\sqrt{\overline{X S}}}(I-\bar{P}) \frac{1}{\sqrt{\overline{X S}}} \widetilde{\delta}_{\mu} .
$$

which are matching (2.11) and (2.12).
To see why the STOCHASTICSTEP outputs $\widetilde{\delta}_{x}, \widetilde{\delta}_{s}$ satisfying (2.11) and (2.12), we note that

$$
p_{\mu}=\sqrt{\widetilde{V}} A^{\top}\left(A \frac{\bar{X}}{\bar{S}} A^{\top}\right)^{-1} A \sqrt{\widetilde{V}} \frac{1}{\sqrt{\overline{X S}}} \widetilde{\delta}_{\mu}=\bar{P} \frac{1}{\sqrt{\overline{X S}}} \widetilde{\delta}_{\mu}
$$

because of Theorem 2.5.1.

Using the explicitly formula, we are ready to bound all quantities we needed in the following two subsubsections.

### 2.4.2.1 Bounding $\widetilde{\delta}_{s}, \widetilde{\delta}_{x}$ and $\widetilde{\delta}_{\mu}$

Lemma 2.4.3. Under the Assumption 2.4.1, the two vectors $\widetilde{\delta}_{x}$ and $\widetilde{\delta}_{s}$ found by STOCHASTICSTEP satisfy :

1. $\left\|\mathbb{E}\left[\bar{s}^{-1} \widetilde{\delta}_{s}\right]\right\|_{2} \leq 2 \epsilon,\left\|\mathbb{E}\left[\bar{x}^{-1} \widetilde{\delta}_{x}\right]\right\|_{2} \leq 2 \epsilon,\left\|\mathbb{E}\left[s^{-1} \widetilde{\delta}_{s}\right]\right\|_{2} \leq 2 \epsilon,\left\|\mathbb{E}\left[x^{-1} \widetilde{\delta}_{x}\right]\right\|_{2} \leq 2 \epsilon,\left\|\mathbb{E}\left[\mu^{-1} \widetilde{\delta}_{\mu}\right]\right\|_{2} \leq$ $4 \epsilon$.
2. $\mathbb{V}\left[\frac{\tilde{s}_{s, i}}{\bar{s}_{i}}\right] \leq \frac{2 \epsilon^{2}}{k}, \mathbb{V}\left[\frac{\tilde{\phi}_{x, i}}{\bar{x}_{i}}\right] \leq \frac{2 \epsilon^{2}}{k}, \mathbb{V}\left[\frac{\tilde{s}_{s, i}}{s_{i}}\right] \leq \frac{2 \epsilon^{2}}{k}, \mathbb{V}\left[\frac{\widetilde{\boldsymbol{\gamma}}_{x, i}}{x_{i}}\right] \leq \frac{2 \epsilon^{2}}{k}, \mathbb{V}\left[\frac{\widetilde{\mathcal{p}}_{\mu, i}}{\mu_{i}}\right] \leq \frac{8 \epsilon^{2}}{k}$.
3. $\left\|\bar{s}^{-1} \widetilde{\delta}_{s}\right\|_{\infty} \leq \frac{0.01}{\log n},\left\|s^{-1} \widetilde{\delta}_{s}\right\|_{\infty} \leq \frac{0.02}{\log n},\left\|\bar{x}^{-1} \widetilde{\delta}_{x}\right\|_{\infty} \leq \frac{0.01}{\log n},\left\|x^{-1} \widetilde{\delta}_{x}\right\|_{\infty} \leq \frac{0.02}{\log n},\left\|\mu^{-1} \widetilde{\delta}_{\mu}\right\|_{\infty} \leq \frac{0.02}{\log n}$.

Remark 2.4.1. For notational simplicity, the $\mathbb{E}$ and $\mathbb{V}$ in the proof are for the case without resampling (Line 12). Since the all the additional terms due to resampling are polynomially bounded and since we can set failure probability to an arbitrarily small inverse polynomial (see Claim 2.4.6), the proof does not change and the result remains the same.

## Proof.

Claim 2.4.4 (Part 1 , bounding the $\ell_{2}$ norm of expectation).

$$
\left\|\mathbb{E}\left[\bar{s}^{-1} \widetilde{\delta}_{s}\right]\right\|_{2} \leq 2 \epsilon,\left\|\mathbb{E}\left[\bar{x}^{-1} \widetilde{\delta}_{x}\right]\right\|_{2} \leq 2 \epsilon,\left\|\mathbb{E}\left[s^{-1} \widetilde{\delta}_{s}\right]\right\|_{2} \leq 2 \epsilon,\left\|\mathbb{E}\left[x^{-1} \widetilde{\delta}_{x}\right]\right\|_{2} \leq 2 \epsilon,\left\|\mathbb{E}\left[\mu^{-1} \widetilde{\delta}_{\mu}\right]\right\|_{2} \leq 4 \epsilon
$$

Proof. For $\left\|\bar{s}^{-1} \widetilde{\delta}_{s}\right\|_{\infty}$, we consider the $i$-th coordinate of the vector

$$
\bar{s}_{i}^{-1} \widetilde{\delta}_{s, i}=\frac{1}{\sqrt{\bar{x}_{i} \bar{s}_{i}}} \sum_{j=1}^{n} \bar{P}_{i, j} \frac{\widetilde{\delta}_{\mu, j}}{\sqrt{\bar{x}_{j} \bar{s}_{j}}}
$$

Then, we have

$$
\mathbb{E}\left[\bar{s}_{i}^{-1} \widetilde{\delta}_{s, i}\right]=\frac{1}{\sqrt{\bar{x}_{i} \bar{s}_{i}}} \sum_{j=1}^{n} \bar{P}_{i, j} \frac{\mathbb{E}\left[\widetilde{\delta}_{\mu, j}\right]}{\sqrt{\bar{x}_{j} \bar{s}_{j}}}=\frac{1}{\sqrt{\bar{x}_{i} \bar{s}_{i}}} \sum_{j=1}^{n} \bar{P}_{i, j} \frac{\delta_{\mu, j}}{\sqrt{\bar{x}_{j} \bar{s}_{j}}}
$$

Since $x s \approx_{0.1} t$ and $\left\|\delta_{\mu}\right\| \leq \epsilon t$, we have $\left\|\frac{\delta_{\mu}}{\sqrt{x s}}\right\|_{2} \leq \frac{1.1 \epsilon t}{\sqrt{t}}$. Since $\bar{P}$ is an orthogonal projection matrix, we have $\left\|\bar{P} \frac{\delta_{\mu}}{\sqrt{\overline{x s}}}\right\|_{2} \leq\left\|\frac{\delta_{\mu}}{\sqrt{\overline{x s}}}\right\|_{2}$. Putting all the above facts and $x s=\overline{x s}$, we can show

$$
\begin{aligned}
\left\|\mathbb{E}\left[\bar{s}^{-1} \widetilde{\delta}_{s}\right]\right\|_{2}^{2} & =\sum_{i=1}^{n}\left(\frac{1}{\sqrt{\bar{x}_{i} \overline{\bar{s}}_{i}}} \sum_{j=1}^{n} \bar{P}_{i, j} \frac{\delta_{\mu, j}}{\sqrt{\bar{x}_{j} \bar{s}_{j}}}\right)^{2}=\sum_{i=1}^{n} \frac{1}{\bar{x}_{i} \bar{s}_{i}}\left(\sum_{j=1}^{n} \bar{P}_{i, j} \frac{\delta_{\mu, j}}{\sqrt{\bar{x}_{j} \bar{s}_{j}}}\right)^{2} \\
& \leq \frac{1}{0.9 t} \sum_{i=1}^{n}\left(\sum_{j=1}^{n} \bar{P}_{i, j} \frac{\delta_{\mu, j}}{\sqrt{\bar{x}_{j} \bar{s}_{j}}}\right)^{2}=\frac{1}{0.9 t}\left\|\bar{P} \frac{\delta_{\mu}}{\sqrt{\overline{x s}}}\right\|_{2}^{2} \\
& \leq \frac{1}{0.9 t}\left\|\frac{\delta_{\mu}}{\sqrt{\overline{x s}}}\right\|_{2}^{2} \leq \frac{(1.1)^{2}}{0.9 t} \cdot \frac{(\epsilon t)^{2}}{t} \leq 1.4 \epsilon^{2}
\end{aligned}
$$

which implies that

$$
\begin{equation*}
\left\|\mathbb{E}\left[\bar{s}^{-1} \widetilde{\delta}_{s}\right]\right\|_{2} \leq 1.2 \epsilon . \tag{2.14}
\end{equation*}
$$

Notice that the proof for $x$ is identical to the proof for $s$ because $(I-\bar{P})$ is also a projection matrix. Since $\bar{s} \approx_{0.1} s$ and $\bar{x} \approx_{0.1} x$, then we can also prove the next two inequalities in the Claim statement.

Now, we are ready to bound $\left\|\mathbb{E}\left[\mu^{-1} \widetilde{\delta}_{\mu}\right]\right\|_{2}$

$$
\left\|\mathbb{E}\left[\mu^{-1} \widetilde{\delta}_{\mu}\right]\right\|_{2}=\left\|\mathbb{E}\left[\bar{s}^{-1} \bar{x}^{-1}\left(\bar{x} \widetilde{\delta}_{s}+\bar{s} \widetilde{\delta}_{x}\right)\right]\right\|_{2} \leq\left\|\mathbb{E}\left[\bar{s}^{-1} \widetilde{\delta}_{s}\right]\right\|_{2}+\left\|\mathbb{E}\left[\bar{x}^{-1} \widetilde{\delta}_{x}\right]\right\|_{2} \leq 4 \epsilon
$$

by using $\mu=x s=\overline{x s}$ and $\bar{x} \widetilde{\delta}_{s}+\bar{s} \widetilde{\delta}_{x}=\widetilde{\delta}_{\mu}$ from (2.7).
Claim 2.4.5 (Part 2, bounding the variance per coordinate).

$$
\mathbb{V}\left[\bar{s}_{i}^{-1} \widetilde{\delta}_{s, i}\right] \leq \frac{2 \epsilon^{2}}{k}, \mathbb{V}\left[\bar{x}_{i}^{-1} \widetilde{\delta}_{x, i}\right] \leq \frac{2 \epsilon^{2}}{k}, \mathbb{V}\left[s_{i}^{-1} \widetilde{\delta}_{s, i}\right] \leq \frac{2 \epsilon^{2}}{k}, \mathbb{V}\left[x_{i}^{-1} \widetilde{\delta}_{x, i}\right] \leq \frac{2 \epsilon^{2}}{k}, \mathbb{V}\left[\mu_{i}^{-1} \widetilde{\delta}_{\mu, i}\right] \leq \frac{8 \epsilon^{2}}{k}
$$

Proof. Consider the $i$-th coordinate of the vector

$$
\bar{s}_{i}^{-1} \widetilde{\delta}_{s, i}=\frac{1}{\sqrt{\bar{x}_{i} \bar{s}_{i}}} \sum_{j=1}^{n} \bar{P}_{i, j} \frac{\widetilde{\delta}_{\mu, j}}{\sqrt{\bar{x}_{j} \bar{s}_{j}}}
$$

For variance of $\bar{s}_{i}^{-1} \widetilde{\delta}_{s, i}$, we have

$$
\begin{array}{rlr}
\mathbb{V}\left[\bar{s}_{i}^{-1} \widetilde{\delta}_{s, i}\right] & =\frac{1}{\bar{x}_{i} \bar{s}_{i}} \sum_{j=1}^{n} \frac{\bar{P}_{i, j}^{2}}{\bar{x}_{j} \bar{s}_{j}} \mathbb{V}\left[\widetilde{\delta}_{\mu, j}\right] & \text { by all } \widetilde{\delta}_{\mu, j} \text { are independent } \\
& \leq \frac{1}{\bar{x}_{i} \bar{s}_{i}} \sum_{j=1}^{n} \frac{\bar{P}_{i, j}^{2}}{\bar{x}_{j} \bar{s}_{j}} \frac{1}{k} \frac{\delta_{\mu, j}^{2}}{\frac{\delta_{\mu, j}^{2}}{\sum_{l=1}^{n} \delta_{\mu, l}^{2}}+\frac{1}{n}} & \text { by }(2.6) \\
& \leq \frac{1}{\bar{x}_{i} \bar{s}_{i}} \sum_{j=1}^{n} \frac{\bar{P}_{i, j}^{2}}{\bar{x}_{j} \bar{s}_{j}} \frac{1}{k} \sum_{l=1}^{n} \delta_{\mu, l}^{2} & \\
& \leq \frac{1.3}{t^{2}} \sum_{j=1}^{n} \bar{P}_{i, j}^{2} \frac{1}{k} \sum_{l=1}^{n} \delta_{\mu, l}^{2} \leq \frac{1.3 \epsilon^{2}}{k}, & \text { by } \bar{x}_{i} \bar{s}_{i}=x_{i} s_{i} \approx_{1 / 10} t
\end{array}
$$

where we used that $\sum_{j=1}^{n} \bar{P}_{i, j}^{2}=\bar{P}_{i, i} \leq 1,\left\|\delta_{\mu}\right\|_{2} \leq \epsilon t$ at the end.
The proof for the other three inequalities in the Claim statement are identical to this one. We omit here.

For the variance of $\mu_{i}^{-1} \widetilde{\delta}_{\mu, i}$,

$$
\begin{aligned}
\mathbb{V}\left[\mu_{i}^{-1} \widetilde{\delta}_{\mu, i}\right] & =\mathbb{V}\left[\bar{x}_{i}^{-1} \bar{s}_{i}^{-1}\left(\bar{x}_{i} \widetilde{\delta}_{s, i}+\bar{s}_{i} \widetilde{\delta}_{x, i}\right)\right] \\
& \leq 2 \mathbb{V}\left[\bar{x}_{i}^{-1} \bar{x}_{i} \bar{s}_{i}^{-1} \widetilde{\delta}_{s, i}\right]+2 \mathbb{V}\left[\bar{s}_{i}^{-1} \bar{s}_{i} \bar{x}_{i}^{-1} \widetilde{\delta}_{x, i}\right] \\
& =2 \mathbb{V}\left[\bar{s}_{i}^{-1} \widetilde{\delta}_{s, i}\right]+2 \mathbb{V}\left[\bar{x}_{i}^{-1} \widetilde{\delta}_{x, i}\right] \leq 8 \epsilon^{2} / k
\end{aligned}
$$

where the first step follows by definition of $\mu=x s=\overline{x s}$ and (2.7), the second step follows by triangle inequality and, the last step follows by $\mathbb{V}\left[\bar{s}_{i}^{-1} \widetilde{\delta}_{s, i}\right], \mathbb{V}\left[\bar{x}_{i}^{-1} \widetilde{\delta}_{x, i}\right] \leq 2 \epsilon^{2} / k$

Claim 2.4.6 (Part 3, bounding the probability of success). Without resampling, the following holds with probability $1-2 n \exp \left(-\frac{0.003 k}{\epsilon \sqrt{n} \log n}\right)$. $\left\|\bar{s}^{-1} \widetilde{\delta}_{s}\right\|_{\infty} \leq \frac{0.01}{\log n},\left\|s^{-1} \widetilde{\delta}_{s}\right\|_{\infty} \leq \frac{0.02}{\log n},\left\|\bar{x}^{-1} \widetilde{\delta}_{x}\right\|_{\infty} \leq \frac{0.01}{\log n},\left\|x^{-1} \widetilde{\delta}_{x}\right\|_{\infty} \leq \frac{0.02}{\log n},\left\|\mu^{-1} \widetilde{\delta}_{\mu}\right\|_{\infty} \leq \frac{0.02}{\log n}$. With resampling, it always holds.

Proof. We can write $\bar{s}_{i}^{-1} \widetilde{\delta}_{s, i}-\mathbb{E}\left[\bar{s}_{i}^{-1} \widetilde{\delta}_{s, i}\right]=\sum_{j} Y_{j}$ where $Y_{j}$ are independent random variables defined by

$$
Y_{j}=\frac{1}{\sqrt{\bar{x}_{i} \bar{s}_{i}}} \bar{P}_{i, j} \frac{\widetilde{\delta}_{\mu, j}}{\sqrt{\bar{x}_{j} \bar{s}_{j}}}-\frac{1}{\sqrt{\bar{x}_{i} \bar{s}_{i}}} \bar{P}_{i, j} \frac{\delta_{\mu, j}}{\sqrt{\bar{x}_{j} \bar{s}_{j}}}
$$

We bound the sum using Bernstein inequality (Lemma A.1.2). Note that $Y_{j}$ are mean 0 and that Claim 2.4.5 shows that $\sum_{j=1}^{n} \mathbb{E}\left[Y_{j}^{2}\right]=\mathbb{V}\left[\bar{s}_{i}^{-1} \widetilde{\delta}_{s, i}\right] \leq \frac{2 \epsilon^{2}}{k}$. We also need to give an upper bound for $Y_{j}$

$$
\begin{array}{rlr}
\left|Y_{j}\right| & =\left|\frac{1}{\sqrt{\bar{x}_{i} \bar{s}_{i}}} \bar{P}_{i, j}\left(\frac{\widetilde{\delta}_{\mu, j}-\delta_{\mu, j}}{\sqrt{\bar{x}_{j} \bar{s}_{j}}}\right)\right| & \\
& \leq \frac{1.2}{t}\left|\widetilde{\delta}_{\mu, j}-\delta_{\mu, j}\right| & \text { by }\left|\bar{P}_{i, j}\right| \leq 1, x_{i} s_{i} \approx_{1 / 10} t \\
& \leq \frac{1.2}{t}\left|\delta_{\mu, j} / p_{j}\right| & \text { by } \widetilde{\delta}_{\mu, j} \in\left[0, \delta_{\mu, j} / p_{j}\right]  \tag{2.6}\\
& =\frac{1.2}{t} \frac{1}{k} \frac{1}{\left(\frac{\delta_{\mu, i}}{\sum_{l=1}^{n} \delta_{\mu, l}^{2}}+\frac{1}{n \delta_{\mu, i}}\right)} & \text { by }(2.6) \\
& \leq \frac{0.6}{t} \frac{1}{k}\left(n \sum_{l=1}^{n} \delta_{\mu, l}^{2}\right)^{1 / 2} & \text { by } a^{2}+b^{2} \geq 2 a b \\
& \leq \frac{0.6 \epsilon \sqrt{n}}{k} \stackrel{\text { def }}{=} M . & \text { by }\left\|\delta_{\mu}\right\|_{2} \leq \epsilon t
\end{array}
$$

Now, we can apply Bernstein inequality (Lemma A.1.2)

$$
\begin{aligned}
\operatorname{Pr}\left[\left|\sum_{j=1}^{n} Y_{j}\right|>b\right] & \leq 2 \exp \left(-\frac{b^{2} / 2}{\sum_{j=1}^{n} \mathbb{E}\left[Y_{j}^{2}\right]+M b / 3}\right) \\
& \leq 2 \exp \left(-\frac{b^{2} / 2}{2 \epsilon^{2} / k+(0.6 \epsilon \sqrt{n} / k) \cdot b / 3}\right)
\end{aligned}
$$

We choose $b=\frac{0.005}{\log n}$ and use $\epsilon \leq \frac{1}{400 \log n}$ and $n \geq 10$ to get

$$
\operatorname{Pr}\left[\left|\sum_{j=1}^{n} Y_{j}\right| \geq \frac{0.05}{\log n}\right] \leq 2 \exp \left(-\frac{0.003 k}{\epsilon \sqrt{n} \log n}\right) .
$$

Since $\left\|\mathbb{E}\left[\bar{s}_{i}^{-1} \widetilde{\delta}_{s, i}\right]\right\|_{2} \leq 2 \epsilon \leq \frac{0.005}{\log n}$, we have that $\left|\bar{s}_{i}^{-1} \widetilde{\delta}_{s, i}\right| \leq \frac{0.01}{\log n}$ with probability 1 $2 \exp \left(-\frac{0.003 k}{\epsilon \sqrt{n} \log n}\right)$. Taking a union bound, we have that $\left\|\bar{s}^{-1} \widetilde{\delta}_{s}\right\|_{\infty} \leq \frac{0.01}{\log n}$ with probability $1-2 n \exp \left(-\frac{0.003 k}{\epsilon \sqrt{n} \log n}\right)$. Similarly, this holds for the other 3 terms.

Now, the last term follows by

$$
\left|\mu_{i}^{-1} \widetilde{\delta}_{\mu, i}\right|=\left|\bar{x}_{i}^{-1} \bar{s}_{i}^{-1}\left(\bar{x}_{i} \widetilde{\delta}_{s, i}+\bar{s}_{i} \widetilde{\delta}_{x, i}\right)\right|=\left|\bar{s}_{i}^{-1} \widetilde{\delta}_{s, i}\right|+\left|\bar{x}_{i}^{-1} \widetilde{\delta}_{x, i}\right| \leq \frac{0.02}{\log n} .
$$

### 2.4.2.2 Bounding $\mu^{\text {new }}-\mu$

Lemma 2.4.7. Under the Assumption 2.4.1, the vector $\mu_{i}^{\text {new }} \stackrel{\text { def }}{=}\left(x_{i}+\widetilde{\delta}_{x, i}\right)\left(s_{i}+\widetilde{\delta}_{s, i}\right)$ satisfies

1. $\left\|\mathbb{E}\left[\mu^{-1}\left(\mu^{\text {new }}-\mu-\widetilde{\delta}_{\mu}\right)\right]\right\|_{2} \leq 10 \epsilon_{m p} \cdot \epsilon$ and $\left\|\mathbb{E}\left[\mu^{-1}\left(\mu^{\text {new }}-\mu\right)\right]\right\|_{2} \leq 5 \epsilon$.
2. $\mathbb{V}\left[\mu_{i}^{-1} \mu_{i}^{\text {new }}\right] \leq 50 \epsilon^{2} / k$ for all $i$.
3. $\left\|\mu^{-1}\left(\mu^{\mathrm{new}}-\mu\right)\right\|_{\infty} \leq \frac{0.021}{\log n}$.

Claim 2.4.8 (Part 1 of Lemma 2.4.7).

$$
\left\|\mathbb{E}\left[\mu^{-1}\left(\mu^{\text {new }}-\mu-\widetilde{\delta}_{\mu}\right)\right]\right\|_{2} \leq 10 \epsilon_{m p} \cdot \epsilon, \text { and }\left\|\mathbb{E}\left[\mu^{-1}\left(\mu^{\text {new }}-\mu\right)\right]\right\|_{2} \leq 5 \epsilon
$$

Proof.

$$
\mu^{\mathrm{new}}=\left(x+\widetilde{\delta}_{x}\right)\left(s+\widetilde{\delta}_{s}\right)=\mu+x \widetilde{\delta}_{s}+s \widetilde{\delta}_{x}+\widetilde{\delta}_{x} \widetilde{\delta}_{s}=\mu+\underbrace{\overline{\delta_{\delta}}+\bar{s} \widetilde{\delta}_{x}}_{\widetilde{\delta}_{\mu}}+\underbrace{(x-\bar{x}) \widetilde{\delta}_{s}+(s-\bar{s}) \widetilde{\delta}_{x}+\widetilde{\delta}_{x} \widetilde{\delta}_{s}}_{\epsilon_{\mu}} .
$$

Taking the expectation on both sides, we have

$$
\mathbb{E}\left[\mu^{\text {new }}-\mu-\widetilde{\delta}_{\mu}\right]=(x-\bar{x}) \mathbb{E}\left[\widetilde{\delta}_{s}\right]+(s-\bar{s}) \mathbb{E}\left[\widetilde{\delta}_{x}\right]+\mathbb{E}\left[\widetilde{\delta}_{x} \widetilde{\delta}_{s}\right] .
$$

Hence, we have that

$$
\begin{align*}
& \left\|\mu^{-1} \mathbb{E}\left[\mu^{\text {new }}-\mu-\widetilde{\delta}_{\mu}\right]\right\|_{2} \\
\leq & \left\|\mu^{-1}(x-\bar{x}) s \cdot s^{-1} \mathbb{E}\left[\widetilde{\delta}_{s}\right]\right\|_{2}+\left\|\mu^{-1}(s-\bar{s}) x \cdot x^{-1} \mathbb{E}\left[\widetilde{\delta}_{x}\right]\right\|_{2}+\left\|\mu^{-1} \mathbb{E}\left[\widetilde{\delta}_{x} \widetilde{\delta}_{s}\right]\right\|_{2} \\
\leq & \left\|\mu^{-1}(x-\bar{x}) s\right\|_{\infty} \cdot\left\|s^{-1} \mathbb{E}\left[\widetilde{\delta}_{s}\right]\right\|_{2}+\left\|\mu^{-1}(s-\bar{s}) x\right\|_{\infty} \cdot\left\|x^{-1} \mathbb{E}\left[\widetilde{\delta}_{x}\right]\right\|_{2}+\left\|\mu^{-1} \mathbb{E}\left[\widetilde{\delta}_{x} \widetilde{\delta}_{s}\right]\right\|_{2} \\
\leq & \epsilon_{m p} \cdot\left\|s^{-1} \mathbb{E}\left[\widetilde{\delta}_{s}\right]\right\|_{2}+\epsilon_{m p} \cdot\left\|x^{-1} \mathbb{E}\left[\widetilde{\delta}_{x}\right]\right\|_{2}+\left\|\mu^{-1} \mathbb{E}\left[\widetilde{\delta}_{x} \widetilde{\delta}_{s}\right]\right\|_{2} \\
\leq & 4 \epsilon_{m p} \cdot \epsilon+\left\|\mu^{-1} \mathbb{E}\left[\widetilde{\delta}_{x} \widetilde{\delta}_{s}\right]\right\|_{2} \tag{2.15}
\end{align*}
$$

where the first step follows by triangle inequality, the second step follows by $\|a b\|_{2} \leq\|a\|_{\infty}$. $\|b\|_{2}$, the third step follows by $\left\|\mu^{-1}(x-\bar{x}) s\right\|_{\infty} \leq \epsilon_{m p}$ and $\left\|\mu^{-1}(s-\bar{s}) x\right\|_{\infty} \leq \epsilon_{m p}$ (since $\bar{x} \approx_{\epsilon_{m p}} x, \bar{s} \approx_{\epsilon_{m p}} s$, the last step follows by $\left\|\mathbb{E}\left[s^{-1} \widetilde{\delta}_{s}\right]\right\|_{2} \leq 2 \epsilon$ and $\left\|\mathbb{E}\left[x^{-1} \widetilde{\delta}_{x}\right]\right\|_{2} \leq 2 \epsilon$ (Part 1 of Lemma 2.4.3).

To bound the last term, using $\mathbb{E}\left[\widetilde{\delta}_{s}\right]=\delta_{s}$ and $\mathbb{E}\left[\widetilde{\delta}_{x}\right]=\delta_{x}$, we note that

$$
\mathbb{E}\left[\widetilde{\delta}_{x, i} \widetilde{\delta}_{s, i}\right]=\delta_{x, i} \delta_{s, i}+\mathbb{E}\left[\left(\widetilde{\delta}_{x, i}-\delta_{x, i}\right)\left(\widetilde{\delta}_{s, i}-\delta_{s, i}\right)\right]
$$

Hence, we have

$$
\begin{align*}
\left\|\mu^{-1} \mathbb{E}\left[\widetilde{\delta}_{x} \widetilde{\delta}_{s}\right]\right\|_{2} & \leq\left\|\mu^{-1} \delta_{x} \delta_{s}\right\|_{2}+\left(\sum_{i=1}^{n}\left(\mathbb{E}\left[x_{i}^{-1}\left(\widetilde{\delta}_{x, i}-\delta_{x, i}\right) \cdot s_{i}^{-1}\left(\widetilde{\delta}_{s, i}-\delta_{s, i}\right)\right]\right)^{2}\right)^{1 / 2} \\
& \leq 4 \epsilon^{2}+\frac{1}{2}\left(\sum_{i=1}^{n}\left(\mathbb{V}\left[x_{i}^{-1} \widetilde{\delta}_{x, i}\right]+\mathbb{V}\left[s_{i}^{-1} \widetilde{\delta}_{s, i}\right]\right)^{2}\right)^{1 / 2} \\
& \leq 4 \epsilon^{2}+\frac{1}{2}\left(\sum_{i=1}^{n} 2\left(\mathbb{V}\left[x_{i}^{-1} \widetilde{\delta}_{x, i}\right]\right)^{2}+2\left(\mathbb{V}\left[s_{i}^{-1} \widetilde{\delta}_{s, i}\right]\right)^{2}\right)^{1 / 2} \\
& \leq 4 \epsilon^{2}+2 \sqrt{n \cdot \epsilon^{4} / k^{2}} \leq 4 \epsilon^{2}+2 \epsilon \cdot \epsilon_{m p} \leq 6 \epsilon \cdot \epsilon_{m p} \tag{2.16}
\end{align*}
$$

where the second step follows by $\left\|\mu^{-1} \delta_{x} \delta_{s}\right\|_{2} \leq\left\|x^{-1} \delta_{x}\right\|_{2} \cdot\left\|s^{-1} \delta_{s}\right\|_{2} \leq 4 \epsilon^{2}$ (Part 1 of Lemma 2.4.3) and $2 a b \leq a^{2}+b^{2}$, the third step follows by $(a+b)^{2} \leq 2 a^{2}+2 b^{2}$, the fourth step follows by
$\mathbb{V}\left[x_{i}^{-1} \widetilde{\delta}_{x, i}\right] \leq 2 \epsilon^{2} / k$ and $\mathbb{V}\left[s_{i}^{-1} \widetilde{\delta}_{s, i}\right] \leq 2 \epsilon^{2} / k$ (Part 2 of Lemma 2.4.3), the last step follows by $k \geq \frac{\epsilon \sqrt{n}}{\epsilon_{m p}}$.

Combining (2.15) and (2.16), we have that

$$
\left\|\mu^{-1}\left(\mathbb{E}\left[\mu^{\text {new }}-\mu-\widetilde{\delta}_{\mu}\right]\right)\right\|_{2} \leq 4 \epsilon_{m p} \cdot \epsilon+\left\|\mu^{-1} \mathbb{E}\left[\widetilde{\delta}_{x} \widetilde{\delta}_{s}\right]\right\|_{2} \leq 10 \epsilon_{m p} \cdot \epsilon
$$

where we used $\epsilon \leq \epsilon_{m p}$.
From Part 1 of Lemma 2.4.3, we know that $\left\|\mu^{-1} \mathbb{E}\left[\widetilde{\delta}_{\mu}\right]\right\|_{2} \leq 4 \epsilon$. Thus using triangle inequality, we know

$$
\left\|\mu^{-1}\left(\mathbb{E}\left[\mu^{\text {new }}-\mu\right]\right)\right\|_{2} \leq 10 \epsilon_{m p} \cdot \epsilon+4 \epsilon \leq 5 \epsilon
$$

Claim 2.4.9 (Part 2 of Lemma 2.4.7). $\mathbb{V}\left[\mu_{i}^{-1} \mu_{i}^{\text {new }}\right] \leq 50 \epsilon^{2} / k$ for all $i$.

Proof. Recall that

$$
\mu^{\mathrm{new}}=\mu+\widetilde{\delta}_{\mu}+(x-\bar{x}) \widetilde{\delta}_{s}+(s-\bar{s}) \widetilde{\delta}_{x}+\widetilde{\delta}_{x} \widetilde{\delta}_{s}
$$

We can upper bound the variance of $\mu_{i}^{-1} \mu_{i}^{\text {new }}$,

$$
\begin{aligned}
\mathbb{V}\left[\mu_{i}^{-1} \mu_{i}^{\mathrm{new}}\right] & \leq 4 \mathbb{V}\left[\mu_{i}^{-1} \widetilde{\delta}_{\mu, i}\right]+4 \mathbb{V}\left[\mu_{i}^{-1}\left(x_{i}-\bar{x}_{i}\right) \widetilde{\delta}_{s, i}\right]+4 \mathbb{V}\left[\mu_{i}^{-1}\left(s_{i}-\bar{s}_{i}\right) \widetilde{\delta}_{x, i}\right]+4 \mathbb{V}\left[\mu_{i}^{-1} \widetilde{\delta}_{x, i} \widetilde{\delta}_{s, i}\right] \\
& \leq 32 \frac{\epsilon^{2}}{k}+4 \frac{\epsilon^{2}}{k}+4 \frac{\epsilon^{2}}{k}+\mathbb{V}\left[\mu_{i}^{-1} \widetilde{\delta}_{x, i} \widetilde{\delta}_{s, i}\right] \\
& =40 \frac{\epsilon^{2}}{k}+\mathbb{V}\left[x_{i}^{-1} \widetilde{\delta}_{x, i} \cdot s_{i}^{-1} \widetilde{\delta}_{s, i}\right] \\
& \leq 40 \frac{\epsilon^{2}}{k}+2 \operatorname{Sup}\left[\left(x_{i}^{-1} \widetilde{\delta}_{x, i}\right)^{2}\right] \cdot \mathbb{V}\left[s_{i}^{-1} \widetilde{\delta}_{s, i}\right]+2 \operatorname{Sup}\left[\left(y_{i}^{-1} \widetilde{\delta}_{y, i}\right)^{2}\right] \cdot \mathbb{V}\left[x_{i}^{-1} \widetilde{\delta}_{x, i}\right] \\
& \leq 40 \frac{\epsilon^{2}}{k}+2 \cdot\left(\frac{0.02}{\log n}\right)^{2} \cdot \frac{\epsilon^{2}}{k}+2 \cdot\left(\frac{0.02}{\log n}\right)^{2} \cdot \frac{\epsilon^{2}}{k} \leq 50 \frac{\epsilon^{2}}{k} .
\end{aligned}
$$

where the second step follows by $\mathbb{V}\left[\mu_{i}^{-1} \widetilde{\delta}_{\mu, i}\right] \leq 8 \epsilon^{2} / k$ (Part 2 of Lemma 2.4.3),

$$
\mathbb{V}\left[\mu_{i}^{-1}\left(x_{i}-\bar{x}_{i}\right) \widetilde{\delta}_{s, i}\right]=\mathbb{V}\left[x_{i}^{-1}\left(x_{i}-\bar{x}_{i}\right) s_{i}^{-1} \widetilde{\delta}_{s, i}\right] \leq 2 \epsilon_{m p}^{2} \mathbb{V}\left[s_{i}^{-1} \widetilde{\delta}_{s, i}\right] \leq \epsilon^{2} / k .
$$

and a similar inequality $\mathbb{V}\left[\mu_{i}^{-1}\left(s_{i}-\bar{s}_{i}\right) \widetilde{\delta}_{x, i}\right] \leq \epsilon^{2} / k$, the third step follows by $\mu=x s$, the fourth step follows by $\mathbb{V}[x y] \leq 2 \boldsymbol{S u p}\left[x^{2}\right] \mathbb{V}[y]+2 \boldsymbol{S u p}\left[y^{2}\right] \mathbb{V}[x]$ (Lemma 2.6.1) with Sup denoting the deterministic maximum of the random variable, the fifth step follows by $\mathbb{V}\left[s_{i}^{-1} \widetilde{\delta}_{s, i}\right] \leq 2 \epsilon^{2} / k$ and $\mathbb{V}\left[x_{i}^{-1} \widetilde{\delta}_{x, i}\right] \leq 2 \epsilon^{2} / k$ (Part 2 of Lemma 2.4.3).

Claim 2.4.10 (Part 3 of Lemma 2.4.7). $\left\|\mu^{-1}\left(\mu^{\text {new }}-\mu\right)\right\|_{\infty} \leq \frac{0.021}{\log n}$.

Proof. We again note that

$$
\mu^{\mathrm{new}}=\mu+\widetilde{\delta}_{\mu}+(x-\bar{x}) \widetilde{\delta}_{s}+(s-\bar{s}) \widetilde{\delta}_{x}+\widetilde{\delta}_{x} \widetilde{\delta}_{s} .
$$

Hence, we have

$$
\begin{aligned}
& \left|\mu_{i}^{-1}\left(\mu_{i}^{\mathrm{new}}-\mu_{i}-\widetilde{\delta}_{\mu, i}\right)\right| \\
\leq & \left|(x-\bar{x})_{i} \mu_{i}^{-1} \widetilde{\delta}_{s, i}\right|+\left|(s-\bar{s})_{i} \mu_{i}^{-1} \widetilde{\delta}_{x, i}\right|+\left|\mu_{i}^{-1} \widetilde{\delta}_{x, i} \widetilde{\delta}_{s, i}\right| \\
= & \left|(x-\bar{x})_{i} x_{i}^{-1}\right| \cdot\left|s_{i}^{-1} \widetilde{\delta}_{s, i}\right|+\left|(s-\bar{s})_{i} s_{i}^{-1}\right| \cdot\left|x_{i}^{-1} \widetilde{\delta}_{x, i}\right|+\left|x_{i}^{-1} \widetilde{\delta}_{x, i}\right| \cdot\left|s_{i}^{-1} \widetilde{\delta}_{s, i}\right| \\
\leq & \epsilon_{m p}\left|s_{i}^{-1} \widetilde{\delta}_{s, i}\right|+\epsilon_{m p}\left|x_{i}^{-1} \widetilde{\delta}_{x, i}\right|+\left|s_{i}^{-1} \widetilde{\delta}_{s, i}\right|\left|x_{i}^{-1} \widetilde{\delta}_{x, i}\right| \\
\leq & \epsilon_{m p} \cdot \frac{0.2}{\log n}+\epsilon_{m p} \cdot \frac{0.02}{\log n}+\left(\frac{0.02}{\log n}\right)^{2} \\
\leq & \frac{1}{1000 \log n},
\end{aligned}
$$

where the first step follows by triangle inequality, the second step follows by $\mu_{i}=x_{i} s_{i}$, the third step follows by $x \approx_{\epsilon_{m p}} \bar{x}$ and $s \approx_{\epsilon_{m p}} \bar{s}$, the fifth step follows by $\left|s_{i}^{-1} \widetilde{\delta}_{s, i}\right| \leq \frac{0.02}{\log n}$ and $\left|x_{i}^{-1} \widetilde{\delta}_{x, i}\right| \leq \frac{0.02}{\log n}$ (Part 3 of Lemma 2.4.3).

Since we know that $\left|\mu_{i}^{-1} \widetilde{\delta}_{\mu, i}\right| \leq \frac{0.02}{\log n}$ (Part 3 of Lemma 2.4.3), we have

$$
\left|\mu_{i}^{-1}\left(\mu_{i}^{\mathrm{new}}-\mu_{i}\right)\right| \leq \frac{1}{1000 \log n}+\frac{0.02}{\log n} \leq \frac{0.021}{\log n}
$$

### 2.4.3 Stochastic central path

Now, we are ready to prove $x_{i} s_{i} \approx_{0.1} t$ during the whole algorithm. As explained in the proof outline (see Section 2.4.1), we will prove this bound by analyzing the potential $\Phi_{\lambda}(\mu / t-1)$ where $\Phi_{\lambda}(r)=\sum_{i=1}^{n} \cosh \left(\lambda r_{i}\right)$.

First, we give some basic properties of $\Phi_{\lambda}$.
Lemma 2.4.11 (Basic properties of potential function). Let $\Phi_{\lambda}(r)=\sum_{i=1}^{n} \cosh \left(\lambda r_{i}\right)$ for some $\lambda>0$. For any vector $r \in \mathbb{R}^{n}$,

1. For any vector $\|v\|_{\infty} \leq 1 / \lambda$, we have that

$$
\Phi_{\lambda}(r+v) \leq \Phi_{\lambda}(r)+\left\langle\nabla \Phi_{\lambda}(r), v\right\rangle+2\|v\|_{\nabla^{2} \Phi_{\lambda}(r)}^{2}
$$

2. $\left\|\nabla \Phi_{\lambda}(r)\right\|_{2} \geq \frac{\lambda}{\sqrt{n}}\left(\Phi_{\lambda}(r)-n\right)$.
3. $\left(\sum_{i=1}^{n} \lambda^{2} \cosh ^{2}\left(\lambda r_{i}\right)\right)^{1 / 2} \leq \lambda \sqrt{n}+\left\|\nabla \Phi_{\lambda}(r)\right\|_{2}$.

Proof. For each $i \in[n]$, we use $r_{i}$ to denote the $i$-th coordinate of vector $r$.
Proof of Part 1. We have that

$$
\cosh \left(\lambda\left(r_{i}+v_{i}\right)\right)=\cosh \left(\lambda r_{i}\right)+\lambda \sinh \left(\lambda r_{i}\right) v_{i}+\frac{\lambda^{2}}{2} \cosh \left(\zeta_{i}\right) v_{i}^{2}
$$

where $\zeta_{i}$ is between $\lambda r_{i}$ and $\lambda\left(r_{i}+v_{i}\right)$. By definition of cosh and the assumption that $\|v\|_{\infty} \leq \frac{1}{2 \lambda}$, we have that

$$
\cosh \left(\zeta_{i}\right)=\frac{1}{2} \exp \left(\zeta_{i}\right)+\frac{1}{2} \exp \left(-\zeta_{i}\right) \leq \exp (1) \cdot \frac{1}{2}\left(\exp \left(\lambda r_{i}\right)+\exp \left(-\lambda r_{i}\right)\right) \leq 3 \cosh \left(\lambda r_{i}\right)
$$

Hence, we have

$$
\cosh \left(\lambda\left(r_{i}+v_{i}\right)\right) \leq \cosh \left(\lambda r_{i}\right)+\lambda \sinh \left(\lambda r_{i}\right) v_{i}+2 \lambda^{2} \cosh \left(\lambda r_{i}\right) v_{i}^{2}
$$

Summing over all the coordinates gives

$$
\begin{aligned}
& \sum_{i=1}^{n} \cosh \left(\lambda\left(r_{i}+v_{i}\right)\right) \leq \sum_{i=1}^{n}\left[\cosh \left(\lambda r_{i}\right)+2 \lambda \sinh \left(\lambda r_{i}\right) v_{i}+\lambda^{2} \cosh \left(\lambda r_{i}\right) v_{i}^{2}\right] \\
\Longrightarrow & \Phi_{\lambda}(r+v) \leq \Phi_{\lambda}(r)+\left\langle\nabla \Phi_{\lambda}(r), v\right\rangle+2\|v\|_{\nabla^{2} \Phi_{\lambda}(r) .}^{2}
\end{aligned}
$$

Proof of Part 2. Since $\Phi_{\lambda}(r)=\sum_{i=1}^{n} \cosh \left(\lambda r_{i}\right)$, then

$$
\nabla \Phi_{\lambda}(r)=\left[\begin{array}{llll}
\lambda \sinh \left(\lambda r_{1}\right) & \lambda \sinh \left(\lambda r_{2}\right) & \cdots & \lambda \sinh \left(\lambda r_{n}\right)
\end{array}\right]^{\top} .
$$

Thus, we can lower bound $\left\|\nabla \Phi_{\lambda}(r)\right\|_{2}$ in the following way,

$$
\begin{array}{rlr}
\left\|\nabla \Phi_{\lambda}(r)\right\|_{2} & =\left(\sum_{i=1}^{n} \lambda^{2} \sinh ^{2}\left(\lambda r_{i}\right)\right)^{1 / 2} & \\
& =\left(\sum_{i=1}^{n} \lambda^{2}\left(\cosh ^{2}\left(\lambda r_{i}\right)-1\right)\right)^{1 / 2} & \text { by } \cosh ^{2}(y)-\sinh ^{2}(y)=1, \forall y \\
& \geq \frac{\lambda}{\sqrt{n}} \sum_{i=1}^{n} \sqrt{\cosh ^{2}\left(\lambda r_{i}\right)-1} & \text { by }\|\cdot\|_{2} \geq \frac{1}{\sqrt{n}}\|\cdot\|_{1} \\
& \geq \frac{\lambda}{\sqrt{n}} \sum_{i=1}^{n}\left(\cosh \left(\lambda r_{i}\right)-1\right) & \\
& =\frac{\lambda}{\sqrt{n}}\left(\Phi_{\lambda}(r)-n\right) . & \text { by } \cosh \left(\lambda r_{i}\right) \geq 1
\end{array}
$$

Proof of Part 3. We have

$$
\begin{aligned}
\left(\sum_{i=1}^{n} \lambda^{2} \cosh ^{2}\left(\lambda r_{i}\right)\right)^{1 / 2} & =\left(\sum_{i=1}^{n} \lambda^{2}+\lambda^{2} \sinh ^{2}\left(\lambda r_{i}\right)\right)^{1 / 2} \\
& \leq\left(n \lambda^{2}\right)^{1 / 2}+\left(\sum_{i=1}^{n} \lambda^{2} \sinh ^{2}\left(\lambda r_{i}\right)\right)^{1 / 2} \\
& =\lambda \sqrt{n}+\left\|\nabla \Phi_{\lambda}(r)\right\|_{2}
\end{aligned}
$$

where the first step follows from $\cosh ^{2}(y)-\sinh ^{2}(y)=1, \forall y$.

The following lemma shows that the potential $\Phi$ is decreasing in expectation when $\Phi$ is large.

Lemma 2.4.12. Under the Assumption 2.4.1, we have

$$
\mathbb{E}\left[\Phi_{\lambda}\left(\frac{\mu^{\mathrm{new}}}{t^{\mathrm{new}}}-1\right)\right] \leq \Phi_{\lambda}\left(\frac{\mu}{t}-1\right)-\frac{\lambda \epsilon}{15 \sqrt{n}}\left(\Phi_{\lambda}\left(\frac{\mu}{t}-1\right)-10 n\right) .
$$

Proof. Let $\epsilon_{\mu}=\mu^{\text {new }}-\mu-\widetilde{\delta}_{\mu}$. From the definition, we have

$$
\mu^{\text {new }}-t^{\text {new }}=\mu+\widetilde{\delta}_{\mu}+\epsilon_{\mu}-t^{\text {new }}
$$

which implies

$$
\begin{align*}
\frac{\mu^{\text {new }}}{t^{\text {new }}}-1 & =\frac{\mu}{t^{\text {new }}}+\frac{1}{t^{\text {new }}}\left(\widetilde{\delta}_{\mu}+\epsilon_{\mu}\right)-1 \\
& =\frac{\mu}{t} \frac{t}{t^{\text {new }}}+\frac{1}{t^{\text {new }}}\left(\widetilde{\delta}_{\mu}+\epsilon_{\mu}\right)-1 \\
& =\frac{\mu}{t}+\frac{\mu}{t}\left(\frac{t}{t^{\text {new }}}-1\right)+\frac{1}{t^{\text {new }}}\left(\widetilde{\delta}_{\mu}+\epsilon_{\mu}\right)-1 \\
& =\frac{\mu}{t}-1+\underbrace{\frac{\mu}{t}\left(\frac{t}{t^{\text {new }}}-1\right)+\frac{1}{t^{\text {new }}}\left(\widetilde{\delta}_{\mu}+\epsilon_{\mu}\right)}_{v} \tag{2.17}
\end{align*}
$$

To apply Lemma 2.4.11 with $r=\mu / t-1$ and $r+v=\mu^{\text {new }} / t^{\text {new }}-1$, we first compute the expectation of $v$

$$
\begin{align*}
\mathbb{E}[v] & =\frac{\mu}{t}\left(\frac{t}{t^{\text {new }}}-1\right)+\frac{1}{t^{\text {new }}}\left(\mathbb{E}\left[\widetilde{\delta}_{\mu}\right]+\mathbb{E}\left[\epsilon_{\mu}\right]\right) \\
& =\frac{\mu}{t}\left(\frac{t}{t^{\text {new }}}-1\right)+\frac{1}{t^{\text {new }}}\left(\delta_{\mu}+\mathbb{E}\left[\epsilon_{\mu}\right]\right) \\
& =\frac{\mu}{t}\left(\frac{t}{t^{\text {new }}}-1\right)+\frac{1}{t^{\text {new }}}\left(\left(\left(\frac{t^{\text {new }}}{t}-1\right) \mu-\frac{\epsilon}{2} t^{\text {new }} \frac{\nabla \Phi_{\lambda}(\mu / t-1)}{\left\|\nabla \Phi_{\lambda}(\mu / t-1)\right\|_{2}}\right)+\mathbb{E}\left[\epsilon_{\mu}\right]\right) \\
& =-\frac{\epsilon}{2} \frac{\nabla \Phi_{\lambda}(\mu / t-1)}{\left\|\nabla \Phi_{\lambda}(\mu / t-1)\right\|_{2}}+\frac{1}{t^{\text {new }}} \mathbb{E}\left[\epsilon_{\mu}\right], \tag{2.18}
\end{align*}
$$

where the third step follows by definition of $\delta_{\mu}$.
Next, we bound the $\|v\|_{\infty}$ as follows

$$
\begin{aligned}
\|v\|_{\infty} & \leq\left\|\frac{\mu}{t}\left(\frac{t}{t^{\text {new }}}-1\right)\right\|_{\infty}+\left\|\frac{1}{t^{\text {new }}}\left(\widetilde{\delta}_{\mu}+\epsilon_{\mu}\right)\right\|_{\infty} \leq \frac{\epsilon}{\sqrt{n}}+\frac{\left\|\mu^{-1}\left(\mu^{\text {new }}-\mu\right)\right\|_{\infty}}{0.9} \\
& \leq \frac{\epsilon}{\sqrt{n}}+\frac{0.021}{0.9 \log n} \leq \frac{1}{\lambda} .
\end{aligned}
$$

where we used Part 3 of Lemma 2.4.7 and $\epsilon \leq \frac{1}{400 \log n}$.
Since $\|v\|_{\infty} \leq \frac{1}{\lambda}$, we can apply Part 1 of Lemma 2.4.11 and get

$$
\begin{aligned}
& \mathbb{E}\left[\Phi_{\lambda}(\mu / t+v-1)\right] \\
\leq & \Phi_{\lambda}(\mu / t-1)+\left\langle\nabla \Phi_{\lambda}(\mu / t-1), \mathbb{E}[v]\right\rangle+2 \mathbb{E}\left[\|v\|_{\nabla^{2} \Phi_{\lambda}(\mu / t+v-1)}^{2}\right] \\
= & \Phi_{\lambda}(\mu / t-1)-\frac{\epsilon}{2}\left\|\nabla \Phi_{\lambda}(\mu / t-1)\right\|_{2}+\frac{t}{t^{\text {new }}}\left\langle\nabla \Phi_{\lambda}(\mu / t-1), \mathbb{E}\left[t^{-1} \epsilon_{\mu}\right]\right\rangle+2 \mathbb{E}\left[\|v\|_{\nabla^{2} \Phi_{\lambda}(\mu / t-1)}^{2}\right] \\
\leq & \Phi_{\lambda}(\mu / t-1)-\frac{\epsilon}{2}\left\|\nabla \Phi_{\lambda}(\mu / t-1)\right\|_{2}+\frac{t}{t^{n \mathrm{ew}}}\left\|\nabla \Phi_{\lambda}(\mu / t-1)\right\|_{2} \cdot\left\|\mathbb{E}\left[t^{-1} \epsilon_{\mu}\right]\right\|_{2}+2 \mathbb{E}\left[\|[v]\|_{\nabla^{2} \Phi_{\lambda}(\mu / t-1)}^{2}\right] \\
\leq & \Phi_{\lambda}(\mu / t-1)-\frac{\epsilon}{2}\left\|\nabla \Phi_{\lambda}(\mu / t-1)\right\|_{2}+10 \epsilon_{m p} \cdot \epsilon\left\|\nabla \Phi_{\lambda}(\mu / t-1)\right\|_{2}+2 \mathbb{E}\left[\|v\|_{\nabla^{2} \Phi_{\lambda}(\mu / t-1)}^{2}\right],
\end{aligned}
$$

where the second step follows by substituting $\mathbb{E}[v]$ by (2.18), the third step follows by $\langle a, b\rangle \leq$ $\|a\|_{2} \cdot\|b\|_{2}$, the fourth step follows by $\left\|\mathbb{E}\left[t^{-1} \epsilon_{\mu}\right]\right\|_{2} \leq 10 \epsilon_{m p} \cdot \epsilon$ (from Part 1 of Lemma 2.4.7 and $\mu \approx_{0.1} t$.

We still need to bound $\mathbb{E}\left[\|v\|_{\nabla^{2} \Phi_{\lambda}(\mu / t-1)}^{2}\right]$. Before bounding it, we first bound $\mathbb{E}\left[v_{i}^{2}\right]$,

$$
\begin{align*}
\mathbb{E}\left[v_{i}^{2}\right] & \leq 2 \mathbb{E}\left[\left(\frac{\mu_{i}}{t}\left(\frac{t}{t^{\text {new }}}-1\right)\right)^{2}\right]+2 \mathbb{E}\left[\left(\frac{1}{t^{\text {new }}}\left(\widetilde{\delta}_{\mu, i}+\widehat{\delta}_{\mu, i}\right)\right)^{2}\right] \\
& \leq \epsilon^{2} / n+2.5 \mathbb{E}\left[\left(\left(\mu_{i}^{\text {new }}-\mu_{i}\right) / \mu_{i}\right)^{2}\right] \\
& =\epsilon^{2} / n+2.5 \mathbb{V}\left[\left(\mu_{i}^{\text {new }}-\mu_{i}\right) / \mu_{i}\right]+2.5\left(\mathbb{E}\left[\left(\mu_{i}^{\text {new }}-\mu_{i}\right) / \mu_{i}\right]\right)^{2} \\
& \leq \epsilon^{2} / n+125 \epsilon^{2} / k+2.5\left(\mathbb{E}\left[\left(\mu_{i}^{\text {new }}-\mu_{i}\right) / \mu_{i}\right]\right)^{2} \\
& \leq 126 \epsilon^{2} / k+3\left(\mathbb{E}\left[\left(\mu_{i}^{\text {new }}-\mu_{i}\right) / \mu_{i}\right]\right)^{2}, \tag{2.19}
\end{align*}
$$

where the first step follows by definition of $v$ (see (2.17)), the second step follows by $\mu \approx_{0.1} t$ and $\left(t / t^{\text {new }}-1\right)^{2} \leq \epsilon^{2} /(4 n)$, the third step follows by $\mathbb{E}\left[x^{2}\right]=\mathbb{V}[x]+(\mathbb{E}[x])^{2}$, the fourth step follows by Part 2 of Lemma 2.4.7, and the last step follows by $n \geq k$.

Now, we are ready to bound $\mathbb{E}\left[\|v\|_{\nabla^{2} \Phi_{\lambda}(\mu / t-1)}^{2}\right]$

$$
\begin{aligned}
& \mathbb{E}\left[\|v\|_{\nabla^{2} \Phi_{\lambda}(\mu / t-1)}^{2}\right] \\
= & \lambda^{2} \sum_{i=1}^{n} \mathbb{E}\left[\Phi_{\lambda}(\mu / t-1)_{i} v_{i}^{2}\right] \\
\leq & \lambda^{2} \sum_{i=1}^{n} \Phi_{\lambda}(\mu / t-1)_{i} \cdot\left(126 \epsilon^{2} / k+3\left(\mathbb{E}\left[\left(\mu_{i}^{\mathrm{new}}-\mu_{i}\right) / \mu_{i}\right]\right)^{2}\right) \\
= & 126 \frac{\lambda^{2} \epsilon^{2}}{k} \Phi_{\lambda}(\mu / t-1)+3 \lambda^{2} \sum_{i=1}^{n} \Phi_{\lambda}(\mu / t-1)_{i} \cdot\left(\mathbb{E}\left[\left(\mu_{i}^{\mathrm{new}}-\mu_{i}\right) / \mu_{i}\right]\right)^{2} \\
\leq & 126 \frac{\lambda^{2} \epsilon^{2}}{k} \Phi_{\lambda}(\mu / t-1)+3 \lambda\left(\sum_{i=1}^{n} \lambda^{2} \Phi_{\lambda}(\mu / t-1)_{i}^{2}\right)^{1 / 2} \cdot\left\|\mathbb{E}\left[\mu^{-1}\left(\mu^{\mathrm{new}}-\mu\right)\right]\right\|_{4}^{2} \\
\leq & 126 \frac{\lambda^{2} \epsilon^{2}}{k} \Phi_{\lambda}(\mu / t-1)+3 \lambda\left(\lambda \sqrt{n}+\left\|\nabla \Phi_{\lambda}(\mu / t-1)\right\|_{2}\right) \cdot(5 \epsilon)^{2},
\end{aligned}
$$

where the first step follows by defining $\Phi_{\lambda}(x)_{i}=\cosh \left(\lambda x_{i}\right)$, the second step follows from (2.19), the fourth step follows from Cauchy-Schwarz inequality, the fifth step follows from Part 3 of Lemma 2.4.11 and the fact that $\left\|\mathbb{E}\left[\mu^{-1}\left(\mu^{\text {new }}-\mu\right)\right]\right\|_{4}^{2} \leq\left\|\mathbb{E}\left[\mu^{-1}\left(\mu^{\text {new }}-\mu\right)\right]\right\|_{2}^{2} \leq(5 \epsilon)^{2}$ (Lemma 2.4.7).

Then,

$$
\begin{aligned}
& \mathbb{E}\left[\Phi_{\lambda}(\mu / t+v-1)\right] \\
\leq & \Phi_{\lambda}(\mu / t-1)-\left(\frac{\epsilon}{2}-10 \epsilon_{m p} \cdot \epsilon\right)\left\|\nabla \Phi_{\lambda}(\mu / t-1)\right\|_{2}+252 \frac{\lambda^{2} \epsilon^{2}}{k} \Phi_{\lambda}(\mu / t-1) \\
& +150 \lambda^{2} \epsilon^{2} \sqrt{n}+150 \lambda \epsilon^{2}\left\|\Phi_{\lambda}(\mu / t-1)\right\|_{2} \\
\leq & \Phi_{\lambda}(\mu / t-1)-\frac{\epsilon}{3}\left\|\nabla \Phi_{\lambda}(\mu / t-1)\right\|_{2}+252 \frac{\lambda^{2} \epsilon^{2}}{k} \Phi_{\lambda}(\mu / t-1)+150 \lambda^{2} \epsilon^{2} \sqrt{n} \\
\leq & \Phi_{\lambda}(\mu / t-1)-\frac{\lambda \epsilon}{3 \sqrt{n}}\left(\Phi_{\lambda}(\mu / t-1)-n\right)+252 \frac{\lambda^{2} \epsilon^{2}}{k} \Phi_{\lambda}(\mu / t-1)+150 \lambda^{2} \epsilon^{2} \sqrt{n} \\
\leq & \Phi_{\lambda}(\mu / t-1)-\frac{\lambda \epsilon}{3 \sqrt{n}}\left(\Phi_{\lambda}(\mu / t-1) / 5-2 n\right),
\end{aligned}
$$

where the second step follows from $1000 \lambda \epsilon \leq 1$ and $1000 \epsilon_{m p} \leq 1$, the third step follows from Part 2 of Lemma 2.4.11, and the last step follows from $1000 \lambda \epsilon_{m p} \leq \log n$ and $k \geq \frac{\sqrt{n} \epsilon \log n}{\epsilon_{m p}}$.

As a corollary, we have the following:

Lemma 2.4.13. During the MAIN algorithm, Assumption 2.4.1 is always satisfied. Furthermore, the CLASSICALSTEP happens with probability $O\left(\frac{1}{n^{2}}\right)$ each step.

Proof. The second and the fourth assumptions simply follow from the choice of $\epsilon_{m p}$ and $k$.
Let $\Phi^{(k)}$ be the potential at the $k$-th iteration of the Main. The ClassicalStep ensures that $\Phi^{(k)} \leq n^{3}$ at the end of each iteration. By the definition of $\Phi$ and the choice of $\lambda$ in Main, we have that

$$
\left\|\frac{x s}{t}-1\right\|_{\infty} \leq \frac{\ln \left(2 n^{3}\right)}{\lambda} \leq 0.1
$$

This proves the first assumption $x s \approx_{0.1} t$ with $t>0$.

For the third assumption, we note that

$$
\begin{aligned}
\left\|\delta_{\mu}\right\|_{2} & =\left\|\left(\frac{t^{\text {new }}}{t}-1\right) x s-\frac{\epsilon}{2} \cdot t^{\text {new }} \cdot \frac{\nabla \Phi_{\lambda}(\mu / t-1)}{\left\|\nabla \Phi_{\lambda}(\mu / t-1)\right\|_{2}}\right\|_{2} \\
& \leq\left|\frac{t^{\text {new }}}{t}-1\right|\|x s\|_{2}+\frac{\epsilon}{2} t^{\text {new }} \\
& \leq \frac{\epsilon}{3 \sqrt{n}} \cdot 1.1 \sqrt{n} t+1.01 \cdot \frac{\epsilon}{2} t \leq \epsilon t,
\end{aligned}
$$

where we used $x s \approx_{0.1} t$ and the formula of $t^{\text {new }}$. Hence, we proved all assumptions in Assumption 2.4.1.

Now, we bound the probability that ClassicalStep happens. In the beginning of the Main, Lemma 3.8 .2 is used to modify the linear program with parameter $\min \left(\frac{\delta}{2}, \frac{1}{\lambda}\right)$. Hence, the initial point $x$ and $s$ satisfies $x s \approx_{1 / \lambda} 1$. Therefore, we have $\Phi^{(0)} \leq 10 n$. Lemma 2.4.12 shows $\mathbb{E}\left[\Phi^{(k+1)}\right] \leq\left(1-\frac{\lambda \epsilon}{15 \sqrt{n}}\right) \mathbb{E}\left[\Phi^{(k)}\right]+\frac{\lambda \epsilon}{15 \sqrt{n}} 10 n$. By induction, we have that $\mathbb{E}\left[\Phi^{(k)}\right] \leq 10 n$ for all $k$. Since the potential is positive, Markov inequality shows that for any $k, \Phi^{(k)} \geq n^{3}$ with probability at most $O\left(\frac{1}{n^{2}}\right)$.

### 2.4.4 Analysis of cost per iteration

To apply the data structure for projection maintenance (Theorem 2.5.1), we need to first prove the input vector $w$ does not change too much for each step.

Lemma 2.4.14. Let $x^{\text {new }}=x+\widetilde{\delta}_{x}$ and $s^{\text {new }}=s+\widetilde{\delta}_{s}$. Let $w=\frac{x}{s}$ and $w^{\text {new }}=\frac{x^{\text {new }}}{s^{\text {new }}}$. Then we have

$$
\sum_{i=1}^{n}\left(\frac{\mathbb{E}\left[w_{i}^{\mathrm{new}}\right]-w_{i}}{w_{i}}\right)^{2} \leq 64 \epsilon^{2}, \sum_{i=1}^{n}\left(\mathbb{E}\left[\left(\frac{w_{i}^{\mathrm{new}}-w_{i}}{w_{i}}\right)^{2}\right]\right)^{2} \leq 1000 \epsilon^{2},\left|\frac{w_{i}^{\mathrm{new}}-w_{i}}{w_{i}}\right| \leq 0.1
$$

Proof. From the definition, we know that

$$
\frac{w_{i}^{\mathrm{new}}}{w_{i}}=\frac{1}{s_{i}^{-1} x_{i}} \frac{x_{i}+\widetilde{\delta}_{x, i}}{s_{i}+\widetilde{\delta}_{s, i}}=\frac{1+x_{i}^{-1} \widetilde{\delta}_{x, i}}{1+s_{i}^{-1} \widetilde{\delta}_{s, i}} .
$$

Part 1. For each $i \in[n]$, we have

$$
\begin{array}{rlr}
\frac{\mathbb{E}\left[w_{i}^{\text {new }}\right]}{w_{i}}-1 & =\mathbb{E}\left[\frac{1+x_{i}^{-1} \widetilde{\delta}_{x, i}}{1+s_{i}^{-1} \widetilde{\delta}_{s, i}}\right]-1 \\
& =\mathbb{E}\left[\frac{x_{i}^{-1} \widetilde{\delta}_{x, i}-s_{i}^{-1} \widetilde{\delta}_{s, i}}{1+s_{i}^{-1} \widetilde{\delta}_{s, i}}\right] \\
& \leq 2\left|\mathbb{E}\left[x_{i}^{-1} \widetilde{\delta}_{x, i}-s_{i}^{-1} \widetilde{\delta}_{s, i}\right]\right| & \text { by }\left|s_{i}^{-1} \widetilde{\delta}_{s, i}\right| \leq 0.2, \\
& \leq 2\left|\mathbb{E}\left[x_{i}^{-1} \widetilde{\delta}_{x, i}\right]\right|+2\left|\mathbb{E}\left[s_{i}^{-1} \widetilde{\delta}_{s, i}\right]\right| . & \text { part } 3 \text { of Lemma 2.4.3 } \\
& \text { by triangle inequality }
\end{array}
$$

Thus, summing over all the coordinates gives

$$
\sum_{i=1}^{n}\left(\frac{\mathbb{E}\left[w_{i}^{\mathrm{new}}\right]-w_{i}}{w_{i}}\right)^{2} \leq \sum_{i=1}^{n} 8\left(\mathbb{E}\left[x_{i}^{-1} \widetilde{\delta}_{x, i}\right]\right)^{2}+8\left(\mathbb{E}\left[s_{i}^{-1} \widetilde{\delta}_{s, i}\right]\right)^{2} \leq 64 \epsilon^{2}
$$

where the first step follows by triangle inequality, the last step follows by $\left\|\mathbb{E}\left[s^{-1} \widetilde{\delta}_{s}\right]\right\|_{2}^{2},\left\|\mathbb{E}\left[x^{-1} \widetilde{\delta}_{x}\right]\right\|_{2}^{2} \leq$ $4 \epsilon^{2}$ (Part 1 of Lemma 2.4.3).

Part 2. For each $i \in[n]$, we have

$$
\begin{aligned}
\mathbb{E}\left[\left(\frac{w_{i}^{\text {new }}}{w_{i}}-1\right)^{2}\right] & =\mathbb{E}\left[\left(\frac{x_{i}^{-1} \widetilde{\delta}_{x, i}-s_{i}^{-1} \widetilde{\delta}_{s, i}}{1+s_{i}^{-1} \widetilde{\delta}_{s, i}}\right)^{2}\right] \\
& \leq 2 \mathbb{E}\left[\left(x_{i}^{-1} \widetilde{\delta}_{x, i}-s_{i}^{-1} \widetilde{\delta}_{s, i}\right)^{2}\right] \\
& \leq 2 \mathbb{E}\left[2\left(x_{i}^{-1} \widetilde{\delta}_{x, i}\right)^{2}+2\left(s_{i}^{-1} \widetilde{\delta}_{s, i}\right)^{2}\right] \\
& =4 \mathbb{E}\left[\left(x_{i}^{-1} \widetilde{\delta}_{x, i}\right)^{2}\right]+4 \mathbb{E}\left[\left(s_{i}^{-1} \widetilde{\delta}_{s, i}\right)^{2}\right] \\
& =4 \mathbb{V}\left[x_{i}^{-1} \widetilde{\delta}_{x, i}\right]+4\left(\mathbb{E}\left[x_{i}^{-1} \widetilde{\delta}_{x, i}\right]\right)^{2}+4 \mathbb{V}\left[s_{i}^{-1} \widetilde{\delta}_{s, i}\right]+4\left(\mathbb{E}\left[s_{i}^{-1} \widetilde{\delta}_{s, i}\right]\right)^{2} \\
& \leq 16 \epsilon^{2} / k+4\left(\mathbb{E}\left[x_{i}^{-1} \widetilde{\delta}_{x, i}\right]\right)^{2}+4\left(\mathbb{E}\left[s_{i}^{-1} \widetilde{\delta}_{s, i}\right]\right)^{2}
\end{aligned}
$$

where the last step follows by $\mathbb{V}\left[x_{i}^{-1} \widetilde{\delta}_{x, i}\right], \mathbb{V}\left[s_{i}^{-1} \widetilde{\delta}_{s, i}\right] \leq 2 \epsilon^{2} / k$ (Part 2 of Lemma 2.4.3).
Thus summing over all the coordinates

$$
\begin{aligned}
\sum_{i=1}^{n}\left(\mathbb{E}\left[\left(\frac{w_{i}^{\text {new }}}{w_{i}}-1\right)^{2}\right]\right)^{2} & \leq \frac{512 n \epsilon^{4}}{k^{2}}+64 \sum_{i=1}^{n}\left(\left(\mathbb{E}\left[x_{i}^{-1} \widetilde{\delta}_{x, i}\right]\right)^{4}+\left(\mathbb{E}\left[s_{i}^{-1} \widetilde{\delta}_{s, i}\right]\right)^{4}\right) \\
& \leq \frac{512 n \epsilon^{4}}{k^{2}}+2048 \epsilon^{4} \\
& \leq 1000 \epsilon^{2}
\end{aligned}
$$

where the last step follows by $\left\|\mathbb{E}\left[s^{-1} \widetilde{\delta}_{s}\right]\right\|_{2}^{2},\left\|\mathbb{E}\left[x^{-1} \widetilde{\delta}_{x}\right]\right\|_{2}^{2} \leq 4 \epsilon^{2}$ and $k \geq \sqrt{n} \epsilon$.
Part 3. For each $i \in[n]$

$$
\left|\frac{w_{i}^{\text {new }}}{w_{i}}-1\right|=\left|\frac{1+x_{i}^{-1} \widetilde{\delta}_{x, i}}{1+s_{i}^{-1} \widetilde{\delta}_{s, i}}-1\right| \leq\left|\frac{1+0.02}{1-0.02}-1\right| \leq 0.1
$$

where the second step follows by $\left|x_{i}^{-1} \widetilde{\delta}_{x, i}\right| \leq 0.02$ and $\left|s_{i}^{-1} \widetilde{\delta}_{s, i}\right| \leq 0.02$ (Part 3 of Lemma 2.4.3).

Now, we analyze the cost per iteration in procedure MaIn. This is a direct application of our projection maintenance result.

Lemma 2.4.15. For $\epsilon \geq \frac{1}{\sqrt{n}}$, each iteration of Main (Algorithm 3.6) takes

$$
n^{1+a+o(1)}+\epsilon \cdot\left(n^{\omega-1 / 2+o(1)}+n^{2-a / 2+o(1)}\right)
$$

expected time per iteration in amortized where $0 \leq a \leq \alpha$ controls the batch size in the data structure and $\alpha$ is the dual exponent of matrix multiplication.

Proof. Lemma 2.4.13 shows that ClassicalStep happens with only $O\left(1 / n^{2}\right)$ probability each step. Since the cost of each step only takes $\widetilde{O}\left(n^{2.5}\right)$, the expected cost is only $\widetilde{O}\left(n^{0.5}\right)$.

Lemma 2.4.14 shows that the conditions in Theorem 2.5.1 holds with the parameter $C_{1}=O(\epsilon), C_{2}=O(\epsilon), \epsilon_{m p}=\Theta(1)$.

In the procedure StochasticStep, Theorem 2.5.1 shows that the amortized time per iteration is mainly dominated by two steps:

1. mp. $\operatorname{Update}(w): O\left(\epsilon \cdot\left(n^{\omega-1 / 2+o(1)}+n^{2-a / 2+o(1)}\right)\right)$.
2. mp. $\operatorname{QUERY}\left(\frac{1}{\sqrt{\overline{X S}}} \widetilde{\delta}_{\mu}\right): O\left(n \cdot\left\|\widetilde{\delta}_{\mu}\right\|_{0}+n^{1+a+o(1)}\right)$.

Combining both running time and using $\mathbb{E}\left[\left\|\widetilde{\delta}_{\mu}\right\|_{0}\right]=O(1+k)=O\left(\epsilon \sqrt{n} \log ^{2} n\right)$ (according to the probability of success in Claim 2.4.6 and matching Assumption 2.4.1), we have the result.

### 2.4.5 Main result

Proof of Theorem 2.2.1. In the beginning of the Main algorithm, Lemma 3.8.2 is called to modify the linear program. Then, we run the stochastic central path method on this modified linear program.

When the algorithm stops, we obtain a vector $x$ and $s$ such that $x s \approx_{0.1} t$ with $t \leq \frac{\delta^{2}}{2 n}$. Hence, the duality gap is bounded by $\sum_{i} x_{i} s_{i} \leq \delta^{2}$. Lemma 3.8.2 shows how to obtain an approximate solution of the original linear program with the guarantee needed using the $x$ and $s$ we just found.

Since $t$ is decreased by $1-\frac{\epsilon}{3 \sqrt{n}}$ factor each iteration, it takes $O\left(\frac{\sqrt{n}}{\epsilon} \cdot \log \left(\frac{n}{\delta}\right)\right)$ iterations in total. In Lemma 2.4.15, we proved that each iteration takes

$$
n^{1+a+o(1)}+\epsilon \cdot\left(n^{\omega-1 / 2+o(1)}+n^{2-a / 2+o(1)}\right)
$$

and hence the total runtime is

$$
O\left(n^{2.5-a / 2+o(1)}+n^{\omega+o(1)}+\frac{n^{1.5+a+o(1)}}{\epsilon}\right) \cdot \log \left(\frac{n}{\delta}\right)
$$

Since $\epsilon=\Theta\left(\frac{1}{\log n}\right)$, the total runtime is

$$
O\left(n^{2.5-a / 2+o(1)}+n^{\omega+o(1)}+n^{1.5+a+o(1)}\right) \cdot \log \left(\frac{n}{\delta}\right)
$$

Finally, we note that the optimal choice of $a$ is $\min \left(\frac{2}{3}, \alpha\right)$, which gives the promised runtime.

Using the same proof, but different choice of the parameters, we can analyze the ultra short step stochastic central path method, where each step involves sampling only polylogarithmic coordinates. As we mentioned before, the runtime is still around $n^{\omega}$.

Corollary 2.4.16. Under the same assumption as Theorem 2.2.1, if we choose $\epsilon=\Theta(1 / \sqrt{n})$ and $a=\min \left(\frac{1}{3}, \alpha\right)$, the expected time of Main (Algorithm 3.6) is

$$
\left(n^{\omega+o(1)}+n^{2.5-\alpha / 2+o(1)}+n^{2+1 / 3+o(1)}\right) \cdot \log \left(\frac{n}{\delta}\right)
$$

### 2.5 Projection Maintenance

The goal of this section is to prove the following theorem:
Theorem 2.5.1 (Projection maintenance). Given a full rank matrix $A \in \mathbb{R}^{d \times n}$ with $n \geq d$ and a tolerance parameter $0<\epsilon_{m p}<1 / 4$. Given any positive number a such that $a \leq \alpha$ where $\alpha$ is the dual exponent of matrix multiplication. There is a deterministic data structure (Algorithm 2.3, 2.4) that approximately maintains the projection matrices

$$
\sqrt{W} A^{\top}\left(A W A^{\top}\right)^{-1} A \sqrt{W}
$$

for positive diagonal matrices $W$ through the following two operations:

1. Update $(w)$ : Output $a$ vector $\widetilde{v}$ such that for all $i$,

$$
\left(1-\epsilon_{m p}\right) \widetilde{v_{i}} \leq w_{i} \leq\left(1+\epsilon_{m p}\right) \widetilde{v_{i}}
$$

2. $\operatorname{QUERy}(h):$ Output $\sqrt{\widetilde{V}} A^{\top}\left(A \widetilde{V} A^{\top}\right)^{-1} A \sqrt{\widetilde{V}} h$ for the $\widetilde{v}$ outputted by the last call to Update.

The data structure takes $n^{2} d^{\omega-2}$ time to initialize and each call of $\operatorname{QUERY}(h)$ takes time

$$
n \cdot\|h\|_{0}+n^{1+a+o(1)}
$$

Furthermore, if the initial vector $w^{(0)}$ and the (random) update sequence $w^{(1)}, \cdots, w^{(T)} \in \mathbb{R}^{n}$ satisfies

$$
\sum_{i=1}^{n}\left(\frac{\mathbb{E}\left[w_{i}^{(k+1)}\right]-w_{i}^{(k)}}{w_{i}^{(k)}}\right)^{2} \leq C_{1}^{2}, \sum_{i=1}^{n}\left(\mathbb{E}\left[\left(\frac{w_{i}^{(k+1)}-w_{i}^{(k)}}{w_{i}^{(k)}}\right)^{2}\right]\right)^{2} \leq C_{2}^{2},\left|\frac{w_{i}^{(k+1)}-w_{i}^{(k)}}{w_{i}^{(k)}}\right| \leq \frac{1}{4}
$$

```
Algorithm 2.3 Projection Maintenance Data Structure - Initial, Query
    datastructure MaintainPRoJECTION
            \(\triangleright\) Theorem 2.5.1
    members
        \(w \in \mathbb{R}^{n} \quad \triangleright\) Target vector
        \(v, \widetilde{v} \in \mathbb{R}^{n} \quad \triangleright\) Approximate vector
        \(A \in \mathbb{R}^{d \times n}\)
        \(M \in \mathbb{R}^{n \times n} \quad \triangleright\) Approximate Projection Matrix
        \(\epsilon_{m p} \in(0,1 / 4) \quad \triangleright\) Tolerance
        \(a \in(0, \alpha] \quad \triangleright\) Batch Size for Update \(\left(n^{a}\right)\)
    end members
    procedure Initialize \(\left(A, w, \epsilon_{m p}, a\right) \quad \triangleright\) Lemma 2.5.2
        \(w \leftarrow w, v \leftarrow w, \epsilon_{m p} \leftarrow \epsilon_{m p}, A \leftarrow A, a \leftarrow a\)
        \(M \leftarrow A^{\top}\left(A V A^{\top}\right)^{-1} A\)
    end procedure
    procedure QUERY \((h)\)
        \(\triangleright\) Lemma 2.5.4
        Let \(\widetilde{S}\) be the indices \(i\) such that \(\left(1-\epsilon_{m p}\right) v_{i} \leq w_{i} \leq\left(1+\epsilon_{m p}\right) v_{i}\) is false.
        return \(\sqrt{\widetilde{V}} \cdot(M \cdot(\sqrt{\widetilde{V}} \cdot h))-\sqrt{\widetilde{V}} \cdot\left(M_{\widetilde{S}} \cdot\left(\left(\widetilde{\Delta}_{\widetilde{S}, \widetilde{S}}^{-1}+M_{\widetilde{S}, \widetilde{S}}\right)^{-1} \cdot\left(M_{\widetilde{S}}^{\top} \sqrt{\widetilde{V}} h\right)\right)\right)\)
    end procedure
    end datastructure
```

with the expectation is conditional on $w_{i}^{(k)}$ for all $k=0,1, \cdots, T-1$. Then, the amortized expected time ${ }^{5}$ per call of $\operatorname{Update}(w)$ is

$$
\left(C_{1} / \epsilon_{m p}+C_{2} / \epsilon_{m p}^{2}\right) \cdot\left(n^{\omega-1 / 2+o(1)}+n^{2-a / 2+o(1)}\right)
$$

Remark 2.5.1. For our linear program algorithm, we have $C_{1}=O(1 / \log n), C_{2}=O(1 / \log n)$ and $\epsilon_{m p}=\Theta(1)$. See Lemma 2.4.14.

[^11]```
Algorithm 2.4 Projection Maintenance Data Structure - Update
    datastructure MaintainProjection \(\triangleright\) Theorem 2.5.1
    procedure Update \(\left(w^{\text {new }}\right)\)
                                    \(\triangleright\) Lemma 2.5.3
        \(y_{i} \leftarrow w_{i}^{\text {new }} / v_{i}-1, \forall i \in[n]\)
        \(r \leftarrow\) the number of indices \(i\) such that \(\left|y_{i}\right| \geq \epsilon_{m p}\).
        if \(r<n^{a}\) then
            \(v^{\text {new }} \leftarrow v\)
            \(M^{\text {new }} \leftarrow M\)
        else
            Let \(\pi:[n] \rightarrow[n]\) be a sorting permutation such that \(\left|y_{\pi(i)}\right| \geq\left|y_{\pi(i+1)}\right|\)
            while \(1.5 \cdot r<n\) and \(\left|y_{\pi(\lceil 1.5 \cdot r\rceil)}\right| \geq(1-1 / \log n)\left|y_{\pi(r)}\right|\) do
                \(r \leftarrow \min (\lceil 1.5 \cdot r\rceil, n)\)
            end while
            \(v_{\pi(i)}^{\text {new }} \leftarrow \begin{cases}w_{\pi(i)}^{\text {new }} & i \in\{1,2, \cdots, r\} \\ v_{\pi(i)} & i \in\{r+1, \cdots, n\}\end{cases}\)
                            \(\triangleright\) Compute \(M^{\text {new }}=A^{\top}\left(A V^{\text {new }} A^{\top}\right)^{-1} A\) via Matrix Woodbury
            \(\Delta \leftarrow \operatorname{diag}\left(v^{\text {new }}-v\right) \quad \triangleright \Delta \in \mathbb{R}^{n \times n}\) and \(\|\Delta\|_{0}=r\)
            Let \(S \leftarrow \pi([r])\) be the first \(r\) indices in the permutation.
            Let \(M_{S} \in \mathbb{R}^{n \times r}\) be the \(r\) columns from \(S\) of \(M\).
            Let \(M_{S, S}, \Delta_{S, S} \in \mathbb{R}^{r \times r}\) be the \(r\) rows and columns from \(S\) of \(M\) and \(\Delta\).
            \(M^{\text {new }} \leftarrow M-M_{S} \cdot\left(\Delta_{S, S}^{-1}+M_{S, S}\right)^{-1} \cdot\left(M_{S}\right)^{\top}\)
            end if
            \(w \leftarrow w^{\text {new }}, v \leftarrow v^{\text {new }}, M \leftarrow M^{\text {new }}\)
            \(\widetilde{v}_{i} \leftarrow \begin{cases}v_{i} & \text { if }\left(1-\epsilon_{m p}\right) v_{i} \leq w_{i} \leq\left(1+\epsilon_{m p}\right) v_{i} \\ w_{i} & \text { otherwise }\end{cases}\)
        return \(\widetilde{v}\)
    end procedure
    end datastructure
```


### 2.5.1 Proof outline

For intuition, we consider the case $C_{1}=\Theta(1), C_{2}=\Theta(1)$, and $\epsilon_{m p}=\Theta(1)$ in this explanation. The correctness of the data structure directly follows from Woodbury matrix identity. The amortized time analysis is based on a potential function that measures the
distance of the approximate vector $v$ and the target vector $w$. We will show that

- The cost to update the projection $M$ is proportional to the decrease of the potential.
- Each call to query increase the potential by a fixed amount.

Combining both together gives the amortized runtime bound of our data structure.
Now, we explain the definition of the potential. Consider the $k$-th round of the algorithm. For all $i \in[n]$, we define $x_{i}^{(k)}=\frac{w_{i}^{(k)}}{v_{i}^{(k)}}-1$. Note that $\left|x_{i}^{(k)}\right|$ measures the relative distance between $w_{i}^{(k)}$ and $v_{i}^{(k)}$. Our algorithm fixes the indices with largest error $x_{i}^{(k)}$. To capture the fact that updating in a larger batch is more efficient, we define the potential as a weighted combination of the error where we put more weight to higher $x_{i}^{(k)}$. Formally, we sort the coordinates of $x^{(k)}$ such that $\left|x_{i}^{(k)}\right| \geq\left|x_{i+1}^{(k)}\right|$ and define the potential by

$$
\Psi_{k}=\sum_{i=1}^{n} g_{i} \cdot \psi\left(x_{i}^{(k)}\right)
$$

where $g_{i}$ are positive decreasing numbers to be chosen and $\psi$ is a symmetric $(\psi(x)=\psi(-x))$ positive function that increases on both sides. For intuition, one can think $\psi(x)$ behaves roughly like $|x|$.

Each iteration we update the projection matrix such that the error of $\left|x_{1}\right|, \cdots,\left|x_{r}\right|$ drops from roughly $\epsilon_{m p}$ to 0 . This decreases the potential of $\psi\left(x_{i}^{(k)}\right)$ by $\Omega\left(\epsilon_{m p}\right)$ from $i=$ $1, \cdots, r$. Therefore, the whole potential decreases by $\Omega\left(\epsilon_{m p} \sum_{i=1}^{r} g_{i}\right)$. To make the term $\sum_{i=1}^{r} g_{i}$ proportional to the time to update a rank $r$ part of the projection matrix, we set

$$
g_{i}= \begin{cases}n^{-a}, & \text { if } i<n^{a}  \tag{2.20}\\ i^{\frac{\omega-2}{1-a}-1} n^{-\frac{a(\omega-2)}{1-a}}, & \text { otherwise } .\end{cases}
$$

where $\omega$ is the exponent of matrix multiplication and $a$ is any positive number less than or equals to the dual exponent of matrix multiplication. Lemma 2.6 .3 shows that $g$ is indeed non-increasing and Lemma 2.5.3 shows that the update time of data-structure is indeed $O\left(r g_{r} n^{2+o(1)}\right)=O\left(\sum_{i=1}^{r} g_{i} n^{2+o(1)}\right)$ for any $r \geq n^{a}$.

Each call to UpDATE, the expectation of the error vector $x^{(k)}$ moves roughly in an unit $\ell_{2}$ ball. Therefore, the changes of the potential is roughly upper bounded $\left(\sum_{i=1}^{n} g_{i}^{2}\right)^{1 / 2} \approx$ $n^{\omega-5 / 2}$. Since it takes us $n^{2+o(1)}$ time to decrease the potential by roughly 1 in the update step, the total time is roughly $n^{\omega-1 / 2}$.

For the case of stochastic central path, we note that the variance of the vector $x$ is quite small. By choosing a smooth potential function $\psi$ (see (2.21)), we can essentially give the same result as if there is no variance.

### 2.5.2 Proof of Theorem 2.5.1

Now, we give the proof of Theorem 2.5.1. We will defer some simple calculations into later sections.

Proof of Theorem 2.5.1.
Proof of Correctness. The definition of $\widetilde{v}$ in Line 20 ensures that $\left(1-\epsilon_{m p}\right) \widetilde{v}_{i} \leq$ $w_{i} \leq\left(1+\epsilon_{m p}\right) \widetilde{v_{i}}$.

Using the Matrix Woodbury formula, one can verify that the update rule in Line 17 correctly maintains $M=A^{\top}\left(A V A^{\top}\right)^{-1} A$. See the deviation of the formula in Lemma 2.5.2. By the same reasoning, the Line 19 outputs the vector $\sqrt{\widetilde{V}} A^{\top}\left(A \widetilde{V} A^{\top}\right)^{-1} A \sqrt{\widetilde{V}} h$. This completes the proof of correctness.

Definition of $x$ and $y$. Consider the $k$-th round of the algorithm. For all $i \in[n]$, we define $x_{i}^{(k)}, x_{i}^{(k+1)}$ and $y_{i}^{(k)}$ as follows:

$$
x_{i}^{(k)}=\frac{w_{i}^{(k)}}{v_{i}^{(k)}}-1, y_{i}^{(k)}=\frac{w_{i}^{(k+1)}}{v_{i}^{(k)}}-1, x_{i}^{(k+1)}=\frac{w_{i}^{(k+1)}}{v_{i}^{(k+1)}}-1 .
$$

Note that the difference between $x_{i}^{(k)}$ and $y_{i}^{(k)}$ is that $w$ is changing. The difference between $y_{i}^{(k)}$ and $x_{i}^{(k+1)}$ is that $v$ is changing. For simplicity, we define $\beta_{i}=\left(\mathbb{E}\left[w_{i}^{(k+1)}\right]-w_{i}^{(k)}\right) / w_{i}^{(k)}$, then one of assumption becomes $\sum_{i=1}^{n} \beta_{i}^{2} \leq C_{1}^{2}$.

Assume sorting. Assume the coordinates of vector $x^{(k)} \in \mathbb{R}^{n}$ are sorted such that $\left|x_{i}^{(k)}\right| \geq\left|x_{i+1}^{(k)}\right|$. Let $\tau$ and $\pi$ are permutations such that $\left|x_{\tau(i)}^{(k+1)}\right| \geq\left|x_{\tau(i+1)}^{(k+1)}\right|$ and $\left|y_{\pi(i)}^{(k)}\right| \geq$ $\left|y_{\pi(i+1)}^{(k)}\right|$.

Definition of Potential function. Let $g$ be defined in (2.20). Let $\psi: \mathbb{R} \rightarrow \mathbb{R}$ be defined by

$$
\psi(x)= \begin{cases}\frac{|x|^{2}}{2 \epsilon_{m p}}, & |x| \in\left[0, \epsilon_{m p}\right]  \tag{2.21}\\ \epsilon_{m p}-\frac{\left(2 \epsilon_{m p}-|x|\right)^{2}}{2 \epsilon_{m p}}, & |x| \in\left(\epsilon_{m p}, 2 \epsilon_{m p}\right] \\ \epsilon_{m p} . & |x| \in\left(2 \epsilon_{m p},+\infty\right)\end{cases}
$$

We define the potential at the $k$-th round by

$$
\Psi_{k}=\sum_{i=1}^{n} g_{i} \cdot \psi\left(x_{\tau_{k}(i)}^{(k)}\right) .
$$

where $\tau_{k}(i)$ is the permutation such that $\left|x_{\tau_{k}(i)}^{(k)}\right| \geq\left|x_{\tau_{k}(i+1)}^{(k)}\right|$.

## Bounding the potential.

We can express $\Psi_{k+1}-\Psi_{k}$ as follows:

$$
\begin{align*}
\Psi_{k+1}-\Psi_{k} & =\sum_{i=1}^{n} g_{i} \cdot\left(\psi\left(x_{\tau(i)}^{(k+1)}\right)-\psi\left(x_{i}^{(k)}\right)\right) \\
& =\sum_{i=1}^{n} g_{i} \cdot \underbrace{\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)\right)}_{w \text { move }}-\sum_{i=1}^{n} g_{i} \cdot \underbrace{\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\tau(i)}^{(k+1)}\right)\right)}_{v \text { move }} . \tag{2.22}
\end{align*}
$$

Now, using Lemma 2.5.5 and 3.6.12, and the fact that $\Psi_{0}=0$ and $\Psi_{T} \geq 0$, with (2.22), we get

$$
\begin{aligned}
0 & \leq \Psi_{T}-\Psi_{0}=\sum_{k=0}^{T-1}\left(\Psi_{k+1}-\Psi_{k}\right) \\
& \leq \sum_{k=0}^{T-1}\left(O\left(C_{1}+C_{2} / \epsilon_{m p}\right) \cdot \sqrt{\log n} \cdot\left(n^{-a / 2}+n^{\omega-5 / 2}\right)-\Omega\left(\epsilon_{m p} r_{k} g_{r_{k}} / \log n\right)\right) \\
& =T \cdot O\left(C_{1}+C_{2} / \epsilon_{m p}\right) \cdot \sqrt{\log n} \cdot\left(n^{-a / 2}+n^{\omega-5 / 2}\right)-\sum_{k=1}^{T} \Omega\left(\epsilon_{m p} r_{k} g_{r_{k}} / \log n\right)
\end{aligned}
$$

where the third step follows by Lemma 2.5.5 and Lemma 3.6.12 and $r_{k}$ is the number of coordinates we update during that iteration.

Therefore, we get,

$$
\sum_{k=1}^{T} r_{k} g_{r_{k}}=O\left(T \cdot\left(C_{1} / \epsilon_{m p}+C_{2} / \epsilon_{m p}^{2}\right) \cdot \log ^{3 / 2} n \cdot\left(n^{\omega-5 / 2}+n^{-a / 2}\right)\right)
$$

Proof of running time. See the Section 2.5.3.

### 2.5.3 Initialization time, update time, query time

To formalize the amortized runtime proof, we first analyze the initialization time (Lemma 2.5.2), update time (Lemma 2.5.3), and query time (Lemma 2.5.4) of our projection maintenance data-structure.

Lemma 2.5.2 (Initialization time). The initialization time of data-structure MaintainProjection (Algorithm 2.3) is $O\left(n^{2} d^{\omega-2}\right)$.

Proof. Given matrix $A \in \mathbb{R}^{d \times n}$ and diagonal matrix $V \in \mathbb{R}^{n \times n}$, computing $A^{\top}\left(A V A^{\top}\right)^{-1} A$ takes $O\left(n^{2} d^{\omega-2}\right)$.

Lemma 2.5.3 (Update time). The update time of data-structure MaintainProjection (Algorithm 2.4) is $O\left(r g_{r} n^{2+o(1)}\right)$ where $r$ is the number of indices we updated in $v$.

Proof. Let $A_{S} \in \mathbb{R}^{d \times r}$ be the $r$ columns from $S$ of $A$. From $k$-th query to $(k+1)$-th query, we have

$$
\begin{aligned}
& A^{\top}\left(A V^{(k+1)} A^{\top}\right)^{-1} A \\
= & A^{\top}\left(A\left(V^{(k)}+\Delta\right) A^{\top}\right)^{-1} A \\
= & A^{\top}\left(\left(A V^{(k)} A^{\top}\right)^{-1}-\left(A V^{(k)} A^{\top}\right)^{-1} A_{S}\left(\Delta_{S, S}^{-1}+A_{S}^{\top}\left(A V^{(k)} A^{\top}\right)^{-1} A_{S}\right)^{-1} A_{S}^{\top}\left(A V^{(k)} A^{\top}\right)^{-1}\right) A \\
= & A^{\top}\left(A V^{(k)} A^{\top}\right)^{-1} A-A^{\top}\left(A V^{(k)} A^{\top}\right)^{-1} A_{S}\left(\Delta_{S, S}^{-1}+A_{S}^{\top}\left(A V^{(k)} A^{\top}\right)^{-1} A_{S}\right)^{-1} A_{S}^{\top}\left(A V^{(k)} A^{\top}\right)^{-1} A \\
= & M^{(k)}-M_{S}^{(k)}\left(\Delta_{S, S}^{-1}+M_{S, S}^{(k)}\right)^{-1}\left(M_{S}^{(k)}\right)^{\top},
\end{aligned}
$$

where the second step follows by Matrix Woodbury Identity and the last step follows by definition of $M^{(k)} \in \mathbb{R}^{n \times n}$.

Thus the update rule of matrix $M^{(k+1)} \in \mathbb{R}^{n \times n}$ can be written as

$$
M^{(k+1)}=M^{(k)}-M_{S}^{(k)}\left(\Delta_{S, S}^{-1}+\left(M^{(k)}\right)_{S, S}\right)^{-1}\left(M_{S}^{(k)}\right)^{\top} .
$$

The updates in round $k$ can be splitted into four parts:

1. Adding two $r \times r$ matrices takes $O\left(r^{2}\right)$ time.
2. Computing the inverse of an $r \times r$ matrix takes $O\left(r^{\omega+o(1)}\right)$ time.
3. Computing the matrix multiplication of a $n \times r$ and $r \times n$ matrix takes $O\left(r g_{r} \cdot n^{2+o(1)}\right)$ time where we used that $r \geq n^{a}$ (Lemma 2.2.3).
4. Adding two $n \times n$ matrices together takes $O\left(n^{2}\right)$ time.

Hence, the total cost is

$$
O\left(r^{2}+r^{\omega+o(1)}+r g_{r} \cdot n^{2+o(1)}+n^{2}\right)=O\left(r^{2}+r^{\omega+o(1)}+r g_{r} \cdot n^{2+o(1)}\right)=O\left(r g_{r} \cdot n^{2+o(1)}\right) .
$$

where the first step follows by $r g_{r} \geq 1$ for all $r \geq n^{a}$ and the last step follows by the calculations.

Lemma 2.5.4 (Query time). The query time of data-structure MaintainProjection (Algorithm 2.3) is $O\left(n \cdot\|h\|_{0}+n^{1+a+o(1)}\right)$.

Proof. Let $\widetilde{\Delta}$ satisfies $\widetilde{V}=V+\widetilde{\Delta}$. Let $\widetilde{S} \subset[n]$ denote the support of $\widetilde{\Delta}$ and then $|\widetilde{S}| \leq n^{a}$. Let $\widetilde{r}$ denote $|\widetilde{S}|$. We abuse the notation here, $\widetilde{\Delta}$ denotes both $n \times n$ diagonal matrix and a length $n$ vector.

Using Matrix Woodbury Identity and definition of $M$, a same proof as Update time (Lemma 2.5.3) shows

$$
A^{\top}\left(A \widetilde{V} A^{\top}\right)^{-1} A=M+M_{\widetilde{S}}\left(\widetilde{\Delta}_{\widetilde{S}, \widetilde{S}}^{-1}+M_{\widetilde{S}, \widetilde{S}}\right)^{-1} M_{\widetilde{S}}^{\top}
$$

where $\widetilde{\Delta}_{\tilde{S} \times \widetilde{S}}$ has size $\widetilde{r} \times \widetilde{r}, M_{\widetilde{S}, \widetilde{S}}$ has size $\widetilde{r} \times \widetilde{r}$ and $M_{\widetilde{S}}$ has size $n \times \widetilde{r}$.
To compute $\sqrt{\widetilde{V}} A^{\top}\left(A \widetilde{V} A^{\top}\right)^{-1} A \sqrt{\widetilde{V}} h$, we just need to compute

$$
\sqrt{\widetilde{V}} M \sqrt{\widetilde{V}} h+\sqrt{\widetilde{V}} M_{\widetilde{S}}\left(\widetilde{\Delta}_{\tilde{S}, \widetilde{S}}^{-1}+M_{\widetilde{S}, \tilde{S}}\right)^{-1} M_{\widetilde{S}}^{\top} \sqrt{\widetilde{V}} h
$$

Note the running time of computing the first term of the above equation only takes $O\left(n \cdot\|h\|_{0}\right)$ time.

Next, we analyze the cost of computing the second term of the above equation. It contains several parts:

1. Computing $\widetilde{M}_{\widetilde{S}}^{\top} \cdot(\sqrt{\widetilde{V}} \cdot h) \in \mathbb{R}^{\widetilde{r}}$ takes $\widetilde{r}\|h\|_{0}$ time.
2. Computing $\left(\widetilde{\Delta}_{\widetilde{S}, \widetilde{S}}^{-1}+M_{\widetilde{S}, \widetilde{S}}\right)^{-1} \in \mathbb{R}^{\widetilde{r} \times \widetilde{r}}$ that is the inverse of a $\widetilde{r} \times \widetilde{r}$ matrix takes $\widetilde{r}^{\omega+o(1)}$ time.
3. Computing matrix-vector multiplication between $\widetilde{r} \times \widetilde{r}$ matrix $\left(\left(\widetilde{\Delta}_{\widetilde{S}, \widetilde{S}}^{-1}+M_{\widetilde{S}, \widetilde{S}}\right)^{-1}\right)$ and $\widetilde{r} \times 1$ vector $\left(\widetilde{M}_{\widetilde{S}}^{\top} \sqrt{\widetilde{V}} h\right)$ takes $O\left(\widetilde{r}^{2}\right)$ time
4. Computing matrix-vector multiplication between $n \times \widetilde{r}$ matrix $\left(M_{\widetilde{S}}\right)$ and $\widetilde{r} \times 1$ vector $\left(\left(\widetilde{\Delta}_{\widetilde{S}, \widetilde{S}}^{-1}+M_{\widetilde{S}, \widetilde{S}}\right)^{-1} M_{\widetilde{S}}^{\top} \sqrt{\widetilde{V}} h\right)$ takes $O(n \widetilde{r})$ time.
5. Computing the entry-wise product of two $n$ vectors takes $O(n)$ time

Thus, overall the running time is

$$
O\left(\widetilde{r}\|h\|_{0}+\widetilde{r}^{\omega+o(1)}+\widetilde{r}^{2}+n \widetilde{r}+n\right)=O\left(\widetilde{r}^{\omega+o(1)}+n \widetilde{r}\right)=O\left(n^{a \cdot \omega+o(1)}+n^{1+a}\right) .
$$

Finally, we note that $\omega \leq 3-\alpha \leq 3-a$ (Lemma 2.6.3) and hence $a \cdot \omega \leq a(3-a) \leq 1+a$. Therefore, the runtime is $n^{1+a+o(1)}$.

### 2.5.4 Bounding $w$ move

The goal of this section is to prove Lemma 2.5.5.

Lemma 2.5.5 ( $w$ move). We have

$$
\sum_{i=1}^{n} g_{i} \cdot \mathbb{E}\left[\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)\right] \leq O\left(C_{1}+C_{2} / \epsilon_{m p}\right) \cdot \sqrt{\log n} \cdot\left(n^{-a / 2}+n^{\omega-5 / 2}\right)
$$

Proof. Observe that since the errors $\left|x_{i}^{(k)}\right|$ are sorted in descending order, and $\psi(x)$ is symmetric and non-decreasing function for $x \geq 0$, thus $\psi\left(x_{i}^{(k)}\right)$ is also in decreasing order. In addition, note that $g$ is decreasing, we have

$$
\begin{equation*}
\sum_{i=1}^{n} g_{i} \psi\left(x_{\pi(i)}^{(k)}\right) \leq \sum_{i=1}^{n} g_{i} \psi\left(x_{i}^{(k)}\right) \tag{2.23}
\end{equation*}
$$

Hence the first term in (2.22) can be upper bounded as follows:

$$
\begin{aligned}
& \mathbb{E}\left[\sum_{i=1}^{n} g_{i} \cdot\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)\right)\right] \\
\leq & \mathbb{E}\left[\sum_{i=1}^{n} g_{i} \cdot\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\pi(i)}^{(k)}\right)\right)\right] \\
= & \sum_{i=1}^{n} g_{i} \cdot \mathbb{E}\left[\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\pi(i)}^{(k)}\right)\right] \\
= & O\left(C_{1}+C_{2} / \epsilon_{m p}\right) \cdot \sqrt{\log n} \cdot\left(n^{-a / 2}+n^{\omega-5 / 2}\right) . \quad \text { by }(2.23)
\end{aligned}
$$

Thus, we complete the proof of $w$ move Lemma.

It remains to prove the following Lemma,

## Lemma 2.5.6.

$$
\sum_{i=1}^{n} g_{i} \cdot \mathbb{E}\left[\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\pi(i)}^{(k)}\right)\right]=O\left(C_{1}+C_{2} / \epsilon_{m p}\right) \cdot \sqrt{\log n} \cdot\left(n^{-a / 2}+n^{\omega-5 / 2}\right)
$$

Proof. Let $I$ be the set of indices such that $\left|x_{i}^{(k)}\right| \leq 1$. We separate the term into two:
$\sum_{i=1}^{n} g_{i} \cdot \mathbb{E}\left[\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\pi(i)}^{(k)}\right)\right]=\sum_{i \in I} g_{\pi^{-1}(i)} \cdot \mathbb{E}\left[\psi\left(y_{i}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)\right]+\sum_{i \in I^{c}} g_{\pi^{-1}(i)} \cdot \mathbb{E}\left[\psi\left(y_{i}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)\right]$.
Case 1: Terms from $I$ Mean value theorem shows that

$$
\begin{aligned}
\psi\left(y_{i}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right) & =\psi^{\prime}\left(x_{i}^{(k)}\right)\left(y_{i}^{(k)}-x_{i}^{(k)}\right)+\frac{1}{2} \psi^{\prime \prime}(\zeta)\left(y_{i}^{(k)}-x_{i}^{(k)}\right)^{2} \\
& \leq \psi^{\prime}\left(x_{i}^{(k)}\right) \frac{w_{i}^{(k+1)}-w_{i}^{(k)}}{v_{i}^{(k)}}+\frac{L_{2}}{2}\left(\frac{w_{i}^{(k+1)}-w_{i}^{(k)}}{v_{i}^{(k)}}\right)^{2}
\end{aligned}
$$

where $L_{2}=\max _{x} \psi^{\prime \prime}(x)$. Taking conditional expectation given $w^{(k)}$ on both sides

$$
\begin{aligned}
\mathbb{E}\left[\psi\left(y_{i}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)\right] & \leq \psi^{\prime}\left(x_{i}^{(k)}\right) \cdot \frac{\mathbb{E}\left[w_{i}^{(k+1)}\right]-w_{i}^{(k)}}{v_{i}^{(k)}}+\frac{L_{2}}{2} \frac{1}{\left(v_{i}^{(k)}\right)^{2}} \mathbb{E}\left[\left(w_{i}^{(k+1)}-w_{i}^{(k)}\right)^{2}\right] \\
& =\psi^{\prime}\left(x_{i}^{(k)}\right) \cdot \frac{w_{i}^{(k)}}{v_{i}^{(k)}} \beta_{i}+\frac{L_{2}}{2} \frac{\left(w_{i}^{(k)}\right)^{2}}{\left(v_{i}^{(k)}\right)^{2}} \gamma_{i}
\end{aligned}
$$

where $\beta_{i}=\frac{\mathbb{E}\left[w_{i}^{(k+1)}\right]-w_{i}^{(k)}}{w_{i}^{(k)}}$ and $\gamma_{i}=\mathbb{E}\left[\left(\frac{w_{i}^{(k+1)}-w_{i}^{(k)}}{w_{i}^{(k)}}\right)^{2}\right]$.
To bound $\sum_{i \in I} g_{\pi^{-1}(i)} \mathbb{E}\left[\psi\left(y_{i}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)\right]$, we need to bound the following two terms,

$$
\begin{equation*}
\sum_{i \in I} g_{\pi^{-1}(i)} \psi^{\prime}\left(x_{i}^{(k)}\right) \frac{w_{i}^{(k)}}{v_{i}^{(k)}} \beta_{i}, \text { and } \sum_{i \in I} g_{\pi^{-1}(i)} \frac{L_{2}}{2} \frac{\left(w_{i}^{(k)}\right)^{2}}{\left(v_{i}^{(k)}\right)^{2}} \gamma_{i} . \tag{2.24}
\end{equation*}
$$

For the term $\frac{w_{i}^{(k)}}{v_{i}^{(k)}}$, we note that for $i \in I$, we have $\left|\frac{w_{i}^{(k)}}{v_{i}^{(k)}}\right| \leq\left|x_{i}^{(k)}\right|+1 \leq 2$. Using this,
we can bound the first term by

$$
\begin{align*}
\sum_{i \in I} g_{\pi^{-1}(i)} \psi^{\prime}\left(x_{i}^{(k)}\right) \frac{w_{i}^{(k)}}{v_{i}^{(k)}} \beta_{i} & \leq\left(\sum_{i \in I}\left(g_{\pi^{-1}(i)} \psi^{\prime}\left(x_{i}^{(k)}\right) \frac{w_{i}^{(k)}}{v_{i}^{(k)}}\right)^{2} \sum_{i \in I} \beta_{i}^{2}\right)^{1 / 2} \\
& \leq O\left(L_{1}\right)\left(\sum_{i=1}^{n} g_{i}^{2} \cdot C_{1}^{2}\right)^{1 / 2} \\
& =O\left(C_{1} L_{1}\|g\|_{2}\right) \tag{2.25}
\end{align*}
$$

where $L_{1}=\max _{x}\left|\psi^{\prime}(x)\right|$, the first step follows by Cauchy-Schwarz inequality and the second step follows by $\left|\psi^{\prime}\left(x_{i}^{(k)}\right) \cdot w_{i}^{(k)} / v_{i}^{(k)}\right| \leq 2 L_{1}$ and $\sum_{i=1}^{n} \beta_{i}^{2} \leq C_{1}^{2}$.

For the second term, we have

$$
\begin{equation*}
\sum_{i \in I} g_{\pi^{-1}(i)} \frac{L_{2}}{2} \frac{\left(w_{i}^{(k)}\right)^{2}}{\left(v_{i}^{(k)}\right)^{2}} \gamma_{i} \leq O\left(L_{2}\right) \cdot \sum_{i=1}^{n} g_{i} \cdot \gamma_{i}=O\left(C_{2} L_{2}\|g\|_{2}\right) \tag{2.26}
\end{equation*}
$$

Now, combining (2.25) and (2.26) and using that $L_{1}=O(1), L_{2}=O\left(1 / \epsilon_{m p}\right)$ (from part 4 of Lemma 2.5.9) and $\|g\|_{2} \leq \sqrt{\log n} \cdot O\left(n^{-a / 2}+n^{\omega-5 / 2}\right)$ (from Lemma 2.5.7), we have that

$$
\sum_{i \in I} g_{\pi^{-1}(i)} \cdot \mathbb{E}\left[\psi\left(y_{i}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)\right] \leq O\left(C_{1}+C_{2} / \epsilon_{m p}\right) \cdot \sqrt{\log n} \cdot\left(n^{-a / 2}+n^{\omega-5 / 2}\right)
$$

## Case 2: Terms from $I^{c}$

For all $i \in I^{c}$, we have $\left|x_{i}^{(k)}\right| \geq 1$. Note that $\psi(x)$ is a constant for $x \geq 2 \epsilon_{m p}$ and that $\epsilon_{m p} \leq 1 / 4$. Therefore, if $\left|y_{i}^{(k)}\right| \geq 1 / 2$, we have that $\psi\left(y_{i}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)=0$. Hence, we only need to consider the $i \in I^{c}$ such that $\left|y_{i}^{(k)}\right|<1 / 2$. For these $i$, we have that

$$
\frac{1}{2}<\left|y_{i}^{(k)}-x_{i}^{(k)}\right|=\left|\frac{w_{i}^{(k+1)}-w_{i}^{(k)}}{v_{i}^{(k)}}\right|=\left|\frac{w_{i}^{(k+1)}}{v_{i}^{(k)}}\right|\left|\frac{w_{i}^{(k+1)}-w_{i}^{(k)}}{w_{i}^{(k+1)}}\right| \leq \frac{3}{2}\left|\frac{w_{i}^{(k+1)}-w_{i}^{(k)}}{w_{i}^{(k+1)}}\right|,
$$

where we used that $\left|y_{i}^{(k)}\right|=\left|\frac{w_{i}^{(k+1)}}{v_{i}^{(k)}}-1\right| \leq 1 / 2$. Hence, we have that $\left|\frac{w_{i}^{(k+1)}-w_{i}^{(k)}}{w_{i}^{(k+1)}}\right|>1 / 3$ and hence $\left|\frac{w_{i}^{(k+1)}-w_{i}^{(k)}}{w_{i}^{(k)}}\right|>1 / 4$, which is impossible.

Hence, we have

$$
\sum_{i \in I^{c}} g_{\pi^{-1}(i)} \cdot \mathbb{E}\left[\psi\left(y_{i}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)\right]=0
$$

Combining both cases, we have the result.

## Lemma 2.5.7.

$$
\left(\sum_{i=1}^{n} g_{i}^{2}\right)^{1 / 2} \leq \sqrt{\log n} \cdot O\left(n^{-a / 2}+n^{\omega-5 / 2}\right)
$$

Proof. Since function $g$ behaves differently when $i \leq n^{a}$ and $i>n^{a}$, we split the sum into two parts.

For the first part, we have

$$
\sum_{i=1}^{n^{a}} g_{i}^{2}=\sum_{i=1}^{n^{a}} n^{-2 a}=n^{-a}
$$

For the second part, we have

$$
\sum_{i=n^{a}}^{n} g_{i}^{2}=\sum_{i=n^{a}}^{n} i^{\frac{2(\omega-2)}{1-a}-2} n^{-\frac{2 a(\omega-2)}{1-a}}=\sum_{i=n^{a}}^{n} \frac{1}{i} \cdot i^{\frac{2(\omega-2)}{1-a}-1} n^{-\frac{2 a(\omega-2)}{1-a}} .
$$

Note that

$$
\max _{i \in\left[n^{a}, n\right]} i^{\frac{2(\omega-2)}{1-a}-1} n^{-\frac{2 a(\omega-2)}{1-a}}=\max \left(n^{a \frac{2(\omega-2)}{1-a}-a} n^{-\frac{2 a(\omega-2)}{1-a}}, n^{\frac{2(\omega-2)}{1-a}-1} n^{-\frac{2 a(\omega-2)}{1-a}}\right)=\max \left(n^{-a}, n^{2 \omega-5}\right)
$$

Thus, the second part is

$$
\sum_{i=n^{a}}^{n} g_{i}^{2}=\sum_{i=n^{a}}^{n} \frac{1}{i} \cdot \max \left(n^{-a}, n^{2 \omega-5}\right)=O(\log n) \cdot \max \left(n^{-a}, n^{2 \omega-5}\right)
$$

Combining the first part and the second part completes the proof.

### 2.5.5 Bounding $v$ move

Lemma 2.5.8 ( $v$ move). We have,

$$
\sum_{i=1}^{n} g_{i} \cdot\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\tau(i)}^{(k+1)}\right)\right) \geq \Omega\left(\epsilon_{m p} r_{k} g_{r_{k}} / \log n\right)
$$

Proof. We split the proof into two cases.
We first understand some simple facts which are useful in the later proof. Note that from definition of $x_{i}^{(k+1)}$, we know that $x^{(k+1)}$ has $r_{k}$ coordinates are 0 . Basically, $\left\|y^{(k)}-x^{(k+1)}\right\|_{0}=r_{k}$. The difference between those vectors is, for the largest $r_{k}$ coordinates in $y^{(k)}$, we erase them in $x^{(k+1)}$. Then for each $i \in\left[n-r_{k}\right], x_{\tau(i)}^{(k+1)}=y_{\pi\left(i+r_{k}\right)}^{(k)}$. For convenience, we define $y_{\pi(n+i)}^{(k)}=0, \forall i \in\left[r_{k}\right]$.

Case 1. We exit the while loop when $1.5 r_{k} \geq n$.
Let $u^{*}$ denote the largest $u$ s.t. $\left|y_{\pi(u)}^{(k)}\right| \geq \epsilon_{m p}$. If $u^{*}=r_{k}$, we have that $\left|y_{\pi\left(r_{k}\right)}^{(k)}\right| \geq \epsilon_{m p} \geq$ $\epsilon_{m p} / 100$. Otherwise, the condition of the loop shows that

$$
\left|y_{\pi\left(r_{k}\right)}^{(k)}\right| \geq(1-1 / \log n)^{\log _{1.5} r_{k}-\log _{1.5} u^{*}}\left|y_{\pi\left(u^{*}\right)}^{(k)}\right| \geq(1-1 / \log n)^{\log _{1.5} n} \epsilon_{m p} \geq \epsilon_{m p} / 100
$$

where we used that $n \geq 4$.
According to definition of $x_{\tau(i)}^{(k+1)}$, we have

$$
\begin{aligned}
\sum_{i=1}^{n} g_{i}\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\tau(i)}^{(k+1)}\right)\right) & =\sum_{i=1}^{n} g_{i}\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(y_{\pi\left(i+r_{k}\right)}^{(k)}\right)\right) \geq \sum_{i=n / 3+1}^{n} g_{i}\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(y_{\pi\left(i+r_{k}\right)}^{(k)}\right)\right) \\
& \geq \sum_{i=n / 3+1}^{n} g_{i}\left(\psi\left(y_{\pi(i)}^{(k)}\right)\right) \geq \sum_{i=n / 3+1}^{2 n / 3} g_{i} \psi\left(\epsilon_{m p} / 100\right) \geq \Omega\left(r_{k} g_{r_{k}} \epsilon_{m p}\right)
\end{aligned}
$$

where the first step follows from $x_{\tau(i)}^{(k+1)}=y_{\pi\left(i+r_{k}\right)}^{(k)}$, the second step follows from $\psi(|x|)$ is non-decreasing (part 2 of Lemma 2.5.9) and $\left|y_{\pi(i)}^{(k)}\right|$ is non-increasing, the third step follows from $1.5 r_{k}>n$ and hence $\psi\left(y_{\pi\left(i+r_{k}\right)}^{(k)}\right)=0$ for $i \geq n / 3+1$, the fourth step follows from $\psi$ is non-decreasing and $\left|y_{\pi(i)}^{(k)}\right| \geq\left|y_{\pi\left(r_{k}\right)}^{(k)}\right| \geq \epsilon_{m p} / 100$ for all $i<2 n / 3$, and the last step follows by $g$ is decreasing and part 3 of Lemma 2.5.9.

Case 2. We exit the while loop when $1.5 r_{k}<n$ and $\left|y_{\pi\left(1.5 r_{k}\right)}^{(k)}\right|<(1-1 / \log n)\left|y_{\pi\left(r_{k}\right)}^{(k)}\right|$.
By the same argument as Case 1, we have that $\left|y_{\pi\left(r_{k}\right)}^{(k)}\right| \geq \epsilon_{m p} / 100$. Part 3 of Lemma 2.5.9 together with the fact

$$
\left|y_{\pi(1.5 r)}^{(k)}\right|<\min \left(\epsilon_{m p},\left|y_{\pi(r)}^{(k)}\right| \cdot(1-1 / \log n)\right)
$$

shows that

$$
\begin{equation*}
\psi\left(\left|y_{\pi(1.5 r)}^{(k)}\right|\right)-\psi\left(\left|y_{\pi(r)}^{(k)}\right|\right)=\Omega\left(\epsilon_{m p} / \log n\right) \tag{2.27}
\end{equation*}
$$

Putting it all together, we have

$$
\begin{align*}
& \sum_{i=1}^{n} g_{i} \cdot\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\tau(i)}^{(k+1)}\right)\right) \\
= & \sum_{i=1}^{n} g_{i} \cdot\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(y_{\pi\left(i+r_{k}\right)}^{(k)}\right)\right) \\
\geq & \sum_{i=r_{k} / 2}^{r_{k}} g_{i} \cdot\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(y_{\pi\left(i+r_{k}\right)}^{(k)}\right)\right. \\
\geq & \sum_{i=r_{k} / 2}^{r_{k}} g_{i} \cdot\left(\psi\left(y_{\pi\left(r_{k}\right)}^{(k)}\right)-\psi\left(y_{\pi\left(1.5 r_{k}\right)}^{(k)}\right)\right. \\
\geq & \sum_{i=r_{k} / 2}^{r_{k}} g_{i} \cdot \Omega\left(\frac{\epsilon_{m p}}{\log n}\right) \\
\geq & \sum_{i=r_{k} / 2}^{r_{k}} g_{r_{k}} \cdot \Omega\left(\frac{\epsilon_{m p}}{\log n}\right) \\
= & \Omega\left(\epsilon_{m p}^{(k+1)} r_{k} g_{r_{k}} / \log n\right), y_{\pi\left(i+r_{k}\right)}^{(k)} \tag{2.27}
\end{align*}
$$

where the third step follows by $\left|y_{\pi(i)}^{(k)}\right|$ is decreasing and $\psi$ is non-decreasing (from part 2 of Lemma 2.5.9).


Figure 2.2: $\psi(x), \psi(x)^{\prime}$ and $\psi(x)^{\prime \prime}$. For $\epsilon_{m p} \in(0,1)$.

### 2.5.6 Potential function $\psi$

Lemma 2.5.9 (Properties of function $\psi$ ). Let function $\psi$ be defined in (2.21). Then function $\psi$ satisfies the following properties:

1. Symmetric $(\psi(-x)=\psi(x))$ and $\psi(0)=0$;
2. $\psi(|x|)$ is non-decreasing;
3. $\left|\psi^{\prime}(x)\right|=\Omega(1), \forall|x| \in\left[0.01 \epsilon_{m p}, \epsilon_{m p}\right]$;
4. $L_{1} \stackrel{\text { def }}{=} \max _{x} \psi^{\prime}(x)=1$ and $L_{2} \stackrel{\text { def }}{=} \max _{x} \psi^{\prime \prime}(x)=1 / \epsilon_{m p}$;
5. $\psi(x)$ is a constant for $|x| \geq 2 \epsilon_{m p}$.

Proof. We can see that
$\psi(x)^{\prime}=\left\{\begin{array}{ll}\frac{|x|}{\epsilon_{m p}}, & |x| \in\left[0, \epsilon_{m p}\right] \\ \frac{2 \epsilon_{m p}-|x|}{\epsilon_{m p}}, & |x| \in\left(\epsilon_{m p}, 2 \epsilon_{m p}\right] \\ 0, & x \in\left(2 \epsilon_{m p},+\infty\right)\end{array} \quad\right.$ and $\quad \psi(x)^{\prime \prime}= \begin{cases}\frac{1}{\epsilon_{m p}}, & x \in\left[0, \epsilon_{m p}\right] \cup\left[-2 \epsilon_{m p},-\epsilon_{m p}\right] \\ -\frac{1}{\epsilon_{m p}}, & x \in\left(\epsilon_{m p}, 2 \epsilon_{m p}\right] \cup\left[-\epsilon_{m p}, 0\right] \\ 0 . & x \in\left(2 \epsilon_{m p},+\infty\right)\end{cases}$
From the $\psi(x)^{\prime}$ and $\psi(x)^{\prime \prime}$, it is not hard to see that $\psi$ satisfies the properties needed.

### 2.6 Omitted Proofs

Lemma 2.6.1. Let $x$ and $y$ are (possibly dependent) random variables such that $|x| \leq c_{x}$ and $|y| \leq c_{y}$ almost surely. Then, we have

$$
\mathbb{V}[x y] \leq 2 c_{x}^{2} \cdot \mathbb{V}[y]+2 c_{y}^{2} \cdot \mathbb{V}[x]
$$

Proof. Recall that $\mathbb{V}[x y] \leq \mathbb{E}\left[(x y-t)^{2}\right]$ for any scalar $t$. Hence,

$$
\begin{aligned}
\mathbb{V}[x y] & \leq \mathbb{E}\left[(x y-\mathbb{E}[x] \mathbb{E}[y])^{2}\right]=\mathbb{E}\left[(x y-x \mathbb{E}[y]+x \mathbb{E}[y]-\mathbb{E}[x] \mathbb{E}[y])^{2}\right] \\
& \leq 2 \mathbb{E}\left[(x y-x \mathbb{E}[y])^{2}\right]+2 \mathbb{E}\left[(x \mathbb{E}[y]-\mathbb{E}[x] \mathbb{E}[y])^{2}\right] \\
& \leq 2 c_{x}^{2} \cdot \mathbb{V}[y]+2 c_{y}^{2} \cdot \mathbb{V}[x] .
\end{aligned}
$$

Lemma 2.6.2 ([Vai89b]). Given a matrix $A \in \mathbb{R}^{d \times n}$, vectors $b \in \mathbb{R}^{d}, c \in \mathbb{R}^{n}$. Suppose $x, s, y \in \mathbb{R}^{n}$ satisfy that $x s \approx_{0.1} t, A x=b$ and $A^{\top} y+s=c$ for some $t>0$. For any $\epsilon \in(0,1 / 2]$, in $\widetilde{O}\left(n^{2.5} \log (n / \epsilon)\right)$ time, we can find vectors $x^{\text {new }}, s^{\text {new }} \in \mathbb{R}^{n}$ and $y^{\text {new }} \in \mathbb{R}^{d}$ such that

$$
\begin{array}{r}
\left\|x^{\text {new }} s^{\text {new }}-t\right\|_{2} \leq \epsilon, \\
A x^{\mathrm{new}}=b, \\
A^{\top} y^{\mathrm{new}}+s=c .
\end{array}
$$

Remark 2.6.1. Instead of using this, one can also run our algorithm with $k=n$ for $O(\sqrt{n} \log n)$ iterations. Since $k=n$, there is no randomness involves and hence $\Phi$ will decrease deterministically to $O(n)$.

Lemma 2.6.3. $\omega \leq 3-\alpha$.

Proof. We consider a $n \times n$ matrix $A$ multiply another $n \times n B$, we split $A$ into $n^{1-\alpha}$ fat matrices where each of them has size $n^{\alpha} \times n$. Since $\omega$ is the best exponent of matrix multiplication, thus we know

$$
n^{\omega+o(1)} \leq n^{1-\alpha} \cdot n^{2+o(1)}
$$

which implies $\omega \leq 3-\alpha$.

Lemma 2.6.4 (Rectangular matrix multiplication). For any $n \geq r$, multiplying an $n \times r$ with an $r \times n$ matrix or $n \times n$ with $n \times r$ takes time

$$
n^{2+o(1)}+r^{\frac{\omega-2}{1-\alpha}} n^{2-\frac{\alpha(\omega-2)}{1-\alpha}+o(1)} .
$$

Proof. The cost for multiplying a $n \times n$ and a $n \times r$ matrix is the same as multiplying a $n \times r$ and a $r \times n$ matrix. So, we focus on the later case.

For the case $r \leq n^{\alpha}$, it follows from the rectangular matrix multiplication result in [LGU18].

For the case $r \geq n^{\alpha}$, we let $k=(n / r)^{\frac{1}{1-\alpha}}$. We can view the problem as multiplying a $k \times k^{\alpha}$ and a $k^{\alpha} \times k$ block matrices and each block has size $\frac{n}{k} \times \frac{n}{k}$ size. Therefore, the total cost is

$$
k^{2+o(1)} \times\left(\frac{n}{k}\right)^{\omega+o(1)}=r^{\frac{\omega-2}{1-\alpha}} n^{2-\frac{\alpha(\omega-2)}{1-\alpha}+o(1)} .
$$

For a matrix $A$, we define $\|A\|_{1}$ to be $\sum_{i, j}\left|A_{i, j}\right|$.

Lemma 2.6.5. Consider a linear program $\min _{A x=b, x \geq 0} c^{\top} x$ with $n$ variables and $d$ constraints. Assume that

1. Diameter of the polytope : For any $x \geq 0$ with $A x=b$, we have that $\|x\|_{\infty} \leq R$.
2. Lipschitz constant of the linear program: $\|c\|_{\infty} \leq L$.

For any $\delta \in(0,1]$, the modified linear program $\min _{\bar{A} \bar{x}=\bar{b}, \bar{x} \geq 0} \bar{c}^{\top} \bar{x}$ with

$$
\bar{A}=\left[\begin{array}{ccc}
A & 0 & \frac{1}{R} b-A 1_{n} \\
1_{n}^{\top} & 1 & 0
\end{array}\right] \in \mathbb{R}^{(d+1) \times(n+2)}, \bar{b}=\left[\begin{array}{c}
\frac{1}{R} b \\
n+1
\end{array}\right] \in \mathbb{R}^{d+1} \text { and, } \bar{c}=\left[\begin{array}{c}
\frac{\delta}{L} \cdot c \\
0 \\
1
\end{array}\right] \in \mathbb{R}^{n+2}
$$

satisfies the following :

1. $\bar{x}=\left[\begin{array}{c}1_{n} \\ 1 \\ 1\end{array}\right], \bar{y}=\left[\begin{array}{c}0_{d} \\ -1\end{array}\right]$ and $\bar{s}=\left[\begin{array}{c}1_{n}+\frac{\delta}{L} \cdot c \\ 1 \\ 1\end{array}\right]$ are feasible primal dual vectors.
2. For any feasible primal dual vectors $(\bar{x}, \bar{y}, \bar{s})$ with duality gap $\leq \delta^{2}$, consider the vector $\widehat{x}=R \cdot \bar{x}_{1: n}\left(\bar{x}_{1: n}\right.$ is the first $n$ coordinates of $\left.x\right)$ is an approximate solution to the original linear program in the following sense

$$
\begin{aligned}
& c^{\top} \widehat{x} \leq \min _{A x=b, x \geq 0} c^{\top} x+L R \cdot \delta, \\
&\|A \widehat{x}-b\|_{1} \leq 2 n \delta \cdot\left(R\|A\|_{1}+\|b\|_{1}\right), \\
& \widehat{x} \geq 0
\end{aligned}
$$

Proof. For the first result, straightforward calculations show that $(\bar{x}, \bar{y}, \bar{s})$ are feasible, i.e.,

$$
\bar{A} \bar{x}=\left[\begin{array}{ccc}
A & 0 & \frac{1}{R} b-A 1_{n} \\
1_{n}^{\top} & 1 & 0
\end{array}\right] \cdot\left[\begin{array}{c}
1_{n} \\
1 \\
1
\end{array}\right]=\left[\begin{array}{c}
\frac{1}{R} b \\
n+1
\end{array}\right]=\bar{b}
$$

and

$$
\begin{aligned}
\bar{A}^{\top} \bar{y}+\bar{s} & =\left[\begin{array}{cc}
A^{\top} & 1_{n} \\
0 & 1 \\
\frac{1}{R} b^{\top}-1_{n}^{\top} A^{\top} & 0
\end{array}\right] \cdot\left[\begin{array}{c}
0_{d} \\
-1
\end{array}\right]+\left[\begin{array}{c}
1_{n}+\frac{\delta}{L} \cdot c \\
1 \\
1
\end{array}\right] \\
& =\left[\begin{array}{c}
-1_{n} \\
-1 \\
0
\end{array}\right]+\left[\begin{array}{c}
1_{n}+\frac{\delta}{L} \cdot c \\
1 \\
1
\end{array}\right] \\
& =\left[\begin{array}{c}
\frac{\delta}{L} \cdot c \\
0 \\
1
\end{array}\right] \\
& =\bar{c}
\end{aligned}
$$

For the second result, we let

$$
\mathrm{OPT}=\min _{A x=b, x \geq 0} c^{\top} x, \quad \text { and, } \quad \overline{\mathrm{OPT}}=\min _{\bar{A} \bar{x}=\bar{b}, \bar{x} \geq 0} \bar{c}^{\top} \bar{x}
$$

For any optimal $x$ in the original LP, we consider the following $\bar{x}$

$$
\bar{x}=\left[\begin{array}{c}
\frac{1}{R} x  \tag{2.28}\\
n+1-\frac{1}{R} \sum_{i} x_{i} \\
0
\end{array}\right]
$$

We want to argue that $\bar{x}$ is feasible in the modified LP. It is obvious that $\bar{x} \geq 0$, it remains to show $\bar{A} \bar{x}=\bar{b}$. We have

$$
\bar{A} \bar{x}=\left[\begin{array}{ccc}
A & 0 & \frac{1}{R} b-A 1_{n} \\
1_{n}^{\top} & 1 & 0
\end{array}\right] \cdot\left[\begin{array}{c}
\frac{1}{R} x \\
n+1-\frac{1}{R} \sum_{i} x_{i} \\
0
\end{array}\right]=\left[\begin{array}{c}
\frac{1}{R} A x \\
n+1
\end{array}\right]=\left[\begin{array}{c}
\frac{1}{R} b \\
n+1
\end{array}\right]=\bar{b},
$$

where the third step follows from $A x=b$, and the last step follows from definition of $\bar{b}$.
Therefore, using the definition of $\bar{x}$ in (2.28) we have that

$$
\overline{\mathrm{OPT}} \leq \bar{c}^{\top} \bar{x}=\left[\begin{array}{lll}
\frac{\delta}{L} \cdot c^{\top} & 0 & 1
\end{array}\right] \cdot\left[\begin{array}{c}
\frac{1}{R} x  \tag{2.29}\\
n+1-\frac{1}{R} \sum_{i} x_{i} \\
0
\end{array}\right]=\frac{\delta}{L R} \cdot c^{\top} x=\frac{\delta}{L R} \cdot \mathrm{OPT} .
$$

where the first step follows from modified program is solving a minimization problem, the second step follows from definition of $\bar{c}$ and (2.28), the last step follows from $x$ is an optimal solution in the original linear program.

Given a feasible $(\bar{x}, \bar{y}, \bar{s})$ with duality gap $\delta^{2}$. Write $\bar{x}=\left[\begin{array}{c}\bar{x}_{1: n} \\ \tau \\ \theta\end{array}\right]$ for some $\tau \geq 0, \theta \geq 0$.
We can compute $\bar{c}^{\top} \bar{x}$ which is $\frac{\delta}{L} \cdot c^{\top} \bar{x}_{1: n}+\theta$. Then, we have

$$
\begin{equation*}
\frac{\delta}{L} \cdot c^{\top} \bar{x}_{1: n}+\theta \leq \overline{\mathrm{OPT}}+\delta^{2} \leq \frac{\delta}{L R} \cdot \mathrm{OPT}+\delta^{2} \tag{2.30}
\end{equation*}
$$

where the first step follows from definition of duality gap, the last step follows from (2.29).
Hence, we can upper bound the OPT of the transformed program as follows:

$$
c^{\top} \widehat{x}=R \cdot c^{\top} \bar{x}_{1: n}=\frac{R L}{\delta} \cdot \frac{\delta}{L} c^{\top} \bar{x}_{1: n} \leq \frac{R L}{\delta}\left(\frac{\delta}{L R} \cdot \mathrm{OPT}+\delta^{2}\right)=\mathrm{OPT}+L R \cdot \delta
$$

where the first step follows by $\widehat{x}=R \cdot \bar{x}_{1: n}$, the third step follows by (2.30).
We can upper bound the $\theta$ in the following sense,

$$
\theta \leq \frac{\delta}{L R} \cdot \mathrm{OPT}+\delta^{2} \leq n \delta+\delta^{2} \leq 2 n \delta
$$

where the first step follows from (2.30) and $\frac{\delta}{L} c^{\top} \bar{x}_{1: n} \geq 0$, the second step follows by OPT $=$ $\min _{A x=b, x \geq 0} c^{\top} x \leq n L R$ (because $\|c\|_{\infty} \leq L$ and $\|x\|_{\infty} \leq R$ ), and the last step follows from $\delta \leq 1 \leq n$.

The constraint in the new polytope shows that

$$
A \bar{x}_{1: n}+\left(\frac{1}{R} b-A 1_{n}\right) \theta=\frac{1}{R} b .
$$

Using $\widehat{x}=R x_{1: n}$, we have

$$
A \frac{1}{R} \widehat{x}+\left(\frac{1}{R} b-A 1_{n}\right) \theta=\frac{1}{R} b .
$$

Rewriting it, we have $A \widehat{x}-b=\left(R A 1_{n}-b\right) \theta$ and hence

$$
\begin{aligned}
\|A \widehat{x}-b\|_{1} & =\left\|\left(R A 1_{n}-b\right) \theta\right\|_{1} \leq \theta\left(\left\|R A 1_{n}\right\|_{1}+\|b\|_{1}\right) \leq \theta \cdot\left(R\|A\|_{1}+\|b\|_{1}\right) \\
& \leq 2 n \delta \cdot\left(R\|A\|_{1}+\|b\|_{1}\right)
\end{aligned}
$$

## Chapter 3

## Empirical Risk Minimization

Many convex problems in machine learning and computer science share the same form:

$$
\min _{x} \sum_{i} f_{i}\left(A_{i} x+b_{i}\right),
$$

where $f_{i}$ are convex functions on $\mathbb{R}^{n_{i}}$ with constant $n_{i}, A_{i} \in \mathbb{R}^{n_{i} \times d}, b_{i} \in \mathbb{R}^{n_{i}}$ and $\sum_{i} n_{i}=n$. This problem generalizes linear programming and includes many problems in empirical risk minimization.

In this chapter, we give an algorithm that runs in time

$$
O^{*}\left(\left(n^{\omega}+n^{2.5-\alpha / 2}+n^{2+1 / 6}\right) \log (n / \delta)\right)
$$

where $\omega$ is the exponent of matrix multiplication, $\alpha$ is the dual exponent of matrix multiplication, and $\delta$ is the relative accuracy. Note that the runtime has only a $\log$ dependence on the condition numbers or other data dependent parameters and these are captured in $\delta$. For the current bound $\omega \sim 2.38$ [Vassilevska Williams'12, Le Gall'14] and $\alpha \sim 0.31$ [Le Gall, Urrutia'18], our runtime $O^{*}\left(n^{\omega} \log (n / \delta)\right)$ matches the current best for solving a dense least squares regression problem, a special case of the problem we consider. Very recently, [Alman'18] proved that all the current known techniques can not give a better $\omega$ below 2.168 which is larger than our $2+1 / 6$.

Our result generalizes the very recent result of solving linear programs in the current matrix multiplication time [Cohen, Lee, Song'19] to a more broad class of problems. Our algorithm proposes two concepts which are different from [Cohen, Lee, Song'19] :

- We give a robust deterministic central path method, whereas the previous one is a stochastic central path which updates weights by a random sparse vector.
- We propose an efficient data-structure to maintain the central path of interior point methods even when the weights update vector is dense.


### 3.1 Introduction

Empirical Risk Minimization (ERM) problem is a fundamental question in statistical machine learning. There are a huge number of papers that have considered this topic [Nes83, Vap92, PJ92, Nes04, BBM05, BB08, NJLS09, MB11, FGRW12, LRSB12, JZ13, Vap13, SSZ13, DB14, DBLJ14, FGKS15, DB16, SLC ${ }^{+} 17$, ZYJ17, ZX17, ZWX+17, GSS17, MS17, NS17, AKK+17, Csi18, JLGJ18] as almost all convex optimization machine learning can be phrased in the ERM framework [SSBD14, Vap92]. While the statistical convergence properties and generalization bounds for ERM are well-understood, a general runtime bound for general ERM is not known although fast runtime bounds do exist for specific instances [AKPS19].

Examples of applications of ERM include linear regression, LASSO [Tib96], elastic net [ZH05], logistic regression [Cox58, HJLS13], support vector machines [CV95], $\ell_{p}$ regression [Cla05, $\mathrm{DDH}^{+}$09, BCLL18, AKPS19], quantile regression [Koe00, KH01, Koe05], AdaBoost [FS97], kernel regression [Nad64, Wat64], and mean-field variational inference [XJR02].

The classical Empirical Risk Minimization problem is defined as

$$
\min _{x} \sum_{i=1}^{m} f_{i}\left(a_{i}^{\top} x+b_{i}\right)
$$

where $f_{i}: \mathbb{R} \rightarrow \mathbb{R}$ is a convex function, $a_{i} \in \mathbb{R}^{d}$, and $b_{i} \in \mathbb{R}, \forall i \in[m]$. Note that this formulation also captures most standard forms of regularization as well.

Letting $y_{i}=a_{i}^{\top} x+b_{i}$, and $z_{i}=f_{i}\left(a_{i}^{\top} x+b_{i}\right)$ allows us to rewrite the original problem
in the following sense,

$$
\begin{array}{ll}
\min _{x, y, z} & \sum_{i=1}^{m} z_{i}  \tag{3.1}\\
\text { s.t. } & A x+b=y \\
& \left(y_{i}, z_{i}\right) \in K_{i}=\left\{\left(y_{i}, z_{i}\right): f_{i}\left(y_{i}\right) \leq z_{i}\right\}, \forall i \in[m]
\end{array}
$$

We can consider a more general version where dimension of $K_{i}$ can be arbitrary, e.g. $n_{i}$. Therefore, we come to study the general $n$-variable form

$$
\min _{x \in \prod_{i=1}^{m} K_{i}, A x=b} c^{\top} x
$$

where $\sum_{i=1}^{m} n_{i}=n$. We state our main result for solving the general model.
Theorem 3.1.1 (Main result, informal version of Theorem 3.7.3). Given a matrix $A \in \mathbb{R}^{d \times n}$, two vectors $b \in \mathbb{R}^{d}, c \in \mathbb{R}^{n}$, and $m$ compact convex sets $K_{1}, K_{2}, \cdots, K_{m}$. Assume that there is no redundant constraints and $n_{i}=O(1), \forall i \in[m]$. There is an algorithm (procedure Main in Algorithm 3.6) that solves

$$
\min _{x \in \prod_{i=1}^{m} K_{i}, A x=b} c^{\top} x
$$

up to $\delta$ precision and runs in expected time

$$
\widetilde{O}\left(\left(n^{\omega+o(1)}+n^{2.5-\alpha / 2+o(1)}+n^{2+1 / 6+o(1)}\right) \cdot \log \left(\frac{n}{\delta}\right)\right)
$$

where $\omega$ is the exponent of matrix multiplication, $\alpha$ is the dual exponent of matrix multiplication.

For the current value of $\omega \sim 2.38$ [Wil12, LG14] and $\alpha \sim 0.31$ [LGU18], the expected time is simply $n^{\omega+o(1)} \widetilde{O}\left(\log \left(\frac{n}{\delta}\right)\right)$.

Remark 3.1.1. More precisely, when $n_{i}$ is super constant, our running time depends polynomially on $\max _{i \in[m]} n_{i}$ (but not exponential dependence).

Also note that our runtime depends on diameter, but logarithmically to the diameter. So, it can be applied to linear program by imposing an artificial bound on the solution.

### 3.1.1 Related Work

First-order algorithms for ERM are well-studied and a long series of accelerated stochastic gradient descent algorithms have been developed and optimized [Nes98, JZ13, XZ14, SSZ14, FGKS15, LMH15, MLF15, AY16, RHS+16, SS16, AH16, SLRB17, MS17, LMH17, LJCJ17, All17b, All17a, All18b, All18a]. However, these rates depend polynomially on the Lipschitz constant of $\nabla f_{i}$ and in order to achieve a $\log (1 / \epsilon)$ dependence, the runtime will also have to depend on the strong convexity of the $\sum_{i} f_{i}$. In this chapter, we want to focus on algorithms that depend logarithmically on diameter/smoothness/strong convexity constants, as well as the error parameter $\epsilon$. Note that gradient descent and a direct application of Newton's method do not belong to these class of algorithms, but for example, interior point method and ellipsoid method does.

Therefore, in order to achieve high-accuracy solutions for non-smooth and non strongly convex case, most convex optimization problems will rely on second-order methods, often under the general interior point method (IPM) or some sort of iterative refinement framework. So, we note that our algorithm is thus optimal in this general setting since second-order methods require at least $n^{\omega}$ runtime for general matrix inversion.

Our algorithm applies the interior point method framework to solve ERM. The most general interior point methods require $O(\sqrt{n})$-iterations of linear system solves [Nes98], requiring a naive runtime bound of $O\left(n^{\omega+1 / 2}\right)$. Using the inverse maintenance technique [Vai89b, CLS19], one can improve the running time for LP to $O\left(n^{\omega}\right)$. This essentially implies that almost all convex optimization problems can be solved, up to subpolynomial factors, as fast as linear regression or matrix inversion!

The specific case of $\ell_{2}$ regression can be solved in $O\left(n^{\omega}\right)$ time since the solution is explicitly given by solving a linear system. In the more general case of $\ell_{p}$ regression, [BCLL18] proposed a $\widetilde{O}_{p}\left(n^{|1 / 2-1 / p|}\right)$-iteration iterative solver with a naive $O\left(n^{\omega}\right)$ system solve at each step. Recently, [AKPS19] improved the runtime to $\widetilde{O}_{p}\left(n^{\max (\omega, 7 / 3)}\right)$, which is current matrix multiplication time as $\omega>7 / 3$. However, both these results depend exponentially on $p$ and fail to be impressive for large $p$. Otherwise, we are unaware of other ERM formulations that have have general runtime bounds for obtaining high-accuracy solutions.

Recently several works [AW18a, AW18b, Alm19] try to show the limitation of current known techniques for improving matrix multiplication time. Alman and Vassilevska Williams [AW18b] proved limitations of using the Galactic method applied to many tensors of interest (including Coppersmith-Winograd tensors [CW87]). More recently, Alman [Alm19] proved that by applying the Universal method on those tensors, we cannot hope to achieve any running time better than $n^{2.168}$ which is already above our $n^{2+1 / 6}$.

### 3.2 Overview of Techniques

In this section, we discuss the key ideas in this chapter. Generalizing the stochastic sparse update approach of [CLS19] to our setting is a natural first step to speeding up the matrix-vector multiplication that is needed in each iteration of the interior point method. In linear programs, maintaining approximate complementary slackness means that we maintain $x, s$ to be close multiplicatively to the central path under some notion of distance. However, the generalized notion of complementary slackness requires a barrier-dependent notion of distance. Specifically, if $\phi(x)$ is a barrier function, then our distance is now defined as our function gradient being small in a norm depending on $\nabla^{2} \phi(x)$. One key fact of the stochastic sparse update is that the variance introduced does not perturb the approximation too much, which requires understanding the second derivative of the distance function. For our setting, this would require bounding the 4th derivative of $\phi(x)$, which may not exist for self-concordant functions. So, the stochastic approach may not work algorithmically (not just in the analysis) if $\phi(x)$ is assumed to be simply self-concordant. Even when assumptions on the 4th derivative of $\phi(x)$ are made, the analysis will become significantly more complicated due to the 4th derivative terms. To avoid these problems, the main contributions of this chapter is to 1 ) introduce a robust version of the central path and 2) exploit the robustness via sketching to apply the desired matrix-vector multiplication fast.

More generally, our main observation is that one can generally speed up an iterative method using sketching if the method is robust in a certain sense. To speed up interior point methods, in Section 3.4 and 3.5, we give a robust version of the interior point method; and in Section 3.6, we give a data structure to maintain the sketch; and in Section 3.7, we show how to combine them together. We provide several basic notations and definitions for
numerical linear algebra in Section 3.3. In Section 3.8, we provide some classical lemmas from the literature of interior point methods. In Section 3.9, we prove some basic properties of the sketching matrix. Now, we first begin with an overview of our robust central path and then proceed with an overview of sketching iterative methods.

### 3.2.1 Central Path Method

We consider the following optimization problem

$$
\begin{equation*}
\min _{x \in \prod_{i=1}^{m} K_{i}, A x=b} c^{\top} x \tag{3.2}
\end{equation*}
$$

where $\prod_{i=1}^{m} K_{i}$ is the direct product of $m$ low-dimensional convex sets $K_{i}$. We let $x_{i}$ be the $i$-th block of $x$ corresponding to $K_{i}$. Interior point methods consider the path of solutions to the following optimization problem:

$$
\begin{equation*}
x(t)=\arg \min _{A x=b} c^{\top} x+t \sum_{i=1}^{m} \phi_{i}\left(x_{i}\right) \tag{3.3}
\end{equation*}
$$

where $\phi_{i}: K_{i} \rightarrow \mathbb{R}$ are self-concordant barrier functions. This parameterized path is commonly known as the central path. Many algorithms solve the original problem (3.2) by following the central path as the path parameter is decreased $t \rightarrow 0$. The rate at which we decrease $t$ and subsequently the runtimes of these path-following algorithms are usually governed by the self-concordance properties of the barrier functions we use.

Definition 3.2.1. We call a function $\phi$ a $\nu$ self-concordant barrier for $K$ if $\operatorname{dom} \phi=K$ and for any $x \in \operatorname{dom} \phi$ and for any $u \in \mathbb{R}^{n}$

$$
\left|D^{3} \phi(x)[u, u, u]\right| \leq 2\|u\|_{x}^{3 / 2} \quad \text { and } \quad\|\nabla \phi(x)\|_{x}^{*} \leq \sqrt{\nu}
$$

where $\|v\|_{x}:=\|v\|_{\nabla^{2} \phi(x)}$ and $\|v\|_{x}^{*}:=\|v\|_{\nabla^{2} \phi(x)^{-1}}$, for any vector $v$.
Remark 3.2.1. It is known that $\nu \geq 1$ for any self-concordant barrier function.

Nesterov and Nemirovsky showed that for any open convex set $K \subset \mathbb{R}^{n}$, there is a $O(n)$ self-concordant barrier function [Nes98]. In this chapter, the convex set $K_{i}$ we
considered has $O(1)$ dimension. While Nesterov and Nemirovsky gave formulas for the universal barrier; in practice, most ERM problems lend themselves to explicit $O(1)$ selfconcordant barriers for majority of the convex functions people use. For example, for the set $\{x:\|x\|<1\}$, we use $-\log \left(1-\|x\|^{2}\right)$; for the set $\{x: x>0\}$, we use $-\log (x)$, and so on. That is the reason why we assume the gradient and hessian can be computed in $O(1)$ time. Therefore, in this chapter, we assume a $\nu_{i}$ self-concordant barrier $\phi_{i}$ is provided and that we can compute $\nabla \phi_{i}$ and $\nabla^{2} \phi_{i}$ in $O(1)$ time. The main result we will use about self-concordance is that the norm $\|\cdot\|_{x}$ is stable when we change $x$.

Theorem 3.2.1 (Theorem 4.1.6 in [Nes98]). If $\phi$ is a self-concordant barrier and if $\|y-x\|_{x}<$ 1, then we have:

$$
\left(1-\|y-x\|_{x}\right)^{2} \nabla^{2} \phi(x) \preceq \nabla^{2} \phi(y) \preceq \frac{1}{\left(1-\|y-x\|_{x}\right)^{2}} \nabla^{2} \phi(x) .
$$

In general, we can simply think of $\phi_{i}$ as a function penalizing any point $x_{i} \notin K_{i}$. It is known how to transform the original problem (3.2) by adding $O(n)$ many variables and constraints so that

- The minimizer $x(t)$ at $t=1$ is explicitly given.
- One can obtain an approximate solution of the original problem using the minimizer at small $t$ in linear time.

For completeness, we show how to do it in Lemma 3.8.2. Therefore, it suffices to study how we can move efficiently from $x(1)$ to $x(\epsilon)$ for some tiny $\epsilon$ where $x(t)$ is again the minimizer of the problem (3.3).

### 3.2.2 Robust Central Path

In the standard interior point method, we use a tight $\ell_{2}$-bound to control how far we can deviate from $x(t)$ during the entirety of the algorithm. Specifically, if we denote $\gamma_{i}^{t}\left(x_{i}\right)$ as the appropriate measure of error (this will be specified later and is often called the Newton Decrement) in each block coordinate $x_{i}$ at path parameter $t$, then as we let $t \rightarrow 0$, the old invariant that we are maintaining is,

$$
\Phi_{\mathrm{old}}^{t}(x)=\sum_{i=1}^{m} \gamma_{i}^{t}\left(x_{i}\right)^{2} \leq O(1) .
$$

It can be shown that a Newton step in the standard direction will allow for us to maintain $\Phi_{\text {old }}^{t}$ to be small even as we decrease $t$ by a multiplicative factor of $O\left(m^{-1 / 2}\right)$ in each iteration, thereby giving a standard $O(\sqrt{m})$ iteration analysis. Therefore, the standard approach can be seen as trying to remain within a small $\ell_{2}$ neighborhood of the central path by centering with Newton steps after making small decreases in the path parameter $t$. Note however that if each $\gamma_{i}$ can be perturbed by an error that is $\Omega\left(m^{-1 / 2}\right)$, $\Phi_{\text {old }}^{t}(x)$ can easily become too large for the potential argument to work.

To make our analysis more robust, we introduce a robust version that maintains the soft-max potential:

$$
\Phi_{\text {new }}^{t}(x)=\sum_{i=1}^{m} \exp \left(\lambda \gamma_{i}^{t}\left(x_{i}\right)\right) \leq O(m)
$$

for some $\lambda=\Theta(\log m)$. The robust central path is simply the region of all $x$ that satisfies our potential inequality. We will specify the right constants later but we always make $\lambda$ large enough to ensure that $\gamma_{i} \leq 1$ for all $x$ in the robust central path. Now note that a $\ell_{\infty}$ perturbation of $\gamma$ translates into a small multiplicative change in $\Phi^{t}$, tolerating errors on each $\gamma_{i}$ of up to $O(1 / \operatorname{poly} \log (n))$.

However, maintaining $\Phi_{\text {new }}^{t}(x) \leq O(m)$ is not obvious because the robust central path is a much wider region of $x$ than the typical $\ell_{2}$-neighborhood around the central path. We will show later how to modify the standard Newton direction to maintain $\Phi_{\text {new }}^{t}(x) \leq O(m)$ as we decrease $t$. Specifically, we will show that a variant of gradient descent of $\Phi_{\text {new }}^{t}$ in the Hessian norm suffices to provide the correct guarantees.

### 3.2.3 Speeding up via Sketching

To motivate our sketching algorithm, we consider an imaginary iterative method

$$
z^{(k+1)} \leftarrow z^{(k)}+P \cdot F\left(z^{(k)}\right)
$$

where $P$ is some dense matrix and $F(z)$ is some simple formula that can be computed efficiently in linear time. Note that the cost per iteration is dominated by multiplying $P$ with a vector, which takes $O\left(n^{2}\right)$ time. To avoid the cost of multiplication, instead of storing the solution explicitly, we store it implicitly by $z^{(k)}=P \cdot u^{(k)}$. Now, the algorithm becomes

$$
u^{(k+1)} \leftarrow u^{(k)}+F\left(P \cdot u^{(k)}\right)
$$

This algorithm is as expensive as the previous one except that we switch the location of $P$. However, if we know the algorithm is robust under perturbation of the $z^{(k)}$ term in $F\left(z^{(k)}\right)$, we can instead do

$$
u^{(k+1)} \leftarrow u^{(k)}+F\left(R^{\top} R P \cdot u^{(k)}\right)
$$

for some random Gaussian matrix $R: \mathbb{R}^{b \times n}$. Note that the matrix $R P$ is fixed throughout the whole algorithm and can be precomputed. Therefore, the cost of per iteration decreases from $O\left(n^{2}\right)$ to $O(n b)$.

For our problem, we need to make two adjustments. First, we need to sketch the change of $z$, that is $F\left(P \cdot u^{(k)}\right)$, instead of $z^{(k)}$ directly because the change of $z$ is smaller and this creates a smaller error. Second, we need to use a fresh random $R$ every iteration to avoid the randomness dependence issue in the proof. For the imaginary iterative process, it
becomes

$$
\begin{aligned}
& \bar{z}^{(k+1)} \leftarrow \bar{z}^{(k)}+R^{(k) \top} R^{(k)} P \cdot F\left(\bar{z}^{(k)}\right), \\
& u^{(k+1)} \leftarrow u^{(k)}+F\left(\bar{z}^{(k)}\right) .
\end{aligned}
$$

After some iterations, $\bar{z}^{(k)}$ becomes too far from $z^{(k)}$ and hence we need to correct the error by setting $z^{(k)}=P \cdot u^{(k)}$, which zeros the error.

Note that the algorithm explicitly maintains the approximate vector $\bar{z}$ while implicitly maintaining the exact vector $z$ by $P u^{(k)}$. This is different from the classical way to sketch Newton method [PW16, PW17], which is to simply run $z^{(k+1)} \leftarrow z^{(k)}+R^{\top} R P \cdot F\left(z^{(k)}\right)$ or use another way to subsample and approximate $P$. Such a scheme relies on the iteration method to fix the error accumulated in the sketch, while we are actively fixing the error by having both the approximate explicit vector $\bar{z}$ and the exact implicit vector $z$.

Without precomputation, the cost of computing $R^{(k)} P$ is in fact higher than that of $P \cdot F\left(z^{(k)}\right)$. The first one involves multiplying multiple vectors with $P$ and the second one involves multiplying 1 vector with $P$. However, we can precompute $\left[R^{(1) \top} ; R^{(2) \top} ; \cdots ; R^{(T) \top}\right]^{\top}$. $P$ by fast matrix multiplication. This decreases the cost of multiplying 1 vector with $P$ to $n^{\omega-1}$ per vector. This is a huge saving from $n^{2}$. In our algorithm, we end up using only $\widetilde{O}(n)$ random vectors in total and hence the total cost is still roughly $n^{\omega}$.

### 3.2.4 Maintaining the Sketch

The matrix $P$ we use in interior point methods is of the form

$$
P=\sqrt{W} A^{\top}\left(A W A^{\top}\right)^{-1} A \sqrt{W}
$$

where $W$ is some block diagonal matrix. [CLS19] showed one can approximately maintain the matrix $P$ with total cost $\widetilde{O}\left(n^{\omega}\right)$ across all iterations of interior point method. However, the cost of applying the dense matrix $P$ with a vector $z$ is roughly $O\left(n\|z\|_{0}\right)$ which is $O\left(n^{2}\right)$ for dense vectors. Since interior point methods takes at least $\sqrt{n}$ iterations in general, this gives a total runtime of $O\left(n^{2.5}\right)$. The key idea in [CLS19] is that one can design a stochastic interior point method such that each step only need to multiply $P$ with a vector of density $\widetilde{O}(\sqrt{n})$. This bypasses the $n^{2.5}$ bottleneck.

In this chapter, we do not have this issue because we only need to compute $R P z$ which is much cheaper than $P z$. We summarize why it suffices to maintain $R P$ throughout the algorithm. In general, for interior point method, the vector $z$ is roughly an unit vector and since $P$ is an orthogonal projection, we have $\|P z\|_{2}=O(1)$. One simple insight we have is that if we multiply a random $\sqrt{n} \times n$ matrix $R$ with values $\pm \frac{1}{\sqrt{n}}$ by $P z$, we have $\|R P z\|_{\infty}=\widetilde{O}\left(\frac{1}{\sqrt{n}}\right)$ (Lemma 3.9.1). Since there are $\widetilde{O}(\sqrt{n})$ iterations in interior point method, the total error is roughly $\widetilde{O}(1)$ in a correctly reweighed $\ell_{\infty}$ norm. In Section 3.5, we showed that this is exactly what interior point method needs for convergence. Furthermore, we note that though each step needs to use a fresh random matrix $R_{l}$ of size $\sqrt{n} \times n$, the random matrices $\left[R_{1}^{\top} ; R_{2}^{\top} ; \cdots ; R_{T}^{\top}\right]^{\top}$ we need can all fit into $\widetilde{O}(n) \times n$ budget. Therefore, throughout the algorithm, we simply need to maintain the matrix $\left[R_{1}^{\top} ; R_{2}^{\top} ; \cdots ; R_{T}^{\top}\right]^{\top} P$ which can be done with total cost $\widetilde{O}\left(n^{\omega}\right)$ across all iterations using idea similar to [CLS19].

The only reason the data structure looks complicated is that when the block matrix $W$ changes in different location in $\sqrt{W} A^{\top}\left(A W A^{\top}\right)^{-1} A \sqrt{W}$, we need to update the matrix $\left[R_{1} ; R_{2} ; \cdots ; R_{T}\right] P$ appropriately. This gives us few simple cases to handle in the algorithm and in the proof. For the intuition on how to maintain $P$ under $W$ change, see [CLS19, Section 2.2 and 5.1].

### 3.2.5 Fast rectangular matrix multiplication

Given two size $n \times n$ matrices, the time of multiplying them is $n^{2.81}<n^{3}$ by applying Strassen's original algorithm [Str69]. The current best running time takes $n^{\omega}$ time where $\omega<2.373$ [Wil12, LG14]. One natural extension of multiplying two square matrices is multiplying two rectangular matrices. What is the running time of multiplying one $n \times n^{a}$ matrix with another $n^{a} \times n$ matrix? Let $\alpha$ denote the largest upper bound of $a$ such that multiplying two rectangular matrices takes $n^{2+o(1)}$ time. The $\alpha$ is called the dual exponent of matrix multiplication, and the state-of-the-art result is $\alpha=0.31$ [LGU18]. We use the similar idea as [CLS19] to delay the low-rank update when the rank is small.

### 3.3 Preliminaries

Given a vector $x \in \mathbb{R}^{n}$ and $m$ compact convex sets $K_{1} \subset \mathbb{R}^{n_{1}}, K_{2} \subset \mathbb{R}^{n_{2}}, \cdots, K_{m} \subset$ $\mathbb{R}^{n_{m}}$ with $\sum_{i=1}^{m} n_{i}=n$. We use $x_{i}$ to denote the $i$-th block of $x$, then $x \in \prod_{i=1}^{m} K_{i}$ if $x_{i} \in K_{i}$, $\forall i \in[m]$.

We say a block diagonal matrix $A \in \oplus_{i=1}^{m} \mathbb{R}^{n_{i} \times n_{i}}$ if $A$ can be written as

$$
A=\left[\begin{array}{llll}
A_{1} & & & \\
& A_{2} & & \\
& & \ddots & \\
& & & A_{m}
\end{array}\right]
$$

where $A_{1} \in \mathbb{R}^{n_{1} \times n_{1}}, A_{2} \in \mathbb{R}^{n_{2} \times n_{2}}$, and $A_{m} \in \mathbb{R}^{n_{m} \times n_{m}}$. For a matrix $A$, we use $\|A\|_{F}$ to denote its Frobenius norm and use $\|A\|$ to denote its operator norm. There are some trivial facts $\|A B\|_{2} \leq\|A\|_{2} \cdot\|B\|_{2}$ and $\|A B\|_{F} \leq\|A\|_{F} \cdot\|B\|_{2}$.

For notation convenience, we assume the number of variables $n \geq 10$ and there are no redundant constraints. In particular, this implies that the constraint matrix $A$ is full rank.

For a positive integer $n$, let $[n]$ denote the set $\{1,2, \cdots, n\}$.
For any function $f$, we define $\widetilde{O}(f)$ to be $f \cdot \log ^{O(1)}(f)$. In addition to $O(\cdot)$ notation, for two functions $f, g$, we use the shorthand $f \lesssim g$ (resp. $\gtrsim$ ) to indicate that $f \leq C g$ (resp. $\geq)$ for some absolute constant $C$. For any function $f$, we use $\operatorname{dom} f$ to denote the domain of function $f$.

For a vector $v$, We denote $\|v\|$ as the standard Euclidean norm of $v$ and for a symmetric PSD matrix $A$, we let $\|v\|_{A}=\left(v^{\top} A v\right)^{1 / 2}$. For a convex function $f(x)$ that is clear from context, we denote $\|v\|_{x}=\|v\|_{\nabla^{2} f(x)}$ and $\|v\|_{x}^{*}=\|v\|_{\nabla^{2} f(x)^{-1}}$.

### 3.4 Robust Central Path

In this section we show how to move move efficiently from $x(1)$ to $x(\epsilon)$ for some tiny $\epsilon$ by staying on a robust version of the central path. Because we are maintaining values that are slightly off-center, we show that our analysis still goes through despite $\ell_{\infty}$ perturbations on the order of $O(1 / \operatorname{poly} \log (n))$.

### 3.4.1 Newton Step

To follow the path $x(t)$, we consider the optimality condition of (3.3):

$$
\begin{array}{r}
s / t+\nabla \phi(x)=0, \\
A x=b, \\
A^{\top} y+s=c
\end{array}
$$

where $\nabla \phi(x)=\left(\nabla \phi_{1}\left(x_{1}\right), \nabla \phi_{2}\left(x_{2}\right), \cdots, \nabla \phi_{m}\left(x_{m}\right)\right)$. To handle the error incurred in the progress, we consider the perturbed central path

$$
\begin{array}{r}
s / t+\nabla \phi(x)=\mu, \\
A x=b, \\
A^{\top} y+s=c
\end{array}
$$

where $\mu$ represent the error between the original central path and our central path. Each iteration, we decrease $t$ by a certain factor. It may increase the error term $\mu$. Therefore, we need a step to decrease the norm of $\mu$. The Newton method to move $\mu$ to $\mu+h$ is given by

$$
\begin{aligned}
\frac{1}{t} \cdot \delta_{s}^{\text {ideal }}+\nabla^{2} \phi(x) \cdot \delta_{x}^{\text {ideal }} & =h \\
A \delta_{x}^{\text {ideal }} & =0 \\
A^{\top} \delta_{y}^{\text {ideal }}+\delta_{s}^{\text {ideal }} & =0
\end{aligned}
$$

where $\nabla^{2} \phi(x)$ is a block diagonal matrix with the $i$-th block is given by $\nabla^{2} \phi_{i}\left(x_{i}\right)$. Letting $W=\left(\nabla^{2} \phi(x)\right)^{-1}$, we can solve this:

$$
\begin{aligned}
& \delta_{y}^{\text {ideal }}=-t \cdot\left(A W A^{\top}\right)^{-1} A W h, \\
& \delta_{s}^{\text {ideal }}=t \cdot A^{\top}\left(A W A^{\top}\right)^{-1} A W h, \\
& \delta_{x}^{\text {ideal }}=W h-W A^{\top}\left(A W A^{\top}\right)^{-1} A W h .
\end{aligned}
$$

We define projection matrix $P \in \mathbb{R}^{n \times n}$ as follows

$$
P=W^{1 / 2} A^{\top}\left(A W A^{\top}\right)^{-1} A W^{1 / 2}
$$

and then we rewrite them

$$
\begin{align*}
& \delta_{x}^{\text {ideal }}=W^{1 / 2}(I-P) W^{1 / 2} \delta_{\mu}  \tag{3.4}\\
& \delta_{s}^{\text {ideal }}=t W^{-1 / 2} P W^{1 / 2} \delta_{\mu} \tag{3.5}
\end{align*}
$$

One standard way to analyze the central path is to measure the error by $\|\mu\|_{\nabla^{2} \phi(x)^{-1}}$ and uses the step induced by $h=-\mu$. One can easily prove that if $\|\mu\|_{\nabla^{2} \phi(x)^{-1}}<\frac{1}{10}$, one step of Newton step decreases the norm by a constant factor. Therefore, one can alternatively decrease $t$ and do a Newton step to follow the path.

### 3.4.2 Robust Central Path Method

In this section, we develop a central path method that is robust under certain $\ell_{\infty}$ perturbations. Due to the $\ell_{\infty}$ perturbation, we measure the error $\mu$ by a soft max instead of the $\ell_{2}$ type potential:

Definition 3.4.1. For each $i \in[m]$, let $\mu_{i}^{t}(x, s) \in \mathbb{R}^{n_{i}}$ and $\gamma_{i}^{t}(x, s) \in \mathbb{R}$ be defined as follows:

$$
\begin{align*}
& \mu_{i}^{t}(x, s)=s_{i} / t+\nabla \phi_{i}\left(x_{i}\right),  \tag{3.6}\\
& \gamma_{i}^{t}(x, s)=\left\|\mu_{i}^{t}(x, s)\right\|_{\nabla^{2} \phi_{i}\left(x_{i}\right)^{-1}} \tag{3.7}
\end{align*}
$$

and we define potential function $\Phi$ as follows:

$$
\Phi^{t}(x, s)=\sum_{i=1}^{m} \exp \left(\lambda \gamma_{i}^{t}(x, s)\right)
$$

where $\lambda=O(\log m)$.

The robust central path is the region $(x, s)$ that satisfies $\Phi^{t}(x, s) \leq O(m)$. To run our convergence argument, we will be setting $\lambda$ appropriately so that staying on the robust central path will guarantee a $\ell_{\infty}$ bound on $\gamma$. Then, we will show how to maintain $\Phi^{t}(x, s)$ to be small throughout the algorithm while decreasing $t$, always staying on the robust central path. This is broken into a two step analysis: the progress step (decreasing $t$ ) and the centering step (moving $x, s$ to decrease $\gamma$ ).

It is important to note that to follow the robust central path, we no longer pick the standard Newton direction by setting $h=-\mu$. To explain how we pick our centering step, suppose we can move $\mu \rightarrow \mu+h$ arbitrarily with the only restriction on the distance
$\|h\|_{\nabla^{2} \phi(x)^{-1}}=\alpha$. Then, the natural step would be

$$
h=\arg \min _{\|h\|_{\nabla^{2} \phi(x)^{-1}=\alpha}}\langle\nabla f(\mu(x, s)), h\rangle
$$

where $f(\mu)=\sum_{i=1}^{m} \exp \left(\lambda\|\mu\|_{\left.\nabla^{2} \phi_{i}\left(x_{i}\right)^{-1}\right)}\right.$. Note that

$$
\nabla f\left(\mu^{t}(x, s)\right)_{i}=\lambda \exp \left(\lambda \gamma_{i}^{t}(x, s)\right) / \gamma_{i}^{t}(x, s) \cdot \nabla^{2} \phi_{i}\left(x_{i}\right)^{-1} \mu_{i}^{t}(x, s)
$$

Therefore, the solution for the minimization problem is

$$
h_{i}^{\text {ideal }}=-\alpha \cdot c_{i}^{t}(x, s)^{\text {ideal }} \mu_{i}^{t}(x, s) \in \mathbb{R}^{n_{i}}
$$

where $\mu_{i}^{t}(x, s) \in \mathbb{R}^{n_{i}}$ is defined as Eq. (3.6) and $c_{i}^{t}(x, s) \in \mathbb{R}$ is defined as

$$
c_{i}^{t}(x, s)^{\text {ideal }}=\frac{\exp \left(\lambda \gamma_{i}^{t}(x, s)\right) / \gamma_{i}^{t}(x, s)}{\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(x, s)\right)\right)^{1 / 2}}
$$

Eq. (3.4) and Eq. (3.5) gives the corresponding ideal step on $x$ and $s$.
Now, we discuss the perturbed version of this algorithm. Instead of using the exact $x$ and $s$ in the formula of $h$, we use a $\bar{x}$ which is approximately close to $x$ and a $\bar{s}$ which is close to $s$. Precisely, we have

$$
\begin{equation*}
h_{i}=-\alpha \cdot c_{i}^{t}(\bar{x}, \bar{s}) \mu_{i}^{t}(\bar{x}, \bar{s}) \tag{3.8}
\end{equation*}
$$

where

$$
c_{i}^{t}(x, s)=\left\{\begin{array}{ll}
\frac{\exp \left(\lambda \gamma_{i}^{t}(x, s)\right) / /_{i}^{t}(x, s)}{\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(x, s)\right)\right)^{1 / 2}} & \text { if } \gamma_{i}^{t}(x, s) \geq 96 \sqrt{\alpha}  \tag{3.9}\\
0 & \text { otherwise }
\end{array} .\right.
$$

Note that our definition of $c_{i}^{t}$ ensures that $c_{i}^{t}(x, s) \leq \frac{1}{96 \sqrt{\alpha}}$ regardless of the value of $\gamma_{i}^{t}(x, s)$. This makes sure we do not move too much in any coordinates and indeed when $\gamma_{i}^{t}$ is small,
it is fine to set $c_{i}^{t}=0$. Furthermore, for the formula on $\delta_{x}$ and $\delta_{s}$, we use some matrix $\widetilde{V}$ that is close to $\left(\nabla^{2} \phi(x)\right)^{-1}$. Precisely, we have

$$
\begin{align*}
\delta_{x} & =\widetilde{V}^{1 / 2}(I-\widetilde{P}) \widetilde{V}^{1 / 2} h  \tag{3.10}\\
\delta_{s} & =t \cdot \widetilde{V}^{-1 / 2} \widetilde{P} \widetilde{V}^{1 / 2} h \tag{3.11}
\end{align*}
$$

where

$$
\widetilde{P}=\widetilde{V}^{1 / 2} A^{\top}\left(A \widetilde{V} A^{\top}\right)^{-1} A \widetilde{V}^{1 / 2}
$$

Here we give a quick summary of our algorithm. (The more detailed of our algorithm can be found in Algorithm 3.5 and 3.6 in Section 3.7.)

- Robustipm $(A, b, c, \phi, \delta)$
$-\lambda=2^{16} \log (m), \alpha=2^{-20} \lambda^{-2}, \kappa=2^{-10} \alpha$.
$-\delta=\min \left(\frac{1}{\lambda}, \delta\right)$.
$-\nu=\sum_{i=1}^{m} \nu_{i}$ where $\nu_{i}$ are the self-concordant parameters of $\phi_{i}$.
- Modify the convex problem and obtain an initial $x$ and $s$ according to Lemma 3.8.2.
$-t=1$.
- While $t>\frac{\delta^{2}}{4 \nu}$
* Find $\bar{x}$ and $\bar{s}$ such that $\left\|\bar{x}_{i}-x_{i}\right\|_{\bar{x}_{i}}<\alpha$ and $\left\|\bar{s}_{i}-s_{i}\right\|_{\bar{x}_{i}}^{*}<t \alpha$ for all $i$.
* Find $\widetilde{V}_{i}$ such that $(1-\alpha)\left(\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)\right)^{-1} \preceq \widetilde{V}_{i} \preceq(1+\alpha)\left(\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)\right)^{-1}$ for all $i$.
* Compute $h=-\alpha \cdot c_{i}^{t}(\bar{x}, \bar{s}) \mu_{i}^{t}(\bar{x}, \bar{s})$ where

$$
c_{i}^{t}(\bar{x}, \bar{s})= \begin{cases}\frac{\exp \left(\lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right) / /_{i}^{t}(\bar{x}, \bar{s})}{\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}} & \text { if } \gamma_{i}^{t}(\bar{x}, \bar{s}) \geq 96 \sqrt{\alpha} \\ 0 & \text { otherwise }\end{cases}
$$

and $\mu_{i}^{t}(\bar{x}, \bar{s})=\bar{s}_{i} / t+\nabla \phi_{i}\left(\bar{x}_{i}\right)$ and $\gamma_{i}^{t}(\bar{x}, \bar{s})=\left\|\mu_{i}^{t}(\bar{x}, \bar{s})\right\|_{\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)^{-1}}$

* Let $\widetilde{P}=\widetilde{V}^{1 / 2} A^{\top}\left(A \widetilde{V} A^{\top}\right)^{-1} A \widetilde{V}^{1 / 2}$.
* Compute $\delta_{x}=\widetilde{V}^{1 / 2}(I-\widetilde{P}) \widetilde{V}^{1 / 2} h$ and $\delta_{s}=t \cdot \widetilde{V}^{-1 / 2} \widetilde{P} \widetilde{V}^{1 / 2} h$.
* Move $x \leftarrow x+\delta_{x}, s \leftarrow s+\delta_{s}$.
* $t^{\text {new }}=\left(1-\frac{\kappa}{\sqrt{\nu}}\right) t$.
- Return an approximation solution of the convex problem according to Lemma 3.8.2.

Theorem 3.4.1 (Robust Interior Point Method). Consider a convex problem $\min _{A x=b, x \in \prod_{i=1}^{m} K_{i}} c^{\top} x$ where $K_{i}$ are compact convex sets. For each $i \in[m]$, we are given a $\nu_{i}$-self concordant barrier function $\phi_{i}$ for $K_{i}$. Let $\nu=\sum_{i=1}^{m} \nu_{i}$. Also, we are given $x^{(0)}=\arg \min _{x} \sum_{i=1}^{m} \phi_{i}\left(x_{i}\right)$. Assume that

1. Diameter of the set: For any $x \in \prod_{i=1}^{m} K_{i}$, we have that $\|x\|_{2} \leq R$.
2. Lipschitz constant of the program: $\|c\|_{2} \leq L$.

Then, the algorithm RobustIPM finds a vector $x$ such that

$$
\begin{aligned}
c^{\top} x & \leq \min _{A x=b, x \in \prod_{i=1}^{m} K_{i}} c^{\top} x+L R \cdot \delta, \\
\|A x-b\|_{1} & \leq 3 \delta \cdot\left(R \sum_{i, j}\left|A_{i, j}\right|+\|b\|_{1}\right) \\
x & \in \prod_{i=1}^{m} K_{i}
\end{aligned}
$$

in $O\left(\sqrt{\nu} \log ^{2} m \log \left(\frac{\nu}{\delta}\right)\right)$ iterations.

Proof. Lemma 3.8.2 shows that the initial $x$ and $s$ satisfies

$$
\|s+\nabla \phi(x)\|_{x}^{*} \leq \delta \leq \frac{1}{\lambda}
$$

where the last inequality is due to our step $\delta \leftarrow \min \left(\frac{1}{\lambda}, \delta\right)$. This implies that $\gamma_{i}^{1}(x, s)=$ $\left\|s_{i}+\nabla \phi_{i}\left(x_{i}\right)\right\|_{x_{i}}^{*} \leq \frac{1}{\lambda}$ and hence $\Phi^{1}(x, s) \leq e \cdot m \leq 80 \frac{m}{\alpha}$ for the initial $x$ and $s$. Apply Lemma 3.5.8 repetitively, we have that $\Phi^{t}(x, s) \leq 80 \frac{m}{\alpha}$ during the whole algorithm. In particular, we have this at the end of the algorithm. This implies that

$$
\left\|s_{i}+\nabla \phi_{i}\left(x_{i}\right)\right\|_{x_{i}}^{*} \leq \frac{\log \left(80 \frac{m}{\alpha}\right)}{\lambda} \leq 1
$$

at the end. Therefore, we can apply Lemma 3.8.3 to show that

$$
\langle c, x\rangle \leq\left\langle c, x^{*}\right\rangle+4 t \nu \leq\left\langle c, x^{*}\right\rangle+\delta^{2}
$$

where we used the stop condition for $t$ at the end. Note that this guarantee holds for the modified convex program. Since the error is $\delta^{2}$, Lemma 3.8.2 shows how to get an approximate solution for the original convex program with error $L R \cdot \delta$.

The number of steps follows from the fact we decrease $t$ by $1-\frac{1}{\sqrt{\nu} \log ^{2} m}$ factor every iteration.

### 3.5 Robust Central Path

The goal of this section is to analyze robust central path. We provide an outline in Section 3.5.1. In Section 3.5.2, we bound the changes in $\mu$ and $\gamma$. In Section 3.5.3, we analyze the changes from $(x, x, s)$ to $\left(x^{\text {new }}, x, s^{\text {new }}\right)$. In Section 3.5.4, we analyze the changes from $\left(x^{\text {new }}, x, s^{\text {new }}\right)$ to $\left(x^{\text {new }}, x^{\text {new }}, s^{\text {new }}\right)$. We bound the changes in $t$ in Section 3.5.5. Finally, we analyze entire changes of potential function in Section 3.5.6.

| Statement | Section | Parameters |
| :--- | :--- | :--- |
| Lemma 3.5.2 | Section 3.5.2 | $\mu_{i}^{t}(x, s) \rightarrow \mu_{i}^{t}\left(x^{\text {new }}, s^{\text {new }}\right)$ |
| Lemma 3.5.4 | Section 3.5.2 | $\gamma_{i}^{t}(x, x, s) \rightarrow \gamma_{i}^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right)$ |
| Lemma 3.5.5 | Section 3.5.3 | $\Phi(x, x, s) \rightarrow \Phi\left(x^{\text {new }}, x, s^{\text {new }}\right)$ |
| Lemma 3.5.6 | Section 3.5.4 | $\Phi\left(x^{\text {new }}, x, s^{\text {new }}\right) \rightarrow \Phi\left(x^{\text {new }}, x^{\text {new }}, s^{\text {new }}\right)$ |
| Lemma 3.5.7 | Section 3.5.5 | $\Phi^{t} \rightarrow \Phi^{t^{\text {new }}}$ |
| Lemma 3.5.8 | Section 3.5.6 | $\Phi^{t}(x, s) \rightarrow \Phi^{\text {new }}\left(x^{\text {new }}, s^{\text {new }}\right)$ |

Table 3.1: Bounding the changes of different variables

### 3.5.1 Outline of Analysis

Basically, the main proof is just a simple calculation on how $\Phi^{t}(x, s)$ changes during 1 iteration. It could be compared to the proof of $\ell_{\infty}$ potential reduction arguments for the convergence of long-step interior point methods, although the main difficulty arises from the perturbations from stepping using $\bar{x}, \bar{s}$ instead of $x, s$.

To organize the calculations, we note that the term $\gamma_{i}^{t}(x, s)=\left\|\mu_{i}^{t}(x, s)\right\|_{\nabla^{2} \phi_{i}\left(x_{i}\right)^{-1}}$ has two terms involving $x$, one in the $\mu$ term and one in the Hessian. Hence, we separate how different $x$ affect the potential by defining

$$
\begin{aligned}
\gamma_{i}^{t}(x, z, s) & =\left\|\mu_{i}^{t}(x, s)\right\|_{\nabla^{2} \phi_{i}\left(z_{i}\right)^{-1}} \\
\Phi^{t}(x, z, s) & =\sum_{i=1}^{m} \exp \left(\lambda \gamma_{i}^{t}(x, z, s)\right) .
\end{aligned}
$$

One difference between our proof and standard $\ell_{2}$ proofs of interior point is that we assume the barrier function is decomposable. We define $\alpha_{i}=\left\|\delta_{x, i}\right\|_{\bar{x}_{i}}$ is the "step" size of the coordinate $i$. One crucial fact we are using is that sum of squares of the step sizes is small.

Lemma 3.5.1. Let $\alpha$ denote the parameter in RobustIPM. For all $i \in[m]$, let $\alpha_{i}=\left\|\delta_{x, i}\right\|_{\bar{x}_{i}}$.

Then,

$$
\sum_{i=1}^{m} \alpha_{i}^{2} \leq 4 \alpha^{2}
$$

Proof. Note that

$$
\sum_{i=1}^{m} \alpha_{i}^{2}=\left\|\delta_{x}\right\|_{\bar{x}}^{2}=h^{\top} \widetilde{V}^{1 / 2}(I-\widetilde{P}) \widetilde{V}^{1 / 2} \nabla^{2} \phi(\bar{x}) \widetilde{V}^{1 / 2}(I-\widetilde{P}) \widetilde{V}^{1 / 2} h
$$

Since $(1-\alpha)\left(\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)\right)^{-1} \preceq \widetilde{V}_{i} \preceq(1+\alpha)\left(\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)\right)^{-1}$, we have that

$$
(1-\alpha)\left(\nabla^{2} \phi(\bar{x})\right)^{-1} \preceq \widetilde{V} \preceq(1+\alpha)\left(\nabla^{2} \phi(\bar{x})\right)^{-1} .
$$

Using $\alpha \leq \frac{1}{10000}$, we have that

$$
\sum_{i=1}^{m} \alpha_{i}^{2} \leq 2 h^{\top} \widetilde{V}^{1 / 2}(I-\widetilde{P})(I-\widetilde{P}) \widetilde{V}^{1 / 2} h \leq 2 h^{\top} \widetilde{V} h
$$

where we used that $I-\widetilde{P}$ is an orthogonal projection at the end. Finally, we note that

$$
\begin{aligned}
h^{\top} \widetilde{V} h & \leq 2 \sum_{i=1}^{m}\left\|h_{i}\right\|_{\bar{x}_{i}}^{* 2} \\
& =2 \alpha^{2} \sum_{i=1}^{m} c_{i}^{t}(\bar{x}, \bar{s})^{2}\left\|\mu_{i}^{t}(\bar{x}, \bar{s})\right\|_{\bar{x}_{i}}^{* 2} \\
& \leq 2 \alpha^{2} \sum_{i=1}^{m} \frac{\exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right) / \gamma_{i}^{t}(\bar{x}, \bar{s})^{2}}{\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)^{1 / 2}}\left\|\mu_{i}^{t}(\bar{x}, \bar{s})\right\|_{\bar{x}_{i}}^{* 2} \\
& =2 \alpha^{2} \frac{\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)}{\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)} \\
& =2 \alpha^{2}
\end{aligned}
$$

where the second step follows from definition of $h_{i}(3.8)$, the third step follows from definition $c_{i}^{t}$ (3.9), the fourth step follows from definition of $\gamma_{i}^{t}$ (3.7).

Therefore, putting it all together, we can show

$$
\sum_{i=1}^{m} \alpha_{i}^{2} \leq 4 \alpha^{2}
$$

### 3.5.2 Changes in $\mu$ and $\gamma$

We provide basic lemmas that bound changes in $\mu, \gamma$ due to the centering steps.

Lemma 3.5.2 (Changes in $\mu$ ). For all $i \in[m]$, let

$$
\mu_{i}^{t}\left(x^{\mathrm{new}}, s^{\mathrm{new}}\right)=\mu_{i}^{t}(x, s)+h_{i}+\epsilon_{i}^{(\mu)} .
$$

Then, $\left\|\epsilon_{i}^{(\mu)}\right\|_{x_{i}}^{*} \leq 10 \alpha \cdot \alpha_{i}$.

Proof. Let $x^{(u)}=u x^{\text {new }}+(1-u) x$ and $\mu_{i}^{\text {new }}=\mu_{i}^{t}\left(x^{\text {new }}, s^{\text {new }}\right)$. The definition of $\mu$ (3.6) shows that

$$
\begin{aligned}
\mu_{i}^{\mathrm{new}} & =\mu_{i}+\frac{1}{t} \delta_{s, i}+\nabla \phi_{i}\left(x_{i}^{\mathrm{new}}\right)-\nabla \phi_{i}\left(x_{i}\right) \\
& =\mu_{i}+\frac{1}{t} \delta_{s, i}+\int_{0}^{1} \nabla^{2} \phi_{i}\left(x_{i}^{(u)}\right) \delta_{x, i} \mathrm{~d} u \\
& =\mu_{i}+\frac{1}{t} \delta_{s, i}+\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right) \delta_{x, i}+\int_{0}^{1}\left(\nabla^{2} \phi_{i}\left(x_{i}^{(u)}\right)-\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)\right) \delta_{x, i} \mathrm{~d} u .
\end{aligned}
$$

By the definition of $\delta_{x}$ and $\delta_{s}(3.10)$ and (3.11), we have that $\frac{1}{t} \delta_{s, i}+\widetilde{V}_{i}^{-1} \delta_{x, i}=h_{i}$. Hence, we have

$$
\mu_{i}^{\mathrm{new}}=\mu_{i}+h_{i}+\epsilon_{i}^{(\mu)}
$$

where

$$
\begin{equation*}
\epsilon_{i}^{(\mu)}=\int_{0}^{1}\left(\nabla^{2} \phi_{i}\left(x_{i}^{(u)}\right)-\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)\right) \delta_{x, i} \mathrm{~d} u+\left(\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)-\widetilde{V}_{i}^{-1}\right) \delta_{x, i} . \tag{3.12}
\end{equation*}
$$

To bound $\epsilon_{i}^{(\mu)}$, we note that

$$
\left\|x_{i}^{(t)}-\bar{x}_{i}\right\|_{\bar{x}_{i}} \leq\left\|x_{i}^{(t)}-x_{i}\right\|_{\bar{x}_{i}}+\left\|x_{i}-\bar{x}_{i}\right\|_{\bar{x}_{i}} \leq\left\|\delta_{x, i}\right\|_{\bar{x}_{i}}+\alpha=\alpha_{i}+\alpha \leq 3 \alpha
$$

where the first step follows from triangle inequality, the third step follows from definition of $\alpha_{i}$ (Lemma 3.5.1), and the last step follows from $\alpha_{i} \leq 2 \alpha$ (Lemma 3.5.1).

Using $\alpha \leq \frac{1}{100}$, Theorem 3.2.1 shows that

$$
-7 \alpha \cdot \nabla^{2} \phi_{i}\left(\bar{x}_{i}\right) \preceq \nabla^{2} \phi_{i}\left(x_{i}^{(u)}\right)-\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right) \preceq 7 \alpha \cdot \nabla^{2} \phi_{i}\left(\bar{x}_{i}\right) .
$$

Equivalently, we have

$$
\left(\nabla^{2} \phi_{i}\left(x_{i}^{(u)}\right)-\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)\right) \cdot\left(\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)\right)^{-1} \cdot\left(\nabla^{2} \phi_{i}\left(x_{i}^{(u)}\right)-\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)\right) \preceq(7 \alpha)^{2} \cdot \nabla^{2} \phi_{i}\left(\bar{x}_{i}\right) .
$$

Using this, we have

$$
\begin{align*}
\left\|\int_{0}^{1}\left(\nabla^{2} \phi_{i}\left(x_{i}^{(u)}\right)-\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)\right) \delta_{x, i} \mathrm{~d} u\right\|_{\bar{x}_{i}}^{*} & \leq \int_{0}^{1}\left\|\left(\nabla^{2} \phi_{i}\left(x_{i}^{(u)}\right)-\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)\right) \delta_{x, i}\right\|_{\bar{x}_{i}}^{*} \mathrm{~d} u \\
& \leq 7 \alpha\left\|\delta_{x, i}\right\|_{\bar{x}_{i}}=7 \alpha \cdot \alpha_{i} \tag{3.13}
\end{align*}
$$

where the last step follows from definition of $\alpha_{i}$ (Lemma 3.5.1).
For the other term in $\epsilon_{i}^{(\mu)}$, we note that

$$
(1-2 \alpha) \cdot\left(\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)\right) \preceq \widetilde{V}_{i}^{-1} \preceq(1+2 \alpha) \cdot\left(\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)\right) .
$$

Hence, we have

$$
\begin{equation*}
\left\|\left(\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)-\widetilde{V}_{i}^{-1}\right) \delta_{x, i}\right\|_{\bar{x}_{i}}^{*} \leq 2 \alpha\left\|\delta_{x, i}\right\|_{\bar{x}_{i}}=2 \alpha \cdot \alpha_{i} . \tag{3.14}
\end{equation*}
$$

Combining (3.12), (3.13) and (3.14), we have

$$
\left\|\epsilon_{i}^{(\mu)}\right\|_{\bar{x}_{i}}^{*} \leq 9 \alpha \cdot \alpha_{i}
$$

Finally, we use the fact that $x_{i}$ and $\bar{x}_{i}$ are $\alpha$ close and hence again by self-concordance, $\left\|\epsilon_{i}^{(\mu)}\right\|_{x_{i}}^{*} \leq 10 \alpha \cdot \alpha_{i}$.

Before bounding the change of $\gamma$, we first prove a helper lemma:

Lemma 3.5.3. For all $i \in[m]$, we have

$$
\left\|\mu_{i}^{t}(x, s)-\mu_{i}^{t}(\bar{x}, \bar{s})\right\|_{x_{i}}^{*} \leq 4 \alpha
$$

Proof. Note that

$$
\left\|\mu_{i}^{t}(x, s)-\mu_{i}^{t}(\bar{x}, \bar{s})\right\|_{\bar{x}_{i}}^{*}=\frac{1}{t}\left\|s_{i}-\bar{s}_{i}\right\|_{\bar{x}_{i}}^{*}+\left\|\nabla \phi_{i}\left(x_{i}\right)-\nabla \phi_{i}\left(\bar{x}_{i}\right)\right\|_{\bar{x}_{i}}^{*} .
$$

For the first term, we have $\left\|s_{i}-\bar{s}_{i}\right\|_{\bar{x}_{i}}^{*} \leq t \alpha$.
For the second term, let $x_{i}^{(u)}=u x_{i}+(1-u) \bar{x}_{i}$. Since $x_{i}$ is close enough to $\bar{x}_{i}$, Theorem 3.2.1 shows that $\nabla^{2} \phi_{i}\left(x_{i}^{(u)}\right) \preceq 2 \cdot \nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)$. Hence, we have

$$
\left\|\nabla \phi_{i}\left(x_{i}\right)-\nabla \phi_{i}\left(\bar{x}_{i}\right)\right\|_{\bar{x}_{i}}^{*}=\left\|\int_{0}^{1} \nabla^{2} \phi_{i}\left(x_{i}^{(u)}\right) \cdot\left(x_{i}-\bar{x}_{i}\right) d u\right\|_{\bar{x}_{i}}^{*} \leq 2\left\|x_{i}-\bar{x}_{i}\right\|_{\bar{x}_{i}}=2 \alpha .
$$

Hence, we have $\left\|\mu_{i}^{t}(x, s)-\mu_{i}^{t}(\bar{x}, \bar{s})\right\|_{\bar{x}_{i}}^{*} \leq 3 \alpha$ and using again $x_{i}$ is close enough to $\bar{x}_{i}$ to get the final result.

Lemma 3.5.4 (Changes in $\gamma$ ). For all $i \in[m]$, let

$$
\gamma_{i}^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right) \leq\left(1-\alpha \cdot c_{i}^{t}(\bar{x}, \bar{s})\right) \gamma_{i}^{t}(x, x, s)+\epsilon_{i}^{(\gamma)} .
$$

then $\epsilon_{i}^{(\gamma)} \leq 10 \alpha \cdot\left(\alpha c_{i}^{t}(\bar{x}, \bar{s})+\alpha_{i}\right)$. Furthermore, we have $\left|\gamma_{i}^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right)-\gamma_{i}^{t}(x, x, s)\right| \leq 3 \alpha$.

Proof. For the first claim, Lemma 3.5.2, the definition of $\gamma(3.7), h(3.8)$ and $c$ (3.9) shows that

$$
\begin{aligned}
\gamma_{i}^{t}\left(x^{\mathrm{new}}, x, s^{\mathrm{new}}\right) & =\left\|\mu_{i}^{t}(x, s)+h_{i}+\epsilon_{i}^{(\mu)}\right\|_{x_{i}}^{*} \\
& =\left\|\left(1-\alpha \cdot c_{i}^{t}(\bar{x}, \bar{s})\right) \mu_{i}^{t}(x, s)+\epsilon_{i}\right\|_{x_{i}}^{*}
\end{aligned}
$$

where $\epsilon_{i}=\alpha \cdot c_{i}^{t}(\bar{x}, \bar{s})\left(\mu_{i}^{t}(x, s)-\mu_{i}^{t}(\bar{x}, \bar{s})\right)+\epsilon_{i}^{(\mu)}$.
From the definition of $c_{i}^{t}$, we have that $c_{i}^{t} \leq \frac{1}{96 \sqrt{\alpha}} \leq \frac{1}{\alpha}$ and hence $0 \leq 1-\alpha \cdot c_{i}^{t}(\bar{x}, \bar{s}) \leq 1$. Therefore, we have

$$
\begin{equation*}
\gamma_{i}^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right) \leq\left(1-\alpha \cdot c_{i}^{t}(\bar{x}, \bar{s})\right) \gamma_{i}^{t}(x, x, s)+\left\|\epsilon_{i}\right\|_{x_{i}}^{*} . \tag{3.15}
\end{equation*}
$$

Now, we bound $\left\|\epsilon_{i}\right\|_{x_{i}}^{*}$ :

$$
\begin{align*}
\left\|\epsilon_{i}\right\|_{x_{i}}^{*} & \leq \alpha c_{i}^{t}(\bar{x}, \bar{s}) \cdot\left\|\mu_{i}^{t}(x, s)-\mu_{i}^{t}(\bar{x}, \bar{s})\right\|_{x_{i}}^{*}+\left\|\epsilon_{i}^{(\mu)}\right\|_{x_{i}}^{*} \\
& \leq 4 \alpha^{2} c_{i}^{t}(\bar{x}, \bar{s})+10 \alpha \cdot \alpha_{i} \tag{3.16}
\end{align*}
$$

where we used Lemma 3.5.3 and Lemma 3.5.2 at the end.
For the second claim, we have

$$
\left|\gamma_{i}^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right)-\gamma_{i}^{t}(x, x, s)\right| \leq\left\|h_{i}+\epsilon_{i}^{(\mu)}\right\|_{x_{i}}^{*} \leq 2 \alpha+10 \alpha \cdot \alpha_{i}
$$

where we used (3.16) and that $\left\|h_{i}\right\|_{x_{i}}^{*} \leq 2\|h\|_{\bar{x}}^{*} \leq 2 \alpha$. From Lemma 3.5.1 and that $\alpha \leq \frac{1}{10000}$, we have $10 \alpha \cdot \alpha_{i} \leq 20 \alpha^{2} \leq \alpha$.

### 3.5.3 Movement from $(x, x, s)$ to $\left(x^{\text {new }}, x, s^{\text {new }}\right)$

In the previous section, we see that $\gamma_{i}$ will be expected to decrease by a factor of $\alpha \cdot c_{i}^{t}$ up to some small perturbations. We show that our potential $\Phi^{t}$ will therefore decrease significantly.

Lemma 3.5.5 (Movement along the first and third parameters). Assume that $\gamma_{i}^{t}(x, x, s) \leq 1$ for all $i$. We have

$$
\Phi^{t}\left(x^{\mathrm{new}}, x, s^{\mathrm{new}}\right) \leq \Phi^{t}(x, x, s)-\frac{\alpha \lambda}{5}\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}+\sqrt{m} \lambda \cdot \exp (192 \lambda \sqrt{\alpha})
$$

Note that $\gamma$ is a function that has three inputs. We use $\gamma(x, s)$ to denote $\gamma(x, x, s)$ for simplicity.

Proof. Let $\Phi^{\text {new }}=\Phi^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right), \Phi=\Phi^{t}(x, x, s)$,

$$
\gamma^{(u)}=u \gamma_{i}^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right)+(1-u) \gamma_{i}^{t}(x, x, s) .
$$

Then, we have that

$$
\Phi^{\mathrm{new}}-\Phi=\sum_{i=1}^{m}\left(e^{\lambda \gamma_{i}^{(1)}}-e^{\lambda \gamma_{i}^{(0)}}\right)=\lambda \sum_{i=1}^{m} e^{\lambda \gamma_{i}^{(\zeta)}}\left(\gamma_{i}^{(1)}-\gamma_{i}^{(0)}\right)
$$

for some $0 \leq \zeta \leq 1$. Let $v_{i}=\gamma_{i}^{(1)}-\gamma_{i}^{(0)}$. Lemma 3.5.4 shows that

$$
v_{i} \leq-\alpha \cdot c_{i}^{t}(\bar{x}, \bar{s}) \cdot \gamma_{i}^{t}(x, x, s)+\epsilon_{i}^{(\gamma)}=-\alpha \cdot c_{i}^{t}(\bar{x}, \bar{s}) \cdot \gamma_{i}^{(0)}+\epsilon_{i}^{(\gamma)}
$$

and hence

$$
\begin{equation*}
\frac{\Phi^{\mathrm{new}}-\Phi}{\lambda} \leq-\alpha \sum_{i=1}^{m} c_{i}^{t}(\bar{x}, \bar{s}) \cdot \gamma_{i}^{(0)} \exp \left(\lambda \gamma_{i}^{(\zeta)}\right)+\sum_{i=1}^{m} \epsilon_{i}^{(\gamma)} \exp \left(\lambda \gamma_{i}^{(\zeta)}\right) \tag{3.17}
\end{equation*}
$$

To bound the first term in (3.17), we first relate $\gamma_{i}^{(0)}, \gamma_{i}^{(\zeta)}$ and $\gamma_{i}^{t}(\bar{x}, \bar{s})$. Lemma 3.5.4 shows that

$$
\begin{equation*}
\left|\gamma_{i}^{(0)}-\gamma_{i}^{(\zeta)}\right| \leq\left|\gamma_{i}^{(0)}-\gamma_{i}^{(1)}\right| \leq 3 \alpha . \tag{3.18}
\end{equation*}
$$

Finally, we have

$$
\begin{align*}
\left|\gamma_{i}^{t}(\bar{x}, \bar{s})-\gamma_{i}^{(0)}\right| & =\left|\gamma_{i}^{t}(\bar{x}, \bar{x}, \bar{s})-\gamma_{i}^{t}(x, x, s)\right| \\
& \leq\left|\gamma_{i}^{t}(\bar{x}, \bar{x}, \bar{s})-\gamma_{i}^{t}(x, \bar{x}, s)\right|+\left|\gamma_{i}^{t}(x, \bar{x}, s)-\gamma_{i}^{t}(x, x, s)\right| \\
& \leq\left\|\mu_{i}^{t}(x, s)-\mu_{i}^{t}(\bar{x}, \bar{s})\right\|_{\bar{x}_{i}}^{*}+\left|\left\|\mu_{i}^{t}(x, s)\right\|_{\bar{x}_{i}}^{*}-\left\|\mu_{i}^{t}(x, s)\right\|_{x_{i}}^{*}\right| \\
& \leq 2\left\|\mu_{i}^{t}(x, s)-\mu_{i}^{t}(\bar{x}, \bar{s})\right\|_{x_{i}}^{*}+\left|\left\|\mu_{i}^{t}(x, s)\right\|_{\bar{x}_{i}}^{*}-\left\|\mu_{i}^{t}(x, s)\right\|_{x_{i}}^{*}\right| \\
& \leq 2\left\|\mu_{i}^{t}(x, s)-\mu_{i}^{t}(\bar{x}, \bar{s})\right\|_{x_{i}}^{*}+2 \alpha\left\|\mu_{i}^{t}(x, s)\right\|_{x_{i}}^{*} \\
& \leq 8 \alpha+2 \alpha=10 \alpha \tag{3.19}
\end{align*}
$$

where the first step follows from definition, the second and third step follows from triangle inequality, the fourth step follows from $\left\|\mu_{i}^{t}(x, s)-\mu_{i}^{t}(\bar{x}, \bar{s})\right\|_{\bar{x}_{i}}^{*} \leq 2\left\|\mu_{i}^{t}(x, s)-\mu_{i}^{t}(\bar{x}, \bar{s})\right\|_{x_{i}}^{*}$, the fifth step follows from self-concordance, the sixth step follows from Lemma 3.5.3 and that $\left\|\mu_{i}^{t}(x, s)\right\|_{x_{i}}^{*}=\gamma_{i}^{t}(x, x, s) \leq 1$ for all $i$

Using (3.18) and (3.19), we have

$$
\begin{align*}
& \sum_{i=1}^{m} c_{i}^{t}(\bar{x}, \bar{s}) \cdot \gamma_{i}^{(0)} \exp \left(\lambda \gamma_{i}^{(\zeta)}\right) \\
= & \sum_{i=1}^{m} c_{i}^{t}(\bar{x}, \bar{s}) \cdot \gamma_{i}^{(0)} \exp \left(\lambda \gamma_{i}^{t}(\bar{x}, \bar{s})-\lambda \gamma_{i}^{t}(\bar{x}, \bar{s})+\lambda \gamma_{i}^{(0)}-\lambda \gamma_{i}^{(0)}+\lambda \gamma_{i}^{(\zeta)}\right) \\
\geq & \sum_{i=1}^{m} c_{i}^{t}(\bar{x}, \bar{s}) \cdot \gamma_{i}^{(0)} \exp \left(\lambda \gamma_{i}^{t}(\bar{x}, \bar{s})-13 \lambda \alpha\right) \\
\geq & \frac{1}{2} \sum_{i=1}^{m} c_{i}^{t}(\bar{x}, \bar{s}) \cdot \gamma_{i}^{(0)} \exp \left(\lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right) \\
\geq & \frac{1}{2} \sum_{i=1}^{m} c_{i}^{t}(\bar{x}, \bar{s}) \cdot \gamma_{i}^{t}(\bar{x}, \bar{s}) \exp \left(\lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)-3 \alpha \sum_{i=1}^{m} c_{i}^{t}(\bar{x}, \bar{s}) \exp \left(\lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right) \tag{3.20}
\end{align*}
$$

where the third step follows from $\exp (-13 \lambda \alpha) \geq 1 / 2$, and the last step follows from (3.18).

For the first term in (3.20), we have

$$
\begin{aligned}
& \sum_{i=1}^{m} c_{i}^{t}(\bar{x}, \bar{s}) \cdot \gamma_{i}^{t}(\bar{x}, \bar{s}) \exp \left(\lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right) \\
= & \sum_{\gamma_{i}^{t}(\bar{x}, \bar{s}) \geq 96 \sqrt{\alpha}} \frac{\exp \left(2 \lambda \cdot \gamma_{i}^{t}(\bar{x}, \bar{s})\right)}{\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}} \\
= & \sum_{i=1}^{m} \frac{\exp \left(2 \lambda \cdot \gamma_{i}^{t}(\bar{x}, \bar{s})\right)}{\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}}-\sum_{\gamma_{i}^{t}(\bar{x}, \bar{s})<96 \sqrt{\alpha}} \frac{\exp \left(2 \lambda \cdot \gamma_{i}^{t}(\bar{x}, \bar{s})\right)}{\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}} \\
\geq & \left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}-\frac{m \cdot \exp (192 \lambda \cdot \sqrt{\alpha})}{\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}} .
\end{aligned}
$$

So, if $\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right) \geq m \cdot \exp (192 \lambda \cdot \sqrt{\alpha})$, we have

$$
\sum_{i=1}^{m} c_{i}^{t}(\bar{x}, \bar{s}) \cdot \gamma_{i}^{t}(\bar{x}, \bar{s}) \exp \left(\lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right) \geq\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}-\sqrt{m} \cdot \exp (192 \lambda \sqrt{\alpha})
$$

Note that if $\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right) \leq m \cdot \exp (192 \lambda \cdot \sqrt{\alpha})$, this is still true because left hand
side is lower bounded by 0 . For the second term in (3.20), we have

$$
\begin{aligned}
\sum_{i=1}^{m} c_{i}^{t}(\bar{x}, \bar{s}) \exp \left(\lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right) & =\sum_{\gamma_{i}^{t}(\bar{x}, \bar{s}) \geq 96 \sqrt{\alpha}} \frac{\exp \left(\lambda \cdot \gamma_{i}^{t}(\bar{x}, \bar{s})\right) / \gamma_{i}^{t}(\bar{x}, \bar{s})}{\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}} \exp \left(\lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right) \\
& \leq \frac{1}{96 \sqrt{\alpha}} \sum_{\gamma_{i}^{t}(\bar{x}, \bar{s}) \geq 96 \sqrt{\alpha}} \frac{\exp \left(2 \lambda \cdot \gamma_{i}^{t}(\bar{x}, \bar{s})\right)}{\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}} \\
& \leq \frac{1}{96 \sqrt{\alpha}} \sum_{i=1}^{m} \frac{\exp \left(2 \lambda \cdot \gamma_{i}^{t}(\bar{x}, \bar{s})\right)}{\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}} \\
& =\frac{1}{96 \sqrt{\alpha}}\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}
\end{aligned}
$$

where the second step follows $\frac{1}{\gamma_{i}^{t}(\overline{x, s})} \leq \frac{1}{96 \sqrt{\alpha}}$, and the third step follows from each term in the summation is non-negative.

Combining the bounds for both first and second term in (3.20), we have

$$
\begin{align*}
\sum_{i=1}^{m} c_{i}^{t}(\bar{x}, \bar{s}) \cdot \gamma_{i}^{(0)} \exp \left(\lambda \gamma_{i}^{(\zeta)}\right) \geq & \frac{1}{2}\left(\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}-\sqrt{m} \cdot \exp (192 \lambda \sqrt{\alpha})\right) \\
& -\frac{3 \alpha}{96 \sqrt{\alpha}}\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2} \\
\geq & \frac{2}{5}\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}-\sqrt{m} \cdot \exp (192 \lambda \sqrt{\alpha}) \tag{3.21}
\end{align*}
$$

where the last step follows from $\frac{1}{2}-\frac{3 \alpha}{96 \sqrt{\alpha}} \geq \frac{1}{2}-\frac{3}{96}=\frac{45}{96} \geq \frac{2}{5}$.
For the second term in (3.17), we note that $\left|\gamma_{i}^{(\zeta)}-\gamma_{i}^{t}(\bar{x}, \bar{s})\right| \leq 13 \alpha \leq \frac{1}{2 \lambda}$ by (3.18) and (3.19). Hence,

$$
\sum_{i=1}^{m} \epsilon_{i}^{(\gamma)} \exp \left(\lambda \gamma_{i}^{(\zeta)}\right) \leq 2 \sum_{i=1}^{m} \epsilon_{i}^{(\gamma)} \exp \left(\lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)
$$

Now, we use $\epsilon_{i}^{(\gamma)} \leq 10 \alpha \cdot\left(\alpha c_{i}^{t}(\bar{x}, \bar{s})+\alpha_{i}\right)($ Lemma 3.5.4) to get

$$
\begin{aligned}
\sum_{i=1}^{m} \epsilon_{i}^{(\gamma)} \exp \left(\lambda \gamma_{i}^{(\zeta)}\right) & \leq 20 \alpha \sum_{i=1}^{m}\left(\alpha c_{i}^{t}(\bar{x}, \bar{s})+\alpha_{i}\right) \cdot \exp \left(\lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right) \\
& \leq 20 \alpha\left(\sum_{i=1}^{m}\left(\alpha c_{i}^{t}(\bar{x}, \bar{s})+\alpha_{i}\right)^{2}\right)^{1 / 2}\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}
\end{aligned}
$$

where the last step follows from Cauchy-Schwarz inequality.
Note that by using Cauchy-Schwarz,

$$
\begin{aligned}
\left(\sum_{i=1}^{m}\left(\alpha c_{i}^{t}(\bar{x}, \bar{s})+\alpha_{i}\right)^{2}\right)^{1 / 2} & \leq \alpha\left(\sum_{i=1}^{m} c_{i}^{t}(\bar{x}, \bar{s})^{2}\right)^{1 / 2}+\left(\sum_{i=1}^{m} \alpha_{i}^{2}\right)^{1 / 2} \\
& \leq \alpha \cdot \frac{1}{96 \sqrt{\alpha}}+2 \alpha \leq \frac{\sqrt{\alpha}}{90}
\end{aligned}
$$

where we used the definition of $c_{i}^{t}$, Lemma 3.5.1 and $\alpha \leq \frac{1}{2^{24}}$. Together, we conclude

$$
\begin{equation*}
\sum_{i=1}^{m} \epsilon_{i}^{(\gamma)} \exp \left(\lambda \gamma_{i}^{(\zeta)}\right) \leq \frac{1}{5} \alpha\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2} \tag{3.22}
\end{equation*}
$$

Combining (3.21) and (3.22) to (3.17) gives

$$
\begin{aligned}
\frac{\Phi^{\text {new }}-\Phi}{\lambda} & \leq-\frac{2}{5} \alpha\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}+\sqrt{m} \cdot \exp (192 \lambda \sqrt{\alpha})+\frac{1}{5} \alpha\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2} \\
& =-\frac{1}{5} \alpha\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, \bar{s})\right)\right)^{1 / 2}+\sqrt{m} \cdot \exp (192 \lambda \sqrt{\alpha})
\end{aligned}
$$

where the last step follows from merging the first term with the third term.

### 3.5.4 Movement from $\left(x^{\text {new }}, x, s^{\text {new }}\right)$ to $\left(x^{\text {new }}, x^{\text {new }}, s^{\text {new }}\right)$

Next, we must analyze the potential change when we change the second term.

Lemma 3.5.6 (Movement along the second parameter). Assume that $\left\|\gamma^{t}(x, x, s)\right\|_{\infty} \leq 1$.
Then we have
$\Phi^{t}\left(x^{\text {new }}, x^{\text {new }}, s^{\text {new }}\right) \leq \Phi^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right)+12 \alpha\left(\left\|\gamma^{t}(x, x, s)\right\|_{\infty}+3 \alpha\right) \lambda\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(x, x, s)\right)\right)^{1 / 2}$.
Proof. We can upper bound $\Phi^{t}\left(x^{\text {new }}, x^{\text {new }}, s^{\text {new }}\right)$ as follows

$$
\begin{aligned}
\Phi^{t}\left(x^{\text {new }}, x^{\mathrm{new}}, s^{\mathrm{new}}\right) & =\sum_{i=1}^{m} \exp \left(\lambda \gamma_{i}^{t}\left(x^{\mathrm{new}}, x^{\text {new }}, s^{\text {new }}\right)\right) \\
& \leq \sum_{i=1}^{m} \exp \left(\lambda \gamma_{i}^{t}\left(x^{\mathrm{new}}, x, s^{\mathrm{new}}\right)\left(1+2 \alpha_{i}\right)\right)
\end{aligned}
$$

where the second step follows from $\gamma_{i}^{t}\left(x^{\text {new }}, x^{\text {new }}, s^{\text {new }}\right) \leq \gamma_{i}^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right) \cdot\left(1+2 \alpha_{i}\right)$ by selfconcordance (Theorem 3.2.1) and $\left\|x_{i}^{\text {new }}-x_{i}\right\|_{x_{i}} \leq 2\left\|x_{i}^{\text {new }}-x_{i}\right\|_{\bar{x}_{i}} \leq 2 \alpha_{i}$.

Now, by Lemma 3.5.4, we note that $\gamma_{i}^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right) \leq \gamma_{i}^{t}(x, x, s)+3 \alpha \leq 1+3 \alpha$ and that $\alpha \leq \frac{1}{100 \lambda}$. Hence, by a simple taylor expansion, we have

$$
\begin{aligned}
& \Phi^{t}\left(x^{\text {new }}, x^{\text {new }}, s^{\text {new }}\right) \\
\leq & \sum_{i=1}^{m} \exp \left(\lambda \gamma_{i}^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right)\right)+3 \sum_{i=1}^{m} \alpha_{i} \exp \left(\lambda \gamma_{i}^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right)\right) \gamma_{i}^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right)
\end{aligned}
$$

Finally, we bound the last term by

$$
\begin{aligned}
& \sum_{i=1}^{m} \exp \left(\lambda \gamma_{i}^{t}\left(x^{\mathrm{new}}, x, s^{\mathrm{new}}\right)\right) \gamma_{i}^{t}\left(x^{\mathrm{new}}, x, s^{\mathrm{new}}\right) \alpha_{i} \\
\leq & \sum_{i=1}^{m} \exp \left(\lambda \gamma_{i}^{t}(x, x, s)+3 \lambda \alpha\right)\left(\gamma_{i}^{t}(x, x, s)+3 \alpha\right) \alpha_{i} \\
\leq & 2\left(\left\|\gamma^{t}(x, x, s)\right\|_{\infty}+3 \alpha\right) \sum_{i=1}^{m} \exp \left(\lambda \gamma_{i}^{t}(x, x, s)\right) \alpha_{i} \\
\leq & 2\left(\left\|\gamma^{t}(x, x, s)\right\|_{\infty}+3 \alpha\right)\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(x, x, s)\right)\right)^{1 / 2}\left(\sum_{i=1}^{m} \alpha_{i}^{2}\right)^{1 / 2} \\
\leq & 4 \alpha\left(\left\|\gamma^{t}(x, x, s)\right\|_{\infty}+3 \alpha\right)\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(x, x, s)\right)\right)^{1 / 2},
\end{aligned}
$$

where the first step follows from $\lambda \gamma_{i}^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right) \leq \exp \left(\lambda \gamma_{i}^{t}(x, x, s)+3 \lambda \alpha\right)$, the second step follows $\exp (3 \lambda \alpha) \leq 2$, the third step follows from Cauchy-Schwarz inequality, the last step follows from $\sum_{i=1}^{m} \alpha_{i}^{2} \leq 4 \alpha^{2}$.

### 3.5.5 Movement of $t$

Lastly, we analyze the effect of setting $t \rightarrow t^{\text {new }}$.

Lemma 3.5.7 (Movement in $t$ ). For any $x, s$ such that $\gamma_{i}^{t}(x, s) \leq 1$ for all $i$, let $t^{\text {new }}=$ $\left(1-\frac{\kappa}{\sqrt{\nu}}\right) t$ where $\nu=\sum_{i=1}^{m} \nu_{i}$, we have

$$
\Phi^{\mathrm{new}}(x, s) \leq \Phi^{t}(x, s)+10 \kappa \lambda\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(x, s)\right)\right)^{1 / 2}
$$

Proof. Note that

$$
\begin{aligned}
\gamma_{i}^{t^{\text {new }}}(x, s) & =\left\|\frac{s}{t^{\text {new }}}+\nabla \phi_{i}\left(x_{i}\right)\right\|_{x_{i}}^{*} \\
& =\left\|\frac{s}{t(1-\kappa / \sqrt{\nu})}+\nabla \phi_{i}\left(x_{i}\right)\right\|_{x_{i}}^{*} \\
& \leq(1+2 \kappa / \sqrt{\nu}) \gamma_{i}^{t}(x, s)+2\left\|(\kappa / \sqrt{\nu}) \nabla \phi_{i}\left(x_{i}\right)\right\|_{x_{i}}^{*} \\
& \leq(1+2 \kappa / \sqrt{\nu}) \gamma_{i}^{t}(x, s)+3 \kappa \sqrt{\nu_{i}} / \sqrt{\nu} \\
& \leq \gamma_{i}^{t}(x, s)+5 \kappa \sqrt{\nu_{i}} / \sqrt{\nu}
\end{aligned}
$$

where the first step follows from definition, the second step follows from $t^{\text {new }}=t(1-\kappa / \sqrt{\nu})$, the second last step follows from the fact that our barriers are $\nu_{i}$-self-concordant and the last step used $\gamma_{i}^{t}(x, s) \leq 1$ and $\nu_{i} \geq 1$. Using that $5 \kappa \leq \frac{1}{10 \lambda}$ and $\gamma_{i}^{t}(x, s) \leq 1$, we have by
simple taylor expansion,

$$
\begin{aligned}
\Phi^{\mathrm{tnew}^{\mathrm{ne}}}(x, s) & \leq \sum_{i=1}^{m} \exp \left(\lambda \gamma_{i}^{t}(x, s)\right)+2 \lambda \sum_{i=1}^{m} \exp \left(\lambda \gamma_{i}^{t}(x, s)\right)\left(5 \kappa \sqrt{\nu_{i} / \nu}\right) \\
& =\sum_{i=1}^{m} \exp \left(\lambda \gamma_{i}^{t}(x, s)\right)+10 \kappa \lambda \sum_{i=1}^{m} \exp \left(\lambda \gamma_{i}^{t}(x, s)\right)\left(\sqrt{\nu_{i} / \nu}\right) \\
& \leq \sum_{i=1}^{m} \exp \left(\lambda \gamma_{i}^{t}(x, s)\right)+10 \kappa \lambda\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(x, s)\right)\right)^{1 / 2}\left(\sum_{i=1}^{m} \frac{\nu_{i}}{\nu}\right)^{1 / 2} \\
& =\sum_{i=1}^{m} \exp \left(\lambda \gamma_{i}^{t}(x, s)\right)+10 \kappa \lambda\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(x, s)\right)\right)^{1 / 2}
\end{aligned}
$$

where the third step follows from Cauchy-Schwarz, and the last step follows from $\sum_{i=1}^{m} \nu_{i}=$ $\nu$.

### 3.5.6 Potential Maintenance

Putting it all together, we can show that our potential $\Phi^{t}$ can be maintained to be small throughout our algorithm.

Lemma 3.5.8 (Potential Maintenance). If $\Phi^{t}(x, s) \leq 80 \frac{m}{\alpha}$, then

$$
\Phi^{t^{\mathrm{new}}}\left(x^{\mathrm{new}}, s^{\mathrm{new}}\right) \leq\left(1-\frac{\alpha \lambda}{40 \sqrt{m}}\right) \Phi^{t}(x, s)+\sqrt{m} \lambda \cdot \exp (192 \lambda \sqrt{\alpha})
$$

In particularly, we have $\Phi^{t^{\text {new }}}\left(x^{\text {new }}, s^{\text {new }}\right) \leq 80 \frac{m}{\alpha}$.

Proof. Let

$$
\zeta(x, s)=\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(x, s)\right)\right)^{1 / 2}
$$

By combining our previous lemmas,

$$
\begin{align*}
& \Phi^{\text {new }}\left(x^{\text {new }}, s^{\text {new }}\right) \\
\leq & \Phi^{t}\left(x^{\text {new }}, s^{\text {new }}\right)+10 \kappa \lambda \cdot \zeta\left(x^{\text {new }}, s^{\text {new }}\right) \\
\leq & \Phi^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right)+12 \alpha \lambda\left(\left\|\gamma^{t}(x, s)\right\|_{\infty}+3 \alpha\right) \cdot \zeta(x, s)+10 \kappa \lambda \cdot \zeta\left(x^{\text {new }}, s^{\text {new }}\right) \\
\leq & \Phi^{t}(x, x, s)-\frac{\alpha \lambda}{5} \zeta(\bar{x}, \bar{s})+\sqrt{m} \lambda \cdot \exp (192 \lambda \sqrt{\alpha}) \\
& +12 \alpha \lambda\left(\left\|\gamma^{t}(x, s)\right\|_{\infty}+3 \alpha\right) \cdot \zeta(x, s)+10 \kappa \lambda \cdot \zeta\left(x^{\text {new }}, s^{\text {new }}\right) \tag{3.23}
\end{align*}
$$

where the first step follows from Lemma 3.5.7, the second step follows from Lemma 3.5.6, and the last step follows from Lemma 3.5.5. We note that in all lemma above, we used that fact that $\left\|\gamma^{t}\right\|_{\infty} \leq 1$ (for different combination of $x, \bar{x}, x^{\text {new }}, s, \bar{s}, s^{\text {new }}$ ) which we will show later.

We can upper bound $\gamma_{i}^{t}\left(x^{\text {new }}, s^{\text {new }}\right)$ in the following sense,

$$
\begin{equation*}
\gamma_{i}^{t}\left(x^{\text {new }}, s^{\text {new }}\right) \leq \gamma_{i}^{t}\left(x^{\text {new }}, x, s^{\text {new }}\right)+2 \alpha \leq \gamma_{i}^{t}(x, x, s)+5 \alpha \tag{3.24}
\end{equation*}
$$

where the first step follows from self-concordance and $\gamma_{i} \leq 1$, the second step follows from Lemma 3.5.4.

Hence, since $\zeta$ changes multiplicatively when $\gamma$ changes additively, $\zeta\left(x^{\text {new }}, s^{\text {new }}\right) \leq$ $2 \zeta(x, s)$.

Lemma 3.5.3 shows that $\left\|\mu_{i}^{t}(x, s)-\mu_{i}^{t}(\bar{x}, \bar{s})\right\|_{x_{i}}^{*} \leq 4 \alpha$ and hence

$$
\begin{align*}
\zeta(\bar{x}, \bar{s}) & \geq \frac{2}{3}\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(\bar{x}, x, \bar{s})\right)\right)^{1 / 2} \\
& \geq \frac{2}{3}\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}(x, x, s)-8 \alpha \lambda\right)\right)^{1 / 2} \\
& \geq \frac{1}{2} \zeta(x, s) \tag{3.25}
\end{align*}
$$

Combining (3.24) and (3.25) into (3.23) gives

$$
\begin{aligned}
& \Phi^{t^{\mathrm{new}}}\left(x^{\mathrm{new}}, s^{\mathrm{new}}\right) \\
\geq & \Phi^{t}(x, s)+\left(12 \alpha \lambda\left(\left\|\gamma^{t}(x, s)\right\|_{\infty}+3 \alpha\right)+20 \kappa \lambda-\frac{\alpha \lambda}{10}\right) \cdot \zeta(x, s)+\sqrt{m} \lambda \cdot \exp (192 \lambda \sqrt{\alpha}) \\
\geq & \Phi^{t}(x, s)+\left(12 \alpha \lambda\left\|\gamma^{t}(x, s)\right\|_{\infty}-\frac{\alpha \lambda}{20}\right) \cdot \zeta(x, s)+\sqrt{m} \lambda \cdot \exp (192 \lambda \sqrt{\alpha})
\end{aligned}
$$

where the last step follows from $\kappa \leq \frac{\alpha}{1000}$ and $\alpha \leq \frac{1}{10000}$.
Finally, we need to bound $\left\|\gamma^{t}(x, s)\right\|_{\infty}$. The bound for other $\left\|\gamma^{t}\right\|_{\infty}$, i.e. for different combination of $x, \bar{x}, x^{\text {new }}, s, \bar{s}, s^{\text {new }}$, are similar. We note that

$$
\Phi^{t}(x, s) \leq 80 \frac{m}{\alpha}
$$

implies that $\left\|\gamma^{t}(x, s)\right\|_{\infty} \leq \frac{\log \left(80 \frac{m}{\alpha}\right)}{\lambda}$. Hence, by our choice of $\lambda$ and $\alpha$, we have that $\lambda \geq$ $480 \log \left(80 \frac{m}{\alpha}\right)$ and hence

$$
12 \alpha \lambda\left\|\gamma^{t}(x, s)\right\|_{\infty} \leq \frac{\alpha \lambda}{40}
$$

Finally, using $\Phi^{t}(x, s) \leq \sqrt{m} \cdot \zeta(x, s)$, we have

$$
\begin{aligned}
\Phi^{t^{\mathrm{new}}}\left(x^{\mathrm{new}}, s^{\mathrm{new}}\right) & \geq \Phi^{t}(x, s)-\frac{\alpha \lambda}{40} \zeta(x, s)+\sqrt{m} \lambda \cdot \exp (192 \lambda \sqrt{\alpha}) \\
& \geq\left(1-\frac{\alpha \lambda}{40 \sqrt{m}}\right) \Phi^{t}(x, s)+\sqrt{m} \lambda \cdot \exp (192 \lambda \sqrt{\alpha})
\end{aligned}
$$

Since $\lambda \leq \frac{1}{400 \sqrt{\alpha}}$, we have $\Phi^{t}(x, s) \leq 80 \frac{m}{\alpha}$ implies $\Phi^{t^{\text {new }}}\left(x^{\text {new }}, s^{\text {new }}\right) \leq 80 \frac{m}{\alpha}$.

| Name | Type | Statement | Algorithm | Input | Output |
| :--- | :--- | :--- | :--- | :--- | :--- |
| InITIALIZE | public | Lemma 3.6.2 | Alg. 3.1 | $A, x, s, \bar{W}, \epsilon_{m p}, a, b$ | $\emptyset$ |
| UPDATE | public | Lemma 3.6.3 | Alg. 3.2 | $\bar{W}$ | $\emptyset$ |
| FULLUPDATE | private | Lemma 3.6.5 | Alg. 3.3 | $\bar{W}$ | $\emptyset$ |
| PARTIALUPDATE | private | Lemma 3.6.4 | Alg. 3.2 | $\bar{W}$ | $\emptyset$ |
| QUERY | public | Lemma 3.6.6 | Alg. 3.1 | $\emptyset$ | $\bar{x}, \bar{s}$ |
| MULTIPLYMOVE | public | Lemma 3.6.9 | Alg. 3.4 | $h, t$ | $\emptyset$ |
| MULTIPLY | private | Lemma 3.6.8 | Alg. 3.4 | $h, t$ | $\emptyset$ |
| MOVE | private | Lemma 3.6.7 | Alg. 3.4 | $\emptyset$ | $\emptyset$ |

Table 3.2: Summary of data structure CentralPathMaintenance

### 3.6 Central Path Maintenance

The goal of this section is to present a data-structure to perform our centering steps in $\widetilde{O}\left(n^{\omega-1 / 2}\right)$ amortized time and prove a theoretical guarantee of it. The original idea of inverse maintenance is from Michael B. Cohen [Lee17], then [CLS19] used it to get faster running time for solving Linear Programs. Because a simple matrix vector product would require $O\left(n^{2}\right)$ time, our speedup comes via a low-rank embedding that provides $\ell_{\infty}$ guarantees, which is unlike the sparse vector approach of [CLS19]. In fact, we are unsure if moving in a sparse direction $h$ can have sufficiently controlled noise to show convergence. Here, we give a stochastic version that is faster for dense direction $h$.

Theorem 3.6.1 (Central path maintenance). Given a full rank matrix $A \in \mathbb{R}^{d \times n}$ with $n \geq d$, a tolerance parameter $0<\epsilon_{m p}<1 / 4$ and a block diagonal structure $n=\sum_{i=1}^{m} n_{i}$. Given any positive number a such $a \leq \alpha$ where $\alpha$ is the dual exponent of matrix multiplication. Given any linear sketch of size b, there is a randomized data structure CentralPathMainte-

NANCE (in Algorithm 3.1, 3.2, 3.4) that approximately maintains the projection matrices

$$
\sqrt{W} A^{\top}\left(A W A^{\top}\right)^{-1} A \sqrt{W}
$$

for positive block diagonal psd matrix $W \oplus_{i} \mathbb{R}^{n_{i} \times n_{i}}$; exactly implicitly maintains central path parameters $(x, s)$ and approximately explicitly maintains path parameters through the following five operations:

1. Initialize $\left(\bar{W}^{(0)}, \cdots\right)$ : Assume $\bar{W}^{(0)} \in \otimes_{i} \mathbb{R}^{n_{i} \times n_{i}}$. Initialize all the parameters in $O\left(n^{\omega}\right)$ time.
2. $\operatorname{UpdAtE}(\bar{W})$ : Assume $\bar{W} \in \oplus_{i} \mathbb{R}^{n_{i} \times n_{i}}$. Output a block diagonal matrix $\widetilde{V} \oplus_{i} \mathbb{R}^{n_{i} \times n_{i}}$ such that

$$
\left(1-\epsilon_{m p}\right) \widetilde{v}_{i} \preceq \bar{w}_{i} \preceq\left(1+\epsilon_{m p}\right) \widetilde{v}_{i} .
$$

3. $\operatorname{Query}()$ : Output $(\bar{x}, \bar{s})$ such that $\|\bar{x}-x\|_{\tilde{V}^{-1}} \leq \epsilon_{m p}$ and $\|\bar{s}-s\|_{\tilde{V}} \leq t \epsilon_{m p}$ where $t$ is the last $t$ used in MultiplyMove, where $\epsilon_{m p}=\alpha \log ^{2}(n T) \frac{n^{1 / 4}}{\sqrt{b}}$ and the success probability is $1-1 / \operatorname{poly}(n T)$. This step takes $O(n)$ time.
4. MultiplyMove $(h, t)$ : It outputs nothing. It implicitly maintains:

$$
x=x+\widetilde{V}^{1 / 2}(I-\widetilde{P}) \widetilde{V}^{1 / 2} h, s=s+t \widetilde{V}^{-1 / 2} \widetilde{P} \widetilde{V}^{1 / 2} h
$$

where $\widetilde{P}=\widetilde{V}^{1 / 2} A^{\top}\left(A \widetilde{V} A^{\top}\right)^{-1} A \widetilde{V}^{1 / 2}$. It also explicitly maintains $\bar{x}, \bar{s}$. Assuming $t$ is decreasing, each call takes $O\left(n b+n^{a \omega+o(1)}+n^{a}\|h\|_{0}+n^{1.5}\right)$ amortized time.

Let $\bar{W}^{(0)}$ be the initial matrix and $\bar{W}^{(1)}, \cdots, \bar{W}^{(T)}$ be the (random) update sequence. Under the assumption that there is a sequence of matrix $W^{(0)}, \cdots, W^{(T)} \in \oplus_{i=1}^{m} \mathbb{R}^{n_{i} \times n_{i}}$ sat-
isfies for all $k$

$$
\begin{aligned}
\left\|w_{i}^{-1 / 2}\left(\bar{w}_{i}-w_{i}\right) w_{i}^{-1 / 2}\right\|_{F} & \leq \epsilon_{m p}, \\
\sum_{i=1}^{m}\left\|\left(w_{i}^{(k)}\right)^{-1 / 2}\left(\mathbb{E}\left[w_{i}^{(k+1)}\right]-w_{i}^{(k)}\right)\left(w_{i}^{(k)}\right)^{-1 / 2}\right\|_{F}^{2} & \leq C_{1}^{2}, \\
\sum_{i=1}^{m}\left(\mathbb{E}\left[\left\|\left(w_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k+1)}-w_{i}^{(k)}\right)\left(w_{i}^{(k)}\right)^{-1 / 2}\right\|_{F}^{2}\right]\right)^{2} & \leq C_{2}^{2} \\
\left\|\left(w_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k+1)}-w_{i}^{(k)}\right)\left(w_{i}^{(k)}\right)^{-1 / 2}\right\|_{F} & \leq \frac{1}{4}
\end{aligned}
$$

where $w_{i}^{(k)}$ is the $i$-th block of $W^{(k)}, \forall i \in[m]$.
Then, the amortized expected time per call of $\operatorname{Update}(w)$ is

$$
\left(C_{1} / \epsilon_{m p}+C_{2} / \epsilon_{m p}^{2}\right) \cdot\left(n^{\omega-1 / 2+o(1)}+n^{2-a / 2+o(1)}\right)
$$

Remark 3.6.1. For our algorithm, we have $C_{1}=O\left(1 / \log ^{2} n\right), C_{2}=O\left(1 / \log ^{4} n\right)$ and $\epsilon_{m p}=$ $O\left(1 / \log ^{2} n\right)$. Note that the input of Update $\bar{W}$ can move a lot. It is working as long as $\bar{W}$ is close to some $W$ that is slowly moving. In our application, our $W$ satisfies $C_{1}, C_{2}$ deterministically. We keep it for possible future applications.

```
Algorithm 3.1 Central Path Maintenance Data Structure - Initial, Query, Move
    datastructure CentralPathMaintenance
                                    \(\triangleright\) Theorem 3.6.1
    private : members
        \(\bar{W} \in \otimes_{i \in[m]} \mathbb{R}^{n_{i} \times n_{i}} \quad \triangleright\) Target vector, \(\bar{W}\) is \(\epsilon_{w}\)-close to \(W\)
        \(V, \tilde{V} \in \otimes_{i \in[m]} \mathbb{R}^{n_{i} \times n_{i}} \quad \triangleright\) Approximate vector
        \(A \in \mathbb{R}^{d \times n} \quad \triangleright\) Constraints matrix
        \(M \in \mathbb{R}^{n \times n} \quad \triangleright\) Approximate Projection Matrix
        \(\epsilon_{m p} \in(0,1 / 4) \quad \triangleright\) Tolerance
        \(a \in(0, \alpha] \quad \triangleright\) Batch Size for Update \(\left(n^{a}\right)\)
        \(b \in \mathbb{Z}_{+} \quad \triangleright\) Sketch size of one sketching matrix
        \(R \in \mathbb{R}^{n^{1+o(1)} \times n} \quad \triangleright\) A list of sketching matrices
                                \(\triangleright\) Sketched matrices
        \(u_{1} \in \mathbb{R}^{n}, F \in \mathbb{R}^{n \times n}, u_{2} \in \mathbb{R}^{n} \quad \triangleright\) Implicit representation of \(x, x=u_{1}+F \cdot u_{2}\)
        \(u_{3} \in \mathbb{R}^{n}, G \in \mathbb{R}^{n \times n}, u_{4} \in \mathbb{R}^{n} \quad \triangleright\) Implicit representation of \(s, s=u_{3}+G \cdot u_{4}\)
        \(\bar{x}, \bar{s} \in \mathbb{R}^{n} \quad \triangleright\) Central path parameters, maintain explicitly
        \(l \in \mathbb{Z}_{+} \quad \triangleright\) Randomness counter, \(R_{l} \in \mathbb{R}^{b \times n}\)
        \(t^{\text {pre }} \in \mathbb{R}_{+} \quad \triangleright\) Tracking the changes of \(t\)
    end members
    public : procedure \(\operatorname{Initialize}\left(A, x, s, W, \epsilon_{m p}, a, b\right) \quad \triangleright\) Lemma 3.6.2
        \(\triangleright\) parameters will never change after initialization
        \(A \leftarrow A, a \leftarrow a, b \leftarrow b, \epsilon_{m p} \leftarrow \epsilon_{m p}\)
        \(\triangleright\) parameters will still change after initialization
        \(\bar{W} \leftarrow W, V \leftarrow W, \widetilde{V} \leftarrow V\)
        Choose \(R_{l} \in \mathbb{R}^{b \times n}\) to be sketching matrix, \(\forall l \in[\sqrt{n}] \quad \triangleright\) Lemma 3.9.1
        \(R \leftarrow\left[R_{1}^{\top}, R_{2}^{\top}, \cdots\right]^{\top} \quad \triangleright\) Batch them into one matrix \(R\)
        \(M \leftarrow A^{\top}\left(A V A^{\top}\right)^{-1} A, Q \leftarrow R \sqrt{\widetilde{V}} M \quad \triangleright\) Initialize projection matrices
        \(u_{1} \leftarrow x, u_{2} \leftarrow 0, u_{3} \leftarrow s, u_{4} \leftarrow 0 \quad \triangleright\) Initialize \(x\) and \(s\)
        \(\bar{x} \leftarrow x, \bar{s} \leftarrow s\)
        \(l \leftarrow 1\)
    end procedure
    public : procedure \(\operatorname{QUERY}() \quad \triangleright\) Lemma 3.6.6
        return \((\bar{x}, \bar{s})\)
    end procedure
36:
37: end datastructure
```

```
Algorithm 3.2 Central Path Maintenance Data Structure - Update and PartialUpdate
    datastructure CentralPathMaintenance \(\triangleright\) Theorem 3.6.1
    public : procedure \(\operatorname{UPDATE}\left(\bar{W}^{\text {new }}\right) \quad \triangleright\) Lemma 3.6.3, \(\bar{W}^{\text {new }}\) is close to \(W^{\text {new }}\)
        \(\bar{y}_{i} \leftarrow v_{i}^{-1 / 2} \bar{w}_{i}^{\text {new }} v_{i}^{-1 / 2}-1, \forall i \in[m]\)
        \(r \leftarrow\) the number of indices \(i\) such that \(\left\|\bar{y}_{i}\right\|_{F} \geq \epsilon_{m p}\)
        if \(r<n^{a}\) then
            PartialUpdate \(\left(\bar{W}^{\text {new }}\right)\)
        else
            FullUpdate \(\left(\bar{W}^{\text {new }}\right) \quad \triangleright\) Algorithm 3.3
        end if
    procedure
    private : procedure Partial \(\operatorname{Pbdate}\left(\bar{W}^{\text {new }}\right)\)
                                    \(\triangleright\) Lemma 3.6.4
        \(\bar{W} \leftarrow \bar{W}^{\text {new }}\)
        \(\widetilde{v}_{i}^{\text {new }} \leftarrow \begin{cases}v_{i} & \text { if }\left(1-\epsilon_{m p}\right) v_{i} \preceq \bar{w}_{i} \preceq\left(1+\epsilon_{m p}\right) v_{i} \\ w_{i} & \text { otherwise }\end{cases}\)
        \(F^{\text {new }} \leftarrow F+\left(\left(\widetilde{V}^{\text {new }}\right)^{1 / 2}-(\widetilde{V})^{1 / 2}\right) M \quad \triangleright\) only takes \(n^{1+a}\) time, instead of \(n^{2}\)
        \(G^{\text {new }} \leftarrow G+\left(\left(\widetilde{V}^{\text {new }}\right)^{-1 / 2}-(\widetilde{V})^{-1 / 2}\right) M\)
        \(u_{1} \leftarrow u_{1}+\left(F-F^{\mathrm{new}}\right) u_{2}, u_{3} \leftarrow u_{3}+\left(G-G^{\text {new }}\right) u_{4}\)
        \(F \leftarrow F^{\text {new }}, G \leftarrow G^{\text {new }}\)
        Let \(\widehat{S}\) denote the blocks where \(\widetilde{V}\) and \(\widetilde{V}^{\text {new }}\) are different
        \(\bar{x}_{\widehat{S}} \leftarrow\left(u_{1}\right)_{\widehat{S}}+\left(F u_{2}\right)_{\widehat{S}}, \bar{s}_{\widehat{S}} \leftarrow\left(u_{3}\right)_{\widehat{S}}+\left(G u_{2}\right)_{\widehat{S}} \triangleright\) make sure \(x\) and \(\bar{x}\) are close, similarly
    for \(s\) and \(\bar{s}\)
    end procedure
    end datastructure
```


### 3.6.1 Proof of Theorem 3.6.1

We follow the proof-sketch as [CLS19]. The proof contains four parts : 1) Definition of $X$ and $Y, 2)$ We need to assume sorting, 3) We provide the definition of potential function, 4) We write the potential function.

```
Algorithm 3.3 Central Path Maintenance Data Structure - Full Update
    datastructure CentralPathMaintenance
                                    \(\triangleright\) Theorem 3.6.1
    private : procedure FULLUPDATE \(\left(\bar{W}^{\text {new }}\right) \quad \triangleright\) Lemma 3.6.5
        \(\bar{y}_{i} \leftarrow v_{i}^{-1 / 2} \bar{w}_{i}^{\text {new }} v_{i}^{-1 / 2}-1, \forall i \in[m]\)
        \(r \leftarrow\) the number of indices \(i\) such that \(\left\|\bar{y}_{i}\right\|_{F} \geq \epsilon_{m p}\)
        Let \(\bar{\pi}:[m] \rightarrow[m]\) be a sorting permutation such that \(\left\|\bar{y}_{\bar{\pi}(i)}\right\|_{F} \geq\left\|\bar{y}_{\bar{\pi}(i+1)}\right\|_{F}\)
        while \(1.5 \cdot r<m\) and \(\left\|\bar{y}_{\bar{\pi}(1.5 r)}\right\|_{F} \geq(1-1 / \log m)\left\|\bar{y}_{\bar{\pi}(r)}\right\|_{F}\)
            \(r \leftarrow \min (\lceil 1.5 \cdot r\rceil, m)\)
        end while
            \(v_{\bar{\pi}(i)}^{\text {new }} \leftarrow \begin{cases}\bar{w}_{\bar{\pi}(i)}^{\text {new }} & i \in\{1,2, \cdots, r\} \\ v_{\bar{\pi}(i)} & i \in\{r+1, \cdots, m\}\end{cases}\)
                \(\triangleright\) Compute \(M^{\text {new }}=A^{\top}\left(A V^{\text {new }} A^{\top}\right)^{-1} A\) via Matrix Woodbury
            \(\Delta \leftarrow V^{\text {new }}-V \quad \triangleright \Delta \in \mathbb{R}^{n \times n}\) and \(\|\Delta\|_{0}=r\)
            \(\Gamma \leftarrow \sqrt{V^{\text {new }}}-\sqrt{V}\)
            Let \(S \leftarrow \bar{\pi}([r])\) be the first \(r\) indices in the permutation
            Let \(M_{*, S} \in \mathbb{R}^{n \times O(r)}\) be the \(r\) column-blocks from \(S\) of \(M\)
            Let \(M_{S, S}, \Delta_{S, S} \in \mathbb{R}^{O(r) \times O(r)}\) be the \(r\) row-blocks and column-blocks from \(S\) of \(M, \Delta\)
            \(M^{\text {new }} \leftarrow M-M_{*, S} \cdot\left(\Delta_{S, S}^{-1}+M_{S, S}\right)^{-1} \cdot\left(M_{*, S}\right)^{\top} \quad \triangleright\) Update \(M\)
            \(Q^{\text {new }} \leftarrow Q+R \cdot\left(\Gamma \cdot M^{\text {new }}\right)+R \cdot \sqrt{V} \cdot\left(M^{\text {new }}-M\right) \quad \triangleright\) Update \(Q\)
            \(\bar{W} \leftarrow \bar{W}^{\text {new }}, V \leftarrow V^{\text {new }}, M \leftarrow M^{\text {new }}, Q \leftarrow Q^{\text {new }} \quad \triangleright\) Update in memory
            \(\widetilde{v}_{i} \leftarrow \begin{cases}v_{i} & \text { if }\left(1-\epsilon_{m p}\right) v_{i} \preceq \bar{w}_{i} \preceq\left(1+\epsilon_{m p}\right) v_{i} \\ w_{i} & \text { otherwise }\end{cases}\)
            \(F^{\text {new }} \leftarrow \sqrt{\widetilde{V}} M, G^{\text {new }} \leftarrow \frac{1}{\sqrt{\widetilde{V}}} M\)
            \(u_{1} \leftarrow u_{1}+\left(F-F^{\text {new }}\right) u_{2}, u_{3} \leftarrow u_{3}+\left(G-G^{\text {new }}\right) u_{4}\)
            \(F \leftarrow F^{\text {new }}, G \leftarrow G^{\text {new }}\)
            Let \(\widehat{S}\) denote the blocks where \(\widetilde{V}\) and \(\widetilde{V}^{\text {new }}\) are different
            \(\bar{x}_{\widehat{S}} \leftarrow\left(u_{1}\right)_{\widehat{S}}+\left(F u_{2}\right)_{\widehat{S}}, \bar{s}_{\widehat{S}} \leftarrow\left(u_{3}\right)_{\widehat{S}}+\left(G u_{2}\right)_{\widehat{S}} \triangleright\) make sure \(x\) and \(\bar{x}\) are close, similarly
    for \(s\) and \(\bar{s}\)
            \(t^{\text {pre }} \leftarrow t\)
    end procedure
28:
    end datastructure
```

```
Algorithm 3.4 Central Path Maintenance Data Structure - Multiply and Move
    datastructure CentralPathMaintenance
                                    \(\triangleright\) Theorem 3.6.1
    public : procedure MultiplyAndMove \((h, t)\)
                            \(\triangleright\) Lemma 3.6.9
        \(\operatorname{Multiply}(h, t)\)
        Move()
    end procedure
    private : procedure Multiply \((h, t) \quad \triangleright\) Lemma 3.6.8
        Let \(\widetilde{S}\) be the indices \(i\) such that \(\left(1-\epsilon_{m p}\right) v_{i} \preceq \bar{w}_{i} \preceq\left(1+\epsilon_{m p}\right) v_{i}\) is false.
        \(\widetilde{\Delta} \leftarrow \widetilde{V}-V\)
        \(\widetilde{\Gamma} \leftarrow \sqrt{\widetilde{V}}-\sqrt{V}\)
        \(\delta_{m} \leftarrow\left(\left(\widetilde{\Delta}_{\tilde{S}, \widetilde{S}}^{-1}+M_{\widetilde{S}, \widetilde{S}}\right)^{-1} \cdot\left(\left(M_{\widetilde{S}, *}\right)^{\top} \sqrt{\widetilde{V}} h\right)\right) \quad \triangleright|\widetilde{S}| \leq n^{a}\)
        \(\triangleright\) Compute \(\widetilde{\delta}_{x}=\widetilde{V}^{1 / 2}\left(I-R^{\top} R \widetilde{P}\right) \widetilde{V}^{1 / 2} h\)
        \(\widetilde{\delta}_{x} \leftarrow \widetilde{V} h-\left(\left(R_{l}^{\top} \cdot\left(\left(Q_{l}+R_{l} \cdot \widetilde{\Gamma} \cdot M\right) \cdot \sqrt{\widetilde{V}} \cdot h\right)\right)-\left(R_{l}^{\top} \cdot\left(\left(Q_{l, \tilde{S}}+R_{l} \cdot \widetilde{\Gamma} \cdot M_{\widetilde{S}, *}\right) \cdot \delta_{m}\right)\right)\right)\)
                                    \(\triangleright\) Compute \(\widetilde{\delta}_{s}=t \widetilde{V}^{-1 / 2} R^{\top} R \widetilde{P} \widetilde{V}^{1 / 2} h\)
        \(\widetilde{\delta}_{s} \leftarrow t \cdot \widetilde{V}^{-1} \cdot\left(\left(R_{l}^{\top} \cdot\left(\left(Q+R_{l} \cdot \widetilde{\Gamma} \cdot M\right) \cdot \sqrt{\widetilde{V}} \cdot h\right)\right)-\left(R_{l}^{\top} \cdot\left(\left(Q_{l, \tilde{S}}+R_{l} \cdot \widetilde{\Gamma} \cdot M_{\widetilde{S}, *}\right) \cdot \delta_{m}\right)\right)\right)\)
        \(l \leftarrow l+1 \quad \triangleright\) Increasing the randomness counter, and using the new randomness next time
                                    \(\triangleright\) Implicitly maintain \(x=x+\widetilde{V}^{1 / 2}(I-\widetilde{P}) \widetilde{V}^{1 / 2} h\)
        \(u_{1} \leftarrow u_{1}+\widetilde{V} h\)
        \(u_{2} \leftarrow u_{2}-\sqrt{\widetilde{V}} h+\mathbf{1}_{\widetilde{S}} \delta_{m}\)
                                    \(\triangleright\) Implicitly maintain \(s=s+t \widetilde{V}^{-1 / 2} \widetilde{P} \widetilde{V}^{1 / 2} h\)
        \(u_{3} \leftarrow u_{3}+0\)
        \(u_{4} \leftarrow u_{4}-t \sqrt{\widetilde{V}} h+t \mathbf{1}_{\widetilde{S}} \delta_{m}\)
    end procedure
    private : procedure \(\operatorname{Move}() \quad \triangleright\) Lemma 3.6.7
        if \(l>\sqrt{n}\) or \(t \geq t^{\text {pre }} / 2 \quad \triangleright\) Variance is large enough
            \(x \leftarrow u_{1}+F u_{2}, s \leftarrow u_{3}+F u_{4}\)
            \(\operatorname{Initialize}\left(A, x, s, \bar{W}, \epsilon_{m p}, a, b\right) \quad \triangleright\) Algorithm 3.1
        else
            \(\bar{x} \leftarrow \bar{x}+\widetilde{\delta}_{x}, \bar{s} \leftarrow \bar{s}+\widetilde{\delta}_{s} \quad \triangleright\) Update \(\bar{x}, \bar{s}\)
        end if
    return \((\bar{x}, \bar{s})\)
    end procedure
    end datastructure
```

Definition of matrices $X$ and $Y$. Let us consider the $k$-th round of the algorithm. For all $i \in[m]$, matrix $\bar{y}_{i}^{(k)} \in \mathbb{R}^{n_{i} \times n_{i}}$ is constructed based on procedure UpDATE (Algorithm 3.2)

$$
\bar{y}_{i}^{(k)}=\frac{\bar{w}_{i}^{(k+1)}}{v_{i}^{(k)}}-I .
$$

and $\bar{\pi}$ is a permutation such that $\left\|\bar{y}_{\bar{\pi}(i)}^{(k)}\right\|_{F} \geq\left\|\bar{y}_{\bar{\pi}(i+1)}^{(k)}\right\|_{F}$.
For the purpose of analysis : for all $i \in[m]$, we define $x_{i}^{(k)}, x_{i}^{(k)}$ and $y_{i}^{(k)} \in \mathbb{R}^{n_{i} \times n_{i}}$ as follows:

$$
x_{i}^{(k)}=\frac{w_{i}^{(k)}}{v_{i}^{(k)}}-I, \quad y_{i}^{(k)}=\frac{w_{i}^{(k+1)}}{v_{i}^{(k)}}-I, \quad x_{i}^{(k+1)}=\frac{w_{i}^{(k+1)}}{v_{i}^{(k+1)}}-I,
$$

where $\frac{w_{i}^{(k)}}{v_{i}^{(k)}}$ denotes $\left(v_{i}^{(k)}\right)^{-1 / 2} w_{i}^{(k)}\left(v_{i}^{(k)}\right)^{-1 / 2}$.
It is not hard to observe the difference between $x_{i}^{(k)}$ and $y_{i}^{(k)}$ is that $w$ is changing. We call it " $w$ move". Similarly, the difference between $y_{i}^{(k)}$ and $x_{i}^{(k+1)}$ is that $v$ is changing. We call it " $v$ move".

For each $i$, we define $\beta_{i}$ as follows

$$
\beta_{i}=\left\|\left(w_{i}^{(k)}\right)^{-1 / 2}\left(\mathbb{E}\left[w_{i}^{(k+1)}\right]-w_{i}^{(k)}\right)\left(w_{i}^{(k)}\right)^{-1 / 2}\right\|_{F},
$$

then one of assumption becomes

$$
\sum_{i=1}^{m} \beta_{i}^{2} \leq C_{1}^{2}
$$

Assume sorting for diagonal blocks. Without loss of generality, we can assume the diagonal blocks of matrix $x^{(k)} \in \oplus_{i=1}^{m} \mathbb{R}^{n_{i} \times n_{i}}$ are sorted such that $\left\|x_{i}^{(k)}\right\|_{F} \geq\left\|x_{i+1}^{(k)}\right\|_{F}$. In [CLS19], $x_{i}^{(k)}$ is a scalar. They sorted the sequence based on absolute value. In our situation, $x_{i}^{(k)}$ is a matrix. We sort the sequence based on Frobenius norm. Let $\tau$ permutation such that $\left\|x_{\tau(i)}^{(k+1)}\right\|_{F} \geq\left\|x_{\tau(i+1)}^{(k+1)}\right\|_{F}$. Let $\pi$ denote the permutation such that $\left\|y_{\pi(i)}^{(k)}\right\|_{F} \geq\left\|y_{\pi(i+1)}^{(k)}\right\|_{F}$.

Definition of Potential function. We define three functions $g, \psi$ and $\Phi_{k}$ here. The definition of $\psi$ is different from [CLS19], since we need to handle matrix. The definitions of $g$ and $\Phi_{k}$ are the same as [CLS19].

For the completeness, we still provide a definition of $g$. Let $g$ be defined as

$$
g_{i}= \begin{cases}n^{-a}, & \text { if } i<n^{a} \\ i^{\frac{\omega-2}{1-a}} n^{-\frac{a(\omega-2)}{1-a}}, & \text { otherwise }\end{cases}
$$

In [CLS19], the input of function $\psi: \mathbb{R} \rightarrow \mathbb{R}$ has to be a number. We allow matrix here. Let $\psi$ : square matrix $\rightarrow \mathbb{R}$ be defined by

$$
\psi(x)= \begin{cases}\frac{\|x\|_{F}^{2}}{2 \epsilon_{m p}}, & \|x\|_{F} \in\left[0, \epsilon_{m p}\right]  \tag{3.26}\\ \epsilon_{m p}-\frac{\left(4 \epsilon_{m p}^{2}-\|x\|_{F}^{2}\right)^{2}}{18 \epsilon_{m p}^{3}}, & \|x\|_{F} \in\left(\epsilon_{m p}, 2 \epsilon_{m p}\right] \\ \epsilon_{m p}, & \|x\|_{F} \in\left(2 \epsilon_{m p},+\infty\right)\end{cases}
$$

where $\|x\|_{F}$ denotes the Frobenius norm of square matrix $x$, and let $L_{1}=\max _{x} D_{x} \psi[h] /\|H\|_{F}$, $L_{2}=\max _{x} D_{x}^{2} \psi[h, h] /\|H\|_{F}^{2}$ where $h$ is the vectorization of matrix $H$.

For the completeness, we define the potential at the $k$-th round by

$$
\Phi_{k}=\sum_{i=1}^{m} g_{i} \cdot \psi\left(x_{\tau_{k}(i)}^{(k)}\right)
$$

where $\tau_{k}(i)$ is the permutation such that $\left\|x_{\tau_{k}(i)}^{(k)}\right\|_{F} \geq\left\|x_{\tau_{k}(i+1)}^{(k)}\right\|_{F}$. (Note that in [CLS19] $\|\cdot\|_{F}$ should be $|\cdot|$.)

Rewriting the potential, and bounding it. Following the ideas in [CLS19], we can rewrite $\Phi_{k+1}-\Phi_{k}$ into two terms: the first term is $w$ move, and the second term is $v$ move.

For the completeness, we still provide a proof.

$$
\begin{aligned}
\Phi_{k+1}-\Phi_{k} & =\sum_{i=1}^{m} g_{i} \cdot\left(\psi\left(x_{\tau(i)}^{(k+1)}\right)-\psi\left(x_{i}^{(k)}\right)\right) \\
& =\sum_{i=1}^{m} g_{i} \cdot \underbrace{\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)\right)}_{W \text { move }}-\sum_{i=1}^{m} g_{i} \cdot \underbrace{\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\tau(i)}^{(k+1)}\right)\right)}_{V \text { move }}
\end{aligned}
$$

Using Lemma 3.6.10, we can bound the first term. Using Lemma 3.6.12, we can bound the second term.

### 3.6.2 Initialization time, update time, query time, move time, multiply time

Remark 3.6.2. In terms of implementing this data-structure, we only need three operations Initialize, Update, and Query. However, in order to make the proof more understoodable, we split Update into many operations: FullUpdate, PartialUpdate, Multiply and Move. We give a list of operations in Table 3.2.

Lemma 3.6.2 (Initialization). The initialization time of data-structure CENTRALPATHMaintenance (Algorithm 3.1) is $O\left(n^{\omega+o(1)}\right)$.

Proof. The running time is mainly dominated by two parts, the first part is computing $A^{\top}\left(A V A^{\top}\right)^{-1} A$, this takes $O\left(n^{2} d^{\omega-2}\right)$ time.

The second part is computing $R \sqrt{\widetilde{V}} M$. This takes $O\left(n^{\omega+o(1)}\right)$ time.

Lemma 3.6.3 (Update time). The update time of data-structure CentralPathMainteNANCE (Algorithm 3.2) is $O\left(r g_{r} n^{2+o(1)}\right)$ where $r$ is the number of indices we updated in $V$.

Proof. It is trivially follows from combining Lemma 3.6.4 and Lemma 3.6.5.

Lemma 3.6.4 (Partial Update time). The partial update time of data-structure CentralPathMaintenance (Algorithm 3.2) is $O\left(n^{1+a}\right)$.

Proof. We first analyze the running time of $F$ update, the update equation of $F$ in algorithm
is

$$
\begin{aligned}
F^{\text {new }} & \leftarrow F+\left(\left(\tilde{V}^{\text {new }}\right)^{1 / 2}-(\tilde{V})^{1 / 2}\right) M \\
F & \leftarrow F^{\text {new }}
\end{aligned}
$$

which can be implemented as

$$
F \leftarrow F+\left(\left(\widetilde{V}^{\text {new }}\right)^{1 / 2}-(\tilde{V})^{1 / 2}\right) M
$$

where we only need to change $n^{a}$ row-blocks of $F$. It takes $O\left(n^{1+a}\right)$ time.
Similarly, for the update time of $G$.
Next we analyze the update time of $u_{1}$, the update equation of $u_{1}$ is

$$
u_{1} \leftarrow u_{1}+\left(F-F^{\text {new }}\right) u_{2}
$$

Note that the difference between $F$ and $F^{\text {new }}$ is only $n^{a}$ row-blocks, thus it takes $n^{1+a}$ time to update.

Finally we analyze the update time of $\bar{x}$. Let $\widehat{S}$ denote the blocks where $\widetilde{V}$ and $\widetilde{V}^{\text {new }}$ are different.

$$
\bar{x}_{\widehat{S}} \leftarrow\left(u_{1}\right)_{\widehat{S}}+\left(F u_{2}\right)_{\widehat{S}}
$$

This also can be done in $n^{1+a}$ time, since $\widehat{S}$ indicates only $n^{a}$ blocks.
Therefore, the overall running time is $O\left(n^{1+a}\right)$.

Lemma 3.6.5 (Full Update time). The full update time of data-structure CENTRALPATHMAINTENANCE (Algorithm 3.3) is $O\left(r g_{r} n^{2+o(1)}\right)$ where $r$ is the number of indices we updated in $V$.

Proof. The update equation we use for $Q$ is

$$
Q^{\text {new }} \leftarrow Q+R \cdot\left(\Gamma \cdot M^{\text {new }}\right)+R \cdot \sqrt{V} \cdot\left(M^{\text {new }}-M\right)
$$

It can be re-written as

$$
Q^{\text {new }} \leftarrow Q+R \cdot\left(\Gamma \cdot M^{\text {new }}\right)+R \cdot \sqrt{V} \cdot\left(-M_{*, S} \cdot\left(\Delta_{S, S}^{-1}+M_{S, S}\right)^{-1} \cdot\left(M_{*, S}\right)^{\top}\right)
$$

The running time of computing second term is multiplying a $n \times r$ matrix with another $r \times n$ matrix. The running time of computing third term is also dominated by multiplying a $n \times r$ matrix with another $r \times n$ matrix.

Thus running time of processing $Q$ update is the same as the processing $M$ update.
For the running time of other parts, it is dominated by the time of updating $M$ and $Q$.

Therefore, the rest of the proof is almost the same as Lemma 5.4 in [CLS19], we omitted here.

Lemma 3.6.6 (Query time). The query time of data-structure CentralPathMaintenance (Algorithm 3.1) is $O(n)$ time.

Proof. This takes only $O(n)$ time, since we stored $\bar{x}$ and $\bar{s}$.

Lemma 3.6.7 (Move time). The move time of data-structure CentralPathMainteNANCE (Algorithm 3.4) is $O\left(n^{\omega+o(1)}\right)$ time in the worst case, and is $O\left(n^{\omega-1 / 2+o(1)}\right)$ amortized cost per iteration.

Proof. In one case, it takes only $O(n)$ time. For the other case, the running time is dominated by Initialize, which takes $n^{\omega+o(1)}$ by Lemma 3.6.5.

Lemma 3.6.8 (Multiply time). The multiply time of data-structure CentralPathMainTENANCE (Algorithm 3.4) is $O\left(n b+n^{1+a+o(1)}\right)$ for dense vector $\|h\|_{0}=n$, and is $O(n b+$ $\left.n^{a \omega+o(1)}+n^{a}\|h\|_{0}\right)$ for sparse vector $h$.

Proof. We first analyze the running time of computing vector $\delta_{m}$, the equation is

$$
\delta_{m} \leftarrow\left(\left(\left(\widetilde{\Delta}_{\widetilde{S}, \widetilde{S}}\right)^{-1}+M_{\widetilde{S}, \widetilde{S}}\right)^{-1} \cdot\left(M_{\widetilde{S}, *}\right)^{\top} \sqrt{\widetilde{V}} h\right)
$$

where $\widetilde{\Delta}=\widetilde{V}-V$. Let $\widetilde{r}=\sum_{i \in \widetilde{S}} n_{i}=O(r)$ where $r$ is the number of blocks are different in $\widetilde{V}$ and $V$.

It contains several parts:

1. Computing $\widetilde{M}_{\widetilde{S}}^{\top} \cdot(\sqrt{\widetilde{V}} h) \in \mathbb{R}^{\widetilde{r}}$ takes $O(\widetilde{r})\|h\|_{0}$.
2. Computing $\left(\widetilde{\Delta}_{\widetilde{S}, \widetilde{S}}^{-1}+M_{\widetilde{S}, \widetilde{S}}\right)^{-1} \in \mathbb{R}^{O(\widetilde{r}) \times O(\widetilde{r})}$ that is the inverse of a $O(\widetilde{r}) \times O(\widetilde{r})$ matrix takes $O\left(\widetilde{r}^{\omega+o(1)}\right)$ time.
3. Computing matrix-vector multiplication between $O(\widetilde{r}) \times O(\widetilde{r})$ matrix $\left(\left(\widetilde{\Delta}_{\widetilde{S}, \widetilde{S}}+\right.\right.$ $\left.\left.M_{\widetilde{S}, \widetilde{S}}\right)^{-1}\right)$ and $O(\widetilde{r}) \times 1$ vector $\left(\left(\widetilde{M}_{\widetilde{S}, *}\right)^{\top} \sqrt{\widetilde{V}} h\right)$ takes $O\left(\widetilde{r}^{2}\right)$ time.

Thus, the running time of computing $\delta_{m}$ is

$$
O\left(\widetilde{r}\|h\|_{0}+\widetilde{r}^{\omega+o(1)}+\widetilde{r}^{2}\right)=O\left(\widetilde{r}\|h\|_{0}+\widetilde{r}^{\omega+o(1)}\right) .
$$

Next, we want to analyze the update equation of $\widetilde{\delta}_{x}$

$$
\widetilde{\delta}_{x} \leftarrow \widetilde{V} h-\left(\left(R_{l}^{\top} \cdot\left(\left(Q_{l}+R_{l} \sqrt{\widetilde{\Delta}} M\right) \cdot \sqrt{\widetilde{V}} \cdot h\right)\right)-\left(R_{l}^{\top} \cdot\left(\left(Q_{l, \widetilde{S}}+R_{l} \widetilde{\Gamma} M_{\widetilde{S}}\right) \cdot \delta_{m}\right)\right)\right)
$$

where $\widetilde{\Gamma}=\sqrt{\widetilde{V}}-\sqrt{V}$ has $O(r)$ non-zero blocks.
It is clear that the running time is dominated by the second term in the equation. We only focus on that term.

1. Computing $R_{l}^{\top} Q_{l} \sqrt{\widetilde{V}} h$ takes $O(b n)$ time, because $Q_{l}, R_{l} \in \mathbb{R}^{b \times n}$.
2. Computing $R_{l}^{\top} R_{l} \sqrt{\widetilde{\Delta}} M \sqrt{\widetilde{V}} h$ takes $O\left(b n+b \widetilde{r}+\widetilde{r}\|h\|_{0}\right)$ time. The reason is, computing $\sqrt{\widetilde{\Delta}} M \sqrt{\widetilde{V}} h$ takes $\widetilde{r}\|h\|_{0}$ time, computing $R_{l} \cdot(\sqrt{\widetilde{\Delta}} M \sqrt{\widetilde{V}} h)$ takes $b \widetilde{r}$, then finally computing $R_{l}^{\top} \cdot\left(R_{l} \sqrt{\widetilde{\Delta}} M \sqrt{\widetilde{V}} h\right)$ takes $n b$.

Last, the update equation of $u_{1}, u_{2}, u_{3}, u_{4}$ only takes the $O(n)$ time.
Finally, we note that $r \leq O\left(n^{a}\right)$ due to the guarantee of FullUpdate and ParTIALUPDATE.

Thus, overall the running time of the Multiply is

$$
\begin{aligned}
O\left(\widetilde{r}\|h\|_{0}+\widetilde{r}^{\omega+o(1)}+\widetilde{r}^{2}+b \widetilde{r}+n b\right) & =O\left(\widetilde{r}\|h\|_{0}+\widetilde{r}^{\omega+o(1)}+n b\right) \\
& =O\left(r\|h\|_{0}+r^{\omega+o(1)}+n b\right) \\
& =O\left(n^{a}\|h\|_{0}+n^{a \omega+o(1)}+n b\right)
\end{aligned}
$$

where the first step follows from $b \widetilde{r} \leq n b$ and $\widetilde{r}^{2} \leq \widetilde{r}^{\omega+o(1)}$, and the second step follows from $\widetilde{r}=O(r)$, and the last step follows from $r=O\left(n^{a}\right)$.

If $h$ is the dense vector, then the overall time is

$$
O\left(n b+n^{1+a}+n^{a \omega+o(1)}\right) .
$$

Based on Lemma 5.5 in [CLS19], we know that $a \omega \leq 1+a$. Thus, it becomes $O\left(n b+n^{1+a+o(1)}\right)$ time.

If $h$ is a sparse vector, then the overall time is

$$
O\left(n b+n^{a}\|h\|_{0}+n^{a \omega+o(1)}\right) .
$$

Lemma 3.6.9 (MultiplyMove). The running time of MultiplyMove (Algorithm 3.6.9) is the Multiply time plus Move time.

### 3.6.3 Bounding $W$ move

The goal of this section is to analyze the movement of $W$. [CLS19] provided a scalar version of $W$ move, here we provide a matrix version.

Lemma 3.6.10 ( $W$ move, matrix version of Lemma 5.7 in [CLS19]).

$$
\sum_{i=1}^{m} g_{i} \cdot \mathbb{E}\left[\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\pi(i)}^{(k)}\right)\right]=O\left(C_{1}+C_{2} / \epsilon_{m p}\right) \cdot \sqrt{\log n} \cdot\left(n^{-a / 2}+n^{\omega-5 / 2}\right)
$$

Proof. In scalar version, [CLS19] used absolute $(|\cdot|)$ to measure each $x_{i}^{(k)}$. In matrix version, we use Frobenius norm $\left(\left\|\|_{F}\right)\right.$ to measure each $x_{i}^{(k)}$. Let $I \subseteq[m]$ be the set of indices such that $\left\|x_{i}^{(k)}\right\|_{F} \leq 1$. We separate the term into two :
$\sum_{i=1}^{m} g_{i} \cdot \mathbb{E}\left[\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\pi(i)}^{(k)}\right)\right]=\sum_{i \in I} g_{\pi^{-1}(i)} \cdot \mathbb{E}\left[\psi\left(y_{i}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)\right]+\sum_{i \in I^{c}} g_{\pi^{-1}(i)} \cdot \mathbb{E}\left[\psi\left(y_{i}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)\right]$.
Case 1. Let us consider the terms from $I$.
Let $\operatorname{vec}\left(y_{i}^{(k)}\right)$ denote the vectorization of matrix $y_{i}^{(k)}$. Similarly, $\operatorname{vec}\left(x_{i}^{(k)}\right)$ denotes the vectorization of $x_{i}^{(k)}$. Mean value theorem shows that

$$
\begin{aligned}
\psi\left(y_{i}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)= & \left\langle\psi^{\prime}\left(x_{i}^{(k)}\right), y_{i}^{(k)}-x_{i}^{(k)}\right\rangle+\frac{1}{2} \operatorname{vec}\left(y_{i}^{(k)}-x_{i}^{(k)}\right)^{\top} \psi^{\prime \prime}(\zeta) \operatorname{vec}\left(y_{i}^{(k)}-x_{i}^{(k)}\right) \\
\leq & \left\langle\psi^{\prime}\left(x_{i}^{(k)}\right), y_{i}^{(k)}-x_{i}^{(k)}\right\rangle+\frac{L_{2}}{2}\left\|y_{i}^{(k)}-x_{i}^{(k)}\right\|_{F}^{2} \\
= & \left\langle\psi^{\prime}\left(x_{i}^{(k)}\right),\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k+1)}-w_{i}^{(k)}\right)\left(v_{i}^{(k)}\right)^{-1 / 2}\right\rangle \\
& +\frac{L_{2}}{2}\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k+1)}-w_{i}^{(k)}\right)\left(v_{i}^{(k)}\right)^{-1 / 2}\right\|_{F}^{2}
\end{aligned}
$$

where the second step follows from definition of $L_{2}$ (see Part 4 of Lemma 3.6.14).

Taking conditional expectation given $w^{(k)}$ on both sides

$$
\begin{align*}
\mathbb{E}\left[\psi\left(y_{i}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)\right] \leq & \left\langle\psi^{\prime}\left(x_{i}^{(k)}\right),\left(v_{i}^{(k)}\right)^{-1 / 2}\left(\mathbb{E}\left[w_{i}^{(k+1)}\right]-w_{i}^{(k)}\right)\left(v_{i}^{(k)}\right)^{-1 / 2}\right\rangle \\
& +\frac{L_{2}}{2} \mathbb{E}\left[\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k+1)}-w_{i}^{(k)}\right)\left(v_{i}^{(k)}\right)^{-1 / 2}\right\|_{F}^{2}\right] \\
\leq & L_{1}\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(\mathbb{E}\left[w_{i}^{(k+1)}\right]-w_{i}^{(k)}\right)\left(v_{i}^{(k)}\right)^{-1 / 2}\right\|_{F} \\
& +\frac{L_{2}}{2} \mathbb{E}\left[\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k+1)}-w_{i}^{(k)}\right)\left(v_{i}^{(k)}\right)^{-1 / 2}\right\|_{F}^{2}\right] \\
\leq & L_{1}\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k)}\right)^{1 / 2}\right\|^{2} \cdot\left\|\left(w_{i}^{(k)}\right)^{-1 / 2}\left(\mathbb{E}\left[w_{i}^{(k+1)}\right]-w_{i}^{(k)}\right)\left(w_{i}^{(k)}\right)^{-1 / 2}\right\|_{F} \\
& +\frac{L_{2}}{2}\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k)}\right)^{1 / 2}\right\|^{4} \cdot \mathbb{E}\left[\left\|\left(w_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k+1)}-w_{i}^{(k)}\right)\left(w_{i}^{(k)}\right)^{-1 / 2}\right\|_{F}^{2}\right] \\
= & L_{1}\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k)}\right)^{1 / 2}\right\|^{2} \cdot \beta_{i}+\frac{L_{2}}{2}\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k)}\right)^{1 / 2}\right\|^{4} \cdot \gamma_{i} \tag{3.27}
\end{align*}
$$

where the second step follows from definition of $L_{2}$ (see Part 4 of Lemma 3.6.14), the third step follows from $\|A B\|_{F} \leq\|A\|_{F} \cdot\|B\|$, and the last step follows from defining $\beta_{i}$ and $\gamma_{i}$ as follows:

$$
\begin{aligned}
& \beta_{i}=\left\|\left(w_{i}^{(k)}\right)^{-1 / 2}\left(\mathbb{E}\left[w_{i}^{(k+1)}\right]-w_{i}^{(k)}\right)\left(w_{i}^{(k)}\right)^{-1 / 2}\right\|_{F} \\
& \gamma_{i}=\mathbb{E}\left[\left\|\left(w_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k+1)}-w_{i}^{(k)}\right)\left(w_{i}^{(k)}\right)^{-1 / 2}\right\|_{F}^{2}\right] .
\end{aligned}
$$

To upper bound $\sum_{i \in I} g_{\pi^{-1}(i)} \mathbb{E}\left[\psi\left(y_{i}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)\right]$, we need to bound the following two terms,

$$
\begin{equation*}
\sum_{i \in I} g_{\pi^{-1}(i)} L_{1}\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k)}\right)^{1 / 2}\right\|^{2} \beta_{i}, \text { and } \sum_{i \in I} g_{\pi^{-1}(i)} \frac{L_{2}}{2}\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k)}\right)^{1 / 2}\right\|^{4} \gamma_{i} \tag{3.28}
\end{equation*}
$$

For the first term (which is related to $\beta$ ) in Eq. (3.28), we have

$$
\begin{align*}
\sum_{i \in I} g_{\pi^{-1}(i)} L_{1}\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k)}\right)^{1 / 2}\right\|^{2} \beta_{i} & \leq\left(\sum_{i \in I}\left(g_{\pi^{-1}(i)} L_{1}\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k)}\right)^{1 / 2}\right\|^{2}\right)^{2} \sum_{i \in I} \beta_{i}^{2}\right)^{1 / 2} \\
& \leq O\left(L_{1}\right)\left(\sum_{i=1}^{n} g_{i}^{2} \cdot C_{1}^{2}\right)^{1 / 2} \\
& =O\left(C_{1} L_{1}\|g\|_{2}\right) \tag{3.29}
\end{align*}
$$

where the first step follows from Cauchy-Schwarz inequality, the second step follows from $n_{i}=O(1)$ and $\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k)}\right)^{1 / 2}\right\|^{2}=O(1)$.

For the second term (which is related to $\gamma$ ) in Eq. (3.28), we have

$$
\begin{equation*}
\sum_{i \in I} g_{\pi^{-1}(i)} \frac{L_{2}}{2}\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k)}\right)^{1 / 2}\right\|^{4} n_{i} \gamma_{i} \leq O\left(L_{2}\right) \cdot \sum_{i=1}^{m} g_{i} \cdot \gamma_{i}=O\left(C_{2} L_{2}\|g\|_{2}\right) \tag{3.30}
\end{equation*}
$$

Putting Eq. (3.27), Eq. (3.29) and Eq. (3.30) together, and using several facts $L_{1}=$ $O(1), L_{2}=O\left(1 / \epsilon_{m p}\right)\left(\right.$ from part 4 of Lemma 3.6.14) and $\|g\|_{2} \leq \sqrt{\log n} \cdot O\left(n^{-a / 2}+n^{\omega-5 / 2}\right)$ (from Lemma 3.6.11) gives us

$$
\sum_{i \in I} g_{\pi^{-1}(i)} \cdot \mathbb{E}\left[\psi\left(y_{i}^{k)}\right)-\psi\left(x_{i}^{(k)}\right)\right] \leq O\left(C_{1}+C_{2} / \epsilon_{m p}\right) \cdot \sqrt{\log n} \cdot\left(n^{-a / 2}+n^{\omega-5 / 2}\right)
$$

(Note that, the above Equation is the same as [CLS19].)
Case 2. Let us consider the terms from $I^{c}$.
For each $i \in I^{c}$, we know $\left\|x_{i}^{(k)}\right\|_{F} \geq 1$. We observe that $\psi(x)$ is constant for $\|x\|_{F}^{2} \geq$ $\left(2 \epsilon_{m p}\right)^{2}$, where $\epsilon_{m p} \leq 1 / 4$. If $\left\|y_{i}^{(k)}\right\|_{F} \geq 1 / 2$, then $\psi\left(y_{i}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)=0$. Therefore, we only need to focus on the $i \in I^{c}$ such that $\left\|y_{i}^{(k)}\right\|_{F}<1 / 2$.

For each $i \in I^{c}$ with $\left\|y_{i}^{(k)}\right\|_{F}<1 / 2$, we have

$$
\begin{align*}
\frac{1}{2} & <\left\|y_{i}^{(k)}-x_{i}^{(k)}\right\|_{F} \\
& =\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k+1)}-w_{i}^{(k)}\right)\left(v_{i}^{(k)}\right)^{-1 / 2}\right\|_{F} \\
& =\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k)}\right)^{1 / 2} \cdot\left(w_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k+1)}-w_{i}^{(k)}\right)\left(w_{i}^{(k)}\right)^{-1 / 2} \cdot\left(w_{i}^{(k)}\right)^{1 / 2}\left(v_{i}^{(k)}\right)^{-1 / 2}\right\|_{F} \\
& \leq\left\|\left(v_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k)}\right)^{1 / 2}\right\| \cdot\left\|\left(w_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k+1)}-w_{i}^{(k)}\right)\left(w_{i}^{(k)}\right)^{-1 / 2}\right\|_{F} \cdot\left\|\left(w_{i}^{(k)}\right)^{1 / 2}\left(v_{i}^{(k)}\right)^{-1 / 2}\right\| \\
& =\left\|\left(v_{i}^{(k)}\right)^{-1 / 2} w_{i}^{(k)}\left(v_{i}^{(k)}\right)^{-1 / 2}\right\| \cdot\left\|\left(w_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k+1)}-w_{i}^{(k)}\right)\left(w_{i}^{(k)}\right)^{-1 / 2}\right\|_{F} \\
& \leq \frac{3}{2}\left\|\left(w_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k+1)}-w_{i}^{(k)}\right)\left(w_{i}^{(k)}\right)^{-1 / 2}\right\|_{F} \tag{3.31}
\end{align*}
$$

where the last step follows from $\left\|y_{i}^{(k)}\right\|_{F}=\left\|\frac{w_{i}^{(k+1)}}{v_{i}^{(k)}}-I\right\|_{F} \leq 1 / 2$.
It is obvious that Eq. (3.31) implies

$$
\left\|\left(w_{i}^{(k)}\right)^{-1 / 2}\left(w_{i}^{(k+1)}-w_{i}^{(k)}\right)\left(w_{i}^{(k)}\right)^{-1 / 2}\right\|_{F}>1 / 3>1 / 4
$$

But this is impossible, since we assume it is $\leq 1 / 4$.
Thus, we have

$$
\sum_{i \in I^{c}} g_{\pi^{-1}(i)} \cdot \mathbb{E}\left[\psi\left(y_{i}^{(k)}\right)-\psi\left(x_{i}^{(k)}\right)\right]=0
$$

We state a Lemma that was proved in previous work [CLS19].

Lemma 3.6.11 (Lemma 5.8 in [CLS19]).

$$
\left(\sum_{i=1}^{n} g_{i}^{2}\right)^{1 / 2} \leq \sqrt{\log n} \cdot O\left(n^{-a / 2}+n^{\omega-5 / 2}\right)
$$

### 3.6.4 Bounding $V$ move

In previous work, [CLS19] only handled the movement of $V$ in scalar version. Here, the goal of is to understand the movement of $V$ in matrix version. We start to give some definitions about block diagonal matrices.

Definition 3.6.1. We define block diagonal matrices $X^{(k)}, Y^{(k)}, X^{(k+1)}$ and $\bar{Y}^{(k)} \otimes_{i \in[m]} \mathbb{R}^{n_{i} \times n_{i}}$ as follows

$$
x_{i}^{(k)}=\frac{w_{i}^{(k)}}{v_{i}^{(k)}}-I, \quad y_{i}^{(k)}=\frac{w_{i}^{(k+1)}}{v_{i}^{(k)}}-I, \quad x_{i}^{(k+1)}=\frac{w_{i}^{(k+1)}}{v_{i}^{(k+1)}}-I, \quad \bar{y}_{i}^{(k)}=\frac{\bar{w}_{i}^{(k+1)}}{v_{i}^{(k)}}-I
$$

Let $\epsilon_{w}$ denote the error between $W$ and $\bar{W}$

$$
\left\|W_{i}^{-1 / 2}\left(\bar{W}_{i}-W_{i}\right) W_{i}^{-1 / 2}\right\|_{F} \leq \epsilon_{w} .
$$

Lemma 3.6.12 ( $V$ move, matrix version of Lemma 5.9 in [CLS19]). We have,

$$
\sum_{i=1}^{n} g_{i} \cdot\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\tau(i)}^{(k+1)}\right)\right) \geq \Omega\left(\epsilon_{m p} r_{k} g_{r_{k}} / \log n\right)
$$

Proof. To prove the Lemma, similarly as [CLS19], we will split the proof into two cases.
Before getting into the details of each case, let us first understand several simple facts which are useful in the later proof. Note that from the definition of the algorithm, we only change the block if $\left\|\bar{y}_{i}^{(k)}\right\|_{F}$ is larger than the error between $w_{i}$ and $\bar{w}_{i}$. Hence, all the changes only decreases the norm, namely $\psi\left(y_{i}^{(k)}\right) \geq \psi\left(x_{i}^{(k+1)}\right)$ for all $i$. So is their sorted version $\psi\left(y_{\pi(i)}^{(k)}\right) \geq \psi\left(x_{\tau(i)}\right)^{(k+1)}$ for all $i$.

Case 1. The procedure exits the while loop when $1.5 r_{k} \geq n$.
Let $u^{*}$ denote the largest $u$ such that $\left\|\bar{y}_{\bar{\pi}(u)}^{(k)}\right\|_{F} \geq \epsilon_{m p}$.
If $u^{*}=r_{k}$, we have that

$$
\left\|\bar{y}_{\bar{\pi}\left(r_{k}\right)}^{(k)}\right\|_{F} \geq \epsilon_{m p} \geq \epsilon_{m p} / 100
$$

If $u^{*} \neq r_{k}$, using the condition of the loop, we have that

$$
\begin{aligned}
\left\|\bar{y}_{\bar{\pi}\left(r_{k}\right)}^{(k)}\right\|_{F} & \geq(1-1 / \log n)^{\log _{1.5} r_{k}-\log _{1.5} u^{*}} \cdot\left\|\bar{y}_{\bar{\pi}\left(u^{*}\right)}^{(k)}\right\|_{F} \\
& \geq(1-1 / \log n)^{\log _{1.5} n} \cdot \epsilon_{m p} \\
& \geq \epsilon_{m p} / 100 .
\end{aligned}
$$

where the last step follows from $n \geq 4$.
Recall the definition of $x_{\tau(i)}^{(k+1)}$. We can lower bound the LHS in the Lemma statement in the following sense,

$$
\begin{aligned}
\sum_{i=1}^{n} g_{i} \cdot\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\tau(i)}^{(k+1)}\right)\right) & \geq \sum_{i=n / 3+1}^{2 n / 3} g_{i} \cdot\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\tau(i)}^{(k+1)}\right)\right) \\
& \geq \sum_{i=n / 3+1}^{2 n / 3} g_{i} \cdot\left(\Omega\left(\epsilon_{m p}\right)-O\left(\epsilon_{w}\right)\right) \\
& \geq \sum_{i=n / 3+1}^{2 n / 3} g_{i} \cdot \Omega\left(\epsilon_{m p}\right) \\
& =\Omega\left(r_{k} g_{r_{k}} \epsilon_{m p}\right) .
\end{aligned}
$$

where the second step follows from $\left\|y_{\pi(i)}^{(k)}\right\|_{F} \geq\left\|y_{\pi\left(r_{k}\right)}^{(k)}\right\|_{F} \geq\left(1-O\left(\epsilon_{w}\right)\right)\left\|\bar{y}_{\bar{\pi}\left(r_{k}\right)}^{(k)}\right\|_{F} \geq \epsilon_{m p} / 200$ for all $i<2 n / 3$.

Case 2. The procedure exits the while loop when $1.5 r_{k}<n$ and $\left\|\bar{y}_{\bar{\pi}\left(1.5 r_{k}\right)}^{(k)}\right\|_{F}<(1-$ $1 / \log n)\left\|\bar{y}_{\bar{\pi}\left(r_{k}\right)}^{(k)}\right\|_{F}$.

Using the same argument as Case 1, we have

$$
\left\|\bar{y}_{\bar{\pi}\left(r_{k}\right)}^{(k)}\right\|_{F} \geq \epsilon_{m p} / 100
$$

Using Part 3 of Lemma 3.6.14 and the following fact

$$
\left\|\bar{y}_{\bar{\pi}(1.5 r)}^{(k)}\right\|_{F}<\min \left(\epsilon_{m p},\left\|\bar{y}_{\bar{\pi}(r)}^{(k)}\right\|_{F} \cdot(1-1 / \log n)\right)
$$

we can show that

$$
\begin{equation*}
\psi\left(\bar{y}_{\bar{\pi}(1.5 r)}^{(k)}\right)-\psi\left(\bar{y}_{\bar{\pi}(r)}^{(k)}\right)=\Omega\left(\epsilon_{m p} / \log n\right) . \tag{3.32}
\end{equation*}
$$

Now the question is, how to relax $\psi\left(\bar{y}_{\bar{\pi}(1.5 r)}^{(k)}\right)$ to $\psi\left(y_{\pi(1.5 r)}^{(k)}\right)$ and how to relax $\psi\left(\bar{y}_{\bar{\pi}(r)}^{(k)}\right)$ to $\psi\left(y_{\pi(r)}^{(k)}\right)$

Note that $\left\|y_{i}^{(k)}\right\|_{F} \geq\left\|x_{i}^{(k+1)}\right\|_{F}$ for all $i$. Hence, we have $\psi\left(y_{\pi(i)}^{(k)}\right) \geq \psi\left(x_{\tau(i)}^{(k+1)}\right)$ for all $i$.
Recall the definition of $y, \bar{y}, \pi$ and $\bar{\pi}$,

$$
y_{i}^{(k)}=\frac{w_{i}^{(k+1)}}{v_{i}^{(k)}}-I, \quad \bar{y}_{i}^{(k)}=\frac{\bar{w}_{i}^{(k+1)}}{v_{i}^{(k)}}-I .
$$

and $\pi$ and $\bar{\pi}$ denote the permutations such that $\left\|y_{\pi(i)}^{(k)}\right\|_{F} \geq\left\|y_{\pi(i+1)}^{(k)}\right\|_{F}$ and $\left\|\bar{y}_{\bar{\pi}(i)}^{(k)}\right\|_{F} \geq$ $\left\|\bar{y}_{\bar{\pi}(i+1)}^{(k)}\right\|_{F}$.

Using Fact 3.6.13 and $\|\cdot\|_{2}=\Theta(1)\|\cdot\|_{F}$ when the matrix has constant dimension

$$
\left\|y_{\pi(i)}^{(k)}-\bar{y}_{\bar{\pi}(i)}^{(k)}\right\|_{F} \leq O\left(\epsilon_{w}\right)
$$

where $\epsilon_{w}$ is the error between $W$ and $\bar{W}$.

Next, $\forall i$, we have

$$
\begin{equation*}
\psi\left(y_{\pi(i)}^{(k)}\right)=\psi\left(\bar{y}_{\bar{\pi}(i)}^{(k)}\right) \pm O\left(\epsilon_{w} \epsilon_{m p}\right) \tag{3.33}
\end{equation*}
$$

Next, we note that all the blocks the algorithm updated must lies in the range $i=1, \cdots, \frac{3 r_{k}}{2}-$ 1. After the update, the error of $r_{k}$ of these block becomes so small that its rank will much higher than $r_{k}$. Hence, $r_{k} / 2$ of the unchanged blocks in the range $i=1, \cdots, \frac{3 r_{k}}{2}$ will move earlier in the rank. Therefore, the $r_{k} / 2$-th element in $x^{(k+1)}$ must be larger than the $\frac{3}{2} r_{k}$-th element in $y^{(k)}$. In short, we have $\psi\left(x_{\tau(i)}^{(k+1)}\right) \leq \psi\left(y_{\pi\left(1.5 r_{k}\right)}^{(k)}\right)$ for all $i \geq r_{k} / 2$.

Putting it all together, we have

$$
\begin{align*}
& \sum_{i=1}^{n} g_{i} \cdot\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\tau(i)}^{(k+1)}\right)\right) \\
& \geq \sum_{i=r_{k} / 2}^{r_{k}} g_{i} \cdot\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(x_{\tau(i)}^{(k+1)}\right)\right) \\
& \geq \sum_{i=r_{k} / 2}^{r_{k}} g_{i} \cdot\left(\psi\left(y_{\pi(i)}^{(k)}\right)-\psi\left(y_{\pi\left(1.5 r_{k}\right)}^{(k+1)}\right)\right) \\
& \geq \sum_{i=r_{k} / 2}^{r_{k}} g_{i} \cdot\left(\psi\left(\bar{y}_{\bar{\pi}(i)}^{(k)}\right)-\psi\left(\bar{y}_{\bar{\pi}\left(1.5 r_{k}\right)}^{(k+1)}\right)-O\left(\epsilon_{w} \epsilon_{m p}\right)\right) \\
& \geq \sum_{i=r_{k} / 2}^{r_{k}} g_{i} \cdot\left(\psi\left(\bar{y}_{\bar{\pi}\left(r_{k}\right)}^{(k)}\right)-\psi\left(\bar{y}_{\bar{\pi}\left(1.5 r_{k}\right)}^{(k+1)}\right)-O\left(\epsilon_{w} \epsilon_{m p}\right)\right)  \tag{3.33}\\
& \geq \sum_{i=r_{k} / 2}^{r_{k}} g_{r_{k}} \cdot\left(\Omega\left(\frac{\epsilon_{m p}}{\log n}\right)-O\left(\epsilon_{w} \epsilon_{m p}\right)\right) \\
& \geq \quad \text { by } \psi\left(x_{\tau(i)}^{(k+1)}\right) \leq \psi\left(y_{\pi\left(1.5 r_{k}\right)}^{(k)}\right), \forall i \geq r_{k} / 2  \tag{3.32}\\
& \geq \sum_{i=r_{k} / 2}^{r_{k}} g_{r_{k}} \cdot \Omega\left(\frac{\epsilon_{m p}}{\log n}\right) \geq \psi\left(\bar{y}_{\bar{\pi}\left(r_{k}\right)}^{(k)}\right), \forall i \in\left[r_{k} / 2, r_{k}\right] \\
&= \Omega\left(\epsilon_{m p} r_{k} g_{r_{k}} / \log n\right) . \quad \text { by }(3.33) \\
& \quad \text { by } \epsilon_{w}<O(1 / \log n)
\end{align*}
$$

Therefore, we complete the proof.

Fact 3.6.13. Given two length $n$ positive vectors $a, b$. Let $a$ be sorted such that $a_{i} \geq a_{i+1}$. Let $\pi$ denote the permutation such that $b_{\pi(i)} \geq b_{\pi(i+1)}$. If for all $i \in[n],\left|a_{i}-b_{i}\right| \leq \epsilon a_{i}$. Then for all $i \in[n],\left|a_{i}-b_{\pi(i)}\right| \leq \epsilon a_{i}$.

Proof. Case 1. $\pi(i)=i$. This is trivially true.
Case 2. $\pi(i)<i$. We have

$$
b_{\pi(i)} \geq b_{i} \geq(1-\epsilon) a_{i}
$$

Since $\pi(i)<i$, we know that there exists a $j>i$ such that $\pi(j)<\pi(i)$. Then we have

$$
b_{\pi(i)} \leq b_{\pi(j)} \leq(1+\epsilon) a_{j} \leq(1+\epsilon) a_{i}
$$

Combining the above two inequalities, we have $(1-\epsilon) a_{i} \leq b_{\pi(i)} \leq(1+\epsilon) a_{i}$.
Case 3. $\pi(i)>i$. We have

$$
b_{\pi(i)} \leq b_{i} \leq(1+\epsilon) a_{i}
$$

Since $\pi>i$, we know that there exists $j<i$ such that $\pi(j)>\pi(i)$. Then we have

$$
b_{\pi(i)} \geq b_{\pi(j)} \geq(1-\epsilon) a_{j} \geq(1-\epsilon) a_{i} .
$$

Combining the above two inequalities gives us $(1-\epsilon) a_{i} \leq b_{\pi(i)} \leq(1+\epsilon) a_{i}$.
Therefore, putting all the three cases together completes the proof.

### 3.6.5 Potential function $\psi$

[CLS19] used a scalar version potential function. Here, we generalize it to the matrix version.

Lemma 3.6.14 (Matrix version of Lemma 5.10 in [CLS19]). Let function $\psi:$ square matrix $\rightarrow$ $\mathbb{R}$ (defined as Eq. (3.26)) satisfies the following properties :

1. Symmetric $(\psi(x)=\psi(-x))$ and $\psi(0)=0$
2. If $\|x\|_{F} \geq\|y\|_{F}$, then $\psi(x) \geq \psi(y)$
3. $\left|f^{\prime}(x)\right|=\Omega\left(1 / \epsilon_{m p}\right), \forall x \in\left[\left(0.01 \epsilon_{m p}\right)^{2}, \epsilon_{m p}^{2}\right]$
4. $L_{1} \stackrel{\text { def }}{=} \max _{x} \frac{D_{x} \psi[H]}{\|H\|_{F}}=2$ and $L_{2} \stackrel{\text { def }}{=} \max _{x} \frac{D_{x}^{2} \psi[H, H]}{\|H\|_{F}^{2}}=10 / \epsilon_{m p}$
5. $\psi(x)$ is a constant for $\|x\|_{F} \geq 2 \epsilon_{m p}$

Proof. Let $f: \mathbb{R}_{+} \rightarrow \mathbb{R}$ be defined as

$$
f(x)= \begin{cases}\frac{x^{2}}{2 \epsilon_{m p}^{3}}, & x \in\left[0, \epsilon_{m p}^{2}\right] \\ \epsilon_{m p}-\frac{\left(4 \epsilon_{m p}^{2}-x\right)^{2}}{18 \epsilon_{m p}^{3}}, & x \in\left(\epsilon_{m p}^{2}, 4 \epsilon_{m p}^{2}\right] \\ \epsilon_{m p}, & x \in\left(4 \epsilon_{m p}^{2},+\infty\right)\end{cases}
$$

We can see that

$$
f(x)^{\prime}= \begin{cases}\frac{x}{\epsilon_{m p}^{3}}, & x \in\left[0, \epsilon_{m p}^{2}\right] ; \\
\frac{4 \epsilon_{m p}^{2}-x}{9 \epsilon_{m p}^{3}}, & x \in\left(\epsilon_{m p}^{2}, 4 \epsilon_{m p}^{2}\right] ; \quad \text { and } f(x)^{\prime \prime}=\left\{\begin{array}{ll}
\frac{1}{\epsilon_{m p}^{3}}, & x \in\left[0, \epsilon_{m p}^{2}\right] ; \\
0, & x \in\left(4 \epsilon_{m p}^{2},+\infty\right) .
\end{array} \quad x \in\left(\epsilon_{m p}^{2}, 4 \epsilon_{m p}^{2}\right] ;\right. \\
0, & |x| \in\left(4 \epsilon_{m p}^{2},+\infty\right)\end{cases}
$$

It implies that $\max _{x}\left|f(x)^{\prime}\right| \leq \frac{1}{\epsilon_{m p}}$ and $\max _{x}\left|f(x)^{\prime \prime}\right| \leq \frac{1}{\epsilon_{m p}^{3}}$. Let $\psi(x)=f\left(\|X\|_{F}^{2}\right)$.
Proof of Part 1,2 and 5. These proofs are pretty standard from definition of $\psi$.
Proof of Part 3. This is trivially following from definition of scalar function $f$.

Proof of Part 4. By chain rule, we have

$$
\begin{aligned}
D_{x} \psi[h] & =2 f^{\prime}\left(\|X\|_{F}^{2}\right) \cdot \operatorname{tr}[X H] \\
D_{x}^{2} \psi[h, h] & =2 f^{\prime \prime}\left(\|X\|_{F}^{2}\right) \cdot(\operatorname{tr}[X H])^{2}+2 f^{\prime}\left(\|x\|_{F}^{2}\right) \cdot \operatorname{tr}\left[H^{2}\right]
\end{aligned}
$$

where $x$ is the vectorization of matrix $X$ and $h$ is the vectorization of matrix $H$. We can upper bound

$$
\left|D_{x} \psi[h]\right| \leq 2\left|f^{\prime}\left(\|X\|_{F}^{2}\right)\right| \cdot|\operatorname{tr}[X H]| \leq 2\left|f^{\prime}\left(\|X\|_{F}^{2}\right)\right| \cdot\|X\|_{F} \cdot\|H\|_{F}
$$

Then, we have

$$
\left|f^{\prime}\left(\|X\|_{F}^{2}\right)\right| \cdot\|X\|_{F}= \begin{cases}\|X\|_{F}^{3} / \epsilon_{m p}^{3} \leq 1, & \|X\|_{F} \in\left[0, \epsilon_{m p}\right] \\ \left(4 \epsilon_{m p}^{2}-\|X\|_{F}^{2}\right)\|X\|_{F} / 9 \epsilon_{m p} \leq 2 / 3, & \|X\|_{F} \in\left(\epsilon_{m p}, 2 \epsilon_{m p}\right] \\ 0, & \|X\|_{F} \in\left(2 \epsilon_{m p},+\infty\right)\end{cases}
$$

It implies that $\left|D_{x} \psi[h]\right| \leq 2\|H\|_{F}, \forall x$.
By case analysis, we have

$$
\left|f^{\prime \prime}\left(\|X\|_{F}^{2}\right)\right| \cdot\|X\|_{F}^{2} \leq \begin{cases}\frac{1}{\epsilon_{m p}}\|X\|_{F}^{2} \leq 4 / \epsilon_{m p}, & \|X\|_{F}^{2} \in\left[0,4 \epsilon_{m p}^{2}\right] \\ 0, & \|X\|_{F}^{2} \in\left(4 \epsilon_{m p},+\infty\right)\end{cases}
$$

We can also upper bound

$$
\begin{aligned}
\left|D_{x}^{2} \psi[h, h]\right| & \leq 2\left|f^{\prime \prime}\left(\|X\|_{F}^{2}\right)\right| \cdot(\operatorname{tr}[X H])^{2}+2\left|f^{\prime}\left(\|X\|_{F}^{2}\right)\right| \cdot \operatorname{tr}\left[H^{2}\right] \\
& \leq 2\left|f^{\prime \prime}\left(\|X\|_{F}^{2}\right)\right| \cdot\left(\|X\|_{F}\|H\|_{F}\right)^{2}+2\left|f^{\prime}\left(\|X\|_{F}^{2}\right)\right| \cdot\|H\|_{F}^{2} \\
& \leq 2 \cdot \frac{4}{\epsilon_{m p}}\|H\|_{F}^{2}+2 \cdot \frac{1}{\epsilon_{m p}}\|H\|_{F}^{2} \\
& =\frac{10}{\epsilon_{m p}}\|H\|_{F}^{2} .
\end{aligned}
$$

### 3.6.6 $\quad x$ and $\bar{x}$ are close

Lemma 3.6.15 ( $x$ and $\bar{x}$ are close in term of $\tilde{V}^{-1}$ ). With probability $1-\delta$ over the randomness of sketching matrix $R \in \mathbb{R}^{b \times n}$, we have

$$
\left\|\bar{x}_{i}-x_{i}\right\|_{\tilde{V}_{i}^{-1}} \leq \epsilon_{x}
$$

$\epsilon_{x}=O\left(\alpha \log ^{2}(n / \delta) \cdot \frac{n^{1 / 4}}{\sqrt{b}}\right), b$ is the size of sketching matrix.
Proof. Recall the definition of $\widetilde{\delta}_{x}$ and $\delta_{x}$, we have

$$
\widetilde{\delta}_{x, i}-\delta_{x, i}=\widetilde{V}_{i}^{1 / 2}\left(I-R^{\top} R \widetilde{P}\right) \widetilde{V}^{1 / 2} h-\widetilde{V}_{i}^{1 / 2}(I-\widetilde{P}) \widetilde{V}^{1 / 2} h=\widetilde{V}_{i}^{1 / 2}\left(\widetilde{P}-R^{\top} R \widetilde{P}\right) \widetilde{V}^{1 / 2} h
$$

For iteration $t$, the definition should be

$$
\widetilde{\delta}_{x, i}^{(t)}-\delta_{x, i}^{(t)}=\left(\widetilde{V}_{i}^{(t)}\right)^{1 / 2}\left(\widetilde{P}^{(t)}-\left(R^{(t)}\right)^{\top} R^{(t)} \widetilde{P}^{(t)}\right)\left(\widetilde{V}^{(t)}\right)^{1 / 2} h
$$

For any $i$, let $k$ be the current iteration, $k_{i}$ be the last when we changed the $\widetilde{V}_{i}$. Then, we have that

$$
x_{i}^{(k)}-\bar{x}_{i}^{(k)}=\sum_{t=k_{i}}^{k} \widetilde{\delta}_{x, i}^{(t)}-\delta_{x, i}^{(t)}
$$

because we have $x_{i}^{\left(k_{i}\right)}=\bar{x}_{i}^{\left(k_{i}\right)}$ (guaranteed by our algorithm). Since $\widetilde{V}_{i}^{(t)}$ did not change during iteration $k_{i}$ to $k$ for the block $i$. (However, the whole other parts of matrix $\widetilde{V}$ could change). We consider

$$
\begin{aligned}
\left(x_{i}^{(k)}-\bar{x}_{i}^{(k)}\right)^{\top} \cdot\left(\widetilde{V}_{i}^{(k)}\right)^{-1} \cdot\left(x_{i}^{(k)}-\bar{x}_{i}^{(k)}\right) & =\left(\sum_{t=k_{i}}^{k} \widetilde{\delta}_{x, i}^{(t)}-\delta_{x, i}^{(t)}\right)^{\top} \cdot\left(\widetilde{V}_{i}^{(k)}\right)^{-1} \cdot\left(\sum_{t=k_{i}}^{k} \widetilde{\delta}_{x, i}^{(t)}-\delta_{x, i}^{(t)}\right) \\
& =\left\|\sum_{t=k_{i}}^{k}\left(\left(I-\left(R^{(t)}\right)^{\top} R^{(t)}\right) \widetilde{P}^{(t)}\left(\widetilde{V}^{(t)}\right)^{1 / 2} h^{(t)}\right)_{i}\right\|_{2}^{2}
\end{aligned}
$$

We consider block $i$ and a coordinate $j \in$ block $i$. We define random vector $X_{t} \in \mathbb{R}^{n_{i}}$ as follows:

$$
X_{t}=\left(\left(I-R^{(t) \top} R^{(t)}\right) \widetilde{P}^{(t)}\left(\widetilde{V}^{(t)}\right)^{1 / 2} h^{(t)}\right)_{i}
$$

Let $\left(X_{t}\right)_{j}$ denote the $j$-th coordinate of $X_{t}$, for each $j \in\left[n_{i}\right]$.
By Lemma 3.9.1 in Section 3.9, we have for each $t$,

$$
\mathbb{E}\left[X_{t}\right]=0, \quad \text { and } \quad \mathbb{E}\left[\left(X_{t}\right)_{j}^{2}\right]=\frac{1}{b}\left\|\left(\widetilde{P}^{(t)}\left(\widetilde{V}^{(t)}\right)^{1 / 2} h^{(t)}\right)_{i}\right\|_{2}^{2}
$$

and with probability $1-\delta$,

$$
\left|\left(X_{t}\right)_{j}\right| \leq\left\|\left(\widetilde{P}^{(t)}\left(\widetilde{V}^{(t)}\right)^{1 / 2} h^{(t)}\right)_{i}\right\|_{2} \frac{\log (n / \delta)}{\sqrt{b}}:=M
$$

Now, we apply Bernstein inequality (Lemma A.1.2),

$$
\operatorname{Pr}\left[\sum_{t}\left(X_{t}\right)_{j}>\tau\right] \leq \exp \left(-\frac{\tau^{2} / 2}{\sum_{t} \mathbb{E}\left[\left(X_{t}\right)_{j}^{2}\right]+M \tau / 3}\right)
$$

Choosing $\tau=10^{3} \frac{\sqrt{T}}{\sqrt{b}} \log ^{2}(n / \delta) \cdot\left\|\left(\widetilde{P}^{(t)}\left(\widetilde{V}^{(t)}\right)^{1 / 2} h^{(t)}\right)_{i}\right\|_{2}$

$$
\left.\begin{array}{rl} 
& \operatorname{Pr}\left[\sum_{t}\left(X_{t}\right)_{j}>10^{3} \frac{\sqrt{T}}{\sqrt{b}} \log ^{2}(n / \delta) \cdot\left\|\left(\widetilde{P}^{(t)}\left(\widetilde{V}^{(t)}\right)^{1 / 2} h^{(t)}\right)_{i}\right\|_{2}\right] \\
\leq & \exp \left(-\frac{10^{6} \frac{T}{b} \log ^{4}(n / \delta) \cdot\left\|\left(\widetilde{P}^{(t)}\left(\widetilde{V}^{(t)}\right)^{1 / 2} h^{(t)}\right)_{i}\right\|_{2}^{2} / 2}{\frac{T}{b}\left\|\left(\widetilde{P}^{(t)}\left(\widetilde{V}^{(t)}\right)^{1 / 2} h^{(t)}\right)_{i}\right\|_{2}^{2}+10^{3} \frac{\sqrt{T}}{b} \log ^{3}(n / \delta)\left\|\left(\widetilde{P}^{(t)}\left(\widetilde{V}^{(t)}\right)^{1 / 2} h^{(t)}\right)_{i}\right\|_{2}^{2} / 3}\right.
\end{array}\right)
$$

Now, taking a union, we have

$$
\begin{aligned}
\left\|\sum_{t=k_{i}}^{k}\left(\left(I-\left(R^{(t)}\right)^{\top} R^{(t)}\right) \widetilde{P}^{(t)}\left(\widetilde{V}^{(t)}\right)^{1 / 2} h^{(t)}\right)_{i}\right\|_{2} & =O\left(\frac{\sqrt{T}}{\sqrt{b}} \log ^{2}(n / \delta)\left\|\left(\widetilde{P}^{(t)}\left(\widetilde{V}^{(t)}\right)^{1 / 2} h^{(t)}\right)_{i}\right\|_{2}\right) \\
& \leq O\left(\frac{\sqrt{T}}{\sqrt{b}} \log ^{2}(n / \delta) \alpha\right)
\end{aligned}
$$

where we use that $\left\|\left(\widetilde{P}^{(t)}\left(\widetilde{V}^{(t)}\right)^{1 / 2} h^{(t)}\right)_{i}\right\|_{2} \leq\left\|\left(\left(\widetilde{V}^{(t)}\right)^{1 / 2} h^{(t)}\right)_{i}\right\|_{2}=O(\alpha), n_{i}=O(1)$.
Finally, we use the fact that the algorithm reset $\bar{x}=x, \bar{s}=s$ in less than $\sqrt{n}$ iterations.

### 3.6.7 $s$ and $\bar{s}$ are close

Lemma 3.6.16 ( $s$ and $\bar{s}$ are close). With probability $1-\delta$ over the randomness of sketching matrix $R \in \mathbb{R}^{b \times n}$, we have

$$
t^{-1}\left\|\bar{s}_{i}-s_{i}\right\|_{\tilde{v}_{i}} \leq \epsilon_{s},
$$

$\epsilon_{s}=O\left(\alpha \log ^{2}(n / \delta) \cdot \frac{n^{1 / 4}}{\sqrt{b}}\right.$, and $b$ is the size of sketching matrix.
Proof. Recall the definition of $\widetilde{\delta}_{s}$, $\delta_{s}$, we have

$$
\widetilde{\delta}_{s, i}-\delta_{s, i}=t \widetilde{V}_{i}^{-1 / 2}\left(R^{\top} R-I\right) \widetilde{P} \widetilde{V}^{1 / 2} h
$$

The rest of the proof is identical to Lemma 3.6.15 except we use also the fact we make $\bar{s}=s$ whenever our $t$ changed by a constant factor. We omitted the details here.

### 3.6.8 Data structure is maintaining $(x, s)$ implicitly over all the iterations

Lemma 3.6.17. Over all the iterations, $u_{1}+F u_{2}$ is always maintaining $x$ implicitly, $u_{3}+G u_{4}$ is always maintaining s implicitly.

Proof. We only focus on the PartialUpdate. The FullUpdate is trivial, we ignore the proof.

For $x$.

Note that $M$ is not changing. Let's assume that $u_{1}+F u_{2}=x$, we want to show that

$$
u_{1}^{\text {new }}+F^{\mathrm{new}} u_{2}^{\text {new }}=x^{\text {new }} .
$$

which is equivalent to prove

$$
u_{1}^{\text {new }}+F^{\text {new }} u_{2}^{\text {new }}-\left(u_{1}+F u_{2}\right)=\delta_{x}
$$

Let $\Delta u_{1}=u_{1}^{\text {new }}-u_{1}$ be the change of $u_{1}$ over iteration $t$, then

$$
\Delta u_{1}=\widetilde{V}^{\text {new }} h+\left(F-F^{\text {new }}\right) u_{2}
$$

Let $\Delta u_{2}=u_{2}^{\text {new }}-u_{2}$ be the change of $u_{2}$ over iteration $t$, then

$$
\Delta u_{2}=-\left(\widetilde{V}^{\text {new }}\right)^{1 / 2} h+\mathbf{1}_{\widetilde{S}}\left(\widetilde{\Delta}_{\widetilde{S}, \widetilde{S}}^{-1}+M_{\widetilde{S}, \widetilde{S}}\right)^{-1} M_{\widetilde{S}}^{\top}\left(\widetilde{V}^{\text {new }}\right)^{1 / 2} h .
$$

By definition of $\delta_{x}$ at iteration $t$, we have

$$
\delta_{x}=\widetilde{V}^{\text {new }} h-\left(\sqrt{\widetilde{V}^{\text {new }}} M \sqrt{\widetilde{V}^{\text {new }}} h-\sqrt{\widetilde{V}^{\text {new }}} M_{\widetilde{S}}\left(\widetilde{\Delta}_{\widetilde{S}, \widetilde{S}}^{-1}+M_{\widetilde{S}, \widetilde{S}}\right)^{-1}\left(M_{\widetilde{S}}\right)^{\top} \sqrt{\widetilde{V}^{\text {new }}} h\right) .
$$

We can compute

$$
\begin{aligned}
& u_{1}^{\text {new }}+F^{\text {new }} u_{2}^{\text {new }}-\left(u_{1}+F u_{2}\right) \\
= & \Delta u_{1}+\left(F^{\text {new }} u_{2}^{\text {new }}-F u_{2}\right) \\
= & \widetilde{V}^{\text {new }} h+\left(F-F^{\text {new }}\right) u_{2}+\left(F^{\text {new }} u_{2}^{\text {new }}-F u_{2}\right) \\
= & \widetilde{V}^{\text {new }} h+F^{\text {new }}\left(u_{2}^{\text {new }}-u_{2}\right) \\
= & \widetilde{V}^{\text {new }} h+F^{\text {new }} \Delta u_{2} \\
= & \widetilde{V}^{\text {new }} h-F^{\text {new }} \sqrt{\widetilde{V}^{\text {new }}} h+F^{\text {new }} 1_{\widetilde{S}}\left(\widetilde{\Delta}_{\widetilde{S}, \widetilde{S}}^{-1}+M_{\widetilde{S}, \widetilde{S}}\right)^{-1}\left(M_{\widetilde{S}}\right)^{\top} \sqrt{\widetilde{V}^{\text {new }}} h \\
= & \delta_{x}
\end{aligned}
$$

where we used $F^{\text {new }}=\sqrt{\widetilde{V}^{\text {new }}} M$ in the last step.

## For $s$.

We have

$$
G^{\text {new }}=\frac{1}{\sqrt{\widetilde{V}^{\text {new }}}} M, G=\frac{1}{\sqrt{\widetilde{V}}} M
$$

Let $\Delta u_{3}=u_{3}^{\text {new }}-u_{3}$ be the change of $u_{3}$ over iteration $t$, then

$$
\Delta u_{3}=\left(G-G^{\mathrm{new}}\right) u_{4}
$$

Let $\Delta u_{4}=u_{4}^{\text {new }}-u_{4}$ be the change of $u_{4}$ over iteration $t$, then

$$
\Delta u_{4}=t \cdot \Delta u_{2}
$$

By definition of $\delta_{s}$ in iteration $t$,

$$
\delta_{s}=\left(\frac{1}{\sqrt{\widetilde{V}^{\text {new }}}} M \sqrt{\widetilde{V}^{\text {new }}}(t h)-\frac{1}{\sqrt{\widetilde{V}^{\text {new }}}} M_{\widetilde{S}}\left(\widetilde{\Delta}_{\widetilde{S}, \widetilde{S}}^{-1}+M_{\widetilde{S}, \widetilde{S}}\right)^{-1}\left(M_{\widetilde{S}}\right)^{\top} \sqrt{\widetilde{V}^{\text {new }}}(t h)\right)
$$

We can compute

$$
\begin{aligned}
\left(u_{3}^{\text {new }}+G^{\mathrm{new}} u_{4}^{\mathrm{new}}\right)-\left(u_{3}+G u_{4}\right) & =\Delta u_{3}+\left(G^{\mathrm{new}} u_{4}^{\mathrm{new}}-G u_{4}\right) \\
& =\left(G-G^{\mathrm{new}}\right) u_{4}+\left(G^{\mathrm{new}} u_{4}^{\mathrm{new}}-G u_{4}\right) \\
& =G^{\mathrm{new}}\left(u_{4}^{\mathrm{new}}-u_{4}\right) \\
& =G^{\mathrm{new}} t \Delta u_{2} \\
& =\delta_{s}
\end{aligned}
$$

where the last step follows by definition of $\Delta u_{2}$.

### 3.7 Combining Robust Central Path with Data Structure

The goal of this section is to combine Section 3.5 and Section 3.6.

| Ntn | Choice of Parameter | Statement | Comment |
| :--- | :--- | :--- | :--- |
| $C_{1}$ | $\Theta\left(1 / \log ^{2} n\right)$ | Lem. 3.7.1, Thm. 3.6.1 | $\ell_{2}$ accuracy of $W$ sequence |
| $C_{2}$ | $\Theta\left(1 / \log ^{4} n\right)$ | Lem. 3.7.1, Thm. 3.6.1 | $\ell_{4}$ accuracy of $W$ sequence |
| $\epsilon_{m p}$ | $\Theta\left(1 / \log ^{2} n\right)$ | RoBusTIPM Alg in Sec. 3.5 | accuracy for data structure |
| $T$ | $\Theta\left(\sqrt{n} \log ^{2} n \log (n / \delta)\right)$ | Thm. 3.4.1 | \#iterations |
| $\alpha$ | $\Theta\left(1 / \log ^{2} n\right)$ | RobustIPM Alg in Sec. 3.5 | step size in Hessian norm |
| $b$ | $\Theta\left(\sqrt{n} \log ^{6}(n T)\right.$ | Lem. 3.6.15, 3.6.16, 3.7.2 | sketch size |
| $\epsilon_{x}$ | $\Theta\left(1 / \log ^{3} n\right)$ | Lem. 3.6.15 | accuracy of $\bar{x}$ (respect to $x)$ |
| $\epsilon_{s}$ | $\Theta\left(1 / \log ^{3} n\right)$ | Lem. 3.6.16 | accuracy of $\bar{s}$ (respect to $s)$ |
| $\epsilon_{w}$ | $\Theta\left(1 / \log ^{3} n\right)$ | Lem. 3.7.2 | accuracy of $\bar{W}$ (respect to $W)$ |
| $a$ | $\min \left(2 / 3, \alpha_{m}\right)$ | $\alpha_{m}$ is the dual exponent of MM | batch size |

Table 3.3: Summary of parameters. "Ntn" denotes Notation.

### 3.7.1 Guarantee for $W$ matrices

Lemma 3.7.1 (Guarantee of a sequence of $W$ ). Let $x^{\text {new }}=x+\delta_{x}$. Let $W^{\text {new }}=\left(\nabla^{2} \phi\left(x^{\text {new }}\right)\right)^{-1}$ and $W=\left(\nabla^{2} \phi(x)\right)^{-1}$. Then we have

$$
\begin{aligned}
\sum_{i=1}^{m}\left\|w_{i}^{-1 / 2}\left(w_{i}^{\text {new }}-w_{i}\right) w_{i}^{-1 / 2}\right\|_{F}^{2} & \leq C_{1}^{2} \\
\sum_{i=1}^{m}\left\|w_{i}^{-1 / 2}\left(w_{i}^{\text {new }}-w_{i}\right) w_{i}^{-1 / 2}\right\|_{F}^{4} & \leq C_{2}^{2} \\
\left\|w_{i}^{-1 / 2}\left(w_{i}^{\text {new }}-w_{i}\right) w_{i}^{-1 / 2}\right\|_{F} & \leq \frac{1}{4}
\end{aligned}
$$

where $C_{2}=\Theta\left(\alpha^{2}\right)$ and $C_{1}=\Theta(\alpha)$.

Proof. For each $i \in[m]$, we have

$$
\begin{aligned}
& \left\|W_{i}^{-1 / 2}\left(W_{i}^{\text {new }}-W_{i}\right) W_{i}^{-1 / 2}\right\|_{F}^{2} \\
= & n_{i}\left\|W_{i}^{-1 / 2}\left(W_{i}^{\text {new }}-W_{i}\right) W_{i}^{-1 / 2}\right\|^{2} \\
= & n_{i}\left\|\left(\nabla^{2} \phi\left(x_{i}\right)\right)^{1 / 2}\left(\nabla^{2} \phi\left(x_{i}^{\text {new }}\right)^{-1}-\nabla^{2} \phi\left(x_{i}\right)^{-1}\right)\left(\nabla^{2} \phi\left(x_{i}\right)\right)^{1 / 2}\right\|^{2} \\
\leq & \left(\frac{1}{\left(1-\left\|x_{i}^{\text {new }}-x_{i}\right\|_{\nabla^{2} \phi\left(x_{i}\right)}\right)^{2}}-1\right)^{2} \cdot\left\|\left(\nabla^{2} \phi\left(x_{i}\right)\right)^{1 / 2} \nabla^{2} \phi\left(x_{i}\right)^{-1}\left(\nabla^{2} \phi\left(x_{i}\right)\right)^{1 / 2}\right\|^{2} \\
= & n_{i}\left(\frac{1}{\left(1-\left\|x_{i}^{\text {new }}-x_{i}\right\|_{\nabla^{2} \phi\left(x_{i}\right)}\right)^{2}}-1\right)^{2} \\
\leq & 100 n_{i}\left\|x_{i}^{\text {new }}-x_{i}\right\|_{\nabla^{2} \phi\left(x_{i}\right)}^{2},
\end{aligned}
$$

where the second step follows by Theorem 3.2.1.
In our problem, we assume that $n_{i}=O(1)$. It remains to bound

$$
\left\|x_{i}^{\text {new }}-x_{i}\right\|_{\nabla^{2} \phi\left(x_{i}\right)}^{2}=\left\|\delta_{x, i}\right\|_{\nabla^{2} \phi\left(x_{i}\right)}^{2} \lesssim\left\|\delta_{x, i}\right\|_{\bar{x}_{i}}^{2}=\alpha_{i}^{2}
$$

where the last step follows from definition $\alpha_{i}=\left\|\delta_{x, i}\right\|_{\bar{x}_{i}}$.

Then, we have

$$
\sum_{i=1}^{m}\left\|x_{i}^{\mathrm{new}}-x_{i}\right\|_{\nabla^{2} \phi\left(x_{i}\right)}^{2} \leq \sum_{i=1}^{m} O\left(\alpha_{i}^{2}\right) \leq O\left(\alpha^{2}\right)
$$

where the last step follows by Lemma 3.5.1.

Lemma 3.7.2 (Accuracy of $\bar{W}$ ). Let $x$ and $\bar{x}$ be the vectors maintained by data-structure StochasticProjectionMaintenance. Let $W=\left(\nabla^{2} \phi(x)\right)^{-1}$ and $\bar{W}=\left(\nabla^{2} \phi(\bar{x})\right)^{-1}$.

Then we have

$$
\left\|w_{i}^{-1 / 2}\left(\bar{w}_{i}-w_{i}\right) w_{i}^{-1 / 2}\right\|_{F} \leq \epsilon_{w},
$$

where $\epsilon_{w}=O\left(\alpha \log ^{2}(n T) \cdot \frac{n^{1 / 4}}{\sqrt{b}}\right), b$ is the size of sketching matrix.

Proof. By similar calculation, we have

$$
\left\|w_{i}^{-1 / 2}\left(\bar{w}_{i}-w_{i}\right) w_{i}^{-1 / 2}\right\|_{F}=O(1) \cdot\left\|\bar{x}_{i}-x_{i}\right\|_{\nabla^{2} \phi\left(x_{i}\right)} .
$$

Then, using Lemma 3.6.15 with $\delta=1 / T$

$$
\left\|\bar{x}_{i}-x_{i}\right\|_{\nabla^{2} \phi\left(x_{i}\right)} \leq O\left(\alpha \log ^{2}(n T) \cdot \frac{\sqrt{n^{1 / 4}}}{\sqrt{b}}\right) .
$$

```
Algorithm 3.5 Robust Central Path
    procedure CentralPathStep \((\bar{x}, \bar{s}, t, \lambda, \alpha)\)
        for \(i=1 \rightarrow m\) do \(\quad \triangleright\) Figure out direction \(h\)
            \(\mu_{i}^{t} \leftarrow \bar{s}_{i} / t+\nabla \phi_{i}\left(\bar{x}_{i}\right) \quad \triangleright\) According to Eq. (3.6)
            \(\gamma_{i}^{t} \leftarrow\left\|\mu_{i}^{t}\right\|_{\nabla^{2} \phi_{i}\left(\bar{x}_{i}\right)^{-1}} \quad \triangleright\) According to Eq. (3.7)
            \(c_{i}^{t} \leftarrow \frac{\exp \left(\lambda \gamma_{i}^{t}\right) / \gamma_{i}^{t}}{\left(\sum_{i=1}^{m} \exp \left(2 \lambda \gamma_{i}^{t}\right)\right)^{1 / 2}}\) if \(\gamma_{i}^{t} \geq 96 \sqrt{\alpha}\) and \(c_{i}^{t} \leftarrow 0\) otherwise \(\triangleright\) According to Eq. (3.9)
            \(h_{i} \leftarrow-\alpha \cdot c_{i}^{t} \cdot \mu_{i}^{t} \quad \triangleright\) According to Eq. (3.8)
        end for
        \(\bar{W} \leftarrow\left(\nabla^{2} \phi(\bar{x})\right)^{-1} \quad \triangleright\) Computing block-diagonal matrix \(\bar{W}\)
        return \(h, \bar{W}\)
    end procedure
    procedure RobustCentralPath (mp, \(t, \lambda, \alpha)\)
                                    \(\triangleright\) Lemma 3.5.8
13: \(\triangleright\) Standing at \((x, s)\) implicitly via data-structure
                                    \(\triangleright\) Standing at \((\bar{x}, \bar{s})\) explicitly via data-structure
        \((\bar{x}, \bar{s}) \leftarrow \operatorname{mp} \cdot \operatorname{QUERY}() \quad \triangleright\) Algorithm 3.1, Lemma 3.6.6
        \(h, \bar{W} \leftarrow \operatorname{CentraLPathStep}(\bar{x}, \bar{s}, t, \lambda, \alpha)\)
        mp.Update \((\bar{W})\)
                                \(\triangleright\) Algorithm 3.2, Lemma 3.6.3
        mp.MultiplyMove \((h, t) \quad \triangleright\) Algorithm 3.4, Lemma 3.6.8, Lemma 3.6.7
                        \(\triangleright x \leftarrow x+\delta_{x}, s \leftarrow s+\delta_{s}\), achieved by data-structure implicitly
                        \(\triangleright \bar{x} \leftarrow \bar{x}+\widetilde{\delta}_{x}, \bar{s} \leftarrow \bar{s}+\widetilde{\delta}_{s}\), achieved by data-structure explicitly
                                    \(\triangleright\) If \(x\) is far from \(\bar{x}\), then \(\bar{x} \leftarrow x\)
    end procedure
```


### 3.7.2 Main result

The goal of this section is to prove our main result.

Theorem 3.7.3 (Main result, formal version of Theorem 3.1.1). Consider a convex problem

$$
\min _{A x=b, x \in \prod_{i=1}^{m} K_{i}} c^{\top} x
$$

where $K_{i}$ are compact convex set. For each $i \in[m]$, we are given a $\nu_{i}$-self concordant barrier

```
Algorithm 3.6 Our main algorithm (More detailed version of RobustIPM in Section 3.4)
    procedure \(\operatorname{Main}(A, b, c, \phi, \delta)\)
                        \(\triangleright\) Theorem 3.1.1, Theorem 3.7.3
        \(\lambda \leftarrow 2^{16} \log (m), \alpha \leftarrow 2^{-20} \lambda^{-2}, \kappa \leftarrow 2^{-10} \alpha\)
        \(\delta \leftarrow \min \left(\frac{1}{\lambda}, \delta\right) \quad \triangleright\) Choose the target accuracy
        \(a \leftarrow \min \left(2 / 3, \alpha_{m}\right) \quad \triangleright\) Choose the batch size
        \(b_{\text {sketch }} \leftarrow 2^{10} \sqrt{\nu} \log ^{6}(n / \delta) \cdot \log \log (1 / \delta) \quad \triangleright\) Choose the size of sketching matrix
        Modify the \(\operatorname{ERM}(A, b, c, \phi)\) and obtain an initial \(x\) and \(s\)
        CentralPathMaintenance mp \(\triangleright\) Algorithm 3.1, Theorem 3.6.1
        \(\operatorname{mp} . \operatorname{Initialize}\left(A, x, s, \alpha, a, b_{\text {sketch }}\right) \quad \triangleright\) Algorithm 3.1, Lemma 3.6.2
        \(\nu \leftarrow \sum_{i=1}^{m} \nu_{i} \quad \triangleright \nu_{i}\) are the self-concordant parameters of \(\phi_{i}\)
        \(t \leftarrow 1\)
        while \(t>\delta^{2} /(4 \nu)\) do
            \(t^{\text {new }} \leftarrow\left(1-\frac{\kappa}{\sqrt{\nu}}\right) t\)
            RobustCentralPath \((\mathrm{mp}, t, \lambda, \alpha) \quad\) Algorithm 3.5
                \(t \leftarrow t^{\text {new }}\)
        end while
        Return an approximate solution of the original ERM according to Section 3.8
    end procedure
```

function $\phi_{i}$ for $K_{i}$. Also, we are given $x^{(0)}=\arg \min _{x} \sum_{i} \phi_{i}\left(x_{i}\right)$. Assume that

1. Diameter of the set: For any $x \in \prod_{i=1}^{m} K_{i}$, we have that $\|x\|_{2} \leq R$.
2. Lipschitz constant of the program: $\|c\|_{2} \leq L$.

Then, the algorithm Main finds a vector $x$ such that

$$
\begin{aligned}
c^{\top} x & \leq \min _{A x=b, x \in \prod_{i=1}^{m} K_{i}} c^{\top} x+L R \cdot \delta, \\
\|A x-b\|_{1} & \leq 3 \delta \cdot\left(R \sum_{i, j}\left|A_{i, j}\right|+\|b\|_{1}\right), \\
x & \in \prod_{i=1}^{m} K_{i} .
\end{aligned}
$$

in time

$$
O\left(n^{\omega+o(1)}+n^{2.5-\alpha / 2+o(1)}+n^{2+1 / 6+o(1)}\right) \cdot \widetilde{O}(\log (n / \delta)) .
$$

where $\omega$ is the exponent of matrix multiplication [Wil12, LG14], and $\alpha$ is the dual exponent of matrix multiplication [LGU18].

Proof. The number of iterations is

$$
O\left(\sqrt{\nu} \log ^{2}(m) \log (\nu / \delta)\right)=O\left(\sqrt{n} \log ^{2}(n) \log (n / \delta)\right)
$$

For each iteration, the amortized cost per iteration is

$$
\begin{aligned}
& O\left(n b+n^{1+a}+n^{1.5}\right)+O\left(C_{1} / \epsilon_{m p}+C_{2} / \epsilon_{m p}^{2}\right) \cdot\left(n^{\omega-1 / 2+o(1)}+n^{2-a / 2+o(1)}\right)+O\left(n^{\omega-1 / 2+o(1)}\right) \\
= & O\left(n b+n^{1+a}+n^{1.5}\right)+O\left(\alpha+\alpha^{2}\right) \cdot\left(n^{\omega-1 / 2+o(1)}+n^{2-a / 2+o(1)}\right)+O\left(n^{\omega-1 / 2+o(1)}\right) \\
= & O\left(n b+n^{1+a}+n^{1.5}\right)+O\left(1 / \log ^{4} n\right) \cdot\left(n^{\omega-1 / 2+o(1)}+n^{2-a / 2+o(1)}\right)+O\left(n^{\omega-1 / 2+o(1)}\right) \\
= & O\left(n^{1.5+o(1)} \log ^{6} \log (1 / \delta)+n^{1+a+o(1)}\right)+O\left(n^{\omega-1 / 2+o(1)}+n^{2-a / 2+o(1)}\right) .
\end{aligned}
$$

where the last step follows from choice of $b$ (see Table 3.3).

Finally, we have
total time
$=\#$ iterations $\cdot$ cost per iteration

$$
\begin{aligned}
& =\underbrace{O\left(\sqrt{n} \log ^{2} n \log (n / \delta)\right)}_{\text {\#iterations }} \cdot \underbrace{O\left(n^{1.5+o(1)} \log ^{6} \log (1 / \delta)+n^{1+a+o(1)}+n^{\omega-1 / 2+o(1)}+n^{2-a / 2+o(1)}\right)}_{\text {cost per iteration }} \\
& =O\left(n^{1.5+a+o(1)}+n^{\omega+o(1)}+n^{2.5-a / 2+o(1)}\right) \cdot \log (n / \delta) \cdot \log ^{6} \log (1 / \delta) \\
& =O\left(n^{2+1 / 6+o(1)}+n^{\omega+o(1)}+n^{2.5-\alpha_{m} / 2+o(1)}\right) \cdot \log (n / \delta) \cdot \log ^{6} \log (1 / \delta)
\end{aligned}
$$

where we pick $a=\min \left(2 / 3, \alpha_{m}\right)$ and $\alpha_{m}$ is the dual exponent of matrix multiplication[LGU18]. Thus, we complete the proof.

Corollary 3.7.4 (Empirical risk minimization). Given convex function $f_{i}(y): \mathbb{R} \rightarrow \mathbb{R}$. Suppose the solution $x^{*} \in \mathbb{R}^{d}$ lies in $\ell_{\infty}-\operatorname{Ball}(0, R)$. Suppose $f_{i}$ is L-Lipschitz in region $\{y:|y| \leq 4 \sqrt{n} \cdot M \cdot R\}$. Given a matrix $A \in \mathbb{R}^{d \times n}$ with $\|A\| \leq M$ and $A$ has no redundant constraints, and a vector $b \in \mathbb{R}^{d}$ with $\|b\|_{2} \leq M \cdot R$. We can find $x \in \mathbb{R}^{d}$ s.t.

$$
\sum_{i=1}^{n} f_{i}\left(a_{i}^{\top} x+b_{i}\right) \leq \min _{x \in \mathbb{R}^{d}} \sum_{i=1}^{n} f_{i}\left(a_{i}^{\top} x+b_{i}\right)+\delta M R
$$

in time

$$
O\left(n^{\omega+o(1)}+n^{2.5-\alpha / 2+o(1)}+n^{2+1 / 6+o(1)}\right) \cdot \widetilde{O}(\log (n / \delta)) .
$$

where $\omega$ is the exponent of matrix multiplication [Wil12, LG14], and $\alpha$ is the dual exponent of matrix multiplication [LGU18].

Proof. It follows from applying Theorem 3.7.3 on convex program (3.1) with an extra constraint $x^{*}$ lies in $\ell_{\infty}-\operatorname{Ball}(0, R)$. Note that in program (3.1), $n_{i}=2$. Thus $m=O(n)$.

### 3.8 Initial Point and Termination Condition

We first need some result about self concordance.

Lemma 3.8.1 (Theorem 4.1.7, Lemma 4.2.4 in [Nes98]). Let $\phi$ be any $\nu$-self-concordant barrier. Then, for any $x, y \in \operatorname{dom} \phi$, we have

$$
\begin{aligned}
\langle\nabla \phi(x), y-x\rangle & \leq \nu \\
\langle\nabla \phi(y)-\nabla \phi(x), y-x\rangle & \geq \frac{\|y-x\|_{x}^{2}}{1+\|y-x\|_{x}} .
\end{aligned}
$$

Let $x^{*}=\arg \min _{x} \phi(x)$. For any $x \in \mathbb{R}^{n}$ such that $\left\|x-x^{*}\right\|_{x^{*}} \leq 1$, we have that $x \in \operatorname{dom} \phi$.

$$
\left\|x^{*}-y\right\|_{x^{*}} \leq \nu+2 \sqrt{\nu}
$$

Lemma 3.8.2. Consider a convex problem $\min _{A x=b, x \in \prod_{i=1}^{m} K_{i}} c^{\top} x$ where $K_{i}$ are compact convex set. For each $i \in[m]$, we are given a $\nu_{i}$-self concordant barrier function $\phi_{i}$ for $K_{i}$. Also, we are given $x^{(0)}=\arg \min _{x} \sum_{i} \phi_{i}\left(x_{i}\right)$. Assume that

1. Diameter of the set: For any $x \in \prod_{i=1}^{m} K_{i}$, we have that $\|x\|_{2} \leq R$.
2. Lipschitz constant of the program: $\|c\|_{2} \leq L$.

For any $\delta>0$, the modified program $\min _{\bar{A} \bar{x}=\bar{b}, \bar{x} \in \prod_{i=1}^{m} K_{i} \times \mathbb{R}_{+}} \bar{c}^{\top} \bar{x}$ with

$$
\bar{A}=\left[A \mid b-A x^{(0)}\right], \bar{b}=b, \text { and } \bar{c}=\left[\begin{array}{c}
\frac{\delta}{L R} \cdot c \\
1
\end{array}\right]
$$

satisfies the following:

1. $\bar{x}=\left[\begin{array}{c}x^{(0)} \\ 1\end{array}\right], \bar{y}=0_{d}$ and $\bar{s}=\left[\begin{array}{c}\frac{\delta}{L R} \cdot c \\ 1\end{array}\right]$ are feasible primal dual vectors with $\| \bar{s}+$ $\nabla \bar{\phi}(\bar{x}) \|_{\bar{x}}^{*} \leq \delta$ where $\bar{\phi}(\bar{x})=\sum_{i=1}^{m} \phi_{i}\left(\bar{x}_{i}\right)-\log \left(\bar{x}_{m+1}\right)$.
2. For any $\bar{x}$ such that $\bar{A} \bar{x}=\bar{b}, \bar{x} \in \prod_{i=1}^{m} K_{i} \times \mathbb{R}_{+}$and $\bar{c}^{\top} \bar{x} \leq \min _{\bar{A} \bar{x}=\bar{b}, \bar{x} \in \prod_{i=1}^{m} K_{i} \times \mathbb{R}_{+}} \bar{c}^{\top} \bar{x}+\delta^{2}$, the vector $\bar{x}_{1: n}\left(\bar{x}_{1: n}\right.$ is the first $n$ coordinates of $\left.\bar{x}\right)$ is an approximate solution to the original convex program in the following sense

$$
\begin{aligned}
c^{\top} \bar{x}_{1: n} & \leq \min _{A x=b, x \in \prod_{i=1}^{m} K_{i}} c^{\top} x+L R \cdot \delta \\
\left\|A \bar{x}_{1: n}-b\right\|_{1} & \leq 3 \delta \cdot\left(R \sum_{i, j}\left|A_{i, j}\right|+\|b\|_{1}\right), \\
\bar{x}_{1: n} & \in \prod_{i=1}^{m} K_{i} .
\end{aligned}
$$

Proof. For the first result, straightforward calculations show that $(\bar{x}, \bar{y}, \bar{s})$ are feasible.
To compute $\|\bar{s}+\nabla \bar{\phi}(\bar{x})\|_{\bar{x}}^{*}$, note that

$$
\|\bar{s}+\nabla \bar{\phi}(\bar{x})\|_{\bar{x}}^{*}=\left\|\frac{\delta}{L R} \cdot c\right\|_{\nabla^{2} \phi\left(x^{(0)}\right)^{-1}} .
$$

Lemma 3.8.1 shows that $x \in \mathbb{R}^{n}$ such that $\left\|x-x^{(0)}\right\|_{x^{(0)}} \leq 1$, we have that $x \in \prod_{i=1}^{m} K_{i}$ because $x^{(0)}=\arg \min _{x} \sum_{i} \phi_{i}\left(x_{i}\right)$. Hence, for any $v$ such that $v^{\top} \nabla^{2} \phi\left(x^{(0)}\right) v \leq 1$, we have that $x^{(0)} \pm v \in \prod_{i=1}^{m} K_{i}$ and hence $\left\|x^{(0)} \pm v\right\|_{2} \leq R$. This implies $\|v\|_{2} \leq R$ for any $v^{\top} \nabla^{2} \phi\left(x^{(0)}\right) v \leq 1$. Hence, $\left(\nabla^{2} \phi\left(x^{(0)}\right)\right)^{-1} \preceq R^{2} \cdot I$. Hence, we have

$$
\|\bar{s}+\nabla \bar{\phi}(\bar{x})\|_{\bar{x}}^{*}=\left\|\frac{\delta}{L R} \cdot c\right\|_{\nabla^{2} \phi\left(x^{(0)}\right)^{-1}} \leq\left\|\frac{\delta}{L} \cdot c\right\|_{2} \leq \delta .
$$

For the second result, we let

$$
\mathrm{OPT}=\min _{A x=b, x \in \prod_{i=1}^{m} K_{i}} c^{\top} x \quad \text { and } \quad \overline{\mathrm{OPT}}=\min _{\bar{A} \bar{x}=b, \bar{x} \in \prod_{i=1}^{m} K_{i} \times \mathbb{R}_{+}} \bar{c}^{\top} \bar{x} .
$$

For any feasible $x$ in the original problem, $\bar{x}=\left[\begin{array}{l}x \\ 0\end{array}\right]$ is a feasible in the modified problem. Therefore, we have that

$$
\overline{\mathrm{OPT}} \leq \frac{\delta}{L R} \cdot c^{\top} x=\frac{\delta}{L R} \cdot \mathrm{OPT}
$$

Given a feasible $\bar{x}$ with additive error $\delta^{2}$. Write $\bar{x}=\left[\begin{array}{c}\bar{x}_{1: n} \\ \tau\end{array}\right]$ for some $\tau \geq 0$. We can compute $\bar{c}^{\top} \bar{x}$ which is $\frac{\delta}{L R} \cdot c^{\top} \bar{x}_{1: n}+\tau$. Then, we have

$$
\begin{equation*}
\frac{\delta}{L R} \cdot c^{\top} \bar{x}_{1: n}+\tau \leq \overline{\mathrm{OPT}}+\delta^{2} \leq \frac{\delta}{L R} \cdot \mathrm{OPT}+\delta^{2} \tag{3.34}
\end{equation*}
$$

Hence, we can upper bound the OPT of the transformed program as follows:

$$
c^{\top} \bar{x}_{1: n}=\frac{L R}{\delta} \cdot \frac{\delta}{L R} c^{\top} \bar{x}_{1: n} \leq \frac{L R}{\delta}\left(\frac{\delta}{L R} \cdot \mathrm{OPT}+\delta^{2}\right)=\mathrm{OPT}+L R \cdot \delta
$$

where the second step follows by (3.34).
For the feasibility, we have that $\tau \leq-\frac{\delta}{L R} \cdot c^{\top} \bar{x}_{1: n}+\frac{\delta}{L R} \cdot \mathrm{OPT}+\delta^{2} \leq \delta+\delta+\delta$ because $\mathrm{OPT}=\min _{A x=b, x \geq 0} c^{\top} x \leq L R$ and that $c^{\top} \bar{x}_{1: n} \leq L R$. The constraint in the new polytope shows that

$$
A \bar{x}_{1: n}+\left(b-A x^{(0)}\right) \tau=b .
$$

Rewriting it, we have $A \bar{x}_{1: n}-b=\left(A x^{(0)}-b\right) \tau$ and hence

$$
\left\|A \bar{x}_{1: n}-b\right\|_{1} \leq\left\|A x^{(0)}-b\right\|_{1} \cdot \tau .
$$

Lemma 3.8.3. Let $\phi_{i}\left(x_{i}\right)$ be a $\nu_{i}$-self-concordant barrier. Suppose we have $\frac{s_{i}}{t}+\nabla \phi_{i}\left(x_{i}\right)=\mu_{i}$ for all $i \in[m], A^{\top} y+s=c$ and $A x=b$. Suppose that $\left\|\mu_{i}\right\|_{x, i}^{*} \leq 1$ for all $i$, we have that

$$
\langle c, x\rangle \leq\left\langle c, x^{*}\right\rangle+4 t \nu
$$

where $x^{*}=\arg \min _{A x=b, x \in \prod_{i=1}^{m} K_{i}} c^{\top} x$ and $\nu=\sum_{i=1}^{m} \nu_{i}$.

Proof. Let $x_{\alpha}=(1-\alpha) x+\alpha x^{*}$ for some $\alpha$ to be chosen. By Lemma 3.8.1, we have that $\left\langle\nabla \phi\left(x_{\alpha}\right), x^{*}-x_{\alpha}\right\rangle \leq \nu$. Hence, we have $\frac{\nu}{1-\alpha} \geq\left\langle\nabla \phi\left(x_{\alpha}\right), x^{*}-x\right\rangle$. Hence, we have

$$
\begin{aligned}
\frac{\nu \alpha}{1-\alpha} & \geq\left\langle\nabla \phi\left(x_{\alpha}\right), x_{\alpha}-x\right\rangle \\
& =\left\langle\nabla \phi\left(x_{\alpha}\right)-\nabla \phi(x), x_{\alpha}-x\right\rangle+\left\langle\mu-\frac{s}{t}, x_{\alpha}-x\right\rangle \\
& \geq \sum_{i=1}^{m} \frac{\left\|x_{\alpha, i}-x_{i}\right\|_{x_{i}}^{2}}{1+\left\|x_{\alpha, i}-x_{i}\right\|_{x_{i}}}+\left\langle\mu, x_{\alpha}-x\right\rangle-\frac{1}{t}\left\langle c-A^{\top} y, x_{\alpha}-x\right\rangle \\
& \geq \sum_{i=1}^{m} \frac{\alpha^{2}\left\|x_{i}^{*}-x_{i}\right\|_{x_{i}}^{2}}{1+\alpha\left\|x_{i}^{*}-x_{i}\right\|_{x_{i}}}-\alpha \sum_{i=1}^{m}\left\|\mu_{i}\right\|_{x_{i}}^{*}\left\|x_{i}^{*}-x_{i}\right\|_{x_{i}}-\frac{\alpha}{t}\left\langle c, x^{*}-x\right\rangle .
\end{aligned}
$$

where we used Lemma 3.8.1 on the second first, $A x_{\alpha}=A x$ on the second inequality. Hence, we have

$$
\frac{\langle c, x\rangle}{t} \leq \frac{\left\langle c, x^{*}\right\rangle}{t}+\frac{\nu}{1-\alpha}+\sum_{i=1}^{m}\left\|\mu_{i}\right\|_{x_{i}}^{*}\left\|x_{i}^{*}-x_{i}\right\|_{x_{i}}-\sum_{i=1}^{m} \frac{\alpha\left\|x_{i}^{*}-x_{i}\right\|_{x_{i}}^{2}}{1+\alpha\left\|x_{i}^{*}-x_{i}\right\|_{x_{i}}}
$$

Using $\left\|\mu_{i}\right\|_{x_{i}}^{*} \leq 1$ for all $i$, we have

$$
\frac{\langle c, x\rangle}{t} \leq \frac{\left\langle c, x^{*}\right\rangle}{t}+\frac{\nu}{1-\alpha}+\sum_{i=1}^{m} \frac{\left\|x_{i}^{*}-x_{i}\right\|_{x_{i}}}{1+\alpha\left\|x_{i}^{*}-x_{i}\right\|_{x_{i}}} \leq \frac{\left\langle c, x^{*}\right\rangle}{t}+\frac{\nu}{1-\alpha}+\frac{m}{\alpha} .
$$

Setting $\alpha=\frac{1}{2}$, we have $\langle c, x\rangle \leq\left\langle c, x^{*}\right\rangle+2 t(\nu+m) \leq\left\langle c, x^{*}\right\rangle+4 t \nu$ because the self-concordance $\nu_{i}$ is always larger than 1.

### 3.9 Basic Properties of Subsampled Hadamard Transform Matrix

This section provides some standard calculations about sketching matrices, it can be found in previous literatures [PSW17]. Usually, the reason for using subsampled randomized Hadamard/Fourier transform [LDFU13] is multiplying the matrix with $k$ vectors only takes $k n \log n$ time. Unfortunately, in our application, the best way to optimize the running is using matrix multiplication directly (without doing any fast Fourier transform [CT65], or more fancy sparse Fourier transform [HIKP12b, HIKP12a, Pri13, IKP14, IK14, PS15, CKPS16, Kap16, Kap17, NSW19a]). In order to have an easy analysis, we still use subsampled randomized Hadamard/Fourier matrix.

### 3.9.1 Properties obtained by random projection

Remark 3.9.1. The Subsampled Randomized Hadamard Transform [LDFU13] can be defined as $R=S H_{n} \Sigma \in \mathbb{R}^{b \times}$, where $\Sigma$ is an $n \times n$ diagonal matrix with i.i.d. diagonal entries $\Sigma_{i, i}$ in which $\Sigma_{i, i}=1$ with probability $1 / 2$, and $\Sigma_{i, i}=-1$ with probability $1 / 2 . H_{n}$ refers to the Hadamard matrix of size $n$, which we assume is a power of 2 . The $b \times n$ matrix $S$ samples $b$ coordinates of $n$ dimensional vector uniformly at random. If we replace the definition of sketching matrix in Lemma 3.9.1 by Subsampled Randomized Hadamard Transform and let $\bar{R}=S H_{n}$, then the same proof will go through.

Lemma 3.9.1 (Expectation, variance, absolute guarantees for sketching a fixed vector). Let $h \in \mathbb{R}^{n}$ be a fixed vector. Let $\bar{R} \in \mathbb{R}^{b \times n}$ denote a random matrix where each entry is i.i.d. sampled from $+1 / \sqrt{b}$ with probability $1 / 2$ and $-1 / \sqrt{b}$ with probability $1 / 2$. Let $\Sigma \in \mathbb{R}^{n \times n}$ denote a diagonal matrix where each entry is 1 with probability $1 / 2$ and -1 with probability $1 / 2$. Let $R=\bar{R} \Sigma$, then we have

$$
\mathbb{E}\left[R^{\top} R h\right]=h, \quad \mathbb{E}\left[\left(R^{\top} R h\right)_{i}^{2}\right] \leq h_{i}^{2}+\frac{1}{b}\|h\|_{2}^{2}, \quad \operatorname{Pr}\left[\left|\left(R^{\top} R h\right)_{i}-h_{i}\right|>\|h\|_{2} \frac{\log (n / \delta)}{\sqrt{b}}\right] \leq \delta
$$

Proof. Let $R_{i, j}$ denote the entry at $i$-th row and $j$-th column in matrix $R \in \mathbb{R}^{b \times n}$. Let $R_{*, i} \in \mathbb{R}^{b}$ denote the vector in $i$-th column of $R$.

We first show expectation,

$$
\begin{aligned}
\mathbb{E}\left[\left(R^{\top} R h\right)_{i}\right] & =\mathbb{E}\left[\left\langle R, R_{*, i} h^{\top}\right\rangle\right] \\
& =\mathbb{E}\left[\sum_{j=1}^{b} \sum_{l=1}^{n} R_{j, l} R_{j, i} h_{l}\right] \\
& =\mathbb{E}\left[\sum_{j=1}^{b} R_{j, i}^{2} h_{i}\right]+\mathbb{E}\left[\sum_{j=1}^{b} \sum_{l \in[n] \backslash i} R_{j, l} R_{j, i} h_{l}\right] \\
& =h_{i}+0 \\
& =h_{i}
\end{aligned}
$$

Secondly, we prove the variance is small

$$
\begin{aligned}
\mathbb{E}\left[\left(R^{\top} R h_{i}\right)^{2}\right]= & \mathbb{E}\left[\left\langle R, R_{*, i} h^{\top}\right\rangle^{2}\right] \\
= & \mathbb{E}\left[\left(\sum_{j=1}^{b} \sum_{l=1}^{n} R_{j, l} R_{j, i} h_{l}\right)^{2}\right] \\
= & \mathbb{E}\left[\left(\sum_{j=1}^{b} R_{j, i}^{2} h_{i}+\sum_{j=1}^{b} \sum_{l \in[n] \backslash i} R_{j, l} R_{j, i} h_{l}\right)^{2}\right] \\
= & \mathbb{E}\left[\left(\sum_{j=1}^{b} R_{j, i}^{2} h_{i}\right)^{2}\right]+2 \mathbb{E}\left[\sum_{j^{\prime}=1}^{b} R_{j^{\prime}, i}^{2} h_{i} \sum_{j=1}^{b} \sum_{l \in[n] \backslash i} R_{j, l} R_{j, i} h_{l}\right] \\
& +\mathbb{E}\left[\left(\sum_{j=1}^{b} \sum_{l \in[n] \backslash i} R_{j, l} R_{j, i} h_{l}\right)^{2}\right] \\
= & C_{1}+C_{2}+C_{3},
\end{aligned}
$$

where the last step follows from defining those terms to be $C_{1}, C_{2}$ and $C_{3}$. For the term $C_{1}$,
we have

$$
C_{1}=h_{i}^{2} \mathbb{E}\left[\left(\sum_{j=1}^{b} R_{j, i}^{2}\right)^{2}\right]=h_{i}^{2} \mathbb{E}\left[\sum_{j=1}^{b} R_{j, i}^{4}+\sum_{j^{\prime} \neq j} R_{j, i}^{2} R_{j^{\prime}, i}^{2}\right]=h_{i}^{2}\left(b \cdot \frac{1}{b^{2}}+b(b-1) \cdot \frac{1}{b^{2}}\right)=h_{i}^{2}
$$

For the second term $C_{2}$,

$$
C_{2}=0 .
$$

For the third term $C_{3}$,

$$
\begin{aligned}
C_{3} & =\mathbb{E}\left[\left(\sum_{j=1}^{b} \sum_{l \in[n] \backslash i} R_{j, l} R_{j, i} h_{l}\right)^{2}\right] \\
& =\mathbb{E}\left[\sum_{j=1}^{b} \sum_{l \in[n] \backslash i} R_{j, l}^{2} R_{j, i}^{2} h_{l}^{2}\right]+\mathbb{E}\left[\sum_{j=1}^{b} \sum_{l \in[n] \backslash i} R_{j, l} R_{j, i} h_{l} \sum_{j^{\prime} \in[b] \backslash j} \sum_{l^{\prime} \in[n] \backslash i \backslash l} R_{j^{\prime}, l^{\prime}} R_{j^{\prime}, i} h_{l^{\prime}}\right] \\
& =\sum_{j=1}^{b} \sum_{l \in[n] \backslash i} \frac{1}{b} \frac{1}{b} h_{l}^{2}+0 \leq \frac{1}{b}\|h\|_{2}^{2}
\end{aligned}
$$

Therefore, we have

$$
\mathbb{E}\left[\left(R^{\top} R h\right)_{i}^{2}\right] \leq C_{1}+C_{2}+C_{3} \leq h_{i}^{2}+\frac{1}{b}\|h\|_{2}^{2}
$$

Third, we prove the worst case bound with high probability. We can write $\left(R^{\top} R h\right)_{i}-$
$h_{i}$ as follows

$$
\begin{array}{rlrl}
\left(R^{\top} R h\right)_{i}-h_{i} & =\left\langle R, R_{*, i} h^{\top}\right\rangle-h_{i} & \\
& =\sum_{j=1}^{b} \sum_{l=1}^{n} R_{j, l} \cdot R_{j, i} \cdot h_{l}-h_{i} & \\
& =\sum_{j=1}^{b} R_{j, i}^{2} h_{i}-h_{i}+\sum_{j=1}^{b} \sum_{l \in[n] \backslash i} R_{j, l} R_{j, i} \cdot h_{l} & & \\
& =\sum_{j=1}^{b} \sum_{l \in[n] \backslash i} R_{j, l} R_{j, i} \cdot h_{l} & & \text { by } R_{j, i}^{2}=1 / b \\
& =\sum_{l \in[n] \backslash i} h_{l}\left\langle R_{*, l}, R_{*, i}\right\rangle & & \\
& =\sum_{l \in[n] \backslash i} h_{l} \cdot\left\langle\sigma_{l} \bar{R}_{*, l}, \sigma_{i} \bar{R}_{*, i}\right\rangle & & \text { by } R_{*, l}=\sigma_{l} \bar{R}_{*, l}
\end{array}
$$

First, we apply Khintchine's inequality, we have

$$
\operatorname{Pr}\left[\left|\sum_{l \in[n] \backslash i} h_{l} \cdot \sigma_{l} \cdot\left\langle\bar{R}_{*, l}, \sigma_{i} \bar{R}_{*, i}\right\rangle\right| \geq C t\left(\sum_{l \in[n] \backslash i} h_{l}^{2}\left(\left\langle\bar{R}_{*, l}, \sigma_{i} \bar{R}_{*, i}\right\rangle\right)^{2}\right)^{1 / 2}\right] \leq \exp \left(-C^{\prime} t^{2}\right)
$$

and choose $t=\sqrt{\log (n / \delta)}$.
For each $l \neq i$, using [LDFU13] we have

$$
\operatorname{Pr}\left[\left|\left\langle\bar{R}_{*, l}, \bar{R}_{*, i}\right\rangle\right| \geq \frac{\sqrt{\log (n / \delta)}}{\sqrt{b}}\right] \leq \delta / n .
$$

Taking a union bound over all $l \in[n] \backslash i$, we have

$$
\left|\left(R^{\top} R h\right)_{i}-h_{i}\right| \leq\|h\|_{2} \frac{\log (n / \delta)}{\sqrt{b}}
$$

with probability $1-\delta$.

## Chapter 4

## Algorithmic Theory of ODEs

Sampling logconcave functions arising in statistics and machine learning has been a subject of intensive study. Recent developments include analyses for Langevin dynamics and Hamiltonian Monte Carlo (HMC). While both approaches have dimension-independent bounds for the underlying continuous processes under sufficiently strong smoothness conditions, the resulting discrete algorithms have complexity and number of function evaluations growing with the dimension. Motivated by this problem, in this chapter, we give a general algorithm for solving multivariate ordinary differential equations whose solution is close to the span of a known basis of functions (e.g., polynomials or piecewise polynomials). The resulting algorithm has polylogarithmic depth and essentially tight runtime - it is nearly linear in the size of the representation of the solution.

We apply this to the sampling problem to obtain a nearly linear implementation of HMC for a broad class of smooth, strongly logconcave densities, with the number of iterations (parallel depth) and gradient evaluations being polylogarithmic in the dimension (rather than polynomial as in previous work). This class includes the widely-used loss function for logistic regression with incoherent weight matrices and has been subject of much study recently. We also give a faster algorithm with polylogarithmic depth for the more general and standard class of strongly convex functions with Lipschitz gradient. These results are based on (1)
an improved contraction bound for the exact HMC process and (2) logarithmic bounds on the degree of polynomials that approximate solutions of the differential equations arising in implementing HMC.

### 4.1 Introduction

The complexity of sampling a high-dimensional density of the form $e^{-f(x)}$ where $f$ is a convex function is a fundamental problem with many applications [LS90, LS92, LS93, LV06a, LV06b, Dal17, DK17, DRD18, DCWY18]. The focus of this chapter is to give very fast, i.e., nearly linear time algorithms, for a large subclass of such densities. A motivating and important case is the loss function for logistic regression, widely used in machine learning applications [Ber44, Paa00, NJ02, HJLS13, Bou14]:

$$
\sum_{i=1}^{n} \phi_{i}\left(a_{i}^{\top} x\right)
$$

where $\phi_{i}$ are convex functions; a popular choice is $\phi(t)=\log \left(1+e^{-t}\right)$. Sampling according to $e^{-f}$ for this choice of $f$ corresponds to sampling models according to their KL-divergence, a natural and effective choice for classification problems [HJLS13].

A general approach to sampling is by an ergodic Markov chain whose stationary distribution is designed to have the desired density. Traditionally, this is done via a Metropolis filter, which accepts a proposed (random) next step $y$ from the current point $x$ with probability $\min \left\{1, \frac{f(y)}{f(x)}\right\}$. While very general, one downside of this approach is the possibility of high rejection probabilities, which typically force local steps to be very small. Nevertheless, for arbitrary logconcave functions (including nonsmooth ones), this approach has the current best guarantees [LV06b].

Another family of algorithms is derived from an underlying continuous stochastic process with the desired stationary density. A classic example of such a continuous process is Brownian motion. To sample a convex body for example, one could use Brownian motion
with a boundary reflection condition. This is written as the stochastic equation:

$$
\mathrm{d} X_{t}=\mathrm{d} W_{t}
$$

with reflection at the boundary of the domain, and $d W_{t}$ being infinitesimal Brownian motion. To sample from the density proportional to $e^{-f(x)}$, one can use the stochastic differential equation,

$$
\mathrm{d} X_{t}=-\nabla f\left(X_{t}\right) \mathrm{d} t+\sqrt{2} \mathrm{~d} W_{t}
$$

By the classical Fokker-Planck equation, under mild assumptions on $f$, the stationary density of this process is proportional to $e^{-f(x)}$.

How can we turn these continuous processes into algorithms? One approach is to take small rather than infinitesimal steps, and this leads to the Langevin dynamics, of which there are multiple flavors [Dal17, DK17, ZLC17, RRT17, DRD18, CCBJ18, CCAY $\left.{ }^{+} 18, \mathrm{CFM}^{+} 18\right]$. Starting with Dalalyan [Dal17], it has been established that these dynamics converge in polynomial (in dimension) time for strongly logconcave functions, with the underdamped version converging in $O(\sqrt{d})$ iterations (and polynomial dependences on appropriate condition numbers) [CCBJ18]. The dependence on dimension seems unavoidable in the discretized algorithm, even though the continuous process has no such dependence.

Hamiltonian Monte Carlo. HMC is a random process that maintains a position $x$ and velocity pair $v$. To sample according to $e^{-f}$, we define a Hamiltonian $H(x, v)=f(x)+\frac{1}{2}\|v\|^{2}$. At each step $v$ is chosen randomly from $N(0, I)$ and $x$ is updated using the following Ordinary Differential Equation (ODE) for a some fixed time interval.

$$
\frac{\mathrm{d} x(t)}{\mathrm{d} t}=v(t), \quad \frac{\mathrm{d} v(t)}{\mathrm{d} t}=-\nabla f(x(t)) .
$$

This process has the particularly nice property that it conserves the value of $H$, and as a result there is no need to apply a Metropolis filter. HMC has been studied in many works [MS17, MV18, LV18]. Mangoubi and Smith [MS17] gave the following guarantee for strongly logconcave densities.

Theorem 4.1.1 ([MS17]). Let $f$ be a smooth, strongly convex function s.t. for all $y$

$$
m_{2} \cdot I \preceq \nabla^{2} f(y) \preceq M_{2} \cdot I .
$$

Then, HMC converges to the density proportional to $e^{-f}$ in $\widetilde{O}\left(\left(M_{2} / m_{2}\right)^{2}\right)$ iterations with each iteration being the exact solution of an ODE.

For the resulting algorithm presented in [MS17], which needs to approximate the solution of the ODE, the number of function evaluations and overall time grow as square-root of the dimension (and a higher polynomial of the condition number). Table 4.1 summarizes related work on sampling logconcave functions with various structural assumptions. In all these cases, even with higher-order smoothness and incoherence assumptions, the number of gradient/function evaluations grows as a polynomial in $d$. A special case of much interest is Bayesian logistic regression. To address this [MV18] define an incoherence parameter and achieve the previously best dependence on the dimension of $d^{1 / 4}$ for functions with bounded incoherence. They note that this is nearly optimal for the leapfrog implementation of HMC they use. Improving the complexity further, and in particular the dependence on the dimension $d$ is an important open problem. This brings us to our main motivating question:

For what class of functions can we avoid polynomial dependence on dimension (in an algorithm)? Can we do this for the logistic loss function?

| Year | Refs. | Method | Iters. | Gra/Iter | Total time |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 2006 | $[$ LV06b $]$ | B./H. $^{*}$ | $d^{3}, d^{4}$ | 1 | $d^{5}, d^{6}$ |
| 2017 | $[$ Dal17 $]$ | LMC $^{*}$ | $\kappa^{2} d, \kappa^{3} d^{3}$ | 1 | $\kappa^{2} d^{2}, \kappa^{3} d^{4}$ |
| 2017 | $[$ Dal17] | LMCO $^{*}$ | $\kappa^{2} d, \kappa^{2} d^{2.5}$ | 1 | $\kappa^{2} d^{4}, \kappa^{2} d^{5.5}$ |
| 2018 | $[$ CCBJ18] | DL | $\kappa^{2} d^{0.5}$ | 1 | $\kappa^{2} d^{1.5}$ |
| 2018 | $[$ DCWY18] | MALA $^{*}$ | $\kappa d, \kappa d^{2}$ | 1 | $\kappa d^{2}, \kappa d^{3}$ |
| 2017 | $[$ MS17] | HMC $^{*, \dagger}$ | $\kappa^{6.5} d^{0.5}$ | 1 | $\kappa^{6.5} d^{1.5}$ |
| 2018 | $[$ MV18] | HMC $^{*, \dagger}$ | $\kappa^{2.75} d^{0.25}, \kappa^{3.5} d^{0.25}$ | 1 | $\kappa^{2.75} d^{1.25}, \kappa^{3.5} d^{1.25}$ |
| 2018 | [LSV18] <br> (thesis) | HMC <br>  <br> HMC | $\kappa^{1.5}$ <br> $\kappa^{1.5}$ | $\kappa^{0.25} d^{0.5}$ | 1.5 <br> $\kappa^{1.75} d^{1.5}$ |

* have different bounds for warm start and general (cold) start. We stated the runtime for cold start in green color.
${ }^{\dagger}$ make smoothness and incoherence assumptions motivated by and applicable to Bayesian logistic regression.

Table 4.1: Summary of results, $d$ is the dimension, $\kappa$ is the condition number of $\nabla^{2} f$. "Iters" denotes "the number of iterations / parallel depth". "Gra/Iter" denotes the "the number of gradients per iteration". We use the parallel depth of the algorithm as the number of iterations. We suppress polylogarithmic terms and dependence on the error parameter. Ball walk/hit-and-run apply to general logconcave distributions, the rest assume strongly logconcave with Lipschitz gradient and possibly more. "B./H." denotes "Ball Walk/Hit-andrun". "DL" denotes "Damped Langevin". In all previous work, for simplicity, we report the most favorable bounds by making various assumptions such as $\kappa \ll d$.

### 4.1.1 Results

We begin with an informal statement of our result for sampling from a class that includes the logistic loss function.

Theorem 4.1.2 (Informal version of Theorem 4.5.6). Let $A=\left[a_{1} ; a_{2} ; \cdots ; a_{n}\right] \in \mathbb{R}^{n \times d}$, $\phi_{i}: \mathbb{R} \rightarrow \mathbb{R}$ with its $k$-th derivatives bounded by $O(1)^{k}$ and

$$
f(x)=\sum_{i=1}^{n} \phi_{i}\left(a_{i}^{\top} x\right)+\frac{m_{2}}{2}\|x\|^{2} .
$$

Suppose that $\nabla^{2} f$ has condition number $\kappa$ and $\tau=\left\|A A^{\top}\right\|_{\infty \rightarrow \infty}$. Then we can find a random point $X$ whose Wasserstein distance to $Y$ drawn from the density proportional to $e^{-f}$ satisfies

$$
W_{2}(X, Y) \leq \frac{\epsilon}{\sqrt{m_{2}}}
$$

using $\widetilde{O}\left(\kappa^{1.5}+\frac{\tau}{m_{2}}\right)$ iterations, where each iteration takes $\widetilde{O}(d)$ time and $\widetilde{O}(1)$ evaluations of $\nabla f$.

Remark 4.1.1. The $\frac{1}{\sqrt{m_{2}}}$ term in the error is needed to make the statement invariant under scaling of $f$.

For the logistic $\operatorname{loss}^{1} \phi(t)=\log \left(1+e^{-t}\right)$, we have $\phi^{\prime}(t)=-\frac{1}{1+e^{t}}$, and it has Cauchy estimate $M=1$ with radius $r=1$ (See Lemma 4.9.2). The above result has the following application,

Corollary 4.1.3 (Logistic loss sampling). Let $f(x)=\sum_{i=1}^{n} \phi_{i}\left(a_{i}^{\top} x\right)+\frac{m_{2}}{2}\|x\|^{2}$ with $\phi(t)=$ $\log \left(1+e^{-t}\right)$. Let $\tau=\left\|A A^{\top}\right\|_{\infty \rightarrow \infty}$ and suppose that $\nabla^{2} f(x) \preceq M_{2} \cdot I$ for all $x$. Starting at the minimum $x^{(0)}$ of $f$, we can find a random point $X$ whose Wasserstein distance to $Y$ drawn from the density proportional to $e^{-f}$ satisfies

$$
W_{2}(X, Y) \leq \frac{\epsilon}{\sqrt{m_{2}}}
$$

using $\widetilde{O}\left(\frac{\tau}{m_{2}}+\kappa^{1.5}\right)$ iterations with $\kappa=\frac{M_{2}}{m_{2}}$. Each iteration takes $\widetilde{O}(d)$ time and $\widetilde{O}(1)$ matrixvector multiplications for $A$ and $A^{\top}$.

Remark 4.1.2. Lemma 4.6 .1 shows that $\left\|A A^{\top}\right\|_{\infty \rightarrow \infty}=\Theta\left(\lambda_{\max }\left(A A^{\top}\right)\right)$ for sparse enough matrix $A A^{\top}$. Since $\lambda_{\max }\left(A A^{\top}\right)$ usually has the same order as $M_{2}$, the number of iterations is dominated by the $\kappa^{1.5}$ term.

[^12]The above results extend and improve previous work substantially. First, in all previous algorithms, the number of functions calls was polynomial in the dimension $d$, while the dependence here is polylogarithmic. Second, our incoherence assumption for logistic regression is simpler and milder. Third, due to the nature of how we implement each step, the parallel depth of the algorithm is just the number of iterations, i.e., polylogarithmic in the dimension and $\widetilde{O}(1)$ when the condition numbers are bounded. Fourth, the runtime and depth of our algorithm depends polynomially in $\log (1 / \epsilon)$ while all previous (nearly) linear time algorithms depends polynomially in $1 / \epsilon$.

We also give an improved bound on the complexity of sampling from $e^{-f}$ when $f$ is strongly convex and has a Lipschitz gradient (no further smoothness assumptions).

Theorem 4.1.4. (Strongly Convex). Given a function $f$ such that $0 \prec m_{2} \cdot I \preceq \nabla^{2} f(x) \preceq M_{2} \cdot I$ for all $x \in \mathbb{R}^{d}$ and $0<\epsilon<\sqrt{d}$. Starting $x^{(0)}$ at the minimum of $f$, we can find a random point $X$ whose Wasserstein distance to $Y$ drawn from the density proportional to $e^{-f}$ satisfies

$$
W_{2}(X, Y) \leq \frac{\epsilon}{\sqrt{m_{2}}}
$$

using $O\left(\kappa^{1.5} \log \left(\frac{d}{\epsilon}\right)\right)$ iterations where $\kappa=\frac{M_{2}}{m_{2}}$. Each iteration takes $O\left(\frac{\kappa^{\frac{1}{4}} d^{\frac{3}{2}}}{\epsilon} \log \left(\frac{\kappa d}{\epsilon}\right)\right)$ time and $O\left(\frac{\kappa^{\frac{1}{4}} \frac{1}{\epsilon}}{\epsilon} \log \left(\frac{\kappa d}{\epsilon}\right)\right)$ evaluations of $\nabla f$, amortized over all iterations.

The previous best bound was $\kappa^{2} \sqrt{d}$ iterations [CCBJ18]. This result is one of the key surprises of this chapter. Although this problem has been studied extensively with specifically-designed algorithms and analysis, we show how to get a better result by a general ODE algorithm and a general analysis which works for any ODE. Furthermore, our algorithm is the first to achieve polylogarithmic depth dependence on the dimension, which seemed impossible in prior work.

The above results are based on three ingredients: (1) a new contraction rate for HMC of $\kappa^{1.5}$, improving on the previous best bound of $\kappa^{2}(2)$ a proof that a solution to ODE's arising from HMC applied to the above problem are approximated by (piecewise) low-degree polynomials and (3) a fast (nearly linear time and polylog parallel depth) algorithm for solving multivariate second-order ODEs.

We next present the multivariate high-order ODE guarantee. This generalizes and improves on the guarantee from [LV17]. While we state it below for the case of the piecewise polynomial basis of functions, it applies to any basis of functions. This is a general result about solving ODE efficiently, independent of the application to sampling. The only assumptions needed are that the ODE function is Lipschitz and that the solution is close to the span of small number of basis of functions. These natural assumptions suffice to get around the worst-case complexity lower bounds for solving such general ODEs [KF82, Ko83, Ko10, Kaw10, KC12].

Theorem 4.1.5 (Informal version of Theorem 4.2.4 for 1st order ODE). Let $x^{*}(t) \in \mathbb{R}^{d}$ be the solution of the $O D E$

$$
\frac{\mathrm{d}}{\mathrm{~d} t} x(t)=F(x(t), t), x(0)=v
$$

where $F: \mathbb{R}^{d+1} \rightarrow \mathbb{R}^{d}, x(t) \in \mathbb{R}^{d}$ and $v \in \mathbb{R}^{d}$. Given some $L$ and $\epsilon>0$ such that

1. There exists a piece-wise polynomial $q(t)$ such that $q(t)$ on $\left[T_{j-1}, T_{j}\right]$ is a degree $D_{j}$ polynomial with

$$
0=T_{0}<T_{1}<\cdots<T_{n}=T
$$

and that

$$
\left\|q(t)-\frac{\mathrm{d}}{\mathrm{~d} t} x^{*}(t)\right\| \leq \frac{\epsilon}{T}, \forall t \in[0, T]
$$

2. The algorithm knows about the intervals $\left[T_{j-1}, T_{j}\right]$ and the degree $D_{j}$ for all $j \in[n]$.
3. For any $y, z \in \mathbb{R}^{d}$,

$$
\|F(y, t)-F(z, t)\| \leq L\|y-z\|, \forall t \in[0, T] .
$$

Assume $L T \leq 1 / 16000$. Then, we can find a piece-wise polynomial $x(t)$ such that

$$
\max _{t \in[0, T]}\left\|x(t)-x^{*}(t)\right\| \lesssim \epsilon
$$

using $\widetilde{O}\left(\sum_{i=1}^{n}\left(1+D_{i}\right)\right)$ evaluations of $F$ and $\widetilde{O}\left(d \sum_{i=1}^{n}\left(1+D_{i}\right)\right)$ time.

We suspect these methods will be useful in many other settings beyond the focus application of this chapter. Moreover, the result is nearly optimal. Roughly speaking, it says that the complexity of solving the ODE is nearly the same as the complexity of representing the solution. The assumption that $L T<1$ is essential, as otherwise after longer time, the solution can blow up exponentially. Also, the assumption on the piece-wise polynomial approximation has a certain universality since this is how one implicitly represents a function using any iterative method. Finally, each iteration of the ODE algorithm, the collocation method, can be fully parallelized; as a result the parallel time complexity of the sampling algorithms in this chapter are polylogarithmic in the dimension.

### 4.1.2 HMC and improved contraction rate

We give an improved contraction rate for HMC, stated explicitly as Algorithm 4.1. We

Algorithm 4.1 Hamiltonian Monte Carlo Algorithm
procedure $\operatorname{HMC}\left(x^{(0)}, f, \epsilon, h\right)$
$\triangleright$ Theorem 4.3.2
2: $\quad$ Suppose that $f$ is $m_{2}$ strongly convex with $M_{2}$ Lipschitz gradient on $\mathbb{R}^{d}$.
3: $\quad$ Assume that the step size $h \leq \frac{m_{2}^{1 / 4}}{2 M_{2}^{3 / 4}}$.
4: Let the number of iterations $N=\frac{1}{\theta} \cdot \log \left(\frac{4}{\epsilon^{2}}\left(\frac{\left\|\nabla f\left(x^{(0)}\right)\right\|_{2}^{2}}{m_{2}}+d\right)\right)$ with $\theta=\frac{m_{2} h^{2}}{8}$.
For $k=1,2, \cdots, N$ do
Generate a Gaussian random direction $v \sim \mathcal{N}\left(0, I_{d}\right)$.
Let $x(t)$ be the HMC defined by

$$
\frac{\mathrm{d}^{2} x}{\mathrm{~d} t^{2}}=-\nabla f(x), \frac{\mathrm{d} x}{\mathrm{~d} t}(0)=v, x(0)=x^{(k-1)}
$$

8: $\quad$ Find a point $x^{(k)}$ such that
$\triangleright$ Theorem 4.2.4

$$
\left\|x^{(k)}-x(h)\right\|_{2} \leq \bar{\epsilon}:=\frac{\theta \cdot \epsilon}{2 \sqrt{m_{2}}}
$$

end for
return $x^{(N)}$.
end procedure
give two contraction bounds for the ideal HMC. The first bound is $\left(\frac{m_{2}}{M_{2}}\right)^{1.5}$ using $T \sim \frac{m_{2}^{1 / 4}}{M_{2}^{3 / 4}}$. The second bound shows that there is a $T$ that gives the optimal contraction bound $\frac{m_{2}}{M_{2}}$. However, as we will see this part cannot be used to bound the overall mixing time, because the time $T$ depends on the point we use for coupling, which is unknown to the algorithm. We keep this as evidence for a possible $\frac{m_{2}}{M_{2}}$ bound. The improvement is from $\kappa^{2}$ in previous work to $\kappa^{1.5}$.

Lemma 4.1.6 (Contraction bound for HMC). Let $x(t)$ and $y(t)$ be the solution of HMC dynamics on $e^{-f}$ starts at $x(0)$ and $y(0)$ with initial direction $x^{\prime}(0)=y^{\prime}(0)=v$ for some vector $v$. Suppose that $f$ is $m_{2}$ strongly convex with $M_{2}$ Lipschitz gradient., i.e., $m_{2} \cdot I \preceq$ $\nabla^{2} f(x) \preceq M_{2} \cdot I$ for all $x$. Then, for $0 \leq t \leq \frac{m_{2}^{1 / 4}}{2 M_{2}^{3 / 4}}$, we have that

$$
\|x(t)-y(t)\|_{2}^{2} \leq\left(1-\frac{m_{2}}{4} t^{2}\right)\|x(0)-y(0)\|_{2}^{2}
$$

Furthermore, there is $t \geq 0$ depending on $f, x(0), y(0)$, and $v$ such that

$$
\|x(t)-y(t)\|_{2}^{2} \leq\left(1-\frac{1}{16} \frac{m_{2}}{M_{2}}\right)\|x(0)-y(0)\|_{2}^{2}
$$

### 4.1.3 Techniques

In this chapter we give bounds on the Collocation Method for solving ODEs. To ensure the algorithm is applicable to the ODE's that arise in the sampling application, we need to show that the solution of the ODE is close to a low-rank basis. Given only bounds on the Hessian of a function, we do this by approximating the solution of the ODE with a piecewise degree two polynomial. For smooth functions, we can use low-degree polynomials.

The proofs of the degree bounds go via the Cauchy-Kowalevsky method, by showing bounds on all derivatives at the initial point. To do this for multivariate ODE's, we use
the method of majorants, and reduce the problem to bounding the radius of convergence of one-variable ODEs.

The improved convergence guarantees for exact HMC are also based on better analysis of the underlying ODE, showing that a larger step size than previously known is possible.

For many optimization and sampling methods, there are corresponding customized versions that deal with decomposable functions by sampling terms of the functions. These algorithms usually take nearly linear time with the number of iterations being polynomial in the dimension. Often, an improvement in the general case would lead to an improvement in the decomposable case. To limit the length of this chapter, we focus only on results with polylogarithmic depth. Therefore, in Table 4.1, we do not list algorithms that involve sampling each term in decomposable functions [BFR16, $\mathrm{DSM}^{+} 16, \mathrm{DRW}^{+} 16, \mathrm{BFFN} 17$, DK17, $\left.\mathrm{CWZ}^{+} 17, \mathrm{NDH}^{+} 17, \mathrm{CFM}^{+} 18\right]$. We expect our techniques can be further improved for decomposable functions by sampling individual terms.

Outline of chapter. Then we give the main ODE algorithm and guarantees in Section 4.2. We give the proof of the improved convergence bound for HMC in Section 4.3. We use both parts to obtain improved guarantees for sampling strongly logconcave functions with Lipschitz gradients in Section 4.4 and smooth functions, including logistic loss in Section 4.5.

Some preliminaries including standard definitions and well-known theorems about ODEs are in an appendix. Remaining proofs about ODEs are in Appendix 4.7. In Appendix 4.8 , we present some useful tools for Cauchy estimates. Appendix 4.9 shows how to calculate the Cauchy estimates of some function which are extensively used in practice.

### 4.2 ODE Solver for any basis

In this section, we analyze the collocation method for solving ODEs [Ise09]. This method is classical in numerical analysis. The goal of this section is provide an introduction of this method and provide a non-asymptotic bounds for this method. In [LV17, LV18], we applied this method to obtain faster algorithms for sampling on polytopes. Unfortunately, the particular version of collocation method we used assume the solution can be approximated by a low-degree polynomial, which heavily restrict the set of functions we can sample.

To give an intuition for the collocation method, we first consider the following firstorder ODE

$$
\begin{aligned}
\frac{\mathrm{d}}{\mathrm{~d} t} x(t) & =F(x(t), t), \quad \forall 0 \leq t \leq T \\
x(0) & =v
\end{aligned}
$$

where $F: \mathbb{R}^{d+1} \rightarrow \mathbb{R}$. The collocation method is partly inspired by the Picard-Lindelöf theorem, a constructive existence proof for a large class of ODE. Therefore, we will first revisit the proof of Picard-Lindelöf theorem for first-order ODE.

### 4.2.1 Picard-Lindelöf theorem

In general, we can rewritten the first-order ODE as an integral equation

$$
x(t)=v+\int_{0}^{t} F(x(s), s) \mathrm{d} s \quad \text { for all } 0 \leq t \leq T
$$

To simplify the notation, we use $\mathcal{C}\left([0, T], \mathbb{R}^{d}\right)$ to denote $\mathbb{R}^{d}$-valued functions on $[0, T]$. We define the operator $\mathcal{T}$ from $\mathcal{C}\left([0, T], \mathbb{R}^{d}\right)$ to $\mathcal{C}\left([0, T], \mathbb{R}^{d}\right)$ by

$$
\begin{equation*}
\mathcal{T}(x)(t)=v+\int_{0}^{t} F(x(s), s) \mathrm{d} s \quad \text { for all } 0 \leq t \leq T \tag{4.1}
\end{equation*}
$$

Therefore, the integral equation is simply $x=\mathcal{T}(x)$.
Banach fixed point theorem shows that the integral equation $x=\mathcal{T}(x)$ has a unique solution if there is a norm, and $j \in \mathbb{N}$ such that the map $\mathcal{T}^{\circ j}$ has Lipschitz constant less than 1. Recall that $\mathcal{T}^{\circ j}$ is the composition of $j$ many $\mathcal{T}$, i.e.,

$$
\mathcal{T}^{\circ j}(x)=\underbrace{\mathcal{T}(\mathcal{T}(\cdots \mathcal{T}}_{j \text { many } \mathcal{T}}(x) \cdots)) .
$$

Picard-Lindelöf theorem shows that if $F$ is Lipschitz in $x$, then the map $\mathcal{T}^{\circ j}$ has Lipschitz constant less than 1 for some positive integer $j$.

Lemma 4.2.1. Given any norm $\|\cdot\|$ on $\mathbb{R}^{d}$. Let $L$ be the Lipschitz constant of $F$ in $x$, i.e.

$$
\|F(x, s)-F(y, s)\| \leq L\|x-y\| \quad \text { for all } x, y \in \mathbb{R}^{d}, s \in[0, T]
$$

For any $x \in \mathcal{C}\left([0, T], \mathbb{R}^{d}\right)$, we define the corresponding norm

$$
\|x\| \stackrel{\text { def }}{=} \max _{0 \leq t \leq T}\|x(t)\| .
$$

Then, the Lipschitz constant of $\mathcal{T}^{\circ j}$ in this norm is upper bounded by $(L T)^{j} / j$ !.

Proof. We prove this lemma with a stronger induction statement

$$
\left\|\left(\mathcal{T}^{\circ j} x\right)(h)-\left(\mathcal{T}^{\circ j} y\right)(h)\right\| \leq \frac{L^{j} h^{j}}{j!}\|x-y\| \quad \text { for all } 0 \leq h \leq T
$$

The base case $j=0$ is trivial. For the induction case $j$, we can upper bound the term
as follows

$$
\begin{array}{rlr} 
& \left\|\mathcal{T}^{\circ j} x(h)-\mathcal{T}^{\circ j} y(h)\right\| \\
= & \left\|\mathcal{T} \mathcal{T}^{\circ(j-1)} x(h)-\mathcal{T} \mathcal{T}^{\circ(j-1)} y(h)\right\| & \\
\leq & \int_{0}^{h}\left\|F\left(\mathcal{T}^{\circ(j-1)} x(s), s\right)-F\left(\mathcal{T}^{\circ(j-1)} y(s), s\right)\right\| \mathrm{d} s & \\
\leq & L \int_{0}^{h}\left\|\mathcal{T}^{\circ(j-1)} x(s)-\mathcal{T}^{\circ(j-1)} y(s)\right\| \mathrm{d} s & \text { by } f \text { is } L \text { Lipschitz } \\
\leq & L \int_{0}^{h} \frac{L^{j-1} s^{j-1}}{(j-1)!}\|x-y\| \mathrm{d} s & \text { by the induction statement } \\
= & \frac{L^{j}}{j!} h^{j}\|x-y\| . &
\end{array}
$$

This completes the induction.

### 4.2.2 Intuition of collocation method

To make the Picard-Lindelöf theorem algorithmic, we need to discuss how to represent a function in $\mathcal{C}\left([0, T], \mathbb{R}^{d}\right)$. One standard way is to use a polynomial $p_{i}(t)$ in $t$ for each coordinate $i \in[d]$. In this section, we assume that there is a basis $\left\{\varphi_{j}\right\}_{j=1}^{D} \subset \mathcal{C}([0, T], \mathbb{R})$ such that for all $i \in[d], \frac{\mathrm{d} x_{i}}{\mathrm{~d} t}$ is approximated by some linear combination of $\varphi_{j}(t)$.

For example, if $\varphi_{j}(t)=t^{j-1}$ for $j \in[d]$, then our assumption is simply saying $\frac{\mathrm{d} x_{i}}{\mathrm{~d} t}$ is approximated by a degree $D-1$ polynomial. Other possible basis are piecewise-polynomial and Fourier series. By Gram-Schmidt orthogonalization, we can always pick nodes point $\left\{c_{i}\right\}_{j=1}^{D}$ such that

$$
\varphi_{j}\left(c_{i}\right)=\delta_{i, j} \quad \text { for } i, j \in[D] .
$$

The benefit of such basis is that for any $f \in \operatorname{span}\left(\varphi_{j}\right)$, we have that $f(t)=\sum_{j=1}^{D} f\left(c_{j}\right) \varphi_{j}(t)$.

For polynomials, the resulting basis are the Lagrange polynomials

$$
\varphi_{j}(t)=\prod_{i \in[D] \backslash\{j\}} \frac{t-c_{i}}{c_{j}-c_{i}} \quad \text { for } j \in[D] .
$$

The only assumption we make on the basis is that its integral is bounded.

Definition 4.2.1. Given a $D$ dimensional subspace $\mathcal{V} \subset \mathcal{C}([0, T], \mathbb{R})$ and node points $\left\{c_{j}\right\}_{j=1}^{D} \subset[0, T]$. For any $\gamma_{\varphi} \geq 1$, we call a basis $\left\{\varphi_{j}\right\}_{j=1}^{D} \subset \mathcal{V}$ is $\gamma_{\varphi}$ bounded if $\varphi_{j}\left(c_{i}\right)=\delta_{i, j}$ and we have

$$
\sum_{j=1}^{D}\left|\int_{0}^{t} \varphi_{j}(s) \mathrm{d} s\right| \leq \gamma_{\varphi} T \quad \text { for } t \in[0, T]
$$

Note that if the constant function $1 \in \mathcal{V}$, then we have

$$
1=\sum_{j=1}^{D} 1\left(c_{j}\right) \varphi_{j}(t)=\sum_{j=1}^{D} \varphi_{j}(t)
$$

Hence, we have

$$
T=\int_{0}^{T} 1 \mathrm{~d} s \leq \sum_{j=1}^{D}\left|\int_{0}^{T} \varphi_{j}(s) \mathrm{d} s\right| \leq \gamma_{\varphi} T
$$

Therefore, $\gamma_{\varphi} \geq 1$ for most of the interesting basis. This is the reason why we simply put it as an assumption to shorten some formulas.

In the section 4.2.5, we prove that for the space of low degree polynomial and piecewise low degree polynomial, there is a basis on the Chebyshev nodes that is $O(1)$ bounded.

Assuming that $\frac{d x}{d t}$ can be approximated by some element in $\mathcal{V}$, we have that

$$
\frac{\mathrm{d} x}{\mathrm{~d} t}(t) \sim \sum_{j=1}^{D} \frac{\mathrm{~d} x}{\mathrm{~d} t}\left(c_{j}\right) \varphi_{j}(t)=\sum_{j=1}^{D} F\left(x\left(c_{j}\right), c_{j}\right) \varphi_{j}(t)
$$

Integrating both sides, we have

$$
\begin{equation*}
x(t) \approx v+\int_{0}^{t} \sum_{j=1}^{D} F\left(x\left(c_{j}\right), c_{j}\right) \varphi_{j}(s) \mathrm{d} s \tag{4.2}
\end{equation*}
$$

This inspires us to consider the following operator from $\mathcal{C}\left([0, T], \mathbb{R}^{d}\right)$ to $\mathcal{C}\left([0, T], \mathbb{R}^{d}\right)$ :

$$
\begin{equation*}
\mathcal{T}_{\varphi}(x)(t)_{i}=v_{i}+\int_{0}^{t} \sum_{j=1}^{D} F\left(x\left(c_{j}\right), c_{j}\right)_{i} \varphi_{j}(s) \mathrm{d} s \quad \text { for } i \in[d] . \tag{4.3}
\end{equation*}
$$

Equation (4.2) can be written as $x \approx \mathcal{T}_{\varphi}(x)$. To find $x$ to satisfies this, naturally, one can apply the fix point iteration and this is called the collocation method.

### 4.2.3 Collocation method

From the definition of (4.3), we note that $\mathcal{T}_{\varphi}(x)$ depends only on $x(t)$ at $t=c_{j}$. Therefore, we can only need to calculate $\mathcal{T}_{\varphi}(x)(t)$ at $t=c_{j}$. To simplify the notation, for any $x \in \mathcal{C}\left([0, T], \mathbb{R}^{d}\right)$, we define a corresponding matrix $[x] \in \mathbb{R}^{d \times D}$ by $[x]_{i, j}=x_{i}\left(c_{j}\right)$. For any $d \times D$ matrix $X$, we define $F(X, c)$ as an $d \times D$ matrix

$$
\begin{equation*}
F(X, c)_{i, j}=F\left(X_{*, j}, c_{j}\right)_{i} . \tag{4.4}
\end{equation*}
$$

where $X_{*, j}$ is the $j$-th column of $X$. Finally, we define $A_{\varphi}$ as a $D \times D$ matrix

$$
\begin{equation*}
\left(A_{\varphi}\right)_{i, j}=\int_{0}^{c_{j}} \varphi_{i}(s) \mathrm{d} s \tag{4.5}
\end{equation*}
$$

By inspecting the definition of (4.3), (4.4) and (4.5), we have that

$$
\left[\mathcal{T}_{\varphi}(x)\right]=v \cdot 1_{D}^{\top}+F([x], c) A_{\varphi}
$$

where $1_{D}$ is a column of all 1 vector of length $D$. Hence, we can apply the map $\mathcal{T}_{\varphi}$ by simply multiply $F([x], c)$ by a pre-compute $D \times D$ matrix $A_{\varphi}$. For the basis we considered

```
Algorithm 4.2 Collocation Method
    procedure \(\operatorname{CollocationMEthod}(F, v, T, \varphi, c) \quad \triangleright\) Theorem 4.2.3
    Let \(N=\left\lceil\log \left(\frac{T}{\epsilon} \max _{s \in[0, T]}\|F(v, s)\|\right)\right\rceil \quad \triangleright\) Choose number of iterations
    Let \(A_{\varphi}\) be the matrix defined by \(\left(A_{\varphi}\right)_{i, j}=\int_{0}^{c_{j}} \varphi_{i}(s) \mathrm{d} s\).
    \(X^{(0)} \leftarrow v \cdot 1_{D}^{\top}\). \(\quad 1_{D}\) is a column of all 1 vector of length \(D\)
    for \(j=1,2, \cdots, N-1\) do
        \(X^{(j)} \leftarrow v \cdot 1_{D}^{\top}+F\left(X^{(j-1)}, c\right) A_{\varphi} . \quad \triangleright\) Matrix \(F(X, c)\) is defined in Eq. (4.4)
            \(\triangleright\) Note that we evaluate \(D\) many \(F\) every iteration in this matrix notation.
        end for
        \(x^{(N)}(t) \leftarrow v+\int_{0}^{t} \sum_{i=1}^{D} F\left(X_{*, i}^{(N)}, c_{i}\right) \varphi_{i}(s) \mathrm{d} s\)
        return \(x^{(N)}\)
    end procedure
```

in this chapter, each iteration takes only $\widetilde{O}(d D)$ which is nearly linear to the size of our representation of the solution.

We state our guarantee for a first-order ODE (Algorithm 4.2).

Theorem 4.2.2 (First order ODE). Let $x^{*}(t)$ be the solution of an dimensional ODE

$$
x(0)=v, \frac{\mathrm{~d} x(t)}{\mathrm{d} t}=F(x(t), t) \quad \text { for all } 0 \leq t \leq T
$$

We are given a $D$ dimensional subspace $\mathcal{V} \subset \mathcal{C}([0, T], \mathbb{R})$, node points $\left\{c_{j}\right\}_{j=1}^{D} \subset[0, T]$ and a $\gamma_{\varphi}$ bounded basis $\left\{\varphi_{j}\right\}_{j=1}^{D} \subset \mathcal{V}$ (Definition 4.2.1). Given some $L$ and $\epsilon>0$ such that

1. There exists a function $q \in \mathcal{V}$ such that

$$
\left\|q(t)-\frac{\mathrm{d}}{\mathrm{~d} t} x^{*}(t)\right\| \leq \frac{\epsilon}{T}, \forall t \in[0, T] .
$$

2. For any $y, z \in \mathbb{R}^{d}$,

$$
\|F(y, t)-F(z, t)\| \leq L\|y-z\|, \forall t \in[0, T] .
$$

Assume $\gamma_{\varphi} L T \leq 1 / 2$. Then the algorithm CollocationMethod (Algorithm 4.2) outputs a function $x^{(N)} \in \mathcal{V}$ such that

$$
\max _{t \in[0, T]}\left\|x^{(N)}(t)-x^{*}(t)\right\| \leq 20 \gamma_{\varphi} \epsilon .
$$

The algorithm takes $O\left(D \log \left(\frac{T}{\epsilon} \max _{s \in[0, T]}\|F(v, s)\|\right)\right)$ evaluations of $F$.

Next we state the general result for a $k$-th order ODE. We prove this via a reduction from higher order ODE to first-order ODE. See the proof in Appendix 4.7.

Theorem 4.2.3 ( $k$-th order ODE). Let $x^{*}(t) \in \mathbb{R}^{d}$ be the solution of the ODE

$$
\begin{aligned}
& \frac{\mathrm{d}^{k}}{\mathrm{~d} t^{k}} x(t)=F\left(\frac{\mathrm{~d}^{k-1}}{\mathrm{~d} t^{k-1}} x(t), \cdots, x(t), t\right) \\
& \frac{\mathrm{d}^{i}}{\mathrm{~d} t^{i}} x(0)=v_{i}, \forall i \in\{k-1, \cdots, 1,0\}
\end{aligned}
$$

where $F: \mathbb{R}^{k d+1} \rightarrow \mathbb{R}^{d}, x(t) \in \mathbb{R}^{d}$, and $v_{0}, v_{1}, \cdots, v_{k-1} \in \mathbb{R}^{d}$.
We are given a $D$ dimensional subspace $\mathcal{V} \subset \mathcal{C}([0, T], \mathbb{R})$, node points $\left\{c_{j}\right\}_{j=1}^{D} \subset[0, T]$ and a $\gamma_{\varphi}$ bounded basis $\left\{\varphi_{j}\right\}_{j=1}^{D} \subset \mathcal{V}$ (Definition 4.2.1). Given some $L$ and $\epsilon>0$ such that

1. For $i \in[k]$, there exists a function $q^{(i)} \in \mathcal{V}$ such that

$$
\left\|q^{(i)}(t)-\frac{\mathrm{d}^{i}}{\mathrm{~d} t^{i}} x^{*}(t)\right\| \leq \frac{\epsilon}{T^{i}}, \forall t \in[0, T]
$$

2. For any $y, z \in \mathbb{R}^{k d}$,

$$
\|F(y, t)-F(z, t)\| \leq \sum_{i=1}^{k} L_{i}\left\|y_{i}(t)-z_{i}(t)\right\|, \forall t \in[0, T]
$$

Assume $\gamma_{\varphi} L T \leq 1 / 8$ with $L=\sum_{i=1}^{k} L_{i}^{1 / i}$. Then, we can find functions $\left\{q^{(i)}\right\}_{i \in\{0,1, \cdots, k-1\}} \subset$ $\mathcal{V}$ such that

$$
\max _{t \in[0, T]}\left\|q^{(i)}(t)-\frac{\mathrm{d}^{i}}{\mathrm{~d} t^{i}} x^{*}(t)\right\|_{p}=20(1+2 k) \gamma_{\varphi} \frac{\epsilon}{T^{i}}, \forall i \in\{0,1, \cdots, k-1\}
$$

The algorithm takes $O(D \log (C / \epsilon))$ evaluations of $F$ where

$$
C=\left(4 \gamma_{\varphi} T\right)^{k} \cdot \max _{s \in[0, T]}\left\|F\left(v_{k-1}, v_{k-2}, \cdots, v_{0}, s\right)\right\|+\sum_{i=1}^{k-1}\left(4 \gamma_{\varphi} T\right)^{i}\left\|v_{i}\right\|
$$

Note that the statement is a bit awkward. Instead of finding a function whose derivatives are same as the derivatives of $x^{*}$, the algorithm approximates the derivatives of $x^{*}$ individually. This is because we do not know if derivatives/integrals of functions in $\mathcal{V}$ remain in $\mathcal{V}$. For piece-wise polynomials, we can approximate the $j$-th derivative of the solution by taking $(k-j)$-th iterated integral of $q^{(k)}$, which is still a piece-wise polynomial.

In section 4.2.5, we give a basis for piece-wise polynomials (Lemma 4.2.8). Using this basis, we have the following Theorem.

Theorem 4.2.4. ( $k$-th order ODE) Let $x^{*}(t) \in \mathbb{R}^{d}$ be the solution of the ODE

$$
\begin{aligned}
\frac{\mathrm{d}^{k}}{\mathrm{~d} t^{k}} x(t) & =F\left(\frac{\mathrm{~d}^{k-1}}{\mathrm{~d} t^{k-1}} x(t), \cdots, x(t), t\right) \\
\frac{\mathrm{d}^{i}}{\mathrm{~d} t^{i}} x(0) & =v_{i}, \forall i \in\{k-1, \cdots, 1,0\}
\end{aligned}
$$

where $F: \mathbb{R}^{k d+1} \rightarrow \mathbb{R}^{d}, x(t) \in \mathbb{R}^{d}$, and $v_{0}, v_{1}, \cdots, v_{k-1} \in \mathbb{R}^{d}$. Given some $L$ and $\epsilon>0$ such that

1. There exists a piece-wise polynomial $q(t)$ such that $q(t)$ on $\left[T_{j-1}, T_{j}\right]$ is a degree $D_{j}$ polynomial with

$$
0=T_{0}<T_{1}<\cdots<T_{n}=T
$$

and that

$$
\left\|q(t)-\frac{\mathrm{d}^{k}}{\mathrm{~d} t^{k}} x^{*}(t)\right\| \leq \frac{\epsilon}{T^{k}}, \forall t \in[0, T]
$$

2. The algorithm knows about the intervals $\left[T_{j-1}, T_{j}\right]$ and the degree $D_{j}$ for all $j \in[n]$.
3. For any $y, z \in \mathbb{R}^{k d}$,

$$
\|F(y, t)-F(z, t)\| \leq \sum_{i=1}^{k} L_{i}\left\|y_{i}-z_{i}\right\|, \forall t \in[0, T]
$$

Assume $L T \leq 1 / 16000$ with $L=\sum_{i=1}^{k} L_{i}^{1 / i}$. Then, we can find a piece-wise polynomial $x(t)$ such that

$$
\max _{t \in[0, T]}\left\|\frac{\mathrm{d}^{i}}{\mathrm{~d} t^{i}} x(t)-\frac{\mathrm{d}^{i}}{\mathrm{~d} t^{i}} x^{*}(t)\right\|_{p} \lesssim \frac{\epsilon k}{T^{i}}, \forall i \in\{0,1, \cdots, k-1\} .
$$

using $O(D \log (C / \epsilon))$ evaluations of $F$ with the size of basis $D=\sum_{i=1}^{n}\left(1+D_{i}\right)$ and

$$
O\left(d \min \left(\sum_{i=1}^{n}\left(1+D_{i}\right)^{2}, D \log (C D / \epsilon)\right) \log (C / \epsilon)\right)
$$

time where

$$
C=O(T)^{k} \cdot \max _{s \in[0, T]}\left\|F\left(v_{k-1}, v_{k-2}, \cdots, v_{0}, s\right)\right\|+\sum_{i=1}^{k-1} O(T)^{i}\left\|v_{i}\right\|
$$

Remark 4.2.1. The two different runtime come from two different ways to the integrate of basis in Lemma 4.2.8. The first one is an navie method which is good enough for all our application. The second one follows from multipole method which gives an nearly linear time to the size of the basis with an extra $\log$ dependence on the accuracy.

In the rest of this section, we prove the first-order guarantee, Theorem 4.2.2

### 4.2.4 Proof of first order ODE

First, we bound the Lipschitz constant of the map $\mathcal{T}_{\varphi}$. Unlike the Picard-Lindelöf theorem, we are not able to get an improved bound of the Lipschitz constant of the composite of $\mathcal{T}_{\varphi}$. Fortunately, the Lipschitz constant of the map $\mathcal{T}_{\varphi}$ is good enough for all applications in this chapter.

Lemma 4.2.5. Given any norm $\|\cdot\|$ on $\mathbb{R}^{d}$. Let $L$ be the Lipschitz constant of $F$ in x, i.e.

$$
\|F(x, s)-F(y, s)\| \leq L\|x-y\| \quad \text { for all } x, y \in \mathbb{R}^{d}, s \in[0, T] \text {. }
$$

Then, the Lipschitz constant of $\mathcal{T}_{\varphi}$ in this norm is upper bounded by $\gamma_{\varphi} L T$.

Proof. For any $0 \leq t \leq T$,

$$
\begin{aligned}
\left\|\mathcal{T}_{\varphi}(x)(t)-\mathcal{T}_{\varphi}(y)(t)\right\| & =\left\|\int_{0}^{t} \sum_{j=1}^{D} F\left(x\left(c_{j}\right), c_{j}\right) \varphi_{j}(s) \mathrm{d} s-\int_{0}^{t} \sum_{j=1}^{D} F\left(y\left(c_{j}\right), c_{j}\right) \varphi_{j}(s) \mathrm{d} s\right\| \\
& \leq \sum_{j=1}^{D}\left|\int_{0}^{t} \varphi_{j}(s) \mathrm{d} s\right| \cdot \max _{t \in[0, T]}\|F(x(t), t)-F(y(t), t)\| \\
& \leq \gamma_{\varphi} L T \cdot \max _{s \in[0, t]}\|x(t)-y(t)\| \\
& \leq \gamma_{\varphi} L T\|x-y\|
\end{aligned}
$$

where the third step follows by $\sum_{j=1}^{D}\left|\int_{0}^{t} \varphi_{j}(s) \mathrm{d} s\right| \leq \gamma_{\varphi} T$ for all $0 \leq t \leq T$.

For the rest of the proof, let $x_{\varphi}^{*}$ denote the fixed point of $\mathcal{T}_{\varphi}$, i.e., $\mathcal{T}_{\varphi}\left(x_{\varphi}^{*}\right)=x_{\varphi}^{*}$. The Banach fixed point theorem and Lemma 4.2 .5 shows that $x_{\varphi}^{*}$ uniquely exists if $T \leq \frac{1}{L \gamma_{\varphi}}$.

Let $x^{*}$ denote the solution of the ODE, i.e., the fixed point of $\mathcal{T}$, with $\mathcal{T}\left(x^{*}\right)=x^{*}$. Let $x^{(0)}$ denote the initial solution given by $x^{(0)}(t)=v$ and $x^{(N)}$ denote the solution obtained by
applying operator $\mathcal{T}_{\varphi}$ for $N$ times. Note that $x^{(N)}(t)$ is the output of CollocationMethod in Algorithm 4.2.

Let $q \in \mathcal{V}$ denote an approximation of $\frac{\mathrm{d}}{\mathrm{d} t} x^{*}$ such that

$$
\left\|q(t)-\frac{\mathrm{d}}{\mathrm{~d} t} x^{*}(t)\right\| \leq \frac{\epsilon}{T} \cdot \forall t \in[0, T]
$$

The next lemma summarizes how these objects are related and will allow us to prove the main guarantee for first-order ODEs.

Lemma 4.2.6. Let $L^{(j)}$ be the Lipschitz constant of the map $\mathcal{T}_{\varphi}^{\circ j}$. Assume that $L^{(N)} \leq 1 / 2$. Then, we have

$$
\begin{align*}
\left\|x^{(N)}-x^{*}\right\| & \leq L^{(N)}\left\|x^{(0)}-x^{*}\right\|+2\left\|x_{\varphi}^{*}-x^{*}\right\|  \tag{4.6}\\
\left\|x_{\varphi}^{*}-x^{*}\right\| & \leq 2 \cdot\left\|\mathcal{T}_{\varphi}^{\circ N}\left(x^{*}\right)-x^{*}\right\|  \tag{4.7}\\
\left\|x^{*}-\mathcal{T}_{\varphi}^{\circ N}\left(x^{*}\right)\right\| & \leq \sum_{i=0}^{N-1} L^{(i)} \cdot\left\|x^{*}-\mathcal{T}_{\varphi}\left(x^{*}\right)\right\|  \tag{4.8}\\
\left\|x^{*}-\mathcal{T}_{\varphi}\left(x^{*}\right)\right\| & \leq 2 \gamma_{\varphi} \cdot \epsilon . \tag{4.9}
\end{align*}
$$

Proof. We prove the claims in order.

For the first claim,

$$
\begin{array}{rlr}
\left\|x^{(N)}-x^{*}\right\| & \leq\left\|x^{(N)}-x_{\varphi}^{*}\right\|+\left\|x_{\varphi}^{*}-x^{*}\right\| & \quad \text { by triangle inequality } \\
& =\left\|\mathcal{T}_{\varphi}^{\circ N}\left(x^{(0)}\right)-\mathcal{T}_{\varphi}^{\circ N}\left(x_{\varphi}^{*}\right)\right\|+\left\|x_{\varphi}^{*}-x^{*}\right\| \\
& \leq L^{(N)}\left\|x^{(0)}-x_{\varphi}^{*}\right\|+\left\|x_{\varphi}^{*}-x^{*}\right\| \\
& \leq L^{(N)}\left\|x^{(0)}-x^{*}\right\|+L^{(N)}\left\|x^{*}-x_{\varphi}^{*}\right\|+\left\|x_{\varphi}^{*}-x^{*}\right\| \\
& \leq L^{(N)}\left\|x^{(0)}-x^{*}\right\|+2\left\|x_{\varphi}^{*}-x^{*}\right\|
\end{array}
$$

where the last step follows by $L^{(N)} \leq 1$.
For the second claim,

$$
\begin{array}{rlr}
\left\|x_{\varphi}^{*}-x^{*}\right\| & =\left\|\mathcal{T}_{\varphi}^{\circ N}\left(x_{\varphi}^{*}\right)-x^{*}\right\| & \text { by } x_{\varphi}^{*}=\mathcal{T}_{\varphi}^{\circ N}\left(x_{\varphi}^{*}\right) \\
& \leq\left\|\mathcal{T}_{\varphi}^{\circ N}\left(x_{\varphi}^{*}\right)-\mathcal{T}_{\varphi}^{\circ N}\left(x^{*}\right)\right\|+\left\|\mathcal{T}_{\varphi}^{\circ N}\left(x^{*}\right)-x^{*}\right\| & \text { by triangle inequality } \\
& \leq L^{(N)} \cdot\left\|x_{\varphi}^{*}-x^{*}\right\|+\left\|\mathcal{T}_{\varphi}^{\circ N}\left(x^{*}\right)-x^{*}\right\| & \text { by the definition of } L^{(N)} \\
& \leq \frac{1}{2}\left\|x_{\varphi}^{*}-x^{*}\right\|+\left\|\mathcal{T}_{\varphi}^{\circ N}\left(x^{*}\right)-x^{*}\right\| & \text { by } L^{(N)} \leq 1 / 2
\end{array}
$$

For the third claim,

$$
\begin{aligned}
\left\|x^{*}-\mathcal{T}_{\varphi}^{\circ N}\left(x^{*}\right)\right\| & \leq \sum_{i=0}^{N-1}\left\|\mathcal{T}_{\varphi}^{\circ i}\left(x^{*}\right)-\mathcal{T}_{\varphi}^{\circ(i+1)}\left(x^{*}\right)\right\| \\
& \leq \sum_{i=0}^{N-1} L^{(i)} \cdot\left\|x^{*}-\mathcal{T}_{\varphi}\left(x^{*}\right)\right\|
\end{aligned}
$$

For the last claim,

$$
\begin{aligned}
& \left\|x^{*}(t)-\mathcal{T}_{\varphi}\left(x^{*}\right)(t)\right\| \\
= & \left\|\mathcal{T}\left(x^{*}\right)(t)-\mathcal{T}_{\varphi}\left(x^{*}\right)(t)\right\| \\
= & \left\|\int_{0}^{t} F\left(x^{*}(s), s\right) \mathrm{d} s-\int_{0}^{t} \sum_{j=1}^{D} F\left(x^{*}\left(c_{j}\right), c_{j}\right) \varphi_{j}(s) \mathrm{d} s\right\| \\
= & \left\|\int_{0}^{t} \frac{\mathrm{~d}}{\mathrm{~d} t} x^{*}(s) \mathrm{d} s-\int_{0}^{t} \sum_{j=1}^{D} \frac{\mathrm{~d}}{\mathrm{~d} t} x^{*}\left(c_{j}\right) \varphi_{j}(s) \mathrm{d} s\right\| \\
\leq & \left\|\int_{0}^{t}\left(\frac{\mathrm{~d}}{\mathrm{~d} t} x^{*}(s)-q(s)\right) \mathrm{d} s-\int_{0}^{t} \sum_{j=1}^{D}\left(\frac{\mathrm{~d}}{\mathrm{~d} t} x^{*}\left(c_{j}\right)-q\left(c_{j}\right)\right) \varphi_{j}(s) \mathrm{d} s\right\| \\
+ & \left\|\int_{0}^{t} q(s) \mathrm{d} s-\int_{0}^{t} \sum_{j=1}^{D} q\left(c_{j}\right) \varphi_{j}(s) \mathrm{d} s\right\| \\
\leq & \left.\int_{0}^{t}\left\|\frac{\mathrm{~d}}{\mathrm{~d} t} x^{*}(s)-q(s)\right\| \mathrm{d} s+\sum_{j=1}^{D}\left\|\frac{\mathrm{~d}}{\mathrm{~d} t} x^{*}\left(c_{j}\right)-q\left(c_{j}\right)\right\| \int_{0}^{t} \varphi_{j}(s) \mathrm{d} s \right\rvert\,+0 \\
\leq & \left(1+\gamma_{\varphi}\right) \cdot \epsilon+0
\end{aligned}
$$

where the first step follows by $\mathcal{T}\left(x^{*}\right)=x^{*}$, the second step follows by the definition of $\mathcal{T}$ and $\mathcal{T}_{\varphi}$, the third step follows by $x^{*}(t)$ is the solution of ODE, the fourth step follows by triangle inequality, the second last step follows by $q \in \mathcal{V}$, and the last step follows by $\left\|\frac{\mathrm{d}}{\mathrm{d} t} x^{*}-q\right\| \leq \frac{\epsilon}{T}$ and the definition of $\gamma_{\varphi}$.

Now, we are ready to prove Theorem 4.2.2.

Proof. Using Lemma 4.2.6, we have

$$
\begin{align*}
\left\|x^{(N)}-x^{*}\right\| & \leq L^{(N)}\left\|x^{(0)}-x^{*}\right\|+2\left\|x_{\varphi}^{*}-x^{*}\right\| & & \text { by Eq. (4.6) } \\
& \leq L^{(N)}\left\|x^{(0)}-x^{*}\right\|+4\left\|\mathcal{T}_{\varphi}^{\circ N}\left(x^{*}\right)-x^{*}\right\| & & \text { by Eq. (4.7) }  \tag{4.7}\\
& \leq L^{(N)}\left\|x^{(0)}-x^{*}\right\|+4 \sum_{i=0}^{N-1} L^{(i)} \cdot\left\|x^{*}-\mathcal{T}_{\varphi}\left(x^{*}\right)\right\| & & \text { by Eq. (4.8) } \\
& \leq L^{(N)}\left\|x^{(0)}-x^{*}\right\|+8 \sum_{i=0}^{N-1} L^{(i)} \cdot \gamma_{\varphi} \cdot \epsilon . & & \text { by Eq. (4.9) } \tag{4.9}
\end{align*}
$$

Using the assumption that $\gamma_{\varphi} L T \leq \frac{1}{2}$, Lemma 4.2 .5 shows that $L^{(1)} \leq \frac{1}{2}$ and hence $L^{(j)} \leq \frac{1}{2^{j}}$. Therefore, we have

$$
\begin{equation*}
\left\|x^{(N)}-x^{*}\right\| \leq \frac{1}{2^{N}}\left\|x^{(0)}-x^{*}\right\|+16 \gamma_{\varphi} \cdot \epsilon=\frac{1}{2^{N}}\left\|x^{*}-x^{*}(0)\right\|+16 \gamma_{\varphi} \cdot \epsilon \tag{4.10}
\end{equation*}
$$

To bound $\left\|x^{*}-x^{*}(0)\right\|$, for any $0 \leq t \leq T$

$$
x^{*}(t)=x^{*}(0)+\int_{0}^{t} F\left(x^{*}(s), s\right) \mathrm{d} s
$$

Hence, we have that

$$
\begin{aligned}
\left\|x^{*}(t)-x^{*}(0)\right\| & \leq\left\|\int_{0}^{T} F\left(x^{*}(0), s\right) \mathrm{d} s\right\|+\left\|\int_{0}^{t}\left(F\left(x^{*}(s), s\right)-F\left(x^{*}(0), s\right)\right) \mathrm{d} s\right\| \\
& \leq\left\|\int_{0}^{T} F\left(x^{*}(0), s\right) \mathrm{d} s\right\|+L \int_{0}^{t}\left\|x^{*}(s)-x^{*}(0)\right\| \mathrm{d} s
\end{aligned}
$$

Solving this integral inequality (see Lemma 4.6.3), we have that

$$
\left\|x^{*}(t)-x^{*}(0)\right\| \leq e^{L t}\left\|\int_{0}^{T} F\left(x^{*}(0), s\right) \mathrm{d} s\right\| .
$$

Now, we use $L T \leq \frac{1}{2}$ and get

$$
\left\|x^{*}(t)-x^{*}(0)\right\| \leq 2\left\|\int_{0}^{T} F\left(x^{*}(0), s\right) \mathrm{d} s\right\| .
$$

Picking $N=\left\lceil\log _{2}\left(\frac{T}{\epsilon} \max _{s \in[0, T]}\left\|F\left(x^{*}(0), s\right)\right\|\right)\right\rceil$, (4.10) shows that the error is less than $20 \gamma_{\varphi} \epsilon$.

### 4.2.5 A basis for piece-wise polynomials

In this section, we discuss how to construct a bounded basis for low-degree piece-wise polynomials. We are given $n$ intervals $\left\{\left[I_{i-1}, I_{i}\right]\right\}_{i=1}^{n}$ where $I_{0}=0$ and $I_{n}=T$. In the $i^{\text {th }}$ interval $\left[I_{i-1}, I_{i}\right]$, we represent the function by a degree $D_{i}$ polynomial. Formally, we define the function subspace by

$$
\begin{equation*}
\mathcal{V} \stackrel{\text { def }}{=} \bigoplus_{i=1}^{n} \nu_{i} \quad \text { with } \quad \nu_{i} \stackrel{\text { def }}{=}\left\{\left(\sum_{j=0}^{D_{i}} \alpha_{j} t^{j}\right) \cdot 1_{\left[I_{i-1}, I_{i}\right]}: \alpha_{j} \in \mathbb{R}\right\} . \tag{4.11}
\end{equation*}
$$

The following Lemma shows we can construct the basis for $\mathcal{V}$ by concatenating the basis for $\nu_{i}$.

Lemma 4.2.7. For $i \in[n]$, we are given a $\gamma_{i}$ bounded basis $\left\{\varphi_{j, i}\right\}_{j=0}^{D_{i}}$ for the subspace $\mathcal{V}_{i} \subset \mathcal{E}\left(\left[I_{i-1}, I_{i}\right], \mathbb{R}\right)$ on nodes point $\left\{c_{j, i}\right\}_{j=0}^{D_{i}}$. Then, $\left\{\varphi_{j, i}\right\}_{i, j}$ is a $\sum_{i=1}^{n} \gamma_{i}\left(I_{i}-I_{i-1}\right)$ bounded basis for the subspace $\bigoplus_{i=1}^{n} \mathcal{V}_{i} \subset \mathcal{C}\left(\left[I_{0}, I_{n}\right], \mathbb{R}\right)$.

Proof. For any $t \geq 0$, we have

$$
\begin{aligned}
\sum_{i=1}^{n} \sum_{j=0}^{D_{i}}\left|\int_{I_{0}}^{t} \varphi_{i, j}(s) \mathrm{d} s\right| & \leq \sum_{i=1}^{n}\left(\sum_{j=0}^{D_{i}}\left|\int_{I_{i-1}}^{t} \varphi_{i, j}(s) \mathrm{d} s\right| 1_{t \geq I_{i-1}}\right) \\
& \leq \sum_{i=1}^{n} \gamma_{i}\left(I_{i}-I_{i-1}\right)
\end{aligned}
$$

where we used that $\varphi_{i, j}$ is supported on $\left[I_{i-1}, I_{i}\right]$ in the first inequality.

Next, we note that the boundedness for basis is shift and scale invariant. Hence, we will focus on obtaining a basis for $(t-1)$-degree polynomial on $[-1,1]$ for notation convenience.

For $[-1,1]$, we choose the node points $c_{j}=\cos \left(\frac{2 j-1}{2 t} \pi\right)$ and the basis are

$$
\varphi_{j}(x)=\frac{\sqrt{1-c_{j}^{2}} \cos \left(t \cos ^{-1} x\right)}{t\left(x-c_{j}\right)} .
$$

It is easy to see that $\varphi_{j}\left(c_{i}\right)=\delta_{i, j}$. To bound the integral, Lemma 91 in [LV17] shows that

$$
\left|\int_{-1}^{y} \varphi_{j}(x) d x\right| \leq \frac{2000}{t} \text { for all } y \in[-1,1] .
$$

Summing it over $t$ basis functions, we have that $\gamma_{\varphi} \leq 2000$. Together with Lemma 4.2.7, we have the following result:

Lemma 4.2.8. Let $\mathcal{V}$ be a subspace of piecewise polynomials on $[0, T]$ with fixed nodes. Then, there is a 2000 bounded basis $\{\varphi\}$ for $\mathcal{V}$. Furthermore, for any vector $v$, it takes $O\left(\sum_{i=1}^{n}\left(1+D_{i}\right)^{2}\right)$ time to compute $v^{\top} A_{\varphi}$ where $D_{i}$ is the maximum degree of the $i$-th piece.

Alternatively, one can find $u$ such that $\left\|u-v^{\top} A_{\varphi}\right\| \leq \epsilon T\|v\|_{\infty}$ in time

$$
O\left(\operatorname{rank}(\mathcal{V}) \log \left(\frac{\operatorname{rank}(\mathcal{V})}{\epsilon}\right)\right)
$$

Proof. The bound follows from previous discussion. For the computation cost, note that

$$
\left(v^{\top} A_{\varphi}\right)_{(i, j)}=\sum_{i^{\prime}, j^{\prime}} \int_{0}^{c_{(i, j)}} v_{\left(i^{\prime}, j^{\prime}\right)} \varphi_{\left(i^{\prime}, j^{\prime}\right)}(s) \mathrm{d} s
$$

where $(i, j)$ is the $j$-th node at the $i$-th piece. For any $i \neq i^{\prime}$, the support of $\varphi_{\left(i^{\prime}, j^{\prime}\right)}(s)$ is either disjoint from $\left[0, c_{(i, j)}\right]$ (if $i^{\prime}>i$ ) or included in $\left[0, c_{(i, j)}\right]$ (if $i^{\prime}<i$ ). Hence, we have that

$$
\int_{0}^{c_{(i, j)}} \varphi_{\left(i^{\prime}, j^{\prime}\right)}(s) \mathrm{d} s=\left\{\begin{array}{ll}
0 & \text { if } i^{\prime}>i \\
\int_{-\infty}^{\infty} \varphi_{\left(i^{\prime}, j^{\prime}\right)}(s) \mathrm{d} s & \text { if } i^{\prime}<i
\end{array} .\right.
$$

Therefore, we have

$$
\left(v^{\top} A_{\varphi}\right)_{(i, j)}=\sum_{i^{\prime}<i, j^{\prime}} v_{\left(i^{\prime}, j^{\prime}\right)} \cdot \int_{-\infty}^{\infty} \varphi_{\left(i^{\prime}, j^{\prime}\right)}(s) \mathrm{d} s+\sum_{j^{\prime}} v_{\left(i, j^{\prime}\right)} \cdot \int_{0}^{c_{(i, j)}} \varphi_{\left(i, j^{\prime}\right)}(s) \mathrm{d} s
$$

Note that $\int_{0}^{c_{(i, j)}} \varphi_{\left(i^{\prime}, j^{\prime}\right)}(s) \mathrm{d} s$ can precomputed. Since there are $\sum_{i=1}^{n}\left(1+D_{i}\right)$ many pairs of $(i, j)$, the first term can be computed in $\sum_{i=1}^{n}\left(1+D_{i}\right)$ time. Since there are $\sum_{i=1}^{n}\left(1+D_{i}\right)^{2}$ many pairs of $\left(i, j, j^{\prime}\right)$, the second term can be computed in $\sum_{i=1}^{n}\left(1+D_{i}\right)^{2}$ time.

Theorem 4.2.9 gives another way to compute the integration and its runtime is

$$
O\left(\operatorname{rank}(\mathcal{V}) \log \left(\frac{\operatorname{rank}(\mathcal{V})}{\epsilon}\right)\right)
$$

Remark 4.2.2. Experiment seems to suggest the basis we proposed is 1 bounded.

Here is the theorem we used above to compute the Lagrange polynomials.

Theorem 4.2.9 ([DGR96, Section 5]). Let $\phi_{i}$ be the Lagrange basis polynomials on the Chebyshev nodes $c_{j}=\cos \left(\frac{2 j-1}{2 t} \pi\right)$ for $j \in[t]$, namely, $\phi_{i}(s)=\prod_{j \neq i} \frac{s-c_{j}}{c_{i}-c_{j}}$. Given a polynomial $p(s)=\sum_{j=1}^{t} \alpha_{j} \phi_{j}(s)$ represented by $\left\{\alpha_{j}\right\}_{j=1}^{t}$, one can compute $\left\{\ell_{i}\right\}_{i=1}^{t}$ such that

$$
\left|\ell_{i}-\int_{0}^{c_{i}} p(s) d s\right| \leq \epsilon\|\alpha\|_{\infty}
$$

in time $O\left(t \log \left(\frac{t}{\epsilon}\right)\right)$.

### 4.3 Improved Contraction Bound for HMC

In this section, we give an improved contraction bound for HMC (Algorithm 4.1). Each iteration of Algorithm 4.1 solve the HMC dynamics approximately. In the later sections, we will discuss how to solve this ODE.

To give a contraction bound for the noisy HMC, we first analyze the contraction of the ideal HMC. We reduce the problem of bounding the contraction rate of the ideal HMC to the following lemma involving a matrix ODE.

Lemma 4.3.1. Given a symmetric matrix $H(t)$ such that $0 \prec m_{2} \cdot I \preceq H(t) \preceq M_{2} \cdot I$ for all $t \geq 0$. Consider the $O D E$

$$
\begin{aligned}
& u^{\prime \prime}(t)=-H(t) \cdot u(t) \\
& u^{\prime}(0)=0
\end{aligned}
$$

Let $\alpha(t)=\frac{1}{\|u(0)\|_{2}} \int_{0}^{t}(t-s) \cdot\|H(s) u(0)\|_{2} \mathrm{~d} s$. For any $0 \leq T \leq \frac{1}{2 \sqrt{M_{2}}}$ such that $\alpha(T) \leq \frac{1}{8} \sqrt{\frac{m_{2}}{M_{2}}}$, we have that

$$
\|u(T)\|_{2}^{2} \leq\left(1-\max \left(\frac{1}{4} m_{2} T^{2}, \frac{1}{2} \sqrt{\frac{m_{2}}{M_{2}}} \cdot \alpha(T)\right)\right) \cdot\|u(0)\|_{2}^{2}
$$

Using this lemma, we prove both parts of the main contraction bound, Lemma 4.1.6.

Proof of Lemma 4.1.6. Let error function $e(t)=y(t)-x(t)$. The definition of HMC shows that

$$
e^{\prime \prime}(t)=-(\nabla f(y(t))-\nabla f(x(t)))=-H(t) \cdot e(t)
$$

where $H(t)=\int_{0}^{1} \nabla^{2} f(x(t)+s(y(t)-x(t))) \mathrm{d} s$. By the strong convexity and the Lipschitz gradient of $f$, we have that

$$
m_{2} \cdot I \preceq H(t) \preceq M_{2} \cdot I .
$$

Hence, we can apply Lemma 4.3.1.

To get the first bound, we bound the $\alpha(t)$ defined in Lemma 4.3.1 as follows

$$
\alpha(t)=\frac{1}{\|e(0)\|_{2}} \int_{0}^{t}(t-s) \cdot\|H(s) e(0)\|_{2} \mathrm{~d} s \leq M_{2} \int_{0}^{t}(t-s) \mathrm{d} s=M_{2} \frac{t^{2}}{2}
$$

Therefore, for $0 \leq t \leq \frac{m_{2}^{1 / 4}}{2 M_{2}^{3 / 4}}$, we have that $\alpha(t) \leq \frac{1}{8} \sqrt{\frac{m_{2}}{M_{2}}}$ and hence Lemma 4.3.1 gives the first bound.

To get the second bound, we note that $\alpha(t)$ is increasing and hence there is $t$ such that $\alpha(t)=\frac{1}{8} \sqrt{\frac{m_{2}}{M_{2}}}$. Using such $t$ in Lemma 4.3.1 gives the second bound.

Now, we prove the main technical lemma of this section, a contraction estimate for matrix ODE. We note that not all matrix ODEs come from some HMC and hence it might be possible to get a better bound by directly analyzing the HMC.

Proof of Lemma 4.3.1. Let $e_{1}$ denote the basis vector that it is 1 in the first coordinate and 0 everywhere else.

Without loss of generality, we can assume $\|u(0)\|_{2}=1$ and $u(0)=e_{1}$.
The proof involves first getting a crude bound on $\|u(t)\|_{2}$. Then, we boast the bound by splitting the movement of $u(t)$ into one parallel to $e_{1}$ and one orthogonal to $e_{1}$.

Crude bound on $\|u(t)\|_{2}$ :
Integrating both sides of $u^{\prime \prime}(t)=-H(t) \cdot u(t)$ twices and using $u^{\prime}(0)=0$ gives

$$
\begin{equation*}
u(t)=u(0)-\int_{0}^{t}(t-s) H(s) u(s) \mathrm{d} s \tag{4.12}
\end{equation*}
$$

We take the norm on both sides and use $0 \preceq H(s) \preceq M_{2} \cdot I$ to get

$$
\|u(t)\|_{2} \leq 1+M_{2} \cdot \int_{0}^{t}(t-s)\|u(s)\|_{2} \mathrm{~d} s
$$

Applying Lemma 4.6.2 to this equation and using $t \leq \frac{1}{2 \sqrt{M_{2}}}$ gives

$$
\|u(t)\|_{2} \leq \cosh \left(\sqrt{M_{2}} t\right) \leq \frac{6}{5}
$$

Putting it back to (4.12) gives

$$
\left\|u(t)-e_{1}\right\|_{2} \leq \int_{0}^{t}(t-s)\|H(s) \cdot u(s)\|_{2} \mathrm{~d} s=\int_{0}^{t}(t-s) \cdot M_{2} \cdot \frac{6}{5} \mathrm{~d} s=\frac{6}{10} M_{2} t^{2}
$$

In particular, for any $0 \leq t \leq \frac{1}{2 \sqrt{M_{2}}}$, we have that

$$
\begin{equation*}
\frac{5}{6} \leq u_{1}(t) \leq \frac{7}{6} \tag{4.13}
\end{equation*}
$$

## Improved bound on $\|u(t)\|_{2}$ :

Let $P_{1}$ be the orthogonal projection to the first coordinate and $P_{-1}=I-P_{1}$. We write $u(t)=u_{1}(t)+u_{-1}(t)$ with $u_{1}(t)=P_{1} u(t)$ and $u_{-1}(t)=P_{-1} u(t)$, namely, $u_{1}(t)$ is parallel to $e_{1}$ and $u_{-1}(t)$ is orthogonal to $e_{1}$.

Fix any $0 \leq t \leq T$. Let $\beta(t)=e_{1}^{\top} u(t)$. By the definition of $u$, we have

$$
\begin{equation*}
u^{\prime \prime}(t)=-\beta(t) \cdot H(t) e_{1}-H(t) u_{-1}(t) \tag{4.14}
\end{equation*}
$$

Integrating both sides twice and using $u_{-1}(0)=0$, we have

$$
\begin{aligned}
u_{-1}(t) & =\int_{0}^{t}(t-s) P_{-1} u^{\prime \prime}(s) \mathrm{d} s \\
& =-\int_{0}^{t}(t-s) \cdot \beta(s) \cdot P_{-1} H(s) e_{1} \mathrm{~d} s-\int_{0}^{t}(t-s) P_{-1} H(s) u_{-1}(s) \mathrm{d} s
\end{aligned}
$$

Taking norm on both sides and using that $0 \preceq H(t) \preceq M_{2} \cdot I$ and $\frac{5}{6} \leq \beta(s) \leq \frac{7}{6}$, we have that

$$
\begin{align*}
\left\|u_{-1}(t)\right\|_{2} & \leq \int_{0}^{t}(t-s) \cdot \beta(s) \cdot\left\|H(s) e_{1}\right\|_{2} \mathrm{~d} s+M_{2} \cdot \int_{0}^{t}(t-s) \cdot\left\|u_{-1}(s)\right\|_{2} \mathrm{~d} s \\
& \leq \frac{7}{6} \int_{0}^{t}(t-s) \cdot\left\|H(s) e_{1}\right\|_{2} \mathrm{~d} s+M_{2} \cdot \int_{0}^{t}(t-s) \cdot\left\|u_{-1}(s)\right\|_{2} \mathrm{~d} s \\
& \leq \frac{7}{6} \alpha(T)+M_{2} \cdot \int_{0}^{t}(t-s) \cdot\left\|u_{-1}(s)\right\|_{2} \mathrm{~d} s \tag{4.15}
\end{align*}
$$

where $\alpha(T) \stackrel{\text { def }}{=} \int_{0}^{T}(T-s) \cdot\left\|H(s) e_{1}\right\|_{2} \mathrm{~d} s$. Solving this integral inequality (Lemma 4.6.2), we get

$$
\begin{equation*}
\left\|u_{-1}(t)\right\|_{2} \leq \frac{7}{6} \alpha(T) \cdot \cosh \left(\sqrt{M_{2}} t\right) \leq \frac{7}{6} \alpha(T) \cdot \cosh (1 / 2) \leq \frac{4}{3} \alpha(T) \tag{4.16}
\end{equation*}
$$

where the second step follows from $t \leq \frac{1}{2 \sqrt{M_{2}}}$, and the last step follows from $\cosh (1 / 2) \leq \frac{8}{7}$.
Next, we look at the first coordinate of (4.14) and get

$$
\begin{align*}
\beta^{\prime \prime}(t) & =-\beta(t) \cdot e_{1}^{\top} H(t) e_{1}-e_{1}^{\top} H(t) u_{-1}(t) \\
& \leq-\frac{5}{6} e_{1}^{\top} H(t) e_{1}+\left\|H(t) e_{1}\right\|_{2} \cdot\left\|u_{-1}(t)\right\|_{2} \tag{4.17}
\end{align*}
$$

To bound the last term, we note that

$$
\begin{array}{rlrl}
\left\|H(t) e_{1}\right\|_{2} & =\sqrt{e_{1}^{\top} H^{2}(t) e_{1}} & \\
& \leq \sqrt{M_{2} \cdot e_{1}^{\top} H(t) e_{1}} & & \text { by } H^{2}(t) \preceq M_{2} \cdot H(t) \\
& =\sqrt{\frac{M_{2}}{m_{2}}} \sqrt{\frac{m_{2}}{e_{1}^{\top} H(t) e_{1}}} e_{1}^{\top} H(t) e_{1} & \\
& \leq \sqrt{\frac{M_{2}}{m_{2}}} \cdot e_{1}^{\top} H(t) e_{1} & & \text { by } m_{2} \leq e_{1}^{\top} H(t) e_{1} . \tag{4.18}
\end{array}
$$

Using this into (4.17) and $\alpha(T) \leq \frac{1}{8} \sqrt{\frac{m_{2}}{M_{2}}}$, we have

$$
\beta^{\prime \prime}(t) \leq-e_{1}^{\top} H(t) e_{1} \cdot\left(\frac{5}{6}-\frac{4}{3} \sqrt{\frac{M_{2}}{m_{2}}} \alpha(T)\right) \leq-\frac{2}{3} e_{1}^{\top} H(t) e_{1} .
$$

Hence, we have that

$$
\begin{equation*}
\beta(t) \leq 1-\frac{2}{3} \int_{0}^{t}(t-s) \cdot e_{1}^{\top} H(s) e_{1} \mathrm{~d} s \tag{4.19}
\end{equation*}
$$

Using (4.19), $\beta(t) \geq \frac{5}{6}$ and (4.16) gives

$$
\begin{aligned}
\|u(t)\|_{2}^{2} & =\beta^{2}(t)+\left\|u_{-1}(t)\right\|_{2}^{2} \\
& \leq 1-\int_{0}^{t}(t-s) \cdot e_{1}^{\top} H(s) e_{1} \mathrm{~d} s+\left(\frac{4}{3} \int_{0}^{T}(T-s) \cdot\left\|H(s) e_{1}\right\|_{2} \mathrm{~d} s\right)^{2} \\
& \leq 1-\int_{0}^{t}(t-s) \cdot e_{1}^{\top} H(s) e_{1} \mathrm{~d} s+2 \sqrt{\frac{M_{2}}{m_{2}}} \int_{0}^{T}(T-s) \cdot e_{1}^{\top} H(s) e_{1} \mathrm{~d} s \cdot \alpha(T) \\
& =1-\int_{0}^{t}(t-s) \cdot e_{1}^{\top} H(s) e_{1} \mathrm{~d} s\left(1-2 \sqrt{\frac{M_{2}}{m_{2}}} \alpha(T)\right) \\
& \leq 1-\frac{1}{2} \int_{0}^{t}(t-s) \cdot e_{1}^{\top} H(s) e_{1} \mathrm{~d} s
\end{aligned}
$$

where we used (4.18) at the second inequality and $\alpha(T) \leq \frac{1}{8} \sqrt{\frac{m_{2}}{M_{2}}}$ at the end.
Finally, we bound the last term in two way. One way simply uses $e_{1}^{\top} H(s) e_{1} \geq m_{2}$ and get

$$
\int_{0}^{t}(t-s) \cdot e_{1}^{\top} H(s) e_{1} \mathrm{~d} s \geq \frac{m_{2}}{2} t^{2}
$$

which implies

$$
\|u(t)\|_{2}^{2} \leq 1-\frac{m_{2}}{4} t^{2}
$$

For the other way, we apply (4.18) to get that

$$
\int_{0}^{t}(t-s) \cdot e_{1}^{\top} H(s) e_{1} \mathrm{~d} s \geq \int_{0}^{t}(t-s) \cdot\left\|H(t) e_{1}\right\|_{2} \sqrt{\frac{m_{2}}{M_{2}}} \mathrm{~d} s=\sqrt{\frac{m_{2}}{M_{2}}} \cdot \alpha(T)
$$

where the last step follows by the definition of $\alpha(T)$. Thus, we have

$$
\|u(t)\|_{2}^{2} \leq 1-\frac{1}{2} \sqrt{\frac{m_{2}}{M_{2}}} \alpha(T)
$$

Finally, we analyze the contraction of the noisy HMC.

Theorem 4.3.2 (Contraction of noisy HMC). Suppose $f$ is $m_{2}$ strongly convexity with $M_{2}$ Lipschitz gradient. For any step-size $h \leq \frac{m_{2}^{1 / 4}}{2 M_{2}^{3 / 4}}$, let $X \sim \operatorname{HMC}\left(x^{(0)}, f, \epsilon, h\right)$ and $Y \sim e^{-f}$. Then, we have that

$$
W_{2}(X, Y) \leq \frac{\epsilon}{\sqrt{m_{2}}}
$$

In addition, the number of iterations is

$$
N=O\left(\frac{1}{m_{2} h^{2}}\right) \cdot \log \left(\frac{\left\|\nabla f\left(x^{(0)}\right)\right\|_{2}^{2} / m_{2}+d}{\epsilon}\right) .
$$

Proof. To prove the $W_{2}$ distance, we let $x^{(k)}$ be the iterates of the algorithm HMC. Let $y^{(k)}$ be the $k$-th step of the ideal HMC starting from a random point $y^{(0)} \sim e^{-f}$ with the random initial direction identical to the algorithm HMC. Let $x^{*(k)}$ be the 1 step ideal HMC starting from $x^{(k-1)}$ with the same initial direction as $y^{(k)}$. Lemma 4.1.6 shows that

$$
\left\|x^{*(k)}-y^{(k)}\right\|_{2}^{2} \leq\left(1-\frac{m_{2} h^{2}}{4}\right)\left\|x^{(k-1)}-y^{(k-1)}\right\|_{2}^{2}
$$

Let $\theta=\frac{m_{2} h^{2}}{8}$, then $\theta \leq \frac{1 /\left(4 \kappa^{3 / 4}\right)}{8} \leq 1 / 32$. By the assumption of the noise, we have $\| x^{(k)}-$
$x^{*(k)} \|_{2} \leq \frac{\epsilon \theta}{2 \sqrt{m_{2}}}$. Hence, we have

$$
\begin{aligned}
\left\|x^{(k)}-y^{(k)}\right\|_{2}^{2} & =\left\|\left(x^{*(k)}-y^{(k)}\right)+\left(x^{(k)}-x^{*(k)}\right)\right\|_{2}^{2} \\
& \leq(1+\theta)\left\|x^{*(k)}-y^{(k)}\right\|_{2}^{2}+(1+1 / \theta)\left\|x^{(k)}-x^{*(k)}\right\|_{2}^{2} \\
& \leq(1+\theta)(1-2 \theta)\left\|x^{(k-1)}-y^{(k-1)}\right\|_{2}^{2}+(1+1 / \theta)\left\|x^{(k)}-x^{*(k)}\right\|_{2}^{2} \\
& \leq(1-\theta)\left\|x^{(k-1)}-y^{(k-1)}\right\|_{2}^{2}+(1+1 / \theta)\left\|x^{(k)}-x^{*(k)}\right\|_{2}^{2} \\
& \leq(1-\theta)\left\|x^{(k-1)}-y^{(k-1)}\right\|_{2}^{2}+(2 / \theta) \cdot \frac{\epsilon^{2} \theta^{2}}{4 m_{2}}
\end{aligned}
$$

where the second step follows by $(a+b)^{2} \leq(1+\theta) a^{2}+(1+1 / \theta) b^{2}$, the third step follows by $\left\|x^{*(k)}-y^{(k)}\right\|_{2}^{2} \leq(1-2 \theta)\left\|x^{(k-1)}-y^{(k-1)}\right\|_{2}^{2}$, the fourth step follows by $(1+\theta)(1-2 \theta) \leq(1-\theta)$, the fifth step follow by $\theta \leq 1 / 4$ and $\left\|x^{(k)}-x^{*(k)}\right\|_{2} \leq \frac{\epsilon \theta}{2 \sqrt{m_{2}}}$.

Applying this bound iteratively gives

$$
\begin{equation*}
\left\|x^{(k)}-y^{(k)}\right\|_{2}^{2} \leq(1-\theta)^{k}\left\|x^{(0)}-y^{(0)}\right\|_{2}^{2}+\frac{\epsilon^{2}}{2 m_{2}} \tag{4.20}
\end{equation*}
$$

Let $x^{(\min )}$ be the minimum of $f$. Then, we have

$$
\begin{equation*}
\left\|x^{(0)}-y^{(0)}\right\|_{2}^{2} \leq 2\left\|x^{(0)}-x^{(\min )}\right\|_{2}^{2}+2\left\|y^{(0)}-x^{(\min )}\right\|_{2}^{2} \tag{4.21}
\end{equation*}
$$

For the first term, the strong convexity of $f$ shows that

$$
\begin{equation*}
\left\|x^{(0)}-x^{(\min )}\right\|_{2}^{2} \leq \frac{1}{m^{2}}\left\|\nabla f\left(x^{(0)}\right)\right\|_{2}^{2} \tag{4.22}
\end{equation*}
$$

For the second term, Theorem 1 in [DM16] shows that

$$
\begin{equation*}
\mathbb{E}\left[\left\|y^{(0)}-x^{(\min )}\right\|_{2}^{2}\right] \leq \frac{d}{m_{2}} \tag{4.23}
\end{equation*}
$$

Combining (4.20), (4.21), (4.22) and (4.23), we have

$$
\mathbb{E}\left[\left\|x^{(k)}-y^{(k)}\right\|_{2}^{2}\right] \leq(1-\theta)^{k}\left(\frac{2\left\|\nabla f\left(x^{(0)}\right)\right\|_{2}^{2}}{m_{2}^{2}}+\frac{2 d}{m_{2}}\right)+\frac{\epsilon^{2}}{2 m_{2}}
$$

Picking

$$
k=\frac{1}{\theta} \cdot \log \left(\frac{\frac{2\left\|\nabla f\left(x^{(0)}\right)\right\|_{2}^{2}}{m_{2}^{2}}+\frac{2 d}{m_{2}}}{\frac{\epsilon^{2}}{2 m_{2}}}\right)=\frac{1}{\theta} \cdot \log \left(\frac{4}{\epsilon^{2}}\left(\frac{\left\|\nabla f\left(x^{(0)}\right)\right\|_{2}^{2}}{m_{2}}+d\right)\right)
$$

we have that

$$
\mathbb{E}\left[\left\|x^{(k)}-y^{(k)}\right\|_{2}^{2}\right] \leq \frac{\epsilon^{2}}{m_{2}}
$$

This proves that $W_{2}(X, Y) \leq \frac{\epsilon}{\sqrt{m_{2}}}$.

### 4.4 Strongly Convex functions with Lipschitz Gradient

In this section, we give a faster sampling algorithm for strongly convex functions with Lipschitz gradient. The purpose of this section is to illustrate that our contraction bound and ODE theorem are useful even for functions that are not infinitely differentiable. We believe that our bound can be beaten by algorithms designed for this specific setting.

### 4.4.1 Bounding the ODE solution

First, we prove that the HMC dynamic for these functions can be well approximated by piece-wise degree-2 polynomials. Note that this only requires that the Hessian has bounded eigenvalues.

Lemma 4.4.1 (Smoothness implies the existence of degree-2 polynomial approximation). Let $f$ be a twice-differentiable function such that $-M_{2} \cdot I \preceq \nabla^{2} f(x) \preceq M_{2} \cdot I$ for all $x \in \mathbb{R}^{d}$. Let $0 \leq h \leq \frac{1}{2 \sqrt{M_{2}}}$. Consider the HMC dynamic

$$
\begin{aligned}
\frac{\mathrm{d}^{2} x}{\mathrm{~d} t^{2}}(t) & =-\nabla f(x(t)) \text { for } 0 \leq t \leq h \\
\frac{\mathrm{~d} x}{\mathrm{~d} t}(0) & =v_{1} \\
x(0) & =v_{0}
\end{aligned}
$$

For any integer $k$, there is a continuously differentiable $k$-piece degree 2 polynomial $q$ such that $q(0)=v_{0}, \frac{\mathrm{~d} q}{\mathrm{~d} t}(0)=v_{1}$ and $\left\|\frac{\mathrm{d}^{2} q}{\mathrm{~d} t^{2}}(t)-\frac{\mathrm{d}^{2} x}{\mathrm{~d} t^{2}}(t)\right\|_{2} \leq \frac{\epsilon}{h^{2}}$ for $0 \leq t \leq h$ with

$$
\epsilon=\frac{2 M_{2} h^{3}}{k}\left(\left\|v_{1}\right\|_{2}+\left\|\nabla f\left(v_{0}\right)\right\|_{2} \cdot h\right)
$$

Proof. Let $x(t)$ be the solution of the ODE. We define a continuously differentiable $k$-piece
degree- 2 polynomial $q$ by

$$
q(0)=v_{0} \text { and } q^{\prime}(t)=\frac{\mathrm{d} x\left(t_{\text {pre }}\right)}{\mathrm{d} t} \cdot \frac{t_{\text {next }}-t}{t_{\text {next }}-t_{\text {pre }}}+\frac{\mathrm{d} x\left(t_{\text {next }}\right)}{\mathrm{d} t} \cdot \frac{t_{\text {pre }}-t}{t_{\text {next }}-t_{\mathrm{pre}}}
$$

with $t_{\text {pre }}=\left\lfloor\frac{t}{h / k}\right\rfloor \cdot \frac{h}{k}$ and $t_{\text {next }}=t_{\text {pre }}+\frac{h}{k}$. Clearly, we have that $q(0)=x(0)=v_{0}$ and $\frac{\mathrm{d} q}{\mathrm{~d} t}(0)=\frac{\mathrm{d} x}{\mathrm{~d} t}(0)=v_{1}$. Also, we have that

$$
\begin{equation*}
\left\|\frac{\mathrm{d}^{2} q}{\mathrm{~d} t^{2}}(t)-\frac{\mathrm{d}^{2} x}{\mathrm{~d} t^{2}}(t)\right\|_{2}=\left\|\int_{t_{\mathrm{pre}}}^{t_{\text {next }}} \frac{\mathrm{d}^{2} x}{\mathrm{~d} t^{2}}(s) d s-\frac{\mathrm{d}^{2} x}{\mathrm{~d} t^{2}}(t)\right\|_{2} \leq \frac{h}{k} \max _{0 \leq t \leq h}\left\|\frac{\mathrm{~d}^{3} x}{\mathrm{~d} t^{3}}(t)\right\|_{2} \leq \frac{M_{2} h}{k} \max _{0 \leq t \leq h}\left\|\frac{\mathrm{~d} x}{\mathrm{~d} t}(t)\right\| \tag{4.24}
\end{equation*}
$$

where we used that $\frac{\mathrm{d}^{3} x}{\mathrm{~d} t^{3}}(t)=-\nabla^{2} f(x(t)) \frac{\mathrm{d} x}{\mathrm{~d} t}(t)$ at the end.
Therefore, it suffices to bound the term $\max _{0 \leq t \leq h}\left\|\frac{\mathrm{~d} x}{\mathrm{~d} t}(t)\right\|$. Using again that $\frac{\mathrm{d}^{3} x}{\mathrm{~d} t^{2}}(t)=$ $-\nabla^{2} f(x(t)) \frac{\mathrm{d} x}{\mathrm{~d} t}(t)$, we have that

$$
\begin{aligned}
\frac{\mathrm{d} x}{\mathrm{~d} t}(t) & =\frac{\mathrm{d} x}{\mathrm{~d} t}(0)+\frac{\mathrm{d}^{2} x}{\mathrm{~d} t^{2}}(0) \cdot t+\int_{0}^{t}(t-s) \cdot \frac{\mathrm{d}^{3} x}{\mathrm{~d} t^{3}}(s) d s \\
& =v_{1}-\nabla f\left(v_{0}\right) \cdot t-\int_{0}^{t}(t-s) \cdot \nabla^{2} f(x(s)) \frac{\mathrm{d} x}{\mathrm{~d} t}(s) d s .
\end{aligned}
$$

Hence, for $0 \leq t \leq h$, we have that

$$
\left\|\frac{\mathrm{d} x}{\mathrm{~d} t}(t)\right\|_{2} \leq\left\|v_{1}\right\|_{2}+\left\|\nabla f\left(v_{0}\right)\right\|_{2} \cdot h+M_{2} \cdot \int_{0}^{t}(t-s)\left\|\frac{\mathrm{d} x}{\mathrm{~d} t}(t)\right\|_{2} d s
$$

Solving this integral inequality, Lemma 4.6.2 shows that

$$
\max _{0 \leq t \leq h}\left\|\frac{\mathrm{~d} x}{\mathrm{~d} t}(t)\right\| \leq\left(\left\|v_{1}\right\|_{2}+\left\|\nabla f\left(v_{0}\right)\right\|_{2} \cdot h\right) \cdot \cosh \left(\sqrt{M_{2}} \cdot h\right) \leq 2\left(\left\|v_{1}\right\|_{2}+\left\|\nabla f\left(v_{0}\right)\right\|_{2} \cdot h\right)
$$

where we used that $h \leq \frac{1}{2 \sqrt{M_{2}}}$. Applying this in (4.24) gives this result.

The precise runtime depends on how accurate we need to solve the ODE, namely, the parameter $\epsilon$ in Lemma 4.4.1. The term $\left\|v_{1}\right\|_{2}$ can be upper bounded by $O(\sqrt{d})$ with
high probability since $v_{1}$ is sampled from normal distribution. The term $\left\|\nabla f\left(v_{0}\right)\right\|_{2}$ is much harder to bound. Even for a random $v_{0} e^{-f}$, the worst case bound we can give is

$$
\left\|\nabla f\left(v_{0}\right)\right\|_{2}=\left\|\nabla f\left(v_{0}\right)-\nabla f\left(x^{*}\right)\right\|_{2} \leq M_{2}\left\|v_{0}-x^{*}\right\|_{2} \lesssim \sqrt{\kappa M_{2} d}
$$

where we used that $\left\|v_{0}-x^{*}\right\|_{2} \lesssim \sqrt{d / m_{2}}$ for random $v_{0} e^{-f}$ [DM16]. This is not enough for improving existing algorithms, as we would need $\sqrt{\kappa d}$ time per iteration. The crux of this section is to show that $\left\|\nabla f\left(v_{0}\right)\right\|_{2}=O\left(\sqrt{M_{2} d}\right)$ for most of the iterations in the HMC walk if the process starts at the minimum of $f$. This is tight for quadratic $f$.

Lemma 4.4.2 (Smoothness implies expected gradient is upper bounded). Let $f$ be a function such that $-M_{2} \cdot I \preceq \nabla^{2} f(x) \preceq M_{2} \cdot I$ for all $x \in \mathbb{R}^{d}$. Let $x^{(k)}$ be the starting point of the $k^{\text {th }}$ step in $\operatorname{HMC}\left(x^{(0)}, f, \epsilon, h\right)$ (Algorithm 4.1) with step size $h \leq \frac{1}{8 \sqrt{M_{2}}}$. Then, we have that

$$
\frac{1}{N} \mathbb{E}\left[\sum_{k=0}^{N-1}\left\|\nabla f\left(x^{(k)}\right)\right\|_{2}^{2}\right] \leq O\left(\frac{f\left(x^{(0)}\right)-\min _{x} f(x)}{h^{2} N}+M_{2} d+\frac{\bar{\epsilon}^{2}}{h^{4}}\right)
$$

$\bar{\epsilon}$ is the error in solving the HMC defined in Algorithm 4.1.

Proof. Consider one step of the HMC dynamic. Note that

$$
\frac{\mathrm{d}}{\mathrm{~d} t} f(x(t))=\nabla f(x(t))^{\top} \frac{\mathrm{d} x}{\mathrm{~d} t}
$$

Hence, we have

$$
\begin{array}{rlr}
\frac{\mathrm{d}^{2}}{\mathrm{~d} t^{2}} f(x(t)) & =\frac{\mathrm{d} x^{\top}}{\mathrm{d} t} \nabla^{2} f(x(t)) \frac{\mathrm{d} x}{\mathrm{~d} t}+\nabla f(x(t))^{\top} \frac{\mathrm{d}^{2} x}{\mathrm{~d} t^{2}} & \\
& =\frac{\mathrm{d} x}{\mathrm{~d} t} \nabla^{\top} f(x(t)) \frac{\mathrm{d} x}{\mathrm{~d} t}-\|\nabla f(x(t))\|^{2} & \text { by } \frac{\mathrm{d}^{2} x}{\mathrm{~d} t^{2}}=-\nabla f(x(t)) \\
& \leq M_{2} \cdot\left\|\frac{\mathrm{~d} x}{\mathrm{~d} t}\right\|_{2}^{2}-\|\nabla f(x(t))\|^{2} . & \tag{4.25}
\end{array} \quad \text { by } \nabla^{2} f(x(t)) \preceq M_{2} \cdot I
$$

In Lemma 4.4.1, we proved that $\left\|\frac{\mathrm{d} x}{\mathrm{~d} t}\right\|_{2} \leq 2\left(\left\|\frac{\mathrm{~d} x}{\mathrm{~d} t}(0)\right\|_{2}+\|\nabla f(x(0))\|_{2} \cdot h\right)$ for all $0 \leq t \leq h$. Using this, we have that

$$
\begin{aligned}
\|\nabla f(x(t))-\nabla f(x(0))\| & \leq M_{2}\|x(t)-x(0)\|_{2} \\
& \leq M_{2} \int_{0}^{t}\left\|\frac{\mathrm{~d} x}{\mathrm{~d} t}(t)\right\|_{2} \mathrm{~d} t \\
& \leq 2 M_{2} h \cdot\left(\left\|\frac{\mathrm{~d} x}{\mathrm{~d} t}(0)\right\|_{2}+\|\nabla f(x(0))\|_{2} \cdot h\right)
\end{aligned}
$$

which implies

$$
\begin{equation*}
\|\nabla f(x(t))\|=\|\nabla f(x(0))\| \pm 2 M_{2} h\left(\left\|\frac{\mathrm{~d} x}{\mathrm{~d} t}(0)\right\|+\|\nabla f(x(0))\| \cdot h\right) . \tag{4.26}
\end{equation*}
$$

Using our choice of $h$, we have that $\|\nabla f(x(t))\| \geq \frac{1}{2}\|\nabla f(x(0))\|-2 M_{2} h \cdot\left\|\frac{\mathrm{~d} x}{\mathrm{~d} t}(0)\right\|_{2}$. Putting these estimates into (4.25) gives

$$
\begin{aligned}
& \frac{\mathrm{d}^{2}}{\mathrm{~d} t^{2}} f(x(t)) \\
\leq & 2 M_{2}\left(\left\|\frac{\mathrm{~d} x}{\mathrm{~d} t}(0)\right\|^{2}+\|\nabla f(x(0))\|^{2} h^{2}\right)-\frac{1}{4}\|\nabla f(x(0))\|^{2}+2 M_{2} h \cdot\|\nabla f(x(0))\| \cdot\left\|\frac{\mathrm{d} x}{\mathrm{~d} t}(0)\right\| \\
\leq & 2 M_{2}\left(\left\|\frac{\mathrm{~d} x}{\mathrm{~d} t}(0)\right\|^{2}+\|\nabla f(x(0))\|^{2} h^{2}\right)-\frac{1}{4}\|\nabla f(x(0))\|^{2}+2 \cdot\left(\frac{1}{8}\|\nabla f(x(0))\|\right)^{2}+2 \cdot\left(16 M_{2} h\left\|\frac{\mathrm{~d} x}{\mathrm{~d} t}(0)\right\|\right)^{2} \\
= & \left(2 M_{2}+512 M_{2}^{2} h^{2}\right)\left\|\frac{\mathrm{d} x}{\mathrm{~d} t}(0)\right\|^{2}-\left(\frac{1}{4}-2 M_{2} h^{2}-\frac{1}{32}\right)\|\nabla f(x(0))\|^{2} \\
\leq & 10 M_{2} \cdot\left\|\frac{\mathrm{~d} x}{\mathrm{~d} t}(0)\right\|^{2}-\frac{1}{8}\|\nabla f(x(0))\|^{2}
\end{aligned}
$$

where we used that $h \leq \frac{1}{8 \sqrt{M_{2}}}$ at the last two equations.
Since $\frac{\mathrm{d} x}{\mathrm{~d} t}(0)$ is sampled from normal distribution, we have that

$$
\begin{aligned}
\mathbb{E}[f(x(h))] & \leq f(x(0))+\int_{0}^{h}(h-t) \frac{\mathrm{d}^{2}}{\mathrm{~d} t^{2}} f(x(t)) \mathrm{d} t \\
& =f(x(0))+5 M_{2} \cdot h^{2} \cdot \mathbb{E}\left[\left\|\frac{\mathrm{~d} x}{\mathrm{~d} t} x(0)\right\|^{2}\right]-\frac{h^{2}}{16}\|\nabla f(x(0))\|^{2} \\
& =f(x(0))+5 M_{2} \cdot h^{2} \cdot d-\frac{h^{2}}{16}\|\nabla f(x(0))\|^{2}
\end{aligned}
$$

where the last step follows from $\mathbb{E}\left[\left\|\frac{\mathrm{d} x}{\mathrm{~d} t}(0)\right\|^{2}\right]=d$.
For the ODE starting from $x^{(k)}$, we have that

$$
\mathbb{E}[f(x(h))] \leq f\left(x^{(k)}\right)-\frac{h^{2}}{16}\left\|\nabla f\left(x^{(k)}\right)\right\|^{2}+O\left(M_{2} \cdot d \cdot h^{2}\right)
$$

To compare $f(x(h))$ and $f\left(x^{(k+1)}\right)$, we note that the distance between $x^{(k+1)}$ and $x(h)$ is less than $\bar{\epsilon}$ in $\ell_{2}$ norm. Using (4.26) and the fact $\nabla^{2} f \preceq M_{2} I$, the function value changed by at most

$$
\bar{\epsilon} \cdot\left\|\nabla f\left(x^{(k+1)}\right)\right\|+M_{2} \cdot \bar{\epsilon}^{2}=2 \bar{\epsilon} \cdot\left\|\nabla f\left(x^{(k)}\right)\right\|+O\left(\bar{\epsilon} M_{2} \sqrt{d h}+\bar{\epsilon}^{2} M_{2}\right)
$$

with high probability. Hence, we have

$$
\begin{aligned}
\mathbb{E}\left[f\left(x^{(k+1)}\right)\right] & \leq f\left(x^{(k)}\right)+2 \bar{\epsilon}\left\|\nabla f\left(x^{(k)}\right)\right\|-\frac{h^{2}}{16}\left\|\nabla f\left(x^{(k)}\right)\right\|^{2}+O\left(\bar{\epsilon} M_{2} \sqrt{d} h+\bar{\epsilon}^{2} M_{2}+M_{2} d h^{2}\right) \\
& \leq f\left(x^{(k)}\right)+2 \bar{\epsilon}\left\|\nabla f\left(x^{(k)}\right)\right\|-\frac{h^{2}}{16}\left\|\nabla f\left(x^{(k)}\right)\right\|^{2}+O\left(\bar{\epsilon}^{2} M_{2}+M_{2} d h^{2}\right) \\
& \leq f\left(x^{(k)}\right)+2 \bar{\epsilon}\left\|\nabla f\left(x^{(k)}\right)\right\|-\frac{h^{2}}{16}\left\|\nabla f\left(x^{(k)}\right)\right\|^{2}+O\left(\frac{\bar{\epsilon}^{2}}{h^{2}}+M_{2} d h^{2}\right) \\
& \leq f\left(x^{(k)}\right)-\frac{h^{2}}{32}\left\|\nabla f\left(x^{(k)}\right)\right\|^{2}+O\left(\frac{\bar{\epsilon}^{2}}{h^{2}}+M_{2} d h^{2}\right)
\end{aligned}
$$

where the step follows from $2 a b \leq a^{2}+b^{2}$, the third step follows from our choice of $h$, the second last step follows by $2 \bar{\epsilon}\left\|\nabla f\left(x^{(k)}\right)\right\| \leq \frac{h^{2}}{64}\left\|\nabla f\left(x^{(k)}\right)\right\|^{2}+64 \frac{\bar{\epsilon}^{2}}{h^{2}}$.

Summing $k$ from 0 to $N-1$, we have

$$
\sum_{k=0}^{N-1} \mathbb{E}\left[f\left(x^{(k+1)}\right)-f\left(x^{(k)}\right)+\frac{h^{2}}{32}\left\|\nabla f\left(x^{(k)}\right)\right\|^{2}\right] \leq N \cdot O\left(\frac{\bar{\epsilon}^{2}}{h^{2}}+M_{2} d h^{2}\right)
$$

Using $f\left(x^{(N)}\right) \geq \min _{x} f(x)$ and reorganizing the terms gives the desired result.

Remark 4.4.1. Both Lemma 4.4.1 and Lemma 4.4.2 do not need convexity.

### 4.4.2 Sampling

Now, we can apply our algorithm for second-order ODEs to the HMC dynamic. To control the gradient of the initial point in a simple way, we start at the algorithm at a local minimum of $f$.

Theorem 4.1.4. (Strongly Convex). Given a function $f$ such that $0 \prec m_{2} \cdot I \preceq \nabla^{2} f(x) \preceq M_{2} \cdot I$ for all $x \in \mathbb{R}^{d}$ and $0<\epsilon<\sqrt{d}$. Starting $x^{(0)}$ at the minimum of $f$, we can find a random point $X$ whose Wasserstein distance to $Y$ drawn from the density proportional to $e^{-f}$ satisfies

$$
W_{2}(X, Y) \leq \frac{\epsilon}{\sqrt{m_{2}}}
$$

using $O\left(\kappa^{1.5} \log \left(\frac{d}{\epsilon}\right)\right)$ iterations where $\kappa=\frac{M_{2}}{m_{2}}$. Each iteration takes $O\left(\frac{\kappa^{\frac{1}{4}} d^{\frac{3}{2}}}{\epsilon} \log \left(\frac{\kappa d}{\epsilon}\right)\right)$ time and $O\left(\frac{\kappa^{\frac{1}{4}} d^{\frac{1}{2}}}{\epsilon} \log \left(\frac{\kappa d}{\epsilon}\right)\right)$ evaluations of $\nabla f$, amortized over all iterations.

Proof. The number of iterations follows from Theorem 4.3.2 with

$$
\begin{equation*}
h=\frac{m_{2}^{1 / 4}}{16000 M_{2}^{3 / 4}} \tag{4.27}
\end{equation*}
$$

To approximate the HMC dynamic, we apply the ODE algorithm (Theorem 4.2.4). Now, we estimate the parameters in Theorem 4.2.4. Note that $L_{1}=0, L_{2}=M_{2}, L=\sqrt{M_{2}}$, $T=h, L T \leq 1 / 16000, \epsilon_{\mathrm{ODE}}=\frac{\epsilon \cdot \sqrt{m_{2}} h^{2}}{16}$. Hence, if the solution is approximated by a $k$-piece degree 2 polynomial, we can find it in $O(d k) \cdot \log \left(\frac{C}{\epsilon_{\mathrm{ODE}}}\right)$ time and $O(k) \cdot \log \left(\frac{C}{\epsilon_{\mathrm{ODE}}}\right)$ evaluations where

$$
C=O\left(h^{2}\right)\left\|\nabla f\left(v_{0}\right)\right\|+h\left\|v_{1}\right\|
$$

with $v_{0}$ and $v_{1}$ are the initial point and initial velocity of the dynamic.

Finally, Lemma 4.4.1 shows that the ODE can be approximated using a

$$
k \stackrel{\text { def }}{=} \frac{2 \cdot M_{2} h^{3}\left(\left\|v_{1}\right\|_{2}+\left\|\nabla f\left(v_{0}\right)\right\|_{2} \cdot h\right)}{\epsilon /\left(\sqrt{m_{2}} \cdot \kappa^{3 / 2}\right)}
$$

piece degree 2 polynomials where $v_{0}$ is the initial point and $v_{1}$ is the initial velocity. Since $v_{1}$ is sampled from normal distribution, we have that

$$
k \lesssim \frac{1}{\epsilon}\left(\kappa^{\frac{1}{4}} \sqrt{d}+\frac{\left\|\nabla f\left(v_{0}\right)\right\|_{2}}{\sqrt{M_{2}}}\right)
$$

in expectation.
Now, we apply Lemma 4.4.2 with $\bar{\epsilon}=\sqrt{m_{2}} h^{2} \epsilon / 16$ and use the fact that $f\left(x^{(0)}\right)-$ $\min _{x} f(x) \leq \frac{\left\|\nabla f\left(x^{(0)}\right)\right\|_{2}^{2}}{m_{2}}$, we have

$$
\begin{aligned}
\mathbb{E}\left[\frac{1}{N} \sum_{k=0}^{N-1}\left\|\nabla f\left(x^{(k)}\right)\right\|_{2}^{2}\right] & \lesssim \frac{\left\|\nabla f\left(x^{(0)}\right)\right\|_{2}^{2}}{m_{2} h^{2} N}+M_{2} d+\frac{\bar{\epsilon}^{2}}{h^{4}} \\
& \lesssim \frac{\left\|\nabla f\left(x^{(0)}\right)\right\|_{2}^{2}}{m_{2} h^{2} N}+M_{2} d+m_{2} \epsilon^{2} \\
& \lesssim\left\|\nabla f\left(x^{(0)}\right)\right\|_{2}^{2}+M_{2} d+m_{2} \epsilon^{2} \\
& \lesssim\left\|\nabla f\left(x^{(0)}\right)\right\|_{2}^{2}+M_{2} d
\end{aligned}
$$

where the third step follows from $N \geq \frac{1}{m_{2} h^{2}}$, our choice of $h$ and $\epsilon$ (i.e., Eq. (4.27)), and we used that $\epsilon \leq \sqrt{d}$ at the end. Hence, the expected number of evaluations per each HMC iterations (amortized over all HMC iterations) is

$$
\begin{aligned}
O(k) \cdot \log \left(\frac{C}{\epsilon_{\mathrm{ODE}}}\right) & \lesssim \frac{1}{\epsilon}\left(\kappa^{\frac{1}{4}} \sqrt{d}+\frac{\left\|\nabla f\left(x^{(0)}\right)\right\|+\sqrt{M_{2} d}}{\sqrt{M_{2}}}\right) \log \left(\frac{h^{2}\left(\left\|\nabla f\left(x^{(0)}\right)\right\|+\sqrt{M_{2} d}\right)+h \sqrt{d}}{\epsilon_{\mathrm{ODE}}}\right) \\
& \lesssim \frac{1}{\epsilon}\left(\kappa^{\frac{1}{4}} \sqrt{d}+\frac{\left\|\nabla f\left(x^{(0)}\right)\right\|}{\sqrt{M_{2}}}\right) \log \left(\frac{1}{\epsilon}\left(\frac{\left\|\nabla f\left(x^{(0)}\right)\right\|}{\sqrt{m_{2}}}+\kappa^{3 / 4} \sqrt{d}\right)\right)
\end{aligned}
$$

where the last step follows from our choice of $\epsilon_{\mathrm{ODE}}=\frac{\epsilon \cdot \sqrt{m_{2}} h^{2}}{16}$ and our choice of $h=\frac{m_{2}^{1 / 4}}{16000 M_{2}^{3 / 4}}$. Since we start at the minimum of $f,\left\|\nabla f\left(x^{(0)}\right)\right\|=0$ and this gives the expected number of evaluations.

Similarly, we have the bound for expected time. This completes the proof.

### 4.5 Sampling from Incoherent Logistic Loss Functions and More

In this section we prove Theorem 4.5.6. Since the function

$$
f(x)=\sum_{i=1}^{n} \phi_{i}\left(a_{i}^{\top} x\right)+\frac{m_{2}}{2}\|x\|_{2}^{2}
$$

the HMC dynamic for sampling $e^{-f(x)}$ is given by

$$
\frac{\mathrm{d}^{2}}{\mathrm{~d} t^{2}} x(t)=-\nabla f(x(t))=-A^{\top} \phi^{\prime}(A x)-m_{2} x
$$

where the $i^{\text {th }}$ row of $A \in \mathbb{R}^{n \times d}$ is $a_{i}^{\top}, \forall i \in[n]$ and $\phi^{\prime}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n}$ is defined by

$$
\phi^{\prime}(s)=\left(\frac{\mathrm{d}}{\mathrm{~d} s_{1}} \phi_{1}\left(s_{1}\right), \frac{\mathrm{d}}{\mathrm{~d} s_{2}} \phi_{2}\left(s_{2}\right), \cdots, \frac{\mathrm{d}}{\mathrm{~d} s_{n}} \phi_{n}\left(s_{n}\right)\right) .
$$

To simplify the proof, we let $s(t)=A x(t)$. Then, $s$ satisfies the equation

$$
\begin{equation*}
\frac{\mathrm{d}^{2}}{\mathrm{~d} t^{2}} s(t)=F(s(t)) \quad \text { where } \quad F(s)=-A A^{\top} \phi^{\prime}(s)-m_{2} s \tag{4.28}
\end{equation*}
$$

Ignoring the term $m_{2} s, F$ consists of two parts the first part $-A A^{\top}$ is linear and the second part is decoupled in each variable. This structure allows us to study the dynamic $s$ easily. In this section, we discuss how to approximate the solution of (4.28) using the collocation method.

The proof consists of (a) bounding $\|s(t)\|_{\infty}$, (b) bounding Lipschitz constant of $F$ and (c) showing that $s(t)$ can be approximated by a polynomial.

### 4.5.1 $\quad \ell_{\infty}$ bound of the dynamic

Lemma 4.5.1 ( $\ell_{\infty}$ bound of the dynamic). Let $x^{(j)}$ be the $j^{\text {th }}$ iteration of the HMC dynamic defined in Algorithm 4.1 with

$$
f(x)=\sum_{i=1}^{n} \phi_{i}\left(a_{i}^{\top} x\right)+\frac{m_{2}}{2}\|x\|_{2}^{2}
$$

Assume that $m_{2} \cdot I \preceq \nabla^{2} f(x) \preceq M_{2} \cdot I$ for all $x,\left|\phi^{\prime}(s)\right| \leq M$ for all $s$, and that $\tau=$ $\left\|A A^{\top}\right\|_{\infty \rightarrow \infty}$. Let $s^{(j)}=A x^{(j)}$. Suppose that the step size $T \leq \frac{1}{2 \sqrt{m_{2}}}$, we have that

$$
\max _{j \in[N]}\left\|s^{(j)}-s^{(0)}\right\|_{\infty}=O\left(\sqrt{\frac{\tau}{m_{2}}}+\frac{\tau M}{m_{2}}\right) \cdot \log (d N / \eta)
$$

with probability at least $1-\eta$.

Proof. We ignore the index of iteration $(j)$ and focus on how $s$ changes within each ODE first.

For any $0 \leq t \leq T$ and for any $i$, we have that

$$
\begin{equation*}
s_{i}(t)=s_{i}(0)+s_{i}^{\prime}(0) t-\int_{0}^{t}(t-\ell)\left(A A^{\top} \phi^{\prime}(s(\ell))\right)_{i} \mathrm{~d} \ell-m_{2} \int_{0}^{t}(t-\ell) s_{i}(\ell) \mathrm{d} \ell \tag{4.29}
\end{equation*}
$$

Using $\tau=\left\|A A^{\top}\right\|_{\infty \rightarrow \infty}$ and $\left|\phi^{\prime}(s)\right| \leq M$, we have

$$
\begin{aligned}
\left|s_{i}(t)-s_{i}(0)\right| & \leq T \cdot\left|s_{i}^{\prime}(0)\right|+\int_{0}^{T}(T-\ell) \cdot \tau \cdot M \mathrm{~d} \ell+m_{2} \int_{0}^{t}(t-\ell)\left|s_{i}(\ell)-s_{i}(0)\right| \mathrm{d} \ell+\frac{1}{2} m_{2} T^{2}\left|s_{i}(0)\right| \\
& =\frac{1}{2} m_{2} T^{2} \cdot\left|s_{i}(0)\right|+T \cdot\left|s_{i}^{\prime}(0)\right|+\frac{1}{2} T^{2} \cdot \tau M+m_{2} \int_{0}^{t}(t-\ell)\left|s_{i}(\ell)-s_{i}(0)\right| \mathrm{d} \ell
\end{aligned}
$$

Solving this integral inequality (Lemma 4.6.2), we get

$$
\begin{align*}
\left|s_{i}(t)-s_{i}(0)\right| & \leq\left(\frac{1}{2} m_{2} T^{2} \cdot\left|s_{i}(0)\right|+T \cdot\left|s_{i}^{\prime}(0)\right|+\frac{1}{2} T^{2} \cdot \tau M\right) \cdot \cosh \left(\sqrt{m_{2}} t\right) \\
& \leq \frac{1}{4}\left|s_{i}(0)\right|+2 T \cdot\left|s_{i}^{\prime}(0)\right|+T^{2} \cdot \tau M \tag{4.30}
\end{align*}
$$

where we used that $t \leq T \leq \frac{1}{2 \sqrt{m_{2}}}$.

Using this estimate back to (4.29), we have

$$
\begin{aligned}
& \left|s_{i}(T)-s_{i}(0)-s_{i}^{\prime}(0) T+\frac{1}{2} m_{2} T^{2} s_{i}(0)\right| \\
\leq & \int_{0}^{T}(T-\ell)\left|A A^{\top} \phi^{\prime}(s(\ell))_{i}\right| \mathrm{d} \ell+m_{2} \int_{0}^{T}(T-\ell)\left|s_{i}(\ell)-s_{i}(0)\right| \mathrm{d} \ell \\
\leq & \int_{0}^{T}(T-\ell) \tau M \mathrm{~d} \ell+m_{2} \int_{0}^{T}(T-\ell)\left(\frac{1}{4}\left|s_{i}(0)\right|+2 T \cdot\left|s_{i}^{\prime}(0)\right|+T^{2} \cdot \tau M\right) \mathrm{d} \ell \\
= & \frac{1}{2} T^{2} \cdot \tau M+\frac{1}{2} m_{2} T^{2} \cdot\left(\frac{1}{4}\left|s_{i}(0)\right|+2 T \cdot\left|s_{i}^{\prime}(0)\right|+T^{2} \cdot \tau M\right) \\
\leq & \frac{1}{2} T^{2} \cdot \tau M+\frac{1}{2} m_{2} T^{2} \cdot\left(\frac{1}{4}\left|s_{i}(0)\right|+\frac{1}{\sqrt{m_{2}}} \cdot\left|s_{i}^{\prime}(0)\right|+\frac{1}{4 m_{2}} \cdot \tau M\right) \\
= & \frac{5}{8} T^{2} \cdot \tau M+\frac{1}{8} m_{2} T^{2}\left|s_{i}(0)\right|+\frac{1}{2} \sqrt{m_{2}} T^{2} \cdot\left|s_{i}^{\prime}(0)\right|
\end{aligned}
$$

where the second step follows from $\left|A A^{\top} \phi^{\prime}(s(\ell))_{i}\right| \leq \tau M$ and (4.30), the third step follows from $\int_{0}^{T}(T-\ell) \mathrm{d} \ell=\frac{1}{2} T^{2}$, the fourth step follows from $T \leq \frac{1}{2 \sqrt{m_{2}}}$.

Note that we bounded how much each iteration of the HMC dynamic can change the solution. Writing it differently, we have

$$
s_{i}^{(j+1)}=\left(1-\frac{1}{2} m_{2} T^{2}\right) s_{i}^{(j)}+s_{i}^{(j)^{\prime}}(0) \cdot T+\beta
$$

where

$$
|\beta| \leq \frac{1}{4} m_{2} T^{2}\left|s_{i}^{(j)}(0)\right|+\sqrt{m_{2}} T^{2} \cdot\left|s_{i}^{(j)^{\prime}}(0)\right|+\tau M \cdot T^{2}
$$

Note that $s_{i}^{(j)^{\prime}}(0) \sim \mathcal{N}\left(0,\left(A^{\top} A\right)_{i}\right)$ and that

$$
\lambda_{\max }\left(A^{\top} A\right)=\lambda_{\max }\left(A A^{\top}\right) \leq\left\|A A^{\top}\right\|_{\infty \rightarrow \infty}=\tau
$$

Therefore, $s_{i}^{(j)^{\prime}}(0) T \sim \alpha_{i} \cdot \mathcal{N}(0,1)$ with $0 \leq \alpha_{i} \leq \sqrt{\tau} T$. Now, we simplify the dynamic to

$$
s_{i}^{(j+1)}=(1-\delta) s_{i}^{(j)}+\alpha_{i} N^{(j)}+\beta
$$

where $\delta:=\frac{1}{2} m_{2} T^{2}, \alpha_{i} \leq a:=\sqrt{\tau} T$, and

$$
|\beta| \leq \frac{\delta}{2}\left|s_{i}^{(j)}(0)\right|+b\left|N^{(j)}\right|+c
$$

with $b:=\sqrt{m_{2} \tau} T^{2}$ and $c:=\tau M \cdot T^{2}$.
Applying Lemma 4.5.2, we have that

$$
\operatorname{Pr}\left[\max _{j \in[N]}\left|s_{i}^{(j)}-s_{i}^{(0)}\right| \geq C \cdot\left(\frac{a}{\sqrt{\delta}}+\frac{c+b}{\delta}\right) \log (N / \epsilon)\right] \leq \epsilon
$$

for some constant $C$. Taking union bound over $i \in[d]$, the bound follows from the calculation

$$
\frac{a}{\sqrt{\delta}}+\frac{c+b}{\delta}=\frac{\sqrt{2 \tau} T}{\sqrt{m_{2}} T}+\frac{\tau M \cdot T^{2}+\sqrt{m_{2} \tau} \cdot T^{2}}{2 m_{2} T^{2}} \gtrsim \sqrt{\frac{\tau}{m_{2}}}+\frac{\tau M}{m_{2}} .
$$

Lemma 4.5.2 (Bounding the Martingale). Let $X^{(i)}$ be a sequence of random variable such that

$$
X^{(i+1)}=(1-\delta) X^{(i)}+\alpha_{i} N^{(i)}+\beta
$$

where $N^{(i)} \sim \mathcal{N}(0,1)$ are independent, $\alpha_{i} \leq a, \beta \leq \frac{\delta}{2}\left|X^{(i)}\right|+b\left|N^{(i)}\right|+c$ with positive $a, b, c$ and $0<\delta \leq 1$. For some constant universal $C>0$, we have that

$$
\operatorname{Pr}\left[\max _{i \in[k]}\left|X^{(i)}-X^{(0)}\right| \geq C \cdot\left(\frac{a}{\sqrt{\delta}}+\frac{c+b}{\delta}\right) \log (k / \epsilon)\right] \leq \epsilon
$$

for any $0<\epsilon<1$.

Proof. We will first show that $X^{(i)}$ cannot grow to large. The proof of the other direction is similar. Consider the potential $\Phi^{(i)}=\mathbb{E}\left[e^{\lambda X^{(i)}}\right]$. Note that

$$
\begin{aligned}
\Phi^{(i+1)} & \leq \mathbb{E}\left[e^{\lambda\left((1-\delta) X^{(i)}+\alpha_{i} N^{(i)}+b\left|N^{(i)}\right|+c+\frac{\delta}{2}\left|X^{(i)}\right|\right)}\right] \\
& \leq e^{\lambda c} \mathbb{E}\left[e^{\lambda(1-\delta) X^{(i)}+\frac{\delta}{2}\left|X^{(i)}\right|}\right] \mathbb{E}\left[e^{\lambda\left(\alpha_{i} N^{(i)}+b\left|N^{(i)}\right|\right)}\right],
\end{aligned}
$$

where we used that $N^{(i)}$ are independent. Picking $\lambda \leq \frac{1}{a+b}$ and using $N^{(i)} \sim \mathcal{N}(0,1)$ and $\left|\alpha_{i}\right| \leq a$, we have that

$$
\mathbb{E}\left[e^{\lambda\left(\alpha_{i} N^{(i)}+b\left|N^{(i)}\right|\right)}\right] \leq e^{O\left(\lambda b+\lambda^{2} a^{2}\right)}
$$

Therefore, for any $\eta>0$, we have

$$
\begin{aligned}
\Phi^{(i+1)} & \leq e^{O\left(\lambda c+\lambda b+\lambda^{2} a^{2}\right)} \cdot\left(\mathbb{E}\left[e^{\lambda(1-\delta) X^{(i)}+\frac{\delta}{2}\left|X^{(i)}\right|} 1_{X^{(i)} \leq \eta}\right]+\mathbb{E}\left[e^{\lambda\left(1-\frac{\delta}{2}\right) X^{(i)}} 1_{X^{(i)}>\eta}\right]\right) \\
& \leq e^{O\left(\lambda c+\lambda b+\lambda^{2} a^{2}\right)} \cdot\left(e^{\lambda \eta}+\mathbb{E}\left[e^{\lambda X^{(i)}-\frac{\lambda}{2} \delta \eta}\right]\right) \\
& \leq e^{O\left(\lambda c+\lambda b+\lambda^{2} a^{2}+\lambda \eta\right)}+e^{O\left(\lambda c+\lambda b+\lambda^{2} a^{2}\right)-\frac{\lambda}{2} \delta \eta} \Phi^{(i)} .
\end{aligned}
$$

Choose $\eta=\Theta\left(\frac{c+b+\lambda a^{2}}{\delta}\right)$ such that $O\left(\lambda c+\lambda b+\lambda^{2} a^{2}\right)-\frac{\lambda}{2} \delta \eta \leq 0$. Hence, we have

$$
\Phi^{(i+1)} \leq e^{O\left(\lambda c+\lambda b+\lambda^{2} a^{2}+\lambda \eta\right)}+\Phi^{(i)} \leq e^{O\left(\frac{\lambda c+\lambda b+\lambda^{2} a^{2}}{\delta}\right)}+\Phi^{(i)}
$$

Picking an appropriate $\lambda$, we have the result.

### 4.5.2 Lipschitz constant of $F$

Now, we bound the Lipschitz constant of function $F$.

Lemma 4.5.3 (Lipschitz bound of function $F$ ). Let $L_{\phi^{\prime}}$ be the Lipschitz constant of $\phi^{\prime}$, i.e.,

$$
\left\|\phi^{\prime}\left(s_{1}\right)-\phi^{\prime}\left(s_{2}\right)\right\|_{\infty} \leq L_{\phi^{\prime}}\left\|s_{1}-s_{2}\right\|_{\infty}, \forall s_{1}, s_{2} .
$$

The function $F$ defined in (4.28) has Lipschitz constant $\left(L_{\phi^{\prime}} \tau+m_{2}\right)$ in $\ell_{\infty}$ norm where $\tau=\left\|A A^{\top}\right\|_{\infty \rightarrow \infty}$.

Proof. Note that

$$
\begin{aligned}
\left\|F\left(s_{1}\right)-F\left(s_{2}\right)\right\|_{\infty} & =\left\|A A^{\top}\left(\phi^{\prime}\left(s_{1}\right)-\phi^{\prime}\left(s_{2}\right)\right)\right\|_{\infty}+m_{2}\left\|s_{1}-s_{2}\right\|_{\infty} \\
& \leq \tau \cdot\left\|\phi^{\prime}\left(s_{1}\right)-\phi^{\prime}\left(s_{2}\right)\right\|_{\infty}+m_{2}\left\|s_{1}-s_{2}\right\|_{\infty} \\
& \leq\left(\tau \cdot L_{\phi^{\prime}}+m_{2}\right) \cdot\left\|s_{1}-s_{2}\right\|_{\infty}
\end{aligned}
$$

where the second step follows by $\left\|A A^{\top}\right\|_{\infty \rightarrow \infty}=\tau$, and second step follows by $\phi^{\prime}$ is $L_{\phi^{\prime}}$ Lipschitz function.

Remark 4.5.1. If we think of the role of $F$ in our second order ODE (4.28), then $F$ is in fact independent of $\frac{\mathrm{d} s}{\mathrm{~d} t}$. Therefore $L_{1}=0$ and $L_{2}=L_{\phi^{\prime}} \tau+m_{2}$.

### 4.5.3 Existence of low-degree solutions

Next, we establish bounds on the radius up to which the solution to the ODE (4.28) have a low-degree polynomial approximation.

Lemma 4.5.4 (Low-degree polynomial approximation). Assume that $\left\|A A^{\top}\right\|_{\infty \rightarrow \infty}=\tau$ and that for all $i, \phi_{i}^{\prime}$ has Cauchy estimate $M$ with radius r, i.e.,

$$
\forall l \geq 0, \forall a \in \mathbb{R},\left|\left(\phi_{i}\right)^{(l+1)}(a)\right| \leq M \cdot l!\cdot r^{-l} .
$$

Let $s^{*}(t) \in \mathbb{R}^{n}$ denote the solution of the ODE (4.28).
For any

$$
0 \leq T \leq \frac{r}{4}\left(\left(M \tau r+m_{2} r\left(r+\|s(0)\|_{\infty}\right)\right)^{1 / 2}+\left\|s^{\prime}(0)\right\|_{\infty}\right)^{-1}
$$

and any $0<\epsilon<1$, there is a degree $D=2\left\lceil 4+\log _{2}(1 / \epsilon)\right\rceil$ polynomial $q: \mathbb{R} \rightarrow \mathbb{R}^{n}$ such that

$$
q(0)=s^{*}(0), q^{\prime}(0)=s^{* \prime}(0), \text { and }\left\|\frac{\mathrm{d}^{2}}{\mathrm{~d} t^{2}} q(t)-\frac{\mathrm{d}^{2}}{\mathrm{~d} t^{2}} s^{*}(t)\right\|_{\infty} \leq \frac{\epsilon \cdot r}{T^{2}}, \forall t \in[0, T]
$$

First, we verify the condition in Lemma 4.6.6.
Lemma 4.5.5. (Bounding derivatives of $F$ ). Under the same assumptions in Lemma 4.5.4, we have

$$
\left\|D^{k} F(s)\left[\Delta_{1}, \Delta_{2}, \cdots, \Delta_{k}\right]\right\|_{\infty} \leq g^{(k)}(0) \cdot \prod_{j=1}^{k}\left\|\Delta_{j}\right\|_{\infty}, \forall k \geq 0, \forall \Delta_{1}, \cdots, \Delta_{k}
$$

where

$$
g(x)=\frac{\tau \cdot M+m_{2} \cdot\left(r+\|s\|_{\infty}\right)}{1-r^{-1} x}
$$

Proof. Recall that $\phi^{\prime}(x)=\left(\phi_{1}^{\prime}\left(x_{1}\right), \phi_{2}^{\prime}\left(x_{2}\right), \cdots, \phi_{n}^{\prime}\left(x_{n}\right)\right), \forall x \in \mathbb{R}^{n}$. We have

$$
\begin{aligned}
& \left\|A A^{\top} D^{k} \phi^{\prime}(s)\left[\Delta_{1}, \Delta_{2}, \cdots, \Delta_{k}\right]\right\|_{\infty} \\
\leq & \tau \cdot\left\|D^{k} \phi^{\prime}(s)\left[\Delta_{1}, \Delta_{2}, \cdots, \Delta_{k}\right]\right\|_{\infty} \\
\leq & \tau \cdot \max _{i \in[n]}\left|\phi_{i}^{(k+1)}\right| \cdot\left|\Delta_{1, i}\right| \cdot\left|\Delta_{2, i}\right| \cdots\left|\Delta_{k, i}\right| \\
\leq & \tau \cdot \max _{i \in[n]}\left|\phi_{i}^{(k+1)}\right| \prod_{j=1}^{k}\left\|\Delta_{j}\right\|_{\infty} \\
\leq & \tau \cdot M \cdot k!\cdot r^{-k} \cdot \prod_{j=1}^{k}\left\|\Delta_{j}\right\|_{\infty}
\end{aligned}
$$

Next, the derivatives of the $m_{2} s$ term in $F(s)=-A A^{\top} \phi^{\prime}(s)-m_{2} s$ can be bounded by the derivatives of $x \rightarrow m_{2}\left(\|s\|_{\infty}+x\right)$, which then can be bounded by the derivatives of $x \rightarrow m_{2}\left(r+\|s\|_{\infty}\right) /\left(1-r^{-1} x\right)$. This explains the second part of the function $g$.

Proof of Lemma 4.5.4. Theorem 4.6.4 and Lemma 4.5.5 shows that

$$
\begin{equation*}
\left\|s^{*(k)}(0)\right\|_{\infty} \leq \frac{k!\alpha^{k}}{r^{-1}}=r \cdot k!\cdot \alpha^{k} \tag{4.31}
\end{equation*}
$$

where

$$
\alpha=\max \left(\left(\frac{4}{3} \cdot\left(M \tau+m_{2}\left(r+\|s(0)\|_{\infty}\right)\right) \cdot r^{-1}\right)^{1 / 2}, 2 \cdot\left\|s^{\prime}(0)\right\|_{\infty} r^{-1}\right)
$$

Since $s^{*}$ is real analytic at 0 (Theorem 4.6.4), around $t=0$, we have

$$
s^{*}(t)=\sum_{k=0}^{\infty} \frac{s^{*(k)}(0)}{k!} t^{k}
$$

Apply Theorem 4.6.4 repeatedly at every $t$ such that $s^{*}(t)$ is defined, we can show that the above equation holds as long as the right hand side converges.

Let $q(t)=\sum_{k=0}^{D} s^{*(k)}(0) t^{k}$. Then, we have that

$$
\begin{aligned}
\left\|\frac{\mathrm{d}^{2}}{\mathrm{~d} t^{2}} q(t)-\frac{\mathrm{d}^{2}}{\mathrm{~d} t^{2}} s^{*}(t)\right\|_{\infty} & \leq\left\|\sum_{k=D+1}^{\infty} \frac{s^{*(k)}(0)}{(k-2)!} t^{k-2}\right\|_{\infty} \\
& \leq r \cdot \sum_{k=D+1}^{\infty} \frac{k!}{(k-2)!} \alpha^{k} t^{k-2} \\
& \leq \frac{r}{T^{2}} \cdot \sum_{k=D+1}^{\infty} \frac{(k-1) k}{2^{k}} \\
& =\frac{r}{T^{2}} \cdot 2^{-D}\left(D^{2}+3 D+4\right) \\
& \leq 16 \frac{2^{-D / 2} r}{T^{2}}
\end{aligned}
$$

where we used (4.31) at the second inequality, $\alpha T \leq \frac{1}{2}$ at the third inequality, and $D \geq 1$ at the fourth inequality.

### 4.5.4 Main result

Theorem 4.5.6 (Formal version of Theorem 4.1.2). Let $A=\left[a_{1} ; a_{2} ; \cdots ; a_{n}\right] \in \mathbb{R}^{n \times d}, \phi_{i}$ : $\mathbb{R} \rightarrow \mathbb{R}$ be a function, and

$$
f(x)=\sum_{i=1}^{n} \phi_{i}\left(a_{i}^{\top} x\right)+\frac{m_{2}}{2}\|x\|^{2} .
$$

Let $\tau=\left\|A A^{\top}\right\|_{\infty \rightarrow \infty}$ and suppose that

1. $f$ has $M_{2}$ Lipschitz gradient, i.e., $\nabla^{2} f(x) \preceq M_{2} \cdot I$ for all $x$,
2. $\phi_{i}^{\prime}$ has Cauchy estimate $M$ with radius $r$, i.e., $\forall i \in[n], \ell \geq 1, s \in \mathbb{R},\left|\phi_{i}^{(l+1)}(s)\right| \leq$ $M \cdot l!\cdot r^{-l}$.

Starting at $x^{(0)}$, we can output a random point $X$ such that

$$
\underset{Y \propto e^{-f}}{\mathbb{E}}\left[\|X-Y\|_{2}^{2}\right] \leq \frac{\epsilon}{\sqrt{m_{2}}}
$$

using $N \lesssim k \log \left(\frac{k}{\epsilon}\left(\frac{\left\|\nabla f\left(x^{(0)}\right)\right\|^{2}}{m_{2}}+d\right)\right)$ iterations with

$$
k \lesssim \kappa^{1.5}+\frac{M \tau}{m_{2} r} \log (d N)+\frac{\tau}{m_{2} r^{2}} \log ^{2}(d N) \quad \text { with } \quad \kappa=\frac{M_{2}}{m_{2}} .
$$

Each iteration takes $O\left(d \log ^{3}\left(\frac{1}{\delta}\right)\right)$ time and $O\left(\log ^{2}\left(\frac{1}{\delta}\right)\right)$ evaluations to the function $\phi^{\prime}$ and the matrix vector multiplications for $A$ and $A^{\top}$, with

$$
\delta=\Omega\left(\frac{1}{r}\right) \sqrt{\frac{\lambda_{\min }\left(A^{\top} A\right)}{n \cdot m_{2}}} \cdot \frac{\epsilon}{k}
$$

Proof. The proof consists of bounding the cost of each HMC step in Algorithm 4.1 and bounding the number of steps.

## Cost per iteration:

| Parameters | Value | Source |
| :--- | :--- | :--- |
| $k$ | 2 | Eq. (4.28) |
| $L_{1}$ | 0 | Lemma 4.5.3 and Remark 4.5.1 |
| $L_{2}$ | $\frac{M \tau}{r}+m_{2}$ | Lemma 4.5.3 and Remark 4.5.1 |
| $D$ | $O(\log (1 / \delta))$ | Lemma 4.5.4 |
| $C$ | $O(r)$ | Eq. (4.34) |
| $\epsilon_{\text {ODE }}$ | $O(\delta r)$ | Eq. (4.33) |

Table 4.2: Summary of parameters of Theorem 4.2.4

As we discussed in this section, we consider the ODE (4.28) instead. We will use Theorem 4.2.4 to solve the ODE (4.28). Hence, we need to bound the parameters of Theorem 4.2.4, which are summarized in Table 4.2.

Parameters $D$ and $\epsilon_{\text {ODE }}$ : Let $s$ denote its solution. Lemma 4.5.4 shows that if

$$
\begin{equation*}
h \leq \frac{r}{4}\left(\left(M \tau \cdot r+m_{2} r\left(r+\|s(0)\|_{\infty}\right)\right)^{1 / 2}+\left\|s^{\prime}(0)\right\|_{\infty}\right)^{-1}, \tag{4.32}
\end{equation*}
$$

for any $\delta>0$, there is a degree $O(\log (1 / \delta))$ polynomial $q$ such that

$$
\begin{equation*}
q(0)=s(0), \quad q^{\prime}(0)=s^{\prime}(0), \quad \text { and } \quad\left\|\frac{\mathrm{d}^{2}}{\mathrm{~d} t^{2}} q(t)-\frac{\mathrm{d}^{2}}{\mathrm{~d} t^{2}} s(t)\right\|_{\infty} \leq \frac{\delta \cdot r}{T^{2}} \text { for } t \in[0, h] . \tag{4.33}
\end{equation*}
$$

Parameter $C$ : To apply Theorem 4.2.4, we first show parameter $C \leq O(r)$ as follows:

$$
\begin{align*}
C & \lesssim h \cdot\left(h\|F(s(0))\|_{\infty}+\left\|s^{\prime}(0)\right\|_{\infty}\right) \\
& \leq h \cdot\left(h\left(\left\|A A^{\top} \phi^{\prime}(s(0))\right\|_{\infty}+\left\|m_{2} s(0)\right\|_{\infty}\right)+\left\|s^{\prime}(0)\right\|_{\infty}\right) \\
& \leq h^{2} \tau\left\|\phi^{\prime}(s(0))\right\|_{\infty}+h^{2} m_{2}\|s(0)\|_{\infty}+h\left\|s^{\prime}(0)\right\|_{\infty} \\
& \leq h^{2} \tau M+h^{2} m_{2}\|s(0)\|_{\infty}+h\left\|s^{\prime}(0)\right\|_{\infty} \\
& \leq \frac{r}{16}+\frac{r}{16}+\frac{r}{4} \leq r . \tag{4.34}
\end{align*}
$$

where the second step follows from (4.28) and triangle inequality, the third step follows from $\left\|A A^{\top}\right\|_{\infty \rightarrow \infty}=\tau$, the fourth step follows from $\phi^{\prime}$ has Cauchy estimate $M$ with radius $r$, and the last step follows from (4.32).

Now, Theorem 4.2.4 shows that if $h$ satisfies (4.32) and that

$$
\begin{equation*}
h \leq \frac{1}{16000} \frac{1}{L}=\frac{1}{16000} \frac{1}{L_{1}+\sqrt{L_{2}}}=\frac{1}{16000} \sqrt{\frac{r}{M \tau+m_{2} r}}, \tag{4.35}
\end{equation*}
$$

then, we can find $p$ such that

$$
\begin{equation*}
\|s(h)-p\|_{\infty} \leq O(\delta \cdot r) \tag{4.36}
\end{equation*}
$$

using $O\left(\log \left(\frac{1}{\delta}\right) \log \left(\frac{C}{\delta \cdot r}\right)\right)=O\left(\log ^{2}\left(\frac{1}{\delta}\right)\right)$ evaluations of $\phi_{i}^{\prime}$ and $O\left(d \log ^{2}\left(\frac{1}{\delta}\right) \log \left(\frac{C}{\delta \cdot r}\right)\right)=O\left(d \log ^{3}\left(\frac{1}{\delta}\right)\right)$ time.

To understand the condition in Eq. (4.32), we note that $s^{\prime}(0)=A v$ where $v \sim \mathcal{N}(0, I)$. Hence, $s^{* \prime}(0) \sim \mathcal{N}\left(0, A^{\top} A\right)$. Note that $\lambda_{\max }\left(A^{\top} A\right)=\lambda_{\max }\left(A A^{\top}\right) \leq\left\|A A^{\top}\right\|_{\infty \rightarrow \infty}=\tau$. Hence, we have that

$$
\left\|s^{\prime}(0)\right\|_{\infty}=O(1) \cdot \sqrt{\tau \cdot \log (d N / \eta)}
$$

with probability at least $1-\eta$ probability for all $N$ iterations. In Lemma 4.5.1, we proved that

$$
\|s(0)\|_{\infty}=O\left(\sqrt{\frac{\tau}{m_{2}}}+\frac{\tau M}{m_{2}}\right) \cdot \log (d N / \eta)
$$

with probability at least $1-\eta$ for all $N$ iterations.

Putting the bound on $\left\|s^{\prime}(0)\right\|$ and $\|s(0)\|$ into the right hand side of Eq. (4.32) gives

$$
\begin{align*}
& \frac{1}{r}\left(\sqrt{M \tau \cdot r+m_{2} r\left(r+\|s(0)\|_{\infty}\right)}+\left\|s^{\prime}(0)\right\|_{\infty}\right)  \tag{4.37}\\
\lesssim & \sqrt{\frac{M \tau}{r}}+\sqrt{m_{2}}+\sqrt{\frac{m_{2}}{r}\left(\sqrt{\frac{\tau}{m_{2}}}+\frac{\tau M}{m_{2}}\right) \log (d N / \eta)}+\frac{\sqrt{\tau}}{r} \sqrt{\log (d N / \eta)} \\
\leq & \frac{M_{2}^{3 / 4}}{m_{2}^{1 / 4}}+\left(\frac{\left(\tau m_{2}\right)^{1 / 4}}{r^{1 / 2}}+\sqrt{\frac{M \tau}{r}}+\frac{\sqrt{\tau}}{r}\right) \sqrt{\log (d N / \eta)} \\
\leq & \frac{M_{2}^{3 / 4}}{m_{2}^{1 / 4}}+\left(\frac{\sqrt{\tau}}{r} \sqrt{\log (d N / \eta)}+\frac{\sqrt{m_{2}}}{\sqrt{\log (d N / \eta)}}+\sqrt{\frac{M \tau}{r}}+\frac{\sqrt{\tau}}{r}\right) \sqrt{\log (d N / \eta)} \\
\lesssim & \frac{M_{2}^{3 / 4}}{m_{2}^{1 / 4}}+\frac{\sqrt{\tau}}{r} \log (d N / \eta)+\sqrt{\frac{M \tau}{r}} \sqrt{\log (d N / \eta)}
\end{align*}
$$

where the second step follows by Eq. (4.32), the third step follows by $\sqrt{m_{2}} \leq \frac{M_{2}^{3 / 4}}{m_{2}^{1 / 4}}$ and $\sqrt{\log (d N / \eta)} \geq 1$, the fourth step follows by $a b \leq a^{2}+b^{2}$, the fifth step follows by $\sqrt{\log (d N / \eta)} \geq 1$ and $\sqrt{m_{2}} \leq \frac{M_{2}^{3 / 4}}{m_{2}^{1 / 4}}$.

Therefore,

$$
\begin{equation*}
h=\Theta\left(\sqrt{\frac{M \tau}{r}} \sqrt{\log (d N / \eta)}+\frac{\sqrt{\tau}}{r} \log (d N / \eta)+\frac{M_{2}^{3 / 4}}{m_{2}^{1 / 4}}\right)^{-1} \tag{4.38}
\end{equation*}
$$

satisfies the condition in Eq. (4.32) and Eq. (4.35). It also satisfies the condition in Theorem 4.3.2 $\left(h \leq \frac{m_{2}^{1 / 4}}{2 M_{2}^{3 / 4}}\right)$.

Next, we note that the corresponding HMC dynamic $x^{*}(h)$ is given by

$$
x^{*}(h)=\left(A^{\top} A\right)^{-1} A^{\top} s^{*}(h) .
$$

Let $p$ be the approximate of $s^{*}(h)$ we find using Theorem 4.3.2 and $q=\left(A^{\top} A\right)^{-1} A^{\top} p$, then,
we have

$$
\begin{aligned}
\left\|x^{*}(h)-q\right\|_{2} & =\left\|\left(A^{\top} A\right)^{-1} A^{\top} s^{*}(h)-\left(A^{\top} A\right)^{-1} A p\right\|_{2} \\
& \leq\left\|\left(A^{\top} A\right)^{-1} A^{\top}\right\|_{2 \rightarrow 2} \cdot\left\|s^{*}(h)-p\right\|_{2} \\
& \leq\left\|\left(A^{\top} A\right)^{-1} A^{\top}\right\|_{2 \rightarrow 2} \cdot O(\sqrt{n} \cdot r \cdot \delta) \\
& \leq\left(\lambda_{\min }\left(A^{\top} A\right)\right)^{-\frac{1}{2}} \cdot O(\sqrt{n} \cdot r \cdot \delta),
\end{aligned}
$$

where the first step follows by definition of $x^{*}(h)$ and $q$, the second step follows by definition of $\left\|\|_{2 \rightarrow 2}\right.$ norm, the third step follows from $\| s^{*}(h)-p \|_{2} \leq O(\sqrt{n} r \delta)$ (implied by (4.36) and $\left.\left\|\left\|_{2} \leq \sqrt{n}\right\|\right\|_{\infty}\right)$.

Using such $q$ as an approximation of $x^{*}(h)$ in Algorithm 4.1, Theorem 4.3.2 shows that the $W_{2}$ error of the sampled point is bounded by

$$
O\left(\left(\lambda_{\min }\left(A^{\top} A\right)\right)^{-\frac{1}{2}} \cdot \sqrt{n} \cdot r \cdot \delta\right) \leq \frac{\epsilon \cdot \theta}{2 \sqrt{m_{2}}}
$$

where $\theta=\frac{m_{2} h^{2}}{8}$ and the last step follows from picking

$$
\delta=c \cdot \frac{1}{r} \sqrt{\frac{\lambda_{\min }\left(A^{\top} A\right)}{n \cdot m_{2}}} \cdot \epsilon \cdot \theta
$$

for some small enough $c$. The cost of each iteration follows from Theorem 4.2.4 and all parameters we pick.

## Number of iterations:

Theorem 4.3.2 shows that the number of iterations is

$$
O\left(\frac{1}{\theta}\right) \cdot\left(\log \left(\frac{1}{\theta \cdot \epsilon}\right)+\log \left(\frac{\left\|\nabla f\left(x^{(0)}\right)\right\|^{2}}{m_{2}}+d\right)\right) .
$$

Finally, we bound the term $1 / \theta$ as follows

$$
\begin{aligned}
\frac{1}{\theta} & \lesssim \frac{1}{m_{2}}\left(\frac{M \tau}{r} \log \left(\frac{d N}{\eta}\right)+\frac{\tau}{r^{2}} \log ^{2}\left(\frac{d N}{\eta}\right)+\frac{M_{2}^{3 / 2}}{\sqrt{m_{2}}}\right) \\
& =\kappa^{1.5}+\frac{M \tau}{m_{2} r} \log \left(\frac{d N}{\eta}\right)+\frac{\tau}{m_{2} r^{2}} \log ^{2}\left(\frac{d N}{\eta}\right)
\end{aligned}
$$

where we used $\kappa=M_{2} / m_{2}$ and (4.38).

### 4.6 Preliminaries

### 4.6.1 Notation

For any function $f$, we define $\widetilde{O}(f)$ to be $f \cdot \log ^{O(1)}(f)$. In addition to $O(\cdot)$ notation, for two functions $f, g$, we use the shorthand $f \lesssim g$ (resp. $\gtrsim$ ) to indicate that $f \leq C g$ (resp. $\geq)$ for an absolute constant $C$.

Definition 4.6.1 ( $p \rightarrow q$ norm). Given matrix $A \in \mathbb{R}^{n \times d}$, we define $\left\|\|_{p \rightarrow q}\right.$ norm as follows

$$
\|A\|_{p \rightarrow q}=\max _{x \in \mathbb{R}^{d}} \frac{\|A x\|_{q}}{\|x\|_{p}}
$$

$\|A\|_{\infty \rightarrow \infty}$ is a special case where $p=\infty$ and $q=\infty$.

Definition 4.6.2 (Wasserstein distance). The $k$-th Wasserstein distance between two probability measure $\mu$ and $\nu$ is

$$
W_{k}(\mu, \nu)=\left(\inf _{(X, Y) \in \mathbb{C}(\mu, \nu)} \mathbb{E}\left[\|X-Y\|^{k}\right]\right)^{1 / k}
$$

where $\mathcal{C}(\mu, \nu)$ is the set of all couplings of $\mu$ and $\nu$.
Definition 4.6.3 (Cauchy's estimates). We say function $\phi$ has Cauchy estimate $M$ and radius of convergence $r$, if for all $x \in \mathbb{R}$ and for all integers $l \geq 0$

$$
\left|\phi^{(l)}(x)\right| \leq M \cdot l!\cdot r^{-l} .
$$

Lemma 4.6.1. If the number of non-zeros for each row of $A A^{\top}$ is bounded by $s$, then

$$
\lambda_{\max }\left(A A^{\top}\right) \leq\left\|A A^{\top}\right\|_{\infty \rightarrow \infty} \leq \sqrt{s} \cdot \lambda_{\max }\left(A A^{\top}\right)
$$

Proof. Let $v$ be the maximum eigenvalue of $A A^{\top}$. Then, we have that

$$
A A^{\top} v=\lambda_{\max }\left(A A^{\top}\right) \cdot v
$$

Since $\left\|A A^{\top}\right\|_{\infty \rightarrow \infty}=\max _{\|v\|_{\infty}=1}\left\|A A^{\top} v\right\|_{\infty}$, we have that $\lambda_{\max }\left(A A^{\top}\right) \leq\left\|A A^{\top}\right\|_{\infty \rightarrow \infty}$.
For the another direction,

$$
\begin{aligned}
\left\|A A^{\top}\right\|_{\infty \rightarrow \infty} & =\max _{i} \sum_{j}\left|A A^{\top}\right|_{i j} \leq \max _{i} \sqrt{s} \cdot \sqrt{\sum_{j}\left(A A^{\top}\right)_{i j}^{2}} \\
& \leq \sqrt{s} \cdot \max _{i} \max _{\|v\|_{2}=1} e_{i}^{\top} A A^{\top} v=\sqrt{s} \cdot \lambda_{\max }\left(A A^{\top}\right) .
\end{aligned}
$$

### 4.6.2 Simple ODEs

We prove two helper lemmas (Lemma 4.6.2 and 4.6.3) for the later use.

Lemma 4.6.2. Given a continuous function $v(t)$ and positive scalars $\beta, \gamma$ such that

$$
0 \leq v(t) \leq \beta+\gamma \int_{0}^{t}(t-s) v(s) \mathrm{d} s
$$

We have that $v(t) \leq \beta \cosh (\sqrt{\gamma} t)$ for all $t \geq 0$.

Proof. Let $\bar{v}(t)$ be the solution of the integral equation $\bar{v}(t)=\beta+\gamma \int_{0}^{t}(t-s) \bar{v}(s) \mathrm{d} s$. Note that $\bar{v}$ satisfies the ODE

$$
\bar{v}^{\prime \prime}(t)=\gamma \bar{v}(t), \quad \bar{v}^{\prime}(0)=0, \quad \bar{v}(0)=\beta
$$

Solving it, we have $\bar{v}(t)=\beta \cosh (\sqrt{\gamma} t)$. Hence, it suffices to prove that $v(t) \leq \bar{v}(t)$ for all $t \geq 0$.

Fix any $\epsilon>0$. We let $T$ be the supremum such that $(1+\epsilon) \bar{v}(t) \geq v(t)$ for all $0 \leq t \leq T$. Suppose $T<+\infty$. Then, we have that

$$
v(T) \leq \beta+\gamma \int_{0}^{T}(T-s) v(s) \mathrm{d} s \leq \beta+(1+\epsilon) \gamma \int_{0}^{T}(T-s) \bar{v}(s) \mathrm{d} s<(1+\epsilon) \bar{v}(T)
$$

By the continuity of $v$ and $\bar{v}$, we show that $T$ is not the supremum. This is a contradiction. Therefore, $T=+\infty$ for any $\epsilon>0$.

Lemma 4.6.3. Given a continuous function $v(t)$ and positive scalars $\beta, \gamma$ such that

$$
0 \leq v(t) \leq \beta+\gamma \int_{0}^{t} v(s) \mathrm{d} s
$$

We have that $v(t) \leq \beta e^{\gamma t}$ for all $t \geq 0$.

Proof. The proof is identical to Lemma 4.6.2.
Let $\bar{v}(t)$ be the solution of the integral equation $\bar{v}(t)=\beta+\gamma \int_{0}^{t} \bar{v}(s) \mathrm{d} s$. Note that $\bar{v}$ satisfies the ODE

$$
\bar{v}^{\prime \prime}(t)=\gamma \bar{v}(t), \quad \bar{v}^{\prime}(0)=0, \quad \bar{v}(0)=\beta
$$

Solving it, we have $\bar{v}(t)=\beta \exp (\gamma t)$. Hence, it suffices to prove that $v(t) \leq \bar{v}(t)$ for all $t \geq 0$.
Fix any $\epsilon>0$. We let $T$ be the supremum such that $(1+\epsilon) \bar{v}(t) \geq v(t)$ for all $0 \leq t \leq T$. Suppose $T<+\infty$. Then, we have that

$$
v(T) \leq \beta+\gamma \int_{0}^{T} v(s) \mathrm{d} s \leq \beta+(1+\epsilon) \gamma \int_{0}^{T} \bar{v}(s) \mathrm{d} s<(1+\epsilon) \bar{v}(T)
$$

By the continuity of $v$ and $\bar{v}$, we show that $T$ is not the supremum. This is a contradiction. Therefore, $T=+\infty$ for any $\epsilon>0$.

### 4.6.3 Cauchy Estimates and Method of Majorants

In order to prove the solution of the HMC dynamic can be approximated by a low degree polynomial, we first give a general bound the $k^{t h}$ derivative of the second order ODE $u^{\prime \prime}(t)=F(u(t))$, by reducing it to bounding derivatives of a one-variable ODE. The low degree result then follows from: first, we take the Taylor expansion of the original function; second, truncate it at a certain degree; finally we can claim that the low-degree truncation provides a good approximation to the original function.

Theorem 4.6.4. Given vectors $v_{1}, v_{0} \in \mathbb{R}^{d}$ and any norm $\|\cdot\|$ on $\mathbb{R}^{d}$. Let $U \subset \mathbb{R}^{d}$ be a neighborhood of $v_{0}$ and that $F: U \rightarrow \mathbb{R}^{d}$ is real analytic near $v_{0}$. Suppose that

$$
\left\|D^{(k)} F\left(v_{0}\right)\left[\Delta_{1}, \Delta_{2}, \cdots, \Delta_{k}\right]\right\| \leq k!\cdot a \cdot c^{k} \prod_{j=1}^{k}\left\|\Delta_{j}\right\| \text { for all } k \geq 0
$$

Then, the $O D E$

$$
\begin{equation*}
\frac{\mathrm{d}^{2}}{\mathrm{~d} t^{2}} u(t)=F(u(t)), \frac{\mathrm{d}}{\mathrm{~d} t} u(0)=v_{1}, u(0)=v_{0} \tag{4.39}
\end{equation*}
$$

has a unique real analytic solution around $t=0$. Furthermore, we have

$$
\left\|u^{(k)}(0)\right\| \leq \frac{k!\alpha^{k}}{c} \quad \forall k \geq 1
$$

with $\alpha=\max \left(\sqrt{\frac{4}{3} a c}, 2\left\|v_{1}\right\| c\right)$.

Theorem 4.6.4 involves two steps. The first step (Lemma 4.6.6) involves bounding the derivatives of the solution of the multivariate ODE (4.39) by its scalar version. The second step (Lemma 4.6.7) involves bounding the scalar ODE directly.

The first step follows directly from the majorization proof of Cauchy-Kowalevski theorem. See [vdH03] for an introduction of the method of majorants. This theorem usually
stated qualitatively without an explicit bound. For completeness, we include a proof for the first step.

Theorem 4.6.5 (Cauchy-Kowalevski theorem [Cau42, Kow75, Kro83, Nak94]). Given vectors $v_{k-1}, v_{k-2}, \cdots, v_{0} \in \mathbb{R}^{d}$. Let $U \subset \mathbb{R}^{k d+1}$ be a neighborhood of $z \stackrel{\text { def }}{=}\left(v_{k-1}, v_{k-2}, \cdots, v_{0}, 0\right)$ and that $F: U \rightarrow \mathbb{R}^{d}$ is a real analytic near $z$. Then, the $O D E$

$$
\begin{aligned}
\frac{d^{k}}{\mathrm{~d} t^{k}} x(t) & =F\left(\frac{\mathrm{~d}^{k-1}}{\mathrm{~d} t^{k-1}} x(t), \cdots, x(t), t\right) \\
\frac{d^{i}}{\mathrm{~d} t^{i}} x(0) & =v_{i}, \forall i \in\{k-1, \cdots, 1,0\}
\end{aligned}
$$

has a unique real analytic solution around $t=0$.
Lemma 4.6.6. (Bounding multivariate ODE by scalar ODE). Given vectors $v_{1}, v_{0} \in \mathbb{R}^{d}$ and any norm $\|\cdot\|$ on $\mathbb{R}^{d}$. Let $U \subset \mathbb{R}^{d}$ be a neighborhood of $v_{0}$ and that $F: U \rightarrow \mathbb{R}^{d}$ is a real analytic near $v_{0}$. Suppose that

$$
\left\|D^{(k)} F\left(v_{0}\right)\left[\Delta_{1}, \Delta_{2}, \cdots, \Delta_{k}\right]\right\| \leq f^{(k)}(0) \prod_{j=1}^{k}\left\|\Delta_{j}\right\| \text { for all } k \geq 0
$$

for some real analytic function $f$ around 0 . Then the $\operatorname{ODE}$ (4.39) has a unique real analytic solution around $t=0$. Furthermore, for any $b \geq\left\|v_{1}\right\|$, we have

$$
\left\|u^{(k)}(0)\right\| \leq \psi^{(k)}(0) \quad \forall k \geq 1
$$

where $\psi$ is the solution of the ODE

$$
\psi^{\prime \prime}(t)=f(\psi(t)), \psi^{\prime}(0)=b, \psi(0)=0 .
$$

For many functions $F$, including the HMC dynamic or complex analytic $F$, we can pick bound the derivatives of $F$ by the function $f(x)=\frac{a}{1-c x}$ for some $a$ and $c$ in Lemma 4.6.6. Therefore, we only need to give a bound on the scalar ODE for this function $f$.

Lemma 4.6.7. (Bounding scalar ODE) Let $f(x)=\frac{a}{1-c x}$ with positive $a$ and $c$. Let $\psi(t)$ denote the solution of the ODE

$$
\psi^{\prime \prime}(t)=f(\psi(t)), \psi^{\prime}(0)=b, \psi(0)=0
$$

with $b \geq 0$. Then,

$$
\psi^{(k)}(0) \leq \frac{k!\alpha^{k}}{c} \quad \forall k \geq 1
$$

with $\alpha=\max \left(\sqrt{\frac{4}{3} a c}, 2 b c\right)$.
Finally, we note that Theorem 4.6.4 follows from Lemma 4.6.6 and Lemma 4.6.7 with $f(x)=\frac{a}{1-c x}$.

### 4.7 Deferred Proof for ODE (Section 4.2)

### 4.7.1 Proof of general $k$-th order ODE

The goal of this section is to prove Theorem 4.2.3.

Proof. We define $\bar{x}(t) \in \mathbb{R}^{k d}$,

$$
\bar{x}(t)=\left(\bar{x}_{1}(t), \bar{x}_{2}(t), \cdots, \bar{x}_{k}(t)\right)=\left(c_{1} \frac{\mathrm{~d}^{k-1}}{\mathrm{~d} t^{k-1}} x(t), c_{2} \frac{\mathrm{~d}^{k-2}}{\mathrm{~d} t^{k-2}} x(t), \cdots, c_{k} x(t)\right)
$$

where $\bar{x}_{1}(t) \in \mathbb{R}^{d}, \bar{x}_{2}(t) \in \mathbb{R}^{d}, \cdots, \bar{x}_{k}(t) \in \mathbb{R}^{d}$. We define the norm on $\mathbb{R}^{k d}$ by $\|\bar{x}\|=$ $\sum_{i=1}^{k}\left\|\bar{x}_{i}\right\|$.

Then we have

$$
\frac{\mathrm{d}}{\mathrm{~d} t} \bar{x}(t)=\left(c_{1} \frac{\mathrm{~d}^{k}}{\mathrm{~d} t^{k}} x(t), c_{2} \frac{\mathrm{~d}^{k-1}}{\mathrm{~d} t^{k-1}} x(t), \cdots, c_{k} \frac{\mathrm{~d}}{\mathrm{~d} t} x(t)\right)
$$

In order to have $\bar{F}(\bar{x}(t), t)=\frac{\mathrm{d}}{\mathrm{d} t} \bar{x}(t)$, we let

$$
\bar{F}(\bar{x}(t), t)=\left(c_{1} F\left(c_{1}^{-1} \bar{x}_{1}(t), c_{2}^{-1} \bar{x}_{2}(t), \cdots, c_{k}^{-1} \bar{x}_{k}(t), t\right), c_{2} c_{1}^{-1} \bar{x}_{1}(t), c_{3} c_{2}^{-1} \bar{x}_{2}(t), \cdots, c_{k} c_{k-1}^{-1} \bar{x}_{k-1}(t)\right) .
$$

Now, we check that indeed

$$
\begin{aligned}
\frac{\mathrm{d}}{\mathrm{~d} t} \bar{x}(t)= & \left(c_{1} \frac{\mathrm{~d}^{k}}{\mathrm{~d} t^{k}} x(t), c_{2} \frac{\mathrm{~d}^{k-1}}{\mathrm{~d} t^{k-1}} x(t), \cdots, c_{k} \frac{\mathrm{~d}}{\mathrm{~d} t} x(t)\right) \\
= & \left(c_{1} F\left(\frac{\mathrm{~d}^{k-1}}{\mathrm{~d} t^{k-1}} x(t), \frac{\mathrm{d}^{k-2}}{\mathrm{~d} t^{k-2}} x(t), \cdots, x(t), t\right), c_{2} \frac{\mathrm{~d}^{k-1}}{\mathrm{~d} t^{k-1}} x(t), \cdots, c_{k} \frac{\mathrm{~d}}{\mathrm{~d} t} x(t)\right) \\
= & \left(c_{1} F\left(c_{1}^{-1} \bar{x}_{1}(t), c_{2}^{-1} \bar{x}_{2}(t), \cdots, c_{k}^{-1} \bar{x}_{k}(t), t\right), c_{2} c_{1}^{-1} \bar{x}_{1}(t), c_{3} c_{2}^{-1} \bar{x}_{2}(t), \cdots, c_{k} c_{k-1}^{-1} \bar{x}_{k-1}(t)\right) \\
& \bar{F}(\bar{x}(t), t)
\end{aligned}
$$

We bound the Lipschitz constant of function $\bar{F}$ by

$$
\begin{aligned}
\|\bar{F}(y(t), t)-\bar{F}(z(t), t)\| & \leq c_{1} \sum_{i=1}^{k} L_{i} c_{i}^{-1}\left\|y_{i}(t)-z_{i}(t)\right\|+\sum_{i=1}^{k-1} c_{i+1} c_{i}^{-1}\left\|y_{i}(t)-z_{i}(t)\right\| \\
& =\sum_{i=1}^{k-1}\left(c_{1} L_{i} c_{i}^{-1}+c_{i+1} c_{i}^{-1}\right)\left\|y_{i}(t)-z_{i}(t)\right\|+c_{1} L_{k} c_{k}^{-1}\left\|y_{k}(t)-z_{k}(t)\right\|
\end{aligned}
$$

We choose $c_{1}=1, c_{i}=\sum_{j=i}^{k} L_{j}^{(i-1) / j}+\frac{1}{\bar{T}^{i-1}}, \forall i \in\{2, \cdots, k\}$ where $\bar{T}=4 \gamma_{\varphi} T$. Then we can calculate for each $i \in[k-1]$,

$$
\begin{aligned}
c_{1} L_{i} c_{i}^{-1}+c_{i+1} c_{i}^{-1} & =\frac{L_{i}}{\sum_{j=i}^{k} L_{j}^{(i-1) / j}+\frac{1}{\bar{T}^{i-1}}}+\frac{\sum_{j=i+1}^{k} L_{j}^{i / j}+\frac{1}{\bar{T}^{i}}}{\sum_{j=i}^{k} L_{j}^{(i-1) / j}+\frac{1}{\overline{T^{i-1}}}} \\
& =\frac{L_{i}}{L_{i}^{(i-1) / i}+\sum_{j=i+1}^{k} L_{j}^{(i-1) / j}+\frac{1}{\bar{T}^{i-1}}}+\frac{\sum_{j=i+1}^{k} L_{j}^{i / j}+\frac{1}{\bar{T}^{i}}}{\sum_{j=i}^{k} L_{j}^{(i-1) / j}+\frac{1}{\bar{T}^{i-1}}} \\
& \leq L_{i}^{1 / i}+\sum_{j=i+1}^{k} L_{j}^{1 / j}+\frac{1}{\bar{T}} \\
& =\sum_{j=i}^{k} L_{j}^{1 / j}+\frac{1}{\bar{T}}
\end{aligned}
$$

For $i=k$, we have

$$
c_{1} L_{k} c_{k}^{-1}=\frac{L_{k}}{L_{k}^{(k-1) / k}+\frac{1}{\bar{T}^{k-1}}} \leq L_{k}^{1 / k} .
$$

Thus,

$$
\|\bar{F}(y(t), t)-\bar{F}(z(t), t)\| \leq\left(\sum_{j=1}^{k} L_{j}^{1 / j}+\frac{1}{\bar{T}}\right) \cdot \sum_{i=1}^{k}\left\|y_{i}(t)-z_{i}(t)\right\|=\left(\sum_{j=1}^{k} L_{j}^{1 / j}+\frac{1}{\bar{T}}\right)\|y(t)-z(t)\|
$$

It gives the following Claim:
Claim 4.7.1. Function $\bar{F}$ has Lipschitz constant $\bar{L}=\sum_{j=1}^{k} L_{j}^{1 / j}+\frac{1}{\bar{T}}$.

Then, we consider the following first order ODE,

$$
\begin{aligned}
\frac{\mathrm{d}}{\mathrm{~d} t} \bar{x}(t) & =\bar{F}(\bar{x}(t), t), \\
\bar{x}(0) & =\left(c_{1} v_{k-1}, c_{2} v_{k-2}, \cdots, c_{k} v_{0}\right)
\end{aligned}
$$

Let $\bar{x}^{*}(t) \in \mathbb{R}^{k d}$ denote the optimal solution, then

$$
\bar{x}^{*}(t)=\left(c_{1} \frac{\mathrm{~d}^{k-1}}{\mathrm{~d} t^{k-1}} x^{*}(t), c_{2} \frac{\mathrm{~d}^{k-2}}{\mathrm{~d} t^{k-2}} x^{*}(t), \cdots, c_{k} x^{*}(t)\right) .
$$

Now, we prove that $\frac{\mathrm{d}}{\mathrm{d} t} \bar{x}^{*}(t)$ is approximate by some element in $\mathcal{V}^{d}$. Let $\bar{q}: \mathbb{R} \rightarrow \mathbb{R}^{k d}$ be defined as follows

$$
\bar{q}(t)=\left(c_{1} q^{(k)}(t), c_{2} q^{(k-1)}(t), \cdots, c_{k} q^{(1)}(t)\right)
$$

Then,

$$
\begin{aligned}
& \left\|\bar{q}(t)-\frac{\mathrm{d}}{\mathrm{~d} t} \bar{x}^{*}(t)\right\| \\
= & \left\|\left(c_{1} q^{(k)}(t), c_{2} q^{(k-1)}(t), \cdots, c_{k} q^{(1)}(t)\right)-\left(c_{1} \frac{\mathrm{~d}^{k}}{\mathrm{~d} t^{k}} x^{*}(t), c_{2} \frac{\mathrm{~d}^{k-1}}{\mathrm{~d} t^{k-1}} x^{*}(t), \cdots, c_{k} \frac{\mathrm{~d}}{\mathrm{~d} t} x^{*}(t)\right)\right\| \\
= & \sum_{i=1}^{k} c_{i}\left\|q^{(k+1-i)}(t)-\frac{\mathrm{d}^{k+1-i}}{\mathrm{~d} t^{k+1-i}} x^{*}(t)\right\| \\
\leq & \sum_{i=1}^{k} c_{i} \frac{\epsilon}{\bar{T}^{k+1-i}}=\frac{1}{\bar{T}} \underbrace{\sum_{i=1}^{k} \frac{c_{i}}{\bar{T}^{k-i}}}_{\bar{\epsilon}} .
\end{aligned}
$$

By the assumption on $T$, we have that $\gamma_{\varphi} \bar{L} T \leq 1 / 2$ and hence theorem 4.2 .2 finds

$$
\bar{x}^{(N)}(t)=\left(\bar{x}_{1}^{(N)}(t), \bar{x}_{2}^{(N)}(t), \cdots, \bar{x}_{k}^{(N)}(t)\right) \in \mathbb{R}^{k d}
$$

such that

$$
\begin{equation*}
\left\|\bar{x}^{(N)}(t)-\bar{x}^{*}(t)\right\| \leq 20 \gamma_{\varphi} \bar{\epsilon} \tag{4.40}
\end{equation*}
$$

This implies that

$$
\begin{equation*}
\left\|c_{i}^{-1} \bar{x}^{(N)}(t)-\frac{\mathrm{d}^{k-i}}{\mathrm{~d} t^{k-i}} x^{*}(t)\right\| \leq 20 c_{i}^{-1} \gamma_{\varphi} \bar{\epsilon}, \forall i \in[k] . \tag{4.41}
\end{equation*}
$$

To bound the last term, we show the following Claim:
Claim 4.7.2. Let $\bar{\epsilon}=\epsilon \sum_{i=1}^{k} \frac{c_{i}}{\overline{T^{k-i}}}$. If we choose $c_{1}=1$ and $c_{i}=\sum_{j=i}^{k} L_{j}^{(i-1) / j}+\frac{1}{\bar{T}^{i-1}}$, $\forall i \in\{2, \cdots, k\}$, then

$$
c_{i}^{-1} \bar{\epsilon} \leq \frac{\epsilon}{\bar{T}^{k-i}}(2 k+1)
$$

Proof. For each $i \in[k]$,

$$
c_{i}^{-1} \bar{\epsilon}=c_{i}^{-1} \epsilon \sum_{j=1}^{k} \frac{c_{j}}{\bar{T}^{k-j}}=\frac{\epsilon}{\bar{T}^{k-i}} \sum_{j=1}^{k} \frac{c_{j} c_{i}^{-1}}{\bar{T}^{i-j}}
$$

We can lower bound the term $\frac{c_{j} c_{i}^{-1}}{\bar{T}^{i-j}}$ as follows,

$$
\begin{aligned}
\frac{c_{j} c_{i}^{-1}}{\bar{T}^{i-j}} & =\frac{\sum_{l=j}^{k} L_{l}^{(j-1) / l}+\frac{1}{\bar{T}^{j-1}}}{\bar{T}^{i-j}\left(\sum_{l=i}^{k} L_{l}^{(i-1) / l}+\frac{1}{\bar{T}^{i-1}}\right)} \\
& =\frac{\sum_{l=j}^{k} L_{l}^{(j-1) / l}+\frac{1}{\bar{T}^{j-1}}}{\sum_{l=i}^{k} \bar{T}^{i-j} L_{l}^{(i-1) / l}+\frac{1}{\bar{T}^{j-1}}} \\
& \leq \frac{\sum_{l=j}^{k} L_{l}^{(j-1) / l}+\frac{1}{\bar{T}^{j-1}}}{\frac{1}{\bar{T}^{j-1}}} \\
& =1+\sum_{l=j}^{k}\left(L_{l}^{1 / l} \bar{T}\right)^{j-1}
\end{aligned}
$$

Therefore, we have

$$
\begin{aligned}
c_{i}^{-1} \bar{\epsilon} & =\frac{\epsilon}{\bar{T}^{k-i}} \sum_{j=1}^{k} \frac{c_{j} c_{i}^{-1}}{\bar{T}^{i-j}} \\
& \leq \frac{\epsilon}{\bar{T}^{k-i}} \sum_{j=1}^{k}\left(1+\sum_{l=j}^{k}\left(L_{l}^{1 / l} \bar{T}\right)^{j-1}\right) \\
& \leq \frac{\epsilon}{\bar{T}^{k-i}}\left(2 k+\sum_{j=1}^{k}\left(\sum_{l=1}^{k} L_{l}^{1 / l} \bar{T}\right)^{j}\right) \\
& \leq \frac{\epsilon}{\bar{T}^{k-i}}\left(2 k+\sum_{j=1}^{k}(1 / 2)^{j}\right)
\end{aligned}
$$

where we used $\gamma_{\varphi} L T \leq 1 / 8$ at the end. Thus we complete the proof of Claim.

Now, using the claim to (4.41), we have the error is

$$
20(2 k+1) \gamma_{\varphi} \frac{\epsilon}{\bar{T}^{k-i}} \leq 20(2 k+1) \gamma_{\varphi} \frac{\epsilon}{T^{k-i}} .
$$

To bound the number of iterations needed in the log term in Theorem 4.2.2, we note that

$$
\begin{aligned}
\int_{0}^{T}\|\bar{F}(\bar{x}(0), s)\| d s & =c_{1}\left\|\int_{0}^{T} F\left(c_{1}^{-1} \bar{x}_{1}(0), c_{2}^{-1} \bar{x}_{2}(0), \cdots, c_{k}^{-1} \bar{x}_{k}(0), s\right) d s\right\|+T \cdot \sum_{i=1}^{k-1} c_{i+1} c_{i}^{-1}\left\|\bar{x}_{i}(t)\right\| \\
& =\left\|\int_{0}^{T} F\left(\frac{d^{k-1}}{d t^{k-1}} x(0), \frac{d^{k-2}}{d t^{k-2}} x(0), \cdots, x(0), s\right) d s\right\|+T \cdot \sum_{i=1}^{k-1} c_{i+1}\left\|\frac{d^{k-i}}{d t^{k-i}} x(0)\right\|
\end{aligned}
$$

Note that

$$
c_{i+1} \leq \sum_{j=1}^{k} L_{j}^{i / j}+\frac{1}{\bar{T}^{i}} \leq L^{i}+\frac{1}{\bar{T}^{i}} \leq \frac{2}{\bar{T}^{i}}
$$

where we used $L \bar{T} \leq \frac{1}{2}$. Hence, the number of iterations we need is

$$
\begin{aligned}
& O\left(D \log \left(\frac{1}{\bar{\epsilon}}\left(\left\|\int_{0}^{T} F\left(v_{k-1}, v_{k-2}, \cdots, v_{0}, s\right) d s\right\|+\sum_{i=1}^{k-1} \frac{\left\|v_{i}\right\|}{\bar{T}^{k-i-1}}\right)\right)\right) \\
= & O\left(D \log \left(\frac{1}{\epsilon}\left(\bar{T}^{k-1} \cdot\left\|\int_{0}^{T} F\left(v_{k-1}, v_{k-2}, \cdots, v_{0}, s\right) d s\right\|+\sum_{i=1}^{k-1} \bar{T}^{i}\left\|v_{i}\right\|\right)\right)\right)
\end{aligned}
$$

where we used $\bar{\epsilon} \geq \epsilon \cdot c_{k} \geq \frac{\epsilon}{\bar{T}^{k-1}}$.

### 4.8 Deferred Proof for Cauchy Estimates (Section 4.6.3)

In this section, we provide the proofs of some core Lemmas/Claims used for proving Theorem 4.6.4.

Lemma 4.6.6. (Bounding multivariate ODE by scalar ODE). Given vectors $v_{1}, v_{0} \in \mathbb{R}^{d}$ and any norm $\|\cdot\|$ on $\mathbb{R}^{d}$. Let $U \subset \mathbb{R}^{d}$ be a neighborhood of $v_{0}$ and that $F: U \rightarrow \mathbb{R}^{d}$ is a real analytic near $v_{0}$. Suppose that

$$
\left\|D^{(k)} F\left(v_{0}\right)\left[\Delta_{1}, \Delta_{2}, \cdots, \Delta_{k}\right]\right\| \leq f^{(k)}(0) \prod_{j=1}^{k}\left\|\Delta_{j}\right\| \text { for all } k \geq 0
$$

for some real analytic function $f$ around 0 . Then the ODE (4.39) has a unique real analytic solution around $t=0$. Furthermore, for any $b \geq\left\|v_{1}\right\|$, we have

$$
\left\|u^{(k)}(0)\right\| \leq \psi^{(k)}(0) \quad \forall k \geq 1
$$

where $\psi$ is the solution of the ODE

$$
\psi^{\prime \prime}(t)=f(\psi(t)), \psi^{\prime}(0)=b, \psi(0)=0
$$

Proof. Theorem 4.6.5 shows that the solution $u$ uniquely exists and is real analytic around 0 . Therefore, we can take derivatives on both sides of $u^{\prime \prime}(t)=F(u(t))$ and get

$$
\begin{aligned}
u^{(3)}(t) & =D F(u(t))\left[u^{(1)}(t)\right] \\
u^{(4)}(t) & =D F(u(t))\left[u^{(2)}(t)\right]+D^{2} F(u(t))\left[u^{(1)}(t), u^{(1)}(t)\right] \\
u^{(5)}(t) & =D F(u(t))\left[u^{(3)}(t)\right]+2 D^{2} F(u(t))\left[u^{(2)}(t), u^{(1)}(t)\right]+D^{3} F(u(t))\left[u^{(1)}(t), u^{(1)}(t), u^{(1)}(t)\right] \\
\vdots & =\vdots
\end{aligned}
$$

Therefore, we have

$$
\begin{aligned}
\left\|u^{(3)}(0)\right\|= & \left\|D F(u(0))\left[u^{(1)}(0)\right]\right\| \\
\leq & f^{(1)}(0)\left\|u^{(1)}(0)\right\| \\
\left\|u^{(4)}(0)\right\|= & \left\|D F(u(0))\left[u^{(2)}(0)\right]+D^{2} F(u(0))\left[u^{(1)}(0), u^{(1)}(0)\right]\right\| \\
\leq & \left\|D F(u(0))\left[u^{(2)}(0)\right]\right\|+\left\|D^{2} F(u(0))\left[u^{(1)}(0), u^{(1)}(0)\right]\right\| \\
\leq & f^{(1)}(0)\left\|u^{(2)}(0)\right\|+f^{(2)}(0)\left\|u^{(1)}(0)\right\|^{2} \\
\left\|u^{(5)}(0)\right\|= & \| D F(u(0))\left[u^{(3)}(0)\right]+2 D^{2} F(u(0))\left[u^{(2)}(0), u^{(1)}(0)\right] \\
& +D^{3} F(u(0))\left[u^{(1)}(0), u^{(1)}(0), u^{(1)}(0)\right] \| \\
\leq & \left\|D F(u(0))\left[u^{(3)}(0)\right]\right\|+\left\|2 D^{2} F(u(0))\left[u^{(2)}(0), u^{(1)}(0)\right]\right\| \\
& +\left\|D^{3} F(u(0))\left[u^{(1)}(0), u^{(1)}(0), u^{(1)}(0)\right]\right\| \\
\leq & f^{(1)}(0)\left\|u^{(3)}(0)\right\|+2 f^{(2)}(0)\left\|u^{(2)}(0)\right\|\left\|u^{(1)}(0)\right\|+f^{(3)}(0)\left\|u^{(1)}(0)\right\|^{3} \\
\vdots= & \vdots
\end{aligned}
$$

Similarly, Theorem 4.6 .5 shows that the solution $\psi$ uniquely exists and is real analytic around 0 . By expanding $\psi^{\prime \prime}(t)=f(\psi(t))$ at $t=0$, we see that

$$
\begin{aligned}
\psi^{(3)}(0) & =f^{(1)}(0) \psi^{(1)}(0), \\
\psi^{(4)}(0) & =f^{(1)}(0) \psi^{(2)}(0)+f^{(2)}(0)\left(\psi^{(1)}(0)\right)^{2}, \\
\psi^{(5)}(0) & =f^{(1)}(0) \psi^{(3)}(0)+2 f^{(2)} \psi^{(2)}(0) \psi^{(1)}(0)+f^{(3)}(0)\left(\psi^{(1)}(0)\right)^{3}, \\
\vdots & =\vdots
\end{aligned}
$$

Since $\left\|u^{(1)}(0)\right\| \leq b=\psi^{(1)}(0),\left\|u^{(2)}(0)\right\|=\|F(u(0))\| \leq f(0)=\psi^{(2)}(0)$.

For $k=3,4,5, \cdots$, we have

$$
\begin{aligned}
\left\|u^{(3)}(0)\right\| & \leq f^{(1)}(0)\left\|u^{(1)}(0)\right\|=f^{(1)}(0) \psi^{(1)}(0)=\psi^{(3)}(0) \\
\left\|u^{(4)}(0)\right\| & \leq \psi^{(4)}(0) \\
\left\|u^{(5)}(0)\right\| & \leq \psi^{(5)}(0) \\
\quad & \leq \vdots
\end{aligned}
$$

Thus, we have that $\left\|u^{(k)}(0)\right\| \leq \psi^{(k)}(0)$ for all $k \geq 1$.

Lemma 4.6.7. (Bounding scalar ODE) Let $f(x)=\frac{a}{1-c x}$ with positive $a$ and $c$. Let $\psi(t)$ denote the solution of the ODE

$$
\psi^{\prime \prime}(t)=f(\psi(t)), \psi^{\prime}(0)=b, \psi(0)=0
$$

with $b \geq 0$. Then,

$$
\psi^{(k)}(0) \leq \frac{k!\alpha^{k}}{c} \quad \forall k \geq 1
$$

with $\alpha=\max \left(\sqrt{\frac{4}{3} a c}, 2 b c\right)$.
Proof. Let $\widetilde{\psi}(t)=\frac{1}{c}(1-\sqrt{1-\alpha t})$ with $\alpha=\max \left(\sqrt{\frac{4}{3} a c}, 2 b c\right)$. Note that

$$
\widetilde{\psi}^{\prime \prime}(t)=\frac{\alpha^{2}}{4 c(1-\alpha t)^{3 / 2}}=\widetilde{f}(\widetilde{\psi}(t)) \quad \text { with } \quad \widetilde{f}(x)=\frac{\alpha^{2}}{4 c(1-c x)^{3}}
$$

Since $\alpha^{2} \geq \frac{4}{3} a c$, we have that

$$
\widetilde{f}^{(k)}(0)=\frac{a}{3} \cdot \frac{(k+2)!}{2} c^{k} \geq k!\cdot a \cdot c^{k}=f^{(k)}(0)
$$

Also, we have that $\widetilde{\psi}(0)=0$ and $\widetilde{\psi}^{\prime}(0)=\frac{\alpha}{2 c} \geq b$. Hence, Lemma 4.6.6 shows that

$$
\psi^{(k)}(0) \leq \widetilde{\psi}^{(k)}(0)=\frac{\alpha^{k}}{c} \prod_{i=1}^{k} \frac{|2 i-3|}{2} \leq \frac{k!\alpha^{k}}{c}
$$

for all $k \geq 1$.

### 4.9 Cauchy Estimates of Some Functions

We first state a useful tool,
Lemma 4.9.1. Let $f: U \rightarrow \mathbb{C}$ be holomorphic and suppose that $D_{R}=\left\{z:\left|z-z_{0}\right| \leq R\right\} \subset$ $U$. Let $\gamma_{R}=\left\{z:\left|z-z_{0}\right|=R\right\}$ denote the boundary of $D_{R}$. For all $n \geq 0$

$$
\left|f^{(n)}\left(z_{0}\right)\right| \leq \frac{n!}{R^{n}} \max _{z \in \gamma_{R}}|f(z)|
$$

### 4.9.1 Logistic loss function

Lemma 4.9.2 (Property of Logistic loss function). Let $\phi(t)=\log \left(1+e^{-t}\right)$, then we know that $\phi^{\prime}(t)$ has Cauchy estimate $M=1$ with radius $r=1$.

Proof. Given the definition of $\phi(t)$, it is easy to see that

$$
\phi^{\prime}(t)=\frac{-1}{1+e^{t}} .
$$

Let $f(z)=\frac{1}{1+e^{z}}$. Let $z=a+b \mathbf{i}$. We have

$$
|f(z)|=\left|\frac{1}{1+e^{a} \cos b+\mathbf{i} e^{a} \sin b}\right|=\frac{1}{\sqrt{\left(1+e^{a} \cos b\right)^{2}+\left(e^{a} \sin b\right)^{2}}}=\frac{1}{\sqrt{1+2 e^{a} \cos b+e^{2 a}}}
$$

Let $z_{0}=a_{0}+\mathbf{i} b_{0}$. We choose $R=1$, then $\left(a-a_{0}\right)^{2}+\left(b-b_{0}\right)^{2}=1$. Since we only care about real numbers, we have $b_{0}=0$. Then we know that $b \in[-1,1]$, which means $\cos b \in[0.54,1]$. Thus, we have,

$$
\sqrt{1+2 e^{a} \cos b+e^{2 a}} \geq \sqrt{1+e^{a}+e^{2 a}} \geq 1 .
$$

Thus, $\left|\frac{1}{1+e^{z}}\right| \leq 1$. Therefore, using Lemma 4.9.1, for all real $z_{0}$,

$$
\left|f^{(n)}\left(z_{0}\right)\right| \leq n!
$$

### 4.9.2 Pseudo-Huber loss function

Huber function [Hub64] has been extensively studied in a large number of algorithmic questions, e.g. regression [CW15b, CW15a], low-rank approximation [SWZ19b], clustering [BFL16], sparse recovery [NSW19a]. Formally speaking, the Huber loss function can be defined as

$$
f(x)= \begin{cases}\frac{x^{2}}{2 \delta}, & \text { if }|x| \leq \delta \\ |x|-\delta / 2, & \text { otherwise }\end{cases}
$$

where $\delta>0$ is a parameter. For many applications, the Pseudo-Huber loss function can be used as a smooth alternative for the Huber loss function.

Lemma 4.9.3 (Property of Pseudo-Huber function). Fix any $\delta>0$. Let $\phi(x)=\sqrt{x^{2}+\delta^{2}}-$ $\delta$, then we know that $\phi^{\prime}(x)$ has Cauchy estimate $M=1$ with radius $r=\delta / 2$.

Proof. Given the definition of function $\phi(x)$, it is easy to see that

$$
\phi^{\prime}(x)=\frac{x}{\sqrt{x^{2}+\delta^{2}}}
$$

Let $f(z)=\frac{z}{\sqrt{z^{2}+\delta^{2}}}$. Let $z=a+b \mathbf{i}$. We have

$$
|f(z)|=\left|\frac{(a+b \mathbf{i})}{\sqrt{(a+b \mathbf{i})^{2}+\delta^{2}}}\right|=\left|\frac{(a+b \mathbf{i})}{\sqrt{a^{2}-b^{2}+\delta^{2}+2 a b \mathbf{i}}}\right|=\frac{|(a+b \mathbf{i})|}{\left|\sqrt{a^{2}-b^{2}+\delta^{2}+2 a b \mathbf{i}}\right|}
$$

For the numerator, we have $|(a+b \mathbf{i})|=\sqrt{a^{2}+b^{2}}$. For the denominator, we have

$$
\left|\sqrt{a^{2}-b^{2}+\delta^{2}+2 a b \mathbf{i}}\right|=\left(\left(a^{2}-b^{2}+\delta^{2}\right)^{2}+4(a b)^{2}\right)^{1 / 4}
$$

Let $z_{0}=a_{0}+\mathbf{i} b_{0}$. We choose $R=\frac{\delta}{2}$, then $\left(a-a_{0}\right)^{2}+\left(b-b_{0}\right)^{2}=(\delta / 2)^{2}$. Since we are only real $z_{0}$, we have $b_{0}=0$. Then we know that $b \in\left[-\frac{\delta}{2}, \frac{\delta}{2}\right]$. Thus

$$
|f(z)|=\frac{\sqrt{a^{2}+b^{2}}}{\left(\left(a^{2}-b^{2}+\delta^{2}\right)^{2}+4(a b)^{4}\right)^{1 / 4}} \leq \frac{\sqrt{a^{2}+\frac{\delta^{2}}{4}}}{\left(\left(a^{2}-\frac{\delta^{2}}{4}+\delta^{2}\right)^{2}+0\right)^{1 / 4}}=\frac{\sqrt{a^{2}+\frac{\delta^{2}}{4}}}{\sqrt{a^{2}+\frac{3 \delta^{2}}{4}}} \leq 1
$$

Therefore, using Lemma 4.9.1, for all real $z_{0}$,

$$
\left|f^{(n)}\left(z_{0}\right)\right| \leq n!\cdot(2 / \delta)^{n}
$$

## Chapter 5

## Convergence result of Deep Neural Network

Deep neural networks (DNNs) have demonstrated dominating performance in many fields; since AlexNet, networks used in practice are going wider and deeper. On the theoretical side, a long line of works has been focusing on training neural networks with one hidden layer. The theory of multi-layer networks remains largely unsettled.

In this work, we prove why stochastic gradient descent (SGD) can find global minima on the training objective of DNNs in polynomial time. We only make two assumptions: the inputs are non-degenerate and the network is over-parameterized. The latter means the network width is sufficiently large: polynomial in $L$, the number of layers and in $n$, the number of samples.

Our key technique is to derive that, in a sufficiently large neighborhood of the random initialization, the optimization landscape is almost-convex and semi-smooth even with ReLU activations. This implies an equivalence between over-parameterized neural networks and neural tangent kernel (NTK) in the finite (and polynomial) width setting.

As concrete examples, starting from randomly initialized weights, we prove that SGD can attain $100 \%$ training accuracy in classification tasks, or minimize regression loss in linear convergence speed, with running time polynomial in $n, L$. Our theory applies to the widely-used but non-smooth ReLU activation, and to any smooth and possibly non-
convex loss functions. In terms of network architectures, our theory at least applies to fully-connected neural networks, convolutional neural networks (CNN), and residual neural networks (ResNet).

### 5.1 Introduction

Neural networks have demonstrated a great success in numerous machine-learning tasks [KSH12, GMH13, $\left.\mathrm{LHP}^{+} 15, \mathrm{AAA}^{+} 16, \mathrm{HZRS}^{2} 6, \mathrm{SHM}^{+} 16, \mathrm{SSS}^{+} 17\right]$. One of the empirical findings is that neural networks, trained by first-order methods from random initialization, have a remarkable ability to fit training data $\left[\mathrm{ZBH}^{+} 17\right]$.

From an expressibility perspective, this may not be surprising since modern neural networks are often over-parameterized: they have much more parameters than the number of training samples. There certainly exist parameter choices with zero training error as long as data is non-degenerate.

Yet, from an optimization perspective, the fact that randomly-initialized first-order methods can find global minima on the training data is quite non-trivial: neural networks are often equipped with the ReLU activation, making the training objective not only nonconvex, but even non-smooth. Even the general convergence for finding approximate critical points of a non-convex, non-smooth function is not fully-understood [BLO05] and appears to be a challenging question on its own. This is in direct contrast to practice, in which ReLU networks trained by stochastic gradient descent (SGD) from random initialization almost never suffer from non-smoothness or non-convexity, and can avoid local minima for a variety of network architectures (see [GVS15]). A theoretical justification was missing to explain this phenomenon.

There are quite a few papers trying to understand the success of neural networks from optimization perspective. Many of them focus on the case when the inputs are random Gaussian, and work only for two-layer neural networks [BG17, Son17, Tia17, LY17, DLT ${ }^{+}$18,

GLM17, PRSZ18, ZSJ+ 17, ZSD17]. [LL18] show that for a two-layer network with ReLU activation, SGD finds nearly-global optimal (say, 99\% classification accuracy) solutions on the training data, as long as the network is over-parameterized, meaning that the number of neurons is polynomially large comparing to the input size. Moreover, if the data is sufficiently structured (say, coming from mixtures of separable distributions), this accuracy extends also to test data. As a separate note, over-parameterization is suggested as the possible key to avoid bad local minima by [SS18] even for two-layer networks.

There are also results that go beyond two-layer networks with limitations. Some consider deep linear neural networks without any activation functions [HM17, ACGH18, BHL18, Kaw16]. [Dal17] studies multi-layer neural networks but essentially only with respect to the convex task of training the last layer. ${ }^{1}$ [SC16] show that under over-parameterization and under random input perturbation, there is bad local minima for multi-layer neural networks. [JGH18] derive global convergence using neural tangent kernel for infinite-width neural networks.

In this paper, we study the following fundamental questions

Can DNN be trained close to zero training error efficiently under mild assumptions? If so, can the running time depend only polynomially in the network depth and input size?

Motivation In 2012, AlexNet was born with 5 convolutional layers [KSH12]. The later VGG network uses 19 layers [SSZ14], and GoogleNet uses 22 layers [SLJ ${ }^{+}$15]. In practice, we

[^13]cannot go deeper by naively stacking layers together, due to the so-called vanishing/exploding gradient problem. To deal with this issue, networks with residual links (ResNet) were proposed with the capability of handling at least 152 layers [HZRS16]. Compared with practical networks that go much deeper, existing theory has been mostly around two-layer (thus one-hidden-layer) neural networks, even just for the training process alone. Thus,

Can we theoretically justify how the training process has worked for multi-layer neural networks?

In this paper, we extend the over-parameterization theory to multi-layer neural networks.

### 5.1.1 Our Result

We show that over-parameterized neural networks can be trained by vanilla firstorder methods such as gradient descent (GD) or stochastic gradient descent (SGD) to global minima (e.g. zero training error), as long as the data is non-degenerate.

We say that the data is non-degenerate if every pairs of samples are distinct. This is a minimal requirement since a dataset with two identical data points of different labels cannot be trained to zero error. We denote by $\delta$ the minimum (relative) distance between two data points, and by $n$ the number of training samples. Now, consider an $L$-layer fully-connected feedforward neural network, each hidden layer consisting of $m$ neurons equipped with ReLU activation. We show that,

- As long as $m \geq \operatorname{poly}\left(n, L, \delta^{-1}\right)$, starting from random Gaussian initialization, GD/SGD finds an $\epsilon$-error global minimum in $\ell_{2}$ regression using at most $T=\operatorname{poly}\left(n, L, \delta^{-1}\right) \log \frac{1}{\epsilon}$ iterations.
- If the task is multi-label classification, then GD/SGD finds an $100 \%$ accuracy classifier on the training set in $T=\operatorname{poly}\left(n, L, \delta^{-1}\right)$ iterations.
- Our result also applies to other Lipschitz-smooth loss functions, and some other network architectures including convolutional neural networks (CNNs) and residual networks (ResNet).

In contrast, prior work on this task either requires $m$ and $T$ to grow in $e^{O(L)}$ (and essentially only the last layer is trained) [Dal17]; or requires $m=\infty$ [JGH18].

Our Contributions We summarize our technical contributions below.

- For a sufficiently large neighborhood of the random initialization, we prove that the training landscape is almost convex and semi-smooth. This somewhat explains the empirical finding by [GVS15] that GD/SGD will not be trapped in local minima. (See Section 5.4.1.)
- For a sufficiently large neighborhood of the random initialization, we derive an equivalence between neural networks and the neural tangent kernel (NTK) introduced by [JGH18]. Unlike the prior work in which they show the equivalence only for infinitewidth networks (i.e., $m=\infty$ ), here we only need $m=\operatorname{poly}(L)$ for such an equivalence to hold. (See Section 5.4.2.)
- We show that equipped with ReLU activation, neural networks do not suffer from exponential gradient explosion or vanishing. This is the key reason we can avoid exponential dependency on $L$. If one is okay with $e^{O(L)}$ dependency, many proofs shall become trivial. (See Section 5.5.)
- We derive a stability theory of neural networks against small but adversarial perturbations that may be of independent interests. Previous results on this topic either have exponential blowup in $L$ [Dal17] or requires the width to go to infinity [JGH18]. (See Section 5.5.)
- We derive our results by training only hidden layers. This can be more meaningful than training all the layers together, in which if one is not careful with parameter choices, the training process can degenerate as if only the last layer is trained [Dal17].

That is a convex task and may not reflect the true power of deep learning. (Of course, as a simple corollary, our results also apply to training all the layers together.)

Finally, we emphasize that this present paper as a deeply-simplified version of the recurrent neural network (RNN) paper [AZLS18] by the same set of authors. To some extent, DNN is a "special case" of RNN, ${ }^{2}$ thus most of the technical tools were already developed in [AZLS18]. We write this DNN result as a separate paper because: (1) not all the readers can easily derive the DNN result from [AZLS18]; (2) the convergence of DNN can be important on its own; (3) the proof in this paper is much simpler (30 vs 80 pages) and could reach out to a wider audience; (4) the simplicity of this paper allows us to tighten parameters in some non-trivial ways; and (5) the simplicity of this paper allows us to also study convolutional networks, residual networks, as well as different loss functions (all of them were missing from [AZLS18]). We also note that the techniques of this paper can be combined with [AZLS18] to show the global convergence of training over-parameterized deep RNN. We ignore the details so as not to complicate this paper.

Towards Generalization In practice, deeper and wider neural networks generalize better [SGS15, ZK16], so what can we say in theory? Although this paper does not explicitly cover generalization to test data, since a neural network in our parameter regime simulates

[^14]its neural tangent kernel (NTK), it is clear that neural networks provide generalization at least as good as its NTK.

In the PAC-learning language, one may study generalization with respect to concept classes. Follow-up work [ALL18c] shows that three-layer over-parameterized ReLU networks can efficiently (in polynomial time and sample complexity) learn the concept class of threelayer neural networks with smooth activations [ALL18c], and the follow-up work [AL19b] shows stronger results for three-layer ResNet.

It is worth pointing out that the three-layer result [ALL18c] goes beyond the almostconvex regime and thus is not captured by its NTK; more interestingly, the three-layer ResNet result [AL19b] is not achievable (in a provable sense) by any kernel method including any NTK.

A concurrent, independent and beautiful result We acknowledge a concurrent work $\left[\mathrm{DLL}^{+} 19\right]$ by Du, Lee, Li, Wang and Zhai. They proved a similar result.

### 5.1.2 Other Related Works

[LL18] originally prove their result for the cross-entropy loss, together with some test accuracy guarantee. (If data is "well-structured", they prove two-layer over-parameterized neural networks can learn it using SGD with polynomially many samples [LL18].) Later, the "training accuracy" part of [LL18] was extended to the $\ell_{2}$ loss [DZPS19]. From another perspective, [DZPS19] focus on gradient descent.

Linear networks without activation functions are important subjects on its own. Besides the already cited references [HM17, ACGH18, BHL18, Kaw16], there are a number of works that study linear dynamical systems, which can be viewed as the linear version of recurrent neural networks or reinforcement learning. Recent works in this line of research include [HMR18, HSZ17, $\mathrm{HLS}^{+}$18, $\mathrm{DMM}^{+}$17, OO18, AAMM18, $\mathrm{SMT}^{+}$18, MT18, DTMR18, $\mathrm{AHL}^{+}$18].

There is sequence of work about one-hidden-layer (multiple neurons) CNN [BG17, ZSD17, DLT ${ }^{+}$18, GKM18, Oym18]. Whether the patches overlap or not plays a crucial role in analyzing algorithms for such CNN. One category of the results have required the patches to be disjoint [BG17, ZSD17, DLT $\left.{ }^{+} 18\right]$. The other category [GKM18, Oym18] have figured out a weaker assumption or even removed that patch-disjoint assumption. On input data distribution, most relied on inputs being Gaussian [BG17, ZSD17, DLT ${ }^{+}$18, Oym18], and some assumed inputs to be symmetrically distributed with identity covariance and boundedness [GKM18].

As for ResNet, [LY17] proved that SGD learns one-hidden-layer residual neural networks under Gaussian input assumption. The techniques in [ZSJ ${ }^{+}$17, ZSD17] can also be
generalized to one-hidden-layer ResNet under the Gaussian input assumption; they can show that GD starting from good initialization point (via tensor initialization) learns ResNet. [HM17] deep linear residual networks have no spurious local optima.

If no assumption is allowed, neural networks have been shown hard in several different perspectives. Thirty years ago, [BR93] first proved that learning the neural network is NPcomplete. Stronger hardness results have been proved over the last decade [KS09, LSSS14, Dan16, DSS16, GKKT17, SVWX17, MR18].

### 5.2 Preliminaries

We use $\mathcal{N}(\mu, \sigma)$ to denote the Gaussian distribution of mean $\mu$ and variance $\sigma$; and $\mathcal{B}\left(m, \frac{1}{2}\right)$ to denote the binomial distribution with $m$ trials and $1 / 2$ success rate. We use $\|v\|_{2}$ or $\|v\|$ to denote Euclidean norms of vectors $v$, and $\|M\|_{2},\|M\|_{F}$ to denote spectral and Frobenius norms of matrices $M$. For a tuple $\vec{W}=\left(W_{1}, \ldots, W_{L}\right)$ of matrices, we let $\|\vec{W}\|_{2}=\max _{\ell \in[L]}\left\|W_{\ell}\right\|_{2}$ and $\|\vec{W}\|_{F}=\left(\sum_{\ell=1}^{L}\left\|W_{\ell}\right\|_{F}^{2}\right)^{1 / 2}$.

We use $\phi(x)=\max \{0, x\}$ to denote the ReLU function, and extend it to vectors $v \in \mathbb{R}^{m}$ by letting $\phi(v)=\left(\phi\left(v_{1}\right), \ldots, \phi\left(v_{m}\right)\right)$. We use $\mathbb{1}_{\text {event }}$ to denote the indicator function for event.

The training data consist of vector pairs $\left\{\left(x_{i}, y_{i}^{*}\right)\right\}_{i \in[n]}$, where each $x_{i} \in \mathbb{R}^{\mathfrak{d}}$ is the feature vector and $y_{i}^{*}$ is the label of the $i$-th training sample. We assume without loss of generality that data are normalized so that $\left\|x_{i}\right\|=1$ and its last coordinate $\left(x_{i}\right)_{\mathfrak{d}}=\frac{1}{\sqrt{2}}$. ${ }^{3}$ We also assume $\left\|y_{i}^{*}\right\| \leq O(1)$ for notation simplicity. ${ }^{4}$

We make the following separable assumption on the training data (motivated by [LL18]):

Assumption 5.2.1. For every pair $i, j \in[n]$, we have $\left\|x_{i}-x_{j}\right\| \geq \delta$.

[^15]To present the simplest possible proof, the main body of this paper only focuses on depth- $L$ feedforward fully-connected neural networks with an $\ell_{2}$-regression task. Therefore, each $y_{i}^{*} \in \mathbb{R}^{d}$ is a target vector for the regression task. We explain how to extend it to more general settings in Section 5.6 and the Appendix. For notational simplicity, we assume all the hidden layers have the same number of neurons, and our results trivially generalize to each layer having different number of neurons. Specifically, we focus on the following network

$$
\begin{aligned}
g_{i, 0} & =A x_{i} & & h_{i, 0}=\phi\left(A x_{i}\right)
\end{aligned} \begin{array}{ll}
g_{i, \ell} & =W_{\ell} h_{i, \ell-1} \\
y_{i} & =B h_{i, L}
\end{array}
$$

where $A \in \mathbb{R}^{m \times \mathcal{0}}$ is the weight matrix for the input layer, $W_{\ell} \in \mathbb{R}^{m \times m}$ is the weight matrix for the $\ell$-th hidden layer, and $B \in \mathbb{R}^{d \times m}$ is the weight matrix for the output layer. For notational convenience in the proofs, we may also use $h_{i,-1}$ to denote $x_{i}$ and $W_{0}$ to denote A.

Definition 5.2.1 (diagonal sign matrix). For each $i \in[n]$ and $\ell \in\{0,1, \ldots, L\}$, we denote by $D_{i, \ell}$ the diagonal sign matrix where $\left(D_{i, \ell}\right)_{k, k}=\mathbb{1}_{\left(W_{\ell} h_{i, \ell-1}\right)_{k} \geq 0}$ for each $k \in[m]$.

As a result, we have $h_{i, \ell}=D_{i, \ell} W_{\ell} h_{i, \ell-1}=D_{i, \ell} g_{i, \ell}$ and $\left(D_{i, \ell}\right)_{k, k}=\mathbb{1}_{\left(g_{i, \ell}\right)_{k} \geq 0}$. We make the following standard choices of random initialization:

Definition 5.2.2. We say that $\vec{W}=\left(W_{1}, \ldots, W_{L}\right), A$ and $B$ are at random initialization if

- $\left[W_{\ell}\right]_{i, j} \sim \mathcal{N}\left(0, \frac{2}{m}\right)$ for every $i, j \in[m]$ and $\ell \in[L]$;
- $A_{i, j} \sim \mathcal{N}\left(0, \frac{2}{m}\right)$ for every $(i, j) \in[m] \times[\mathfrak{d}]$; and
- $B_{i, j} \sim \mathcal{N}\left(0, \frac{1}{d}\right)$ for every $(i, j) \in[d] \times[m]$.

Assumption 5.2.2. Throughout this paper we assume $m \geq \Omega\left(\operatorname{poly}\left(n, L, \delta^{-1}\right) \cdot d\right)$ for some sufficiently large polynomial. To present the simplest proof, we did not try to improve such polynomial factors. We will also assume $\delta \leq O\left(\frac{1}{L}\right)$ for notation simplicity.

### 5.2.1 Objective and Gradient

Our regression objective is

$$
F(\vec{W}) \stackrel{\text { def }}{=} \sum_{i=1}^{n} F_{i}(\vec{W}) \quad \text { where } \quad F_{i}(\vec{W}) \stackrel{\text { def }}{=} \frac{1}{2}\left\|B h_{i, L}-y_{i}^{*}\right\|^{2} \quad \text { for each } i \in[n]
$$

We also denote by loss $\stackrel{\text { def }}{=} B h_{i, L}-y_{i}^{*}$ the loss vector for sample $i$. For simplicity, we focus on training only hidden weights $\vec{W}$ in this paper and leave $A$ and $B$ at random initialization. Our result naturally extends to the case when $A, B$ and $\vec{W}$ are jointly trained. ${ }^{5}$

Definition 5.2.3. For each $\ell \in\{1,2, \cdots, L\}$, we define Back $_{i, \ell} \stackrel{\text { def }}{=} B D_{i, L} W_{L} \cdots D_{i, \ell} W_{\ell} \in$ $\mathbb{R}^{d \times m}$ and for $\ell=L+1$, we define Back $_{i, \ell}=B \in \mathbb{R}^{d \times m}$.

Using this notation, one can calculate the gradient of $F(\vec{W})$ as follows.
Fact 5.2.3. The gradient with respect to the $k$-th row of $W_{\ell} \in \mathbb{R}^{m \times m}$ is

$$
\nabla_{\left[W_{\ell}\right]_{k}} F(\vec{W})=\sum_{i=1}^{n}\left(\operatorname{Back}_{i, \ell+1}^{\top} \operatorname{loss}_{i}\right)_{k} \cdot h_{i, \ell-1} \cdot \mathbb{1}_{\left\langle\left[W_{\ell}\right]_{k}, h_{i, \ell-1}\right\rangle \geq 0}
$$

The gradient with respect to $W_{\ell}$ is

$$
\nabla_{W_{\ell}} F(\vec{W})=\sum_{i=1}^{n} D_{i, \ell}\left(\operatorname{Back}_{i, \ell+1}^{\top} \operatorname{loss}_{i}\right) h_{i, \ell-1}^{\top}
$$

We denote by $\nabla F(\vec{W})=\left(\nabla_{W_{1}} F(\vec{W}), \ldots, \nabla_{W_{L}} F(\vec{W})\right)$.

[^16]
### 5.3 Our Results and Techniques

To present our result in the simplest possible way, we choose to mainly focus on fullyconnected $L$-layer neural networks with the $\ell_{2}$ regression loss. We shall extend it to more general settings (such as convolutional and residual networks and other losses) in Section 5.6. Our main results can be stated as follows:

Theorem 5.3.1 (gradient descent). Suppose $m \geq \widetilde{\Omega}\left(\operatorname{poly}\left(n, L, \delta^{-1}\right) \cdot d\right)$. Starting from random initialization, with probability at least $1-e^{-\Omega\left(\log ^{2} m\right)}$, gradient descent with learning rate $\eta=\Theta\left(\frac{d \delta}{\text { poly }(n, L) \cdot m}\right)$ finds a point $F(\vec{W}) \leq \epsilon$ in $T=\Theta\left(\frac{\operatorname{poly}(n, L)}{\delta^{2}} \cdot \log \epsilon^{-1}\right)$ iterations.

This is known as the linear convergence rate because $\epsilon$ drops exponentially fast in $T$. We have not tried to improve the polynomial factors in $m$ and $T$, and are aware of several ways to improve these factors (but at the expense of complicating the proof). We note that $\mathfrak{d}$ is the data input dimension and our result is independent of $\mathfrak{d}$.

Theorem 5.3.2 (SGD). Suppose $b \in[n]$ and $m \geq \widetilde{\Omega}\left(\frac{\operatorname{poly}\left(n, L, \delta^{-1}\right) \cdot d}{b}\right)$. Starting from random initialization, with probability at least $1-e^{-\Omega\left(\log ^{2} m\right)}$, SGD with learning rate $\eta=$ $\Theta\left(\frac{b \delta d}{\operatorname{poly}(n, L) m \log ^{2} m}\right)$ and mini-batch size b finds $F(\vec{W}) \leq \epsilon$ in $T=\Theta\left(\frac{\operatorname{poly}(n, L) \cdot \log ^{2} m}{\delta^{2} b} \cdot \log \epsilon^{-1}\right)$ iterations.

This is again a linear convergence rate because $T \propto \log \frac{1}{\epsilon}$. The reason for the additional $\log ^{2} m$ factor comparing to Theorem 5.13 .1 is because we have a $1-e^{-\Omega\left(\log ^{2} m\right)}$ high confidence bound.

Remark 5.3.1. For experts in optimization theory, one may immediately question the accuracy of Theorem 5.14.1, because SGD is known to converge at a slower rate $T \propto \frac{1}{\operatorname{poly}(\epsilon)}$ even
for convex functions. There is no contradiction here. Imaging a strongly convex function $f(x)=\sum_{i=1}^{n} f_{i}(x)$ that has a common minimizer $x^{*} \in \operatorname{argmin}_{x}\left\{f_{i}(x)\right\}$ for every $i \in[n]$, then SGD is known to converge in a linear convergence rate.


Figure 5.1: Landscapes of the CIFAR10 image-classification training objective $F(W)$ near the SGD training trajectory. The blue vertical stick marks the current point $W=W_{t}$ at the current iteration $t$. The $x$ and $y$ axes represent the gradient direction $\nabla F\left(W_{t}\right)$ and the most negatively curved direction of the Hessian after smoothing (approximately found by Oja's method [AL17, AL18]). The $z$ axis represents the objective value. Observation. As far as minimizing objective is concerned, the (negative) gradient direction sufficiently decreases the training objective, and it is not needed to use second-order method to find negative curvature. This is consistent with our findings Theorem 5.11.1 and 5.12.1. Remark 1. Gradient norm does not tend to zero because cross-entropy loss is not strongly convex (see Section 5.6). Remark 2. The task is CIFAR10 (for CIFAR100 or CIFAR10 with noisy label, see Figure 5.2 through 5.7 in appendix). Remark 3. Architecture is VGG19 (for Resnet-32 or ResNet-110, see Figure 5.2 through 5.7 in the later sections). Remark 4. The six plots correspond to epochs $5,40,90,120,130$ and 160 . We start with learning rate 0.1 , and decrease it to 0.01 at epoch 81 , and to 0.001 at epoch 122 . SGD with momentum 0.9 is used. The training code is unchanged from [Yan18] and we only write new code for plotting such landscapes.

### 5.4 Conceptual Messages and Technical Theorems

We highlight two conceptual messages that arise from the proofs of Theorem 5.13.1 and 5.14.1.

### 5.4.1 Objective is Almost Convex and Semi-Smooth

The first message is about the optimization landscape for points that are sufficiently close to the random initialization. It consists of two theorems, Theorem 5.11.1 says that the objective is "almost convex" and Theorem 5.12 .1 says that the objective is "semi-smooth."

Theorem 5.4.1 (no critical point). With probability $\geq 1-e^{-\Omega\left(m / \operatorname{poly}\left(n, L, \delta^{-1}\right)\right)}$ over randomness $\vec{W}^{(0)}, A, B$, it satisfies for every $\ell \in[L]$, every $i \in[n]$, and every $\vec{W}$ with $\left\|\vec{W}-\vec{W}^{(0)}\right\|_{2} \leq$ $\frac{1}{\operatorname{poly}\left(n, L, \delta^{-1}\right)}$,

$$
\|\nabla F(\vec{W})\|_{F}^{2} \leq O\left(F(\vec{W}) \times \frac{L n m}{d}\right) \quad \text { and } \quad\|\nabla F(\vec{W})\|_{F}^{2} \geq \Omega\left(F(\vec{W}) \times \frac{\delta m}{d n^{2}}\right)
$$

The first property above is easy to prove, while the second property above says that as long as the objective is large, the gradient norm is also large. (See also Figure 5.1.) This means, when we are sufficiently close to the random initialization, there is no saddle point or critical point of any order.

Theorem 5.11 .1 gives us hope to find global minima of the objective $F(\vec{W})$, but is not enough. If we follow the negative gradient direction of $F(\vec{W})$, how can we guarantee that the objective truly decreases? Classical optimization theory usually relies on objective's (Lipscthiz) smoothness [Nes04] to derive an objective-decrease guarantee. Unfortunately, smoothness property at least requires the objective to be twice differentiable, but ReLU activation is not. To deal with this issue, we prove the following.

Theorem 5.4.2 (semi-smoothness). With probability at least $1-e^{-\Omega(m / \operatorname{poly}(L, \log m))}$ over the randomness of $\vec{W}^{(0)}, A, B$, we have: for every $\breve{W} \in\left(\mathbb{R}^{m \times m}\right)^{L}$ with $\|\breve{W}-\vec{W}(0)\|_{2} \leq \frac{1}{\operatorname{poly}(L, \log m)}$,
and for every $\vec{W}^{\prime} \in\left(\mathbb{R}^{m \times m}\right)^{L}$ with $\left\|\vec{W}^{\prime}\right\|_{2} \leq \frac{1}{\operatorname{poly}(L, \log m)}$, the following inequality holds
$F\left(\stackrel{\breve{\vec{W}}}{ }+\overrightarrow{W^{\prime}}\right) \leq F(\stackrel{\breve{\vec{W}}}{ })+\left\langle\nabla F(\stackrel{\breve{\vec{W}})}{ }) \vec{W}^{\prime}\right\rangle+\frac{\operatorname{poly}(L) \sqrt{n m \log m}}{\sqrt{d}} \cdot\left\|\overrightarrow{W^{\prime}}\right\|_{2}\left(F\left(\frac{\breve{\vec{W}}))^{1 / 2}+O\left(\frac{n L^{2} m}{d}\right)\left\|\overrightarrow{W^{\prime}}\right\|_{2}^{2}}{}\right.\right.$
Different from classical smoothness, we still have a first-order term $\left\|\overrightarrow{W^{\prime}}\right\|_{2}$ on the right hand side, while classical smoothness only has a second-order term $\left\|\vec{W}^{\prime}\right\|_{2}^{2}$. As one can see in our final proofs, as $m$ goes larger, the effect of the first-order term becomes smaller comparing to the second-order term. This brings Theorem 5.12.1 closer, but still not identical, to the classical Lipschitz smoothness.

Back to Theorem 5.13.1 and 5.14.1 The derivation of Theorem 5.13.1+5.14.1 from Theorem 5.11.1+5.12.1 is quite straightforward, and can be found in Section 5.13 and 5.14. At a high level, we show that GD/SGD can converge fast enough so that the weights stay close to random initialization by spectral norm bound $\frac{1}{\operatorname{poly}\left(n, L, \delta^{-1}\right)}$. This ensures Theorem 5.11.1 and 5.12 .1 both apply. ${ }^{6}$

In practice, one often goes beyond this theory-predicted spectral-norm boundary. However, quite interestingly, we still observe Theorem 5.11.1 and 5.12.1 hold in practice (see Figure 5.1). The gradient is sufficiently large and going in its negative direction can indeed decrease the objective.

[^17]
### 5.4.2 Equivalence to Neural Tangent Kernel

Recall on input $x \in \mathbb{R}^{\boldsymbol{D}}$, the network output $y(\vec{W} ; x) \stackrel{\text { def }}{=} y=B h_{L} \in \mathbb{R}^{d}$ is a function of the weights $\vec{W}$. Let us here focus on $d=1$ for notational simplicity and leave $d>1$ to the appendix. The neural tangent kernel (NTK) [JGH18] is usually referred to as the feature space defined by the network gradient at random initialization. In other words,

- Given two inputs $x, \widetilde{x} \in \mathbb{R}^{\mathfrak{d}}$, the NTK kernel function is given as

$$
K^{\text {ntk }}(x, \widetilde{x}) \stackrel{\text { def }}{=}\left\langle\nabla y\left(\vec{W}^{(0)} ; x\right), \nabla y\left(\vec{W}^{(0)} ; \widetilde{x}\right)\right\rangle
$$

- Given weight matrix tuple $\vec{W}^{\prime}$, the NTK model computes (we call the NTK objective)

$$
y^{\mathrm{ntk}}\left(\vec{W}^{\prime} ; x\right) \stackrel{\text { def }}{=}\left\langle\nabla y\left(\vec{W}^{(0)} ; x\right), \overrightarrow{W^{\prime}}\right\rangle=\sum_{\ell=1}^{L}\left\langle\nabla_{W_{\ell}} y\left(\vec{W}^{(0)} ; x\right), W_{\ell}^{\prime}\right\rangle
$$

In contrast, the dynamic NTK is given by arbitrary weight tuple $\vec{W}=\vec{W}^{(0)}+\vec{W}^{\prime}$ that may not be at random initialization. [JGH18] proved in their original paper that, when $m$ is infinite, dynamic NTK and NTK are identical because during the training process $\lim _{m \rightarrow \infty}\left\|\overrightarrow{W^{\prime}}\right\|_{2}=0$ if $\vec{W}=\vec{W}^{(0)}+\vec{W}^{\prime}$ is output of gradient descent.

In this paper, we complement [JGH18] by showing a polynomial bound on this equivalence, for any point that is within a certain ball of $\vec{W}^{(0)}$. It is a simple corollary of Theorem 5.11.1 and Theorem 5.12.1, but we state it independently here since it may be of additional interest. ${ }^{7}$

[^18]Theorem 5.4.3. Let $\vec{W}^{(0)}, A, B$ be at random initialization. For fixed unit vectors $x, \widetilde{x} \in$ $\mathbb{R}^{\boldsymbol{d}}$, every (small) parameter $\omega \leq \frac{1}{\text { poly }(L, \log m)}$, with probability at least $1-e^{-\Omega\left(m \omega^{2 / 3} L\right)}$ over $\vec{W}^{(0)}, A, B$, we have for all $\vec{W}^{\prime}$ with $\left\|\overrightarrow{W^{\prime}}\right\|_{2} \leq \omega$,
(a) $\left\|\nabla y\left(\vec{W}^{(0)}+\vec{W}^{\prime} ; x\right)-\nabla y^{\text {ntk }}\left(\vec{W}^{\prime} ; x\right)\right\|_{F} \leq \widetilde{O}\left(\omega^{1 / 3} L^{3}\right) \cdot\left\|\nabla y^{\text {ntk }}\left(\vec{W}^{\prime} ; x\right)\right\|_{F}$;
(b) $y\left(\vec{W}^{(0)}+\vec{W}^{\prime} ; x\right)=y\left(\vec{W}^{(0)} ; x\right)+y^{\mathrm{ntk}}\left(\vec{W}^{\prime} ; x\right) \pm \widetilde{O}\left(L^{3} \omega^{4 / 3} \sqrt{m}\right) ;$ and
(c) $\left\langle\nabla y\left(\vec{W}^{(0)}+\vec{W}^{\prime} ; x\right), \nabla y\left(\vec{W}^{(0)}+\vec{W}^{\prime} ; \widetilde{x}\right)\right\rangle=K^{\text {ntk }}(x, \widetilde{x}) \pm \widetilde{O}\left(\omega^{1 / 3} L^{3}\right) \cdot \sqrt{K^{\text {ntk }}(x, x) K^{\text {ntk }}(\widetilde{x}, \widetilde{x})}$.

Theorem 5.15.1a and 5.15.1c says that dynamic NTK and NTK are almost equivalent up to a small multiplicative factor as long as $\omega<\frac{1}{\operatorname{poly}(L, \log m)}$; while Theorem 5.15 .1 b says that the NTK objective is almost exactly the first-order approximation of the neural network output as long as $\omega<\frac{1}{m^{3 / 8} \operatorname{poly}(L)}$.

In comparison, in Theorem 5.13.1 and 5.14.1, GD/SGD outputs $\vec{W}^{\prime}$ satisfying $\omega \leq$ $\frac{\operatorname{poly}\left(n, \delta^{-1}\right)}{\sqrt{m}} \ll \frac{1}{m^{3 / 8} \operatorname{poly}(L)}$ under the assumption $m \geq \operatorname{poly}\left(n, L, \delta^{-1}\right) .{ }^{8}$ Thus, Theorem 5.15 .1 b implies $\vec{W}^{\prime}$ is also a solution to the NTK regression objective.

Remark 5.4.1. If one wishes to have $y\left(\vec{W}^{(0)}+\vec{W}^{\prime} ; x\right) \approx y^{\text {ntk }}\left(\vec{W}^{\prime} ; x\right)$ without the zero-order term $y\left(\vec{W}^{(0)} ; x\right)$, this can be achieved by properly scaling down the random initialization by a factor of the target error $\epsilon<1$. This was used in the follow-up [ALL18c] to achieve small generalization error on over-parameterized neural networks.

[^19]
### 5.5 Proof Overview

Our proof to the Theorem 5.11.1 and 5.12.1 mostly consist of the following steps.

Step 1: properties at random initialization Let $\vec{W}=\vec{W}^{(0)}$ be at random initialization and $h_{i, \ell}$ and $D_{i, \ell}$ be defined with respect to $\vec{W}$. We first show that forward propagation neither explode or vanish. That is,

$$
\left\|h_{i, \ell}\right\| \approx 1 \text { for all } i \in[n] \text { and } \ell \in[L] .
$$

This is basically because for a fixed $y$, we have $\|W y\|^{2}$ is around 2 , and if its signs are sufficiently random, then ReLU activation kills half of the norm, that is $\|\phi(W y)\| \approx 1$. Then applying induction finishes the proof.

Analyzing forward propagation is not enough. We also need spectral norm bounds on the backward matrix and on the intermediate matrix

$$
\begin{equation*}
\left\|B D_{i, L} W_{L} \cdots D_{i, a} W_{a}\right\|_{2} \leq O(\sqrt{m / d}) \quad \text { and } \quad\left\|D_{i, a} W_{a} \cdots D_{i, b} W_{b}\right\|_{2} \leq O(\sqrt{L}) \tag{5.1}
\end{equation*}
$$

for every $a, b \in[L]$. Note that if one naively bounds the spectral norm by induction, then $\left\|D_{i, a} W_{a}\right\|_{2} \approx 2$ and it will exponentially blow up! Our careful analysis ensures that even when $L$ layers are stacked together, there is no exponential blow up in $L$.

The final lemma in this step proves that, as long as $\left\|x_{i}-x_{j}\right\| \geq \delta$, then

$$
\left\|h_{i, \ell}-h_{j, \ell}\right\| \geq \Omega(\delta) \text { for each layer } \ell \in[L] .
$$

Again, if one is willing to sacrifice an exponential factor and prove a lower bound $\delta \cdot 2^{-\Omega(L)}$, this will be easy. What is hard is to derive such lower bound without sacrificing more than a constant factor, but under the condition of $\delta \leq \frac{1}{C L}$. Details are in Section 5.8.

Step 2: stability after adversarial perturbation We show that for every $\vec{W}$ that is "close" to initialization, meaning $\left\|\vec{W}-\vec{W}^{(0)}\right\|_{2} \leq \omega$ for some $\omega \leq \frac{1}{\operatorname{poly}(L)}$, then
(a) the number of sign changes $\left\|D_{i, \ell}-\widetilde{D}_{i, \ell}\right\|_{0}$ is at most $O\left(m \omega^{2 / 3} L\right) \ll m$, and
(b) the perturbation amount $\left\|h_{i, \ell}-\widetilde{h}_{i, \ell}\right\| \leq O\left(\omega L^{5 / 2}\right) \ll 1$.

Since $\omega \leq \frac{1}{\operatorname{poly}(L)}$, both changes above become negligible. We call this "forward stability", and it is the most technical proof of this paper. Intuitively, both "(a) implies (b)" and "(b) implies (a)" are trivial to prove by matrix concentration. ${ }^{9}$ Unfortunately, one cannot apply such derivation by induction, because constants will blow up exponentially in the number of layers. We need some careful double induction introduced by [AZLS18], and details in Section 5.9.1. Another main result in this step is to derive stability for the backward matrix and the intermediate matrix. We show that when $w \leq \operatorname{poly}(L)$, (5.1) remains to hold. Details are in Section 5.9.2 and 5.9.3.

Remark 5.5 .1 . In the final proof, $\vec{W}$ is a point obtained by GD/SGD starting from $\vec{W}^{(0)}$, and thus $\vec{W}$ may depend on the randomness of $\vec{W}^{(0)}$. Since we cannot control how such randomness correlates, we argue for the above stability properties against all possible $\vec{W}$. This is why we call it "stability against adversarial perturbation."

Step 3: gradient bound The hard part of Theorem 5.11.1 is to show gradient lower bound. For this purpose, recall from Fact 5.2.3 that each sample $i \in[n]$ contributes to the

[^20]full gradient matrix by $D_{i, \ell}\left(\operatorname{Back}_{i, \ell+1}^{\top} \operatorname{loss}_{i}\right) h_{i, \ell-1}^{\top}$, where the backward matrix is applied to a loss vector loss ${ }_{i}$. To show this is large, intuitively, one wishes to show ( $\operatorname{Back}_{i, \ell+1}^{\top} \operatorname{loss}_{i}$ ) and $h_{i, \ell-1}$ are both vectors with large Euclidean norm.

Thanks to Step 1 and 2, this is not hard for a single sample $i \in[n]$. For instance, $\left\|\widetilde{h}_{i, \ell-1}\right\| \approx 1$ by Step 1 and we know $\left\|h_{i, \ell-1}-\widetilde{h}_{i, \ell-1}\right\| \leq o(1)$ from Step 2. One can also argue for $\mathrm{Back}_{i, \ell+1}^{\top} \operatorname{loss}_{i}$ but this is a bit harder. Indeed, when moving from random initialization $\vec{W}^{(0)}$ to $\vec{W}$, the loss vector loss $_{i}$ can change completely. Fortunately, loss ${ }_{i} \in \mathbb{R}^{d}$ is a lowdimensional vector, so one can calculate $\left\|\operatorname{Back}_{i, \ell+1}^{\top} u\right\|$ for every fixed $u$ and then apply $\epsilon$-net.

Finally, how to combine the above argument with multiple samples $i \in[n]$ ? These matrices are clearly not independent and may (in principle) sum up to zero. To deal with this, we use $\left\|h_{i, \ell}-h_{j, \ell}\right\| \geq \Omega(\delta)$ from Step 1. In other words, even if the contribution matrix $D_{i, \ell}\left(\operatorname{Back}_{i, \ell+1}^{\top} \operatorname{loss}_{i}\right) h_{i, \ell-1}^{\top}$ with respect to one sample $i$ is fixed, the contribution matrix with respect to other samples $j \in[n] \backslash\{i\}$ are still sufficiently random. Thus, the final gradient matrix will still be large. This idea comes from the prior work [LL18], ${ }^{10}$ and helps us prove Theorem 5.11.1. Details in Section 5.10 and 5.11.

Step 4: semi-smoothness In order to prove Theorem 5.12.1, one needs to argue, if we are currently at $\stackrel{\breve{W}}{W}$ and perturb it by $\vec{W}^{\prime}$, then how much does the objective change in second and higher order terms. This is different from our stability theory in Step 2, because Step 2 is regarding having a perturbation on $\vec{W}^{(0)}$; in contrast, in Theorem 5.12 .1 we need a (small) perturbation $\vec{W}^{\prime}$ on top of $\stackrel{\breve{W}}{ }$, which may already be a point perturbed from $\vec{W}^{(0)}$.

[^21]Nevertheless, we still manage to show that, if $\hbar_{i, \ell}$ is calculated on $\breve{\vec{W}}$ and $h_{i, \ell}$ is calculated on $\stackrel{\breve{W}}{W}+\vec{W}^{\prime}$, then $\left\|h_{i, \ell}-\hbar_{i, \ell}\right\| \leq O\left(L^{1.5}\right)\left\|W^{\prime}\right\|_{2}$. This, along with other properties to prove, ensures semi-smoothness. This explains Theorem 5.12.1 and details are in Section 5.12.

Remark 5.5.2. In other words, the amount of changes to each hidden layer (i.e., $h_{i, \ell}-\hbar_{i, \ell}$ ) is proportional to the amount of perturbation $\left\|W^{\prime}\right\|_{2}$. This may sound familiar to some readers: a ReLU function is Lipschitz continuous $|\phi(a)-\phi(b)| \leq|a-b|$, and composing Lipschitz functions still yield Lipschitz functions. What is perhaps surprising here is that this "composition" does not create exponential blow-up in the Lipschitz continuity parameter, as long as the amount of over-parameterization is sufficient and $\stackrel{\breve{W}}{ }$ is close to initialization.

### 5.6 Notable Extensions

Our Step 1 through Step 4 in Section 5.5 in fact give rise to a general plan for proving the training convergence of any neural network (at least with respect to the ReLU activation). Thus, it is expected that it can be generalized to many other settings. Not only we can have different number of neurons each layer, our theorems can be extended at least in the following three major directions. ${ }^{11}$

Different loss functions There is absolutely no need to restrict only to $\ell_{2}$ regression loss. We prove in Section 5.16 that, for any Lipschitz-smooth loss function $f$ :

Theorem 5.6.1 (arbitrary loss). From random initialization, with probability at least 1 -$e^{-\Omega\left(\log ^{2} m\right)}$, gradient descent with appropriate learning rate satisfy the following.

- If $f$ is nonconvex but $\sigma$-gradient dominant (a.k.a. Polyak-Łojasiewicz), GD finds $\epsilon$ error minimizer in ${ }^{12}$

$$
T=\widetilde{O}\left(\frac{\operatorname{poly}(n, L)}{\sigma \delta^{2}} \cdot \log \frac{1}{\epsilon}\right) \text { iterations }
$$

as long as $m \geq \widetilde{\Omega}\left(\operatorname{poly}\left(n, L, \delta^{-1}\right) \cdot d \sigma^{-2}\right)$.

- If $f$ is convex, then $G D$ finds $\epsilon$-error minimizer in

$$
T=\widetilde{O}\left(\frac{\operatorname{poly}(n, L)}{\delta^{2}} \cdot \frac{1}{\epsilon}\right) \text { iterations }
$$

[^22]as long as $m \geq \widetilde{\Omega}\left(\operatorname{poly}\left(n, L, \delta^{-1}\right) \cdot d \log \epsilon^{-1}\right)$.

- If $f$ is non-convex, then $S G D$ finds a point with $\|\nabla f\| \leq \epsilon$ in at most ${ }^{13}$

$$
T=\widetilde{O}\left(\frac{\operatorname{poly}(n, L)}{\delta^{2}} \cdot \frac{1}{\epsilon^{2}}\right) \text { iterations }
$$

as long as $m \geq \widetilde{\Omega}\left(\operatorname{poly}\left(n, L, \delta^{-1}\right) \cdot d \epsilon^{-1}\right)$.

- If $f$ is cross-entropy for multi-label classification, then GD attains $100 \%$ training accuracy in at most ${ }^{14}$.

$$
T=\widetilde{O}\left(\frac{\operatorname{poly}(n, L)}{\delta^{2}}\right) \text { iterations }
$$

as long as $m \geq \widetilde{\Omega}\left(\operatorname{poly}\left(n, L, \delta^{-1}\right) \cdot d\right)$.

We remark here that the $\ell_{2}$ loss is 1 -gradient dominant so it falls into the above general Theorem 5.6.1. One can also derive similar bounds for (mini-batch) SGD so we do not repeat the statements here.

Convolutional neural networks (CNN) There are lots of different ways to design CNN and each of them may require somewhat different proofs. In Section 5.17, we study the case when $A, W_{1}, \ldots, W_{L-1}$ are convolutional while $W_{L}$ and $B$ are fully connected. We assume for notational simplicity that each hidden layer has $\mathfrak{d}$ points each with $m$ channels. (In vision tasks, a point is a pixel). In the most general setting, these values $\mathfrak{d}$ and $m$ can vary across layers. We prove the following theorem:

[^23]Theorem 5.6.2 (CNN). As long as $m \geq \widetilde{\Omega}\left(\operatorname{poly}\left(n, L, \mathfrak{d}, \delta^{-1}\right) \cdot d\right)$, with high probability, $G D$ and SGD find an $\epsilon$-error solution for $\ell_{2}$ regression in $T=\widetilde{O}\left(\frac{\operatorname{poly}(n, L, \mathfrak{o})}{\delta^{2}} \cdot \log \epsilon^{-1}\right)$ iterations for CNN.

Of course, one can replace $\ell_{2}$ loss with other loss functions in Theorem 5.6.1 to get different types of convergence rates. We do not repeat them here.

Residual neural networks (ResNet) There are lots of different ways to design ResNet and each of them may require somewhat different proofs. In symbols, between two layers, one may study $h_{\ell}=\phi\left(h_{\ell-1}+W h_{\ell-1}\right), h_{\ell}=\phi\left(h_{\ell-1}+W_{2} \phi\left(W_{1} h_{\ell-1}\right)\right)$, or even $h_{\ell}=\phi\left(h_{\ell-1}+\right.$ $\left.W_{3} \phi\left(W_{2} \phi\left(W_{1} h_{\ell-1}\right)\right)\right)$. Since the main purpose here is to illustrate the generality of our techniques but not to attack each specific setting, in Section 5.18, we choose to consider the simplest residual setting $h_{\ell}=\phi\left(h_{\ell-1}+W h_{\ell-1}\right)$ (that was also studied for instance by theoretical work [HM17]). With appropriately chosen random initialization, we prove the following theorem:

Theorem 5.6.3 (ResNet). As long as $m \geq \widetilde{\Omega}\left(\operatorname{poly}\left(n, L, \delta^{-1}\right) \cdot d\right)$, with high probability, GD and SGD find an $\epsilon$-error solution for $\ell_{2}$ regression in $T=\widetilde{O}\left(\frac{\operatorname{poly}(n, L)}{\delta^{2}} \cdot \log \epsilon^{-1}\right)$ iterations for ResNet.

Of course, one can replace $\ell_{2}$ loss with other loss functions in Theorem 5.6.1 to get different types of convergence rates. We do not repeat them here.

### 5.7 Detailed Proofs

- In Section 5.8, we derive network properties at random initialization.
- In Section 5.9, we derive the stability theory against adversarial perturbation.
- In Section 5.10, we gradient upper and lower bounds at random initialization.
- In Section 5.11, we prove Theorem 5.11.1.
- In Section 5.12, we prove Theorem 5.12.1.
- In Section 5.13, we prove Theorem 5.13.1.
- In Section 5.14, we prove Theorem 5.14.1.
- In Section 5.15, we prove Theorem 5.15.1.


### 5.8 Properties at Random Initialization

Throughout this section we assume $\vec{W}, A$ and $B$ are randomly generated according to Definition 5.2.2. The diagonal sign matrices $D_{i, \ell}$ are also determined according to this random initialization.

### 5.8.1 Forward Propagation

Lemma 5.8.1 (forward propagation). If $\epsilon \in(0,1]$, with probability at least $1-O(n L)$. $e^{-\Omega\left(m \epsilon^{2} / L\right)}$ over the randomness of $A \in \mathbb{R}^{m \times \mathfrak{0}}$ and $\vec{W} \in\left(\mathbb{R}^{m \times m}\right)^{L}$, we have

$$
\forall i \in[n], \ell \in\{0,1, \ldots, L\} \quad: \quad\left\|h_{i, \ell}\right\| \in[1-\epsilon, 1+\epsilon] .
$$

Remark 5.8.1. Lemma 5.8 .1 is in fact trivial to prove if the allowed failure probability is instead $e^{-\Omega\left(m \epsilon^{2} / L^{2}\right)}$ (by applying concentration inequality layer by layer).

Before proving Lemma 5.8 .1 we note a simple mathematical fact:
Fact 5.8.2. Let $h, q \in \mathbb{R}^{p}$ be fixed vectors and $h \neq 0, W \in \mathbb{R}^{m \times p}$ be random matrix with i.i.d. entries $W_{i, j} \sim \mathcal{N}\left(0, \frac{2}{m}\right)$, and vector $v \in \mathbb{R}^{m}$ defined as $v_{i}=\phi\left((W h)_{i}\right)=\mathbb{1}_{(W(h+q))_{i} \geq 0}(W h)_{i}$. Then,

- $\left|v_{i}\right|$ follows i.i.d. from the following distribution: with half probability $\left|v_{i}\right|=0$, and with the other half probability $\left|v_{i}\right|$ follows from folded Gaussian distributions $\left|\mathcal{N}\left(0, \frac{2\|h\|^{2}}{m}\right)\right|$.
- $\frac{m\|v\|^{2}}{2\|h\|^{2}}$ is in distribution identical to $\chi_{\omega}^{2}$ (chi-square distribution of order $\omega$ ) where $\omega$ follows from binomial distribution $\mathcal{B}(m, 1 / 2)$.

Proof of Fact 5.8.2. We assume each vector $W_{i}$ is generated by first generating a gaussian vector $g \sim \mathcal{N}\left(0, \frac{2 I}{m}\right)$ and then setting $W_{i}= \pm g$ where the sign is chosen with half-half probability. Now, $\left|\left\langle W_{i}, h\right\rangle\right|=|\langle g, h\rangle|$ only depends on $g$, and is in distribution identical to $\left|\mathcal{N}\left(0, \frac{2\|h\|^{2}}{m}\right)\right|$. Next, after the sign is determined, the indicator $\mathbb{1}_{\left\langle W_{i}, h+q\right\rangle \geq 0}$ is 1 with half probability and 0 with another half. Therefore, $\left|v_{i}\right|$ satisfies the aforementioned distribution. As for $\|v\|^{2}$, letting $\omega \in\{0,1, \ldots, m\}$ be the variable indicator how many indicators are 1 , then $\omega \sim \mathcal{B}(m, 1 / 2)$ and $\frac{m\|v\|^{2}}{2\|h\|^{2}} \sim \chi_{\omega}^{2}$.

Proof of Lemma 5.8.1. We only prove Lemma 5.8 .1 for a fixed $i \in[n]$ and $\ell \in\{0,1,2, \ldots, L\}$ because we can apply union bound at the end. Below, we drop the subscript $i$ for notational convenience, and write $h_{i, \ell}$ and $x_{i}$ as $h_{\ell}$ and $x$ respectively.

Letting $\Delta_{\ell} \xlongequal{\text { def }} \frac{\left\|h_{\ell}\right\|^{2}}{\left\|h_{\ell-1}\right\|^{2}}$, we can write

$$
\log \left\|h_{b-1}\right\|^{2}=\log \|x\|^{2}+\sum_{\ell=0}^{b-1} \log \Delta_{\ell}=\sum_{\ell=0}^{b-1} \log \Delta_{\ell}
$$

According to Fact 5.8.2, fixing any $h_{\ell-1} \neq 0$ and letting $W_{\ell}$ be the only source of randomness, we have $\frac{m}{2} \Delta_{\ell} \sim \chi_{\omega}^{2}$ where $\omega \sim \mathcal{B}(m, 1 / 2)$. For such reason, for each $\Delta_{\ell}$, we can write $\Delta_{\ell}=\Delta_{\ell, \omega}$ where $\frac{m}{2} \Delta_{\ell, \omega} \sim \chi_{\omega}^{2}$ and $\omega \sim \mathcal{B}(m, 1 / 2)$. In the analysis below, we condition on the event that $\omega \in[0.4 m, 0.6 m]$; this happens with probability $\geq 1-e^{-\Omega(m)}$ for each layer $\ell \in[L]$. To simplify our notations, if this event does not hold, we set $\Delta_{\ell}=1$.

Expectation One can verify that $\mathbb{E}\left[\log \Delta_{\ell, \omega} \mid \omega\right]=\log \frac{4}{m}+\psi\left(\frac{\omega}{2}\right)$ where $\psi(h)=\frac{\Gamma^{\prime}(h)}{\Gamma(h)}$ is the digamma function. Using the bound $\log h-\frac{1}{h} \leq \psi(h) \leq \log h-\frac{1}{2 h}$ of digamma function, we have

$$
\log \frac{2 \omega}{m}-\frac{2}{\omega} \leq \mathbb{E}\left[\log \Delta_{\ell, \omega} \mid \omega\right] \leq \log \frac{2 \omega}{m}-\frac{1}{\omega} .
$$

Whenever $\omega \in[0.4 m, 0.6 m]$, we can write

$$
\log \frac{2 \omega}{m}=\log \left(1+\frac{2 \omega-m}{m}\right) \geq \frac{2 \omega-m}{m}-\left(\frac{2 \omega-m}{m}\right)^{2}
$$

It is easy to verify $\mathbb{E}_{\omega}\left[\frac{2 \omega-m}{m}\right]=0$ and $\mathbb{E}_{\omega}\left[\left(\frac{2 \omega-m}{m}\right)^{2}\right]=\frac{1}{m}$. Therefore,

$$
\underset{\omega}{\mathbb{E}}\left[\log \frac{2 \omega}{m}\right] \geq-\frac{1}{m}-\operatorname{Pr}[\omega \notin[0.4 m, 0.6 m]] \cdot \log \frac{2}{m} \geq-\frac{2}{m}
$$

Combining everything together, along with the fact that $\mathbb{E}_{\omega}\left[\log \frac{2 \omega}{m}\right] \leq \log \frac{\mathbb{E}[2 \omega]}{m}=0$, we have (when $m$ is sufficiently larger than a constant)

$$
\begin{equation*}
-\frac{4}{m} \leq \mathbb{E}\left[\log \Delta_{\ell}\right] \leq 0 \tag{5.2}
\end{equation*}
$$

Subgaussian Tail By standard tail bound for chi-square distribution, we know that

$$
\forall t \in[0, \infty): \quad \operatorname{Pr}\left[\left.\left|\frac{m}{2} \Delta_{\ell, \omega}-\omega\right| \leq t \right\rvert\, \omega\right] \geq 1-2 e^{-\Omega\left(t^{2} / \omega\right)}-e^{-\Omega(t)}
$$

Since we only need to focus on $\omega \geq 0.4 m$, this means

$$
\forall t \in[0, m]: \quad \operatorname{Pr}\left[\left.\left|\frac{m}{2} \Delta_{\ell, \omega}-\omega\right| \leq t \right\rvert\, \omega \geq 0.4 m\right] \geq 1-O\left(e^{-\Omega\left(t^{2} / m\right)}\right)
$$

On the other hand, by Chernoff-Hoeffding bound, we also have

$$
\underset{\omega}{\operatorname{Pr}}\left[\left|\omega-\frac{m}{2}\right| \leq t\right] \geq 1-O\left(e^{-\Omega\left(t^{2} / m\right)}\right)
$$

Together, using the definition $\Delta_{\ell}=\Delta_{\ell, \omega}$ (or $\Delta_{\ell}=1$ if $\omega \notin[0.4 m, 0.6 m]$ ), we obtain

$$
\forall t \in[0, m]: \quad \operatorname{Pr}\left[\left|\frac{m}{2} \Delta_{\ell}-\frac{m}{2}\right| \leq t\right] \geq 1-O\left(e^{-\Omega\left(t^{2} / m\right)}\right)
$$

This implies,

$$
\begin{equation*}
\forall t \in\left[0, \frac{m}{4}\right]: \quad \operatorname{Pr}\left[\left|\log \Delta_{\ell}\right| \leq \frac{t}{m}\right] \geq 1-O\left(e^{-\Omega\left(t^{2} / m\right)}\right) \tag{5.3}
\end{equation*}
$$

Now, let us make another simplification: define $\widehat{\Delta}_{\ell}=\Delta_{\ell}$ if $\left|\log \Delta_{\ell}\right| \leq \frac{1}{4}$ and $\widehat{\Delta}_{\ell}=1$ otherwise. In this way, (5.3) implies that $X=\log \widehat{\Delta}_{\ell}$ is an $O(m)$-subgaussian random variable.

Concentration Using martingale concentration on subgaussian variables (see for instance [Sha11]), we have for $\epsilon \in(0,1$ ],

$$
\operatorname{Pr}\left[\left|\sum_{\ell=0}^{b-1} \log \widehat{\Delta}_{\ell}-\mathbb{E}\left[\log \widehat{\Delta}_{\ell}\right]\right|>\epsilon\right] \leq O\left(e^{-\Omega\left(\epsilon^{2} m / L\right)}\right)
$$

Since with probability $\geq 1-L e^{-\Omega(m)}$ it satisfies $\widehat{\Delta}_{\ell}=\Delta_{\ell}$ for all $\ell \in[L]$, combining this with (5.2), we have

$$
\operatorname{Pr}\left[\left|\sum_{\ell=0}^{b-1} \log \Delta_{\ell}\right|>\epsilon\right] \leq O\left(e^{-\Omega\left(\epsilon^{2} m / L\right)}\right) .
$$

In other words, $\left\|h_{b-1}\right\|^{2} \in[1-\epsilon, 1+\epsilon]$ with probability at least $1-O\left(e^{-\Omega\left(\epsilon^{2} m / L\right)}\right)$.

### 5.8.2 Intermediate Layers

Lemma 5.8.3 (intermediate layers). Suppose $m \geq \Omega(n L \log (n L))$. With probability at least $\geq 1-e^{-\Omega(m / L)}$ over the randomness of $\vec{W} \in\left(\mathbb{R}^{m \times m}\right)^{L}$, for all $i \in[n], 1 \leq a \leq b \leq L$,
(a) $\left\|W_{b} D_{i, b-1} W_{b-1} \cdots D_{i, a} W_{a}\right\|_{2} \leq O(\sqrt{L})$.
(b) $\left\|W_{b} D_{i, b-1} W_{b-1} \cdots D_{i, a} W_{a} v\right\| \leq 2\|v\|$ for all vectors $v$ with $\|v\|_{0} \leq O\left(\frac{m}{L \log m}\right)$.
(c) $\left\|u^{\top} W_{b} D_{i, b-1} W_{b-1} \cdots D_{i, a} W_{a}\right\| \leq O(1)\|u\|$ for all vectors $u$ with $\|u\|_{0} \leq O\left(\frac{m}{L \log m}\right)$.

For any integer $s$ with $1 \leq s \leq O\left(\frac{m}{L \log m}\right)$, with probability at least $1-e^{-\Omega(s \log m)}$ over the randomness of $\vec{W} \in\left(\mathbb{R}^{m \times m}\right)^{L}$ :

$$
\text { (d) }\left|u^{\top} W_{b} D_{i, b-1} W_{b-1} \cdots D_{i, a} W_{a} v\right| \leq\|u\|\|v\| \cdot O\left(\frac{\sqrt{s \log m}}{\sqrt{m}}\right) \text { for all vectors } u, v \text { with }\|u\|_{0},\|v\|_{0} \leq
$$ $s$.

Proof. Again we prove the lemma for fixed $i, a$ and $b$ because we can take a union bound at the end. We drop the subscript $i$ for notational convenience.
(a) Let $z_{a-1}$ be any fixed unit vector, and define $z_{\ell}=D_{\ell} W_{\ell} \cdots D_{a} W_{a} z_{a-1}$. According to Fact 5.8.2 again, fixing any $z_{\ell-1}$ and letting $W_{\ell}$ be the only source of randomness, defining $\Delta_{\ell} \xlongequal{\text { def }} \frac{\left\|z_{\ell}\right\|^{2}}{\left\|z_{\ell-1}\right\|^{2}}$, we have that $\frac{m}{2} \Delta_{\ell}$ is distributed according to a $\chi_{\omega}^{2}$ where $\omega \sim$ $\mathcal{B}\left(m, \frac{1}{2}\right)$. Therefore, we have

$$
\log \left\|z_{b-1}\right\|^{2}=\log \left\|z_{a-1}\right\|^{2}+\sum_{\ell=a}^{b-1} \log \Delta_{\ell}=\sum_{\ell=a}^{b-1} \log \Delta_{\ell} .
$$

Using exactly the same proof as Lemma 5.8.1, we have

$$
\left\|z_{b-1}\right\|^{2}=\left\|W_{b} D_{b-1} W_{b-1} \cdots D_{a} W_{a} z_{a-1}\right\|^{2} \in[1-1 / 3,1+1 / 3]
$$

with probability at least $1-e^{-\Omega(m / L)}$. As a result, if we fix a subset $M \subseteq[m]$ of cardinality $|M| \leq O(m / L)$, taking $\epsilon$-net, we know that with probability at least $e^{-\Omega(m / L)}$, it satisfies

$$
\begin{equation*}
\left\|W_{b} D_{b-1} W_{b-1} \cdots D_{a} W_{a} u\right\| \leq 2\|u\| \tag{5.4}
\end{equation*}
$$

for all vectors $u$ whose coordinates are zeros outside $M$. Now, for an arbitrary unit vector $v \in \mathbb{R}^{m}$, we can decompose it as $v=u_{1}+\cdots+u_{N}$ where $N=O(L)$, each $u_{j}$ is non-zero only at $O(m / L)$ coordinates, and the vectors $u_{1}, \ldots, u_{N}$ are non-zeros on different coordinates. We can apply (5.4) for each each such $u_{j}$ and triangle inequality. This gives

$$
\left\|W_{b} D_{b-1} W_{b-1} \cdots D_{a} W_{a} v\right\| \leq 2 \sum_{j=1}^{N}\left\|u_{j}\right\| \leq 2 \sqrt{N}\left(\sum_{j=1}^{N}\left\|u_{j}\right\|^{2}\right)^{1 / 2} \leq O(\sqrt{L}) \cdot\|v\|
$$

(b) The proof of Lemma 5.8.3b is the same as Lemma 5.8.3a, except to take $\epsilon$-net over all $O\left(\frac{m}{L \log m}\right)$-sparse vectors $u$ and then applying union bound.
(c) Similar to the proof of Lemma 5.8.3a, for any fixed vector $v$, we have that with probability at least $1-e^{-\Omega(m / L)}$ (over the randomness of $W_{b-1}, \ldots, W_{1}, A$ ),

$$
\left\|D_{b-1} W_{b-1} \cdots D_{a} W_{a} v\right\| \leq 2\|v\|
$$

Conditioning on this event happens, using the randomness of $W_{b}$, we have for each fixed vector $u \in \mathbb{R}^{m}$, we have

$$
\operatorname{Pr}_{W_{b}}\left[\left|u^{\top} W_{b}\left(D_{b-1} W_{b-1} \cdots D_{a} W_{a} v\right)\right| \geq \frac{4}{\sqrt{L}}\|u\|\|v\|\right] \leq e^{-\Omega(m / L)} .
$$

Now consider the case that $v$ is a sparse vector that is only non-zero over some fixed index set $M \subseteq[m]$ (with $|M| \leq O(m / L))$, and that $u$ is of sparsity $s=O\left(\frac{m}{L \log m}\right)$. Taking $\epsilon$-net over all such possible vectors $u$ and $v$, we have with probability at least $1-e^{-\Omega(m / L)}$, for all vectors $u \in \mathbb{R}^{m}$ with $\|u\|_{0} \leq s$ and all vectors $v \in \mathbb{R}^{m}$ that have non-zeros only in $M$,

$$
\begin{equation*}
\left|u^{\top} W_{b}\left(D_{b-1} W_{b-1} \cdots D_{a} W_{a} v\right)\right| \leq \frac{8}{\sqrt{L}}\|u\|\|v\| \tag{5.5}
\end{equation*}
$$

Back to the case when $v$ is an arbitrary vector, we can partition [ $m$ ] into $N$ index sets $[m]=M_{1} \cup M_{2} \cup \cdots \cup M_{N}$ and write $v=v_{1}+v_{2}+\cdots+v_{N}$, where $N=O(L)$ and each $v_{j}$ is non-zero only in $M_{j}$. By applying (5.5) for $N$ times and using triangle inequality, we have

$$
\begin{aligned}
\left|u^{\top} W_{b}\left(D_{b-1} W_{b-1} \cdots D_{a} W_{a} v\right)\right| & \leq \sum_{j=1}^{N}\left|u^{\top} W_{b}\left(D_{b-1} W_{b-1} \cdots D_{a} W_{a} v_{j}\right)\right| \\
& \leq \frac{8}{\sqrt{L}}\|u\| \times \sum_{j=1}^{N}\left\|v_{j}\right\| \leq O(1) \times\|u\|\|v\|
\end{aligned}
$$

(d) We apply the same proof as Lemma 5.8.3c with minor changes to the parameters. We can show with probability at least $1-e^{-\Omega(m / L)}$ (over the randomness of $W_{b-1}, \ldots, W_{1}, A$ ), for a fixed vector $v \in \mathbb{R}^{m}$ :

$$
\left\|D_{b-1} W_{b-1} \cdots D_{a} W_{a} v\right\| \leq 2\|v\|
$$

Further using the randomness of $W_{b}$, we have that conditioning on the above event, fixing any $u \in \mathbb{R}^{m}$, with probability at least $1-e^{-\Omega(s \log m)}$ over the randomness of $W_{b}$ :

$$
\left|u W_{b}\left(D_{b-1} W_{b-1} \cdots D_{a} W_{a} v\right)\right| \leq\left(\frac{s \log m}{m}\right)^{1 / 2} \times O(\|v\|\|u\|)
$$

Finally, taking $\epsilon$-net over all possible vectors $u, v$ that are $s$ sparse, we have the desired result.

### 5.8.3 Backward Propagation

Lemma 5.8.4 (backward propagation). Suppose $m \geq \Omega(n L \log (n L))$. If $s \geq \Omega\left(\frac{d}{\log m}\right)$ and $s \leq O\left(\frac{m}{L \log m}\right)$, then with probability at least $1-e^{-\Omega(s \log m)}$, for all $i \in[n], a=1,2, \ldots, L+1$, (a) $\left|v^{\top} B D_{i, L} W_{L} \cdots D_{i, a} W_{a} u\right| \leq O\left(\frac{\sqrt{s \log m}}{\sqrt{d}}\right)\|v\|\|u\|$ for all $v \in \mathbb{R}^{d}$ and all $u \in \mathbb{R}^{m}$ with $\|u\|_{0} \leq s$.

With probability at least $\geq 1-e^{-\Omega(m / L)}$, for all $i \in[n], 1 \leq a \leq L$,
(b) $\left\|v^{\top} B D_{i, L} W_{L} \cdots D_{i, a} W_{a}\right\| \leq O(\sqrt{m / d})\|v\|$ for all vectors $u \in \mathbb{R}^{d}$ if $d \leq O\left(\frac{m}{L \log m}\right)$.

Proof. (a) The proof follows the same idea of Lemma 5.8 .3 (but choosing $b=L$ ). Given any fixed vector $u$, we have with probability at least $1-e^{-\Omega(m / L)}$ (over the randomness of $\left.W_{L}, \ldots, W_{1}, A\right)$,

$$
\left\|D_{L} W_{L} \cdots D_{a} W_{a} u\right\| \leq 2\|u\|
$$

Conditioning on this event happens, using the randomness of $B$ (recall each entry of $B$ follows from $\mathcal{N}\left(0, \frac{1}{d}\right)$ ), we have for each fixed vector $u \in \mathbb{R}^{m}$,

$$
\operatorname{Pr}_{B}\left[\left|v^{\top} B\left(D_{L} W_{L} \cdots D_{a} W_{a} u\right)\right| \geq \frac{\sqrt{s \log m}}{\sqrt{d}} \cdot O(\|u\|\|v\|)\right] \leq e^{-\Omega(s \log m)}
$$

Finally, one can take $\epsilon$-net over all $s$-sparse vectors $u \in \mathbb{R}^{m}$ and all vectors $v \in \mathbb{R}^{d}$ and apply union bound.
(b) The proof is identical to Lemma 5.8.3c, except the fact that each entry of $B$ follows from $\mathcal{N}\left(0, \frac{1}{d}\right)$ instead of $\mathcal{N}\left(0, \frac{2}{m}\right)$.

### 5.8.4 $\delta$-Separateness

Lemma 5.8.5 ( $\delta$-separateness). Let $m \geq \Omega\left(\frac{L \log (n L)}{\delta^{6}}\right)$. There exists some constant $C>1$ so that, if $\delta \leq \frac{1}{C L},\left\|x_{1}\right\|=\cdots=\left\|x_{n}\right\|=1$ and $\left\|x_{i}-x_{j}\right\| \geq \delta$ for every pair $i, j \in[n]$, then with probability at least $1-e^{-\Omega\left(\delta^{6} m / L\right)}$, we have :

$$
\forall i \neq j \in[n], \quad \forall \ell \in\{0,1, \ldots, L\}:\left\|\left(I-\frac{h_{i, \ell} h_{i, \ell}^{\top}}{\left\|h_{i, \ell}\right\|^{2}}\right) h_{j, \ell}\right\| \geq \frac{\delta}{2}
$$

Proof of Lemma 5.8.5. We first apply Lemma 5.8 .1 to show that $\left\|h_{i, \ell}\right\| \in\left[1-\delta^{3} / 10,1+\right.$ $\left.\delta^{3} / 10\right]$. Next we prove Lemma 5.8 .5 by induction.

In the base case of $\ell=-1$, since $\left\|x_{i}-x_{j}\right\| \geq \delta$ by our Assumption 5.2.1 and without loss of generality $\left\|x_{i}\right\|=1$ and $\left(x_{i}\right)_{\mathfrak{O}}=\frac{1}{\sqrt{2}}$, we already have

$$
\left\|\left(I-\frac{h_{i, \ell} h_{i, \ell}^{\top}}{\left\|h_{i, \ell}\right\|^{2}}\right) h_{j, \ell}\right\|^{2}=\left\|\left(I-\frac{x_{i} x_{i}^{\top}}{\left\|x_{i}\right\|^{2}}\right) x_{j}\right\|^{2}=\left\|x_{j}-x_{i} \cdot\left\langle x_{i}, x_{j}\right\rangle\right\|^{2}=1-\left(\left\langle x_{i}, x_{j}\right\rangle\right)^{2} \geq \frac{3}{4} \delta^{2}
$$

Suppose $h_{i, \ell-1}$ and $h_{j, \ell-1}$ are fixed and satisfies $\left\|\left(I-\frac{h_{i, \ell-1} h_{i, \ell-1}^{\top}}{\left\|h_{i, \ell-1}\right\|^{2}}\right) h_{j, \ell-1}\right\|^{2} \geq \delta_{\ell-1}^{2}$ for some $\delta_{\ell-1} \geq \delta / 2$. We write $W_{\ell} h_{i, \ell-1}=\vec{g}_{1}$ where $\vec{g}_{1} \sim N\left(0, \frac{2\left\|h_{i, \ell-1}\right\|^{2}}{m} I\right)$.

Denoting by $\widehat{h}=h_{i, \ell-1} /\left\|h_{i, \ell-1}\right\|$, we can write $W_{\ell} h_{j, \ell-1}=W_{\ell} \widehat{h} \widehat{h}^{\top} h_{j, \ell-1}+W_{\ell}(I-$ $\left.\widehat{h} \widehat{h}^{\top}\right) h_{j, \ell-1}$ and the randomness of the two terms are independent. In particular, we can write

$$
\begin{equation*}
W_{\ell} h_{j, \ell-1}=\frac{\left\langle h_{i, \ell-1}, h_{j, \ell-1}\right\rangle}{\left\|h_{i, \ell-1}\right\|^{2}} \cdot \vec{g}_{1}+\left\|\left(I-\widehat{h} \widehat{h}^{\top}\right) h_{j, \ell-1}\right\| \cdot \vec{g}_{2} \tag{5.6}
\end{equation*}
$$

where $\vec{g}_{2} \sim \mathcal{N}\left(0, \frac{2}{m} I\right)$ is independent of $g_{1}$. Applying Claim 5.8.6 for each coordinate $k \in[m]$ (and re-scaling by $\frac{m}{\left\|h_{i, \ell-1}\right\|^{2}}$, we have

$$
\mathbb{E}\left[\left(\phi\left(W_{\ell} h_{i, \ell-1}\right)-\phi\left(W_{\ell} h_{j, \ell-1}\right)\right)_{k}^{2}\right] \geq\left(\frac{\delta_{\ell-1}}{\left\|h_{i, \ell-1}\right\|}\right)^{2}\left(1-\frac{\delta_{\ell-1}}{\left\|h_{i, \ell-1}\right\|}\right) \cdot \frac{\left\|h_{i, \ell-1}\right\|^{2}}{m} \geq \frac{\delta_{\ell-1}^{2}\left(1-O\left(\delta_{\ell-1}\right)\right)}{m}
$$

Applying Chernoff bound (on independent subgaussian random variables), we have with probability at least $1-e^{-\Omega\left(\delta_{\ell-1}^{4} m\right)},{ }^{15}$

$$
\left\|h_{i, \ell}-h_{j, \ell}\right\|^{2}=\left\|\phi\left(W_{\ell} h_{i, \ell-1}\right)-\phi\left(W_{\ell} h_{j, \ell-1}\right)\right\|^{2} \geq \delta_{\ell-1}^{2}\left(1-O\left(\delta_{\ell-1}\right)\right)
$$

Since $\left\|h_{i, \ell}\right\|$ and $\left\|h_{j, \ell}\right\|$ are close to 1 , we have

$$
\begin{aligned}
\left\|\left(I-\frac{h_{i, \ell} h_{i, \ell}^{\top}}{\left\|h_{i, \ell}\right\|^{2}}\right) h_{j, \ell}\right\|^{2} & =\left\|h_{j, \ell}\right\|^{2}-\frac{\left\langle h_{i, \ell}, h_{j, \ell}\right\rangle^{2}}{\left\|h_{i, \ell}\right\|^{2}} \\
& =\left\|h_{j, \ell}\right\|^{2}+\frac{\left\|h_{i, \ell}-h_{j, \ell}\right\|^{2}-\left\|h_{i, \ell}\right\|^{2}-\left\|h_{j, \ell}\right\|^{2}}{2\left\|h_{i, \ell}\right\|^{2}} \geq \delta_{\ell-1}^{2}\left(1-O\left(\delta_{\ell-1}\right)\right)
\end{aligned}
$$

### 5.8.4.1 Auxiliary Claim

The following mathematical fact is needed in the proof of Lemma 5.8.5. Its proof is by carefully integrating the PDF of Gaussian distribution.

Claim 5.8.6. Given $g_{1}, g_{2} \sim \mathcal{N}(0,2)$, constant $\alpha \in \mathbb{R}$ and $\delta \in\left[0, \frac{1}{6}\right]$, we have

$$
\underset{g_{1}, g_{2}}{\mathbb{E}}\left[\left(\phi\left(g_{1}\right)-\phi\left(\alpha g_{1}+\delta g_{2}\right)\right)^{2}\right] \geq \delta^{2}(1-\delta)
$$

Proof of Claim 5.8.6. We first tackle two easy cases.
Suppose $a<\frac{3}{4}$. If so, then with probability at least 0.3 we have $g_{1}>1$. If this happens, then with probability at least $1 / 2$ we have $g_{2}<0$. If both happens, we have

$$
\phi\left(g_{1}\right)-\phi\left(\alpha g_{1}+\delta g_{2}\right)=g_{1}-\phi\left(\alpha g_{1}+\delta g_{2}\right) \geq g_{1}-\alpha g_{1} \geq \frac{1}{4}
$$

[^24]Therefore, we have if $a<\frac{3}{4}$ then the expectation is at least 0.03 . For similar reason, if $a>\frac{5}{4}$ we also have the expectation is at least 0.03 . In the remainder of the proof, we assume $\alpha \in\left[\frac{3}{4}, \frac{5}{4}\right]$.

If $g_{1} \geq 0$, we have

$$
\begin{aligned}
f\left(g_{1}\right) & \stackrel{\text { def }}{=} \underset{g_{2}}{\mathbb{E}}\left[\left(\phi\left(g_{1}\right)-\phi\left(\alpha g_{1}+\delta g_{2}\right)\right)^{2} \mid g_{1} \geq 0\right] \\
& =\int_{0}^{\infty} \frac{\left(x-g_{1}\right)^{2} \exp \left(-\frac{\left(x-\alpha g_{1}\right)^{2}}{4 \delta^{2}}\right)}{\sqrt{4 \pi \delta^{2}}} \mathrm{~d} x \\
& =\frac{(\alpha-2) \delta g_{1} e^{-\frac{\alpha^{2} g_{1}^{2}}{4 \delta^{2}}}}{\sqrt{\pi}}+\frac{1}{2}\left((\alpha-1)^{2} g_{1}^{2}+2 \delta^{2}\right)\left(\operatorname{erf}\left(\frac{\alpha g_{1}}{2 \delta}\right)+1\right) .
\end{aligned}
$$

If $g_{1}<0$, we have

$$
\begin{aligned}
f\left(g_{1}\right) & \stackrel{\text { def }}{=} \underset{g_{2}}{\mathbb{E}}\left[\left(\phi\left(g_{1}\right)-\phi\left(\alpha g_{1}+\delta g_{2}\right)\right)^{2} \mid g_{1}<0\right] \\
& =\int_{0}^{\infty} \frac{x^{2} \exp \left(-\frac{\left(x-\alpha g_{1}\right)^{2}}{4 \delta^{2}}\right)}{\sqrt{4 \pi \delta^{2}}} \mathrm{~d} x \\
& =\frac{1}{2}\left(\alpha^{2} g_{1}^{2}+2 \delta^{2}\right)\left(\operatorname{erf}\left(\frac{\alpha g_{1}}{2 \delta}\right)+1\right)+\frac{\alpha \delta g_{1} e^{-\frac{\alpha^{2} g_{1}^{2}}{4 \delta^{2}}}}{\sqrt{\pi}} .
\end{aligned}
$$

Overall, we have

$$
\begin{aligned}
& \underset{g_{1}, g_{2}}{\mathbb{E}}\left[\left(\phi\left(g_{1}\right)-\phi\left(\alpha g_{1}+\delta g_{2}\right)\right)^{2}\right] \\
= & \int_{0}^{\infty} \frac{f(g) \exp \left(-\frac{g^{2}}{4}\right)}{\sqrt{4 \pi}} \mathrm{~d} g+\int_{-\infty}^{0} \frac{f(g) \exp \left(-\frac{g^{2}}{4}\right)}{\sqrt{4 \pi}} \mathrm{~d} g \\
= & \left(\frac{(\alpha-1)^{2} \alpha \delta}{\pi\left(\alpha^{2}+\delta^{2}\right)}+\frac{(\alpha-2) \delta^{3}}{\pi\left(\alpha^{2}+\delta^{2}\right)}+\frac{1}{2}\left((\alpha-1)^{2}+\delta^{2}\right)+\frac{1}{\pi}\left((\alpha-1)^{2}+\delta^{2}\right) \arctan \left(\frac{\alpha}{\delta}\right)\right) \\
& \quad+\frac{1}{2 \pi}\left(\pi\left(\alpha^{2}+\delta^{2}\right)-2\left(\alpha^{2}+\delta^{2}\right) \arctan \left(\frac{\alpha}{\delta}\right)-2 \alpha \delta\right) \\
= & \frac{\delta\left(-2 \alpha^{2}+\alpha-2 \delta^{2}\right)}{\pi\left(\alpha^{2}+\delta^{2}\right)}+\frac{(1-2 \alpha) \arctan \left(\frac{\alpha}{\delta}\right)}{\pi}+(\alpha-1) \alpha+\delta^{2}+\frac{1}{2} \\
= & \left(\alpha^{2}-2 \alpha+1\right)+\delta^{2}+\frac{2}{\pi} \sum_{k=1}^{\infty}(-1)^{k} \frac{(\alpha+k) \delta^{2 k+1}}{(2 k+1) \alpha^{2 k+1}} .
\end{aligned}
$$

It is easy to see that, as long as $\delta \leq \alpha$, we always have $\frac{(\alpha+k) \delta^{2 k+1}}{(2 k+1) \alpha^{2 k+1}} \geq \frac{(\alpha+k+1) \delta^{2 k+3}}{(2 k+3) \alpha^{2 k+3}}$. Therefore

$$
\underset{g_{1}, g_{2}}{\mathbb{E}}\left[\left(\phi\left(g_{1}\right)-\phi\left(\alpha g_{1}+\delta g_{2}\right)\right)^{2}\right] \geq\left(\alpha^{2}-2 \alpha+1\right)+\delta^{2}-\frac{2}{\pi} \frac{(\alpha+1) \delta^{3}}{3 \alpha^{3}} \geq \delta^{2}(1-\delta)
$$

### 5.9 Stability against Adversarial Weight Perturbations

Let $A, B$ and $\vec{W}^{(0)}=\left(\widetilde{W}_{1}, \ldots, \widetilde{W}_{L}\right)$ be matrices at random initialization (see Definition 5.2.2), and throughout this section, we consider (adversarially) perturbing $\vec{W}$ by $\overrightarrow{W^{\prime}}=\left(W_{1}^{\prime}, \ldots, W_{L}^{\prime}\right)$ satisfying $\left\|\vec{W}^{\prime}\right\|_{2} \leq \omega$ (meaning, $\left\|W_{\ell}^{\prime}\right\|_{2} \leq \omega$ for every $\ell \in[L]$ ). We stick to the following notations in this section

## Definition 5.9.1.

$$
\begin{array}{lll}
\widetilde{g}_{i, 0}=A x_{i} & g_{i, 0}=A x_{i} & \text { for } i \in[n] \\
\widetilde{h}_{i, 0}=\phi\left(A x_{i}\right) & h_{i, 0}=\phi\left(A x_{i}\right) & \text { for } i \in[n] \\
\widetilde{g}_{i, \ell}=\widetilde{W}_{\ell} h_{i, \ell-1} & g_{i, \ell}=\left(\widetilde{W}_{\ell}+W_{\ell}^{\prime}\right) h_{i, \ell-1} & \text { for } i \in[n] \text { and } \ell \in[L] \\
\widetilde{h}_{i, \ell}=\phi\left(\widetilde{W}_{\ell} h_{i, \ell-1}\right) & h_{i, \ell}=\phi\left(\left(\widetilde{W}_{\ell}+W_{\ell}^{\prime}\right) h_{i, \ell-1}\right) & \text { for } i \in[n] \text { and } \ell \in[L]
\end{array}
$$

Define diagonal matrices $\widetilde{D}_{i, \ell} \in \mathbb{R}^{m \times m}$ and $D_{i, \ell} \in \mathbb{R}^{m \times m}$ by letting $\left(\widetilde{D}_{i, \ell}\right)_{k, k}=\mathbb{1}_{\left(\widetilde{g}_{i, \ell}\right)_{k} \geq 0}$ and $\left(D_{i, \ell}\right)_{k, k}=\mathbb{1}_{\left(g_{i, \ell}\right)_{k} \geq 0}, \forall k \in[m]$. Accordingly, we let $g_{i, \ell}^{\prime}=g_{i, \ell}-\widetilde{g}_{i, \ell}, h_{i, \ell}^{\prime}=h_{i, \ell}-\widetilde{h}_{i, \ell}$, and diagonal matrix $D_{i, \ell}^{\prime}=D_{i, \ell}-\widetilde{D}_{i, \ell}$.

### 5.9.1 Forward Perturbation

Lemma 5.9.1 (forward perturbation). Suppose $\omega \leq \frac{1}{C L^{9 / 2} \log ^{3} m}$ for some sufficiently large constant $C>1$. With probability at least $1-e^{-\Omega\left(m \omega^{2 / 3} L\right)}$, for every $\vec{W}^{\prime}$ satisfying $\left\|\vec{W}^{\prime}\right\|_{2} \leq \omega$,
(a) $g_{i, \ell}^{\prime}$ can be written as $g_{i, \ell}^{\prime}=g_{i, \ell, 1}^{\prime}+g_{i, \ell, 2}^{\prime}$ where $\left\|g_{i, \ell, 1}^{\prime}\right\| \leq O\left(\omega L^{3 / 2}\right)$ and $\left\|g_{i, \ell, 2}^{\prime}\right\|_{\infty} \leq$ $O\left(\frac{\omega L^{5 / 2} \sqrt{\log m}}{\sqrt{m}}\right)$
(b) $\left\|D_{i, \ell}^{\prime}\right\|_{0} \leq O\left(m \omega^{2 / 3} L\right)$ and $\left\|D_{i, \ell}^{\prime} g_{i, \ell}\right\| \leq O\left(\omega L^{3 / 2}\right)$.
(c) $\left\|g_{i, \ell}^{\prime}\right\|,\left\|h_{i, \ell}^{\prime}\right\| \leq O\left(\omega L^{5 / 2} \sqrt{\log m}\right)$.

Proof of Lemma 5.9.1. In our proof below, we drop the subscript with respect to $i$ for notational simplicity, and one can always take a union bound over all possible indices $i$ at the end.

Using Lemma 5.8.1, we can first assume that $\left\|\widetilde{h}_{\ell}\right\|,\left\|\widetilde{g}_{\ell}\right\| \in\left[\frac{2}{3}, \frac{4}{3}\right]$ for all $\ell$. This happens with probability at least $1-e^{-\Omega(m / L)}$. We also assume $\left\|\prod_{b=\ell}^{a+1} \widetilde{W}_{b} \widetilde{D}_{b-1}\right\|_{2} \leq c_{1} \sqrt{L}$ where $c_{1}>0$ is the hidden constant in Lemma 5.8.3a.

We shall inductively prove Lemma 5.9.1. In the base case $\ell=0$, we have $g_{\ell}^{\prime}=0$ so all the statements holds. In the remainder of the proof, we assume that Lemma 5.9.1 holds for $\ell-1$ and we shall prove the three statements for layer $\ell$. To help the readers understand how the constants propagate without blowing up, we shall prove $\left\|g_{i, \ell, 1}^{\prime}\right\| \leq 4 c_{1} L^{1.5} \omega$ in Lemma 5.9.1a without the big- $O$ notation, while for all other terms we use big- $O$ to hide polynomial dependency on $c_{1} .{ }^{16}$

[^25]We first carefully rewrite:

$$
\begin{aligned}
g_{\ell}^{\prime} & =\left(\widetilde{W}_{\ell}+W_{\ell}^{\prime}\right)\left(\widetilde{D}_{\ell-1}+D_{\ell-1}^{\prime}\right)\left(\widetilde{g}_{\ell-1}+g_{\ell-1}^{\prime}\right)-\widetilde{W}_{\ell} \widetilde{D}_{\ell-1} \widetilde{g}_{\ell-1} \\
& =W_{\ell}^{\prime}\left(\widetilde{D}_{\ell-1}+D_{\ell-1}^{\prime}\right)\left(\widetilde{g}_{\ell-1}+g_{\ell-1}^{\prime}\right)+\widetilde{W}_{\ell} D_{\ell-1}^{\prime}\left(\widetilde{g}_{\ell-1}+g_{\ell-1}^{\prime}\right)+\widetilde{W}_{\ell} \widetilde{D}_{\ell-1} g_{\ell-1}^{\prime} \\
& =\cdots \\
& =\sum_{a=1}^{\ell}\left(\prod_{b=\ell}^{a+1} \widetilde{W}_{b} \widetilde{D}_{b-1}\right)(\underbrace{W_{a}^{\prime}\left(\widetilde{D}_{a-1}+D_{a-1}^{\prime}\right)\left(\widetilde{g}_{a-1}+g_{a-1}^{\prime}\right)}_{(\diamond)}+\underbrace{\widetilde{W}_{a} D_{a-1}^{\prime}\left(\widetilde{g}_{a-1}+g_{a-1}^{\prime}\right)}_{(\diamond)})
\end{aligned}
$$

For each term in $(\diamond)$, we have

$$
\begin{aligned}
& \left\|\left(\prod_{b=\ell}^{a+1} \widetilde{W}_{b} \widetilde{D}_{b-1}\right)\left(W_{a}^{\prime}\left(\widetilde{D}_{a-1}+D_{a-1}^{\prime}\right)\left(\widetilde{g}_{a-1}+g_{a-1}^{\prime}\right)\right)\right\| \\
\leq & \left\|\prod_{b=\ell}^{a+1} \widetilde{W}_{b} \widetilde{D}_{b-1}\right\|_{2} \cdot\left\|W_{a}^{\prime}\right\|_{2} \cdot\left\|\widetilde{D}_{a-1}+D_{a-1}^{\prime}\right\|_{2} \cdot\left\|\widetilde{g}_{a-1}+g_{a-1}^{\prime}\right\| \\
\stackrel{\otimes}{\leq} & c_{1} \cdot \omega \cdot 1 \cdot\left\|\widetilde{g}_{a-1}+g_{a-1}^{\prime}\right\| \stackrel{(2}{\leq} 2 c_{1} \sqrt{L} \omega+O\left(\omega^{2} L^{3} \sqrt{\log m}\right) .
\end{aligned}
$$

Above, inequality (1) uses Lemma 5.8.3a and $\left\|\widetilde{D}_{a-1}+D_{a-1}^{\prime}\right\|_{2}=\left\|D_{a-1}\right\|_{2} \leq 1$; and inequality (2) has used $\left\|\widetilde{g}_{\ell}\right\| \leq 2$ and our inductive assumption Lemma 5.9.1c. By triangle inequality, we have

$$
g_{\ell}^{\prime}=\overrightarrow{\operatorname{err}}_{1}+\sum_{a=1}^{\ell}\left(\prod_{b=\ell}^{a+1} \widetilde{W}_{b} \widetilde{D}_{b-1}\right)(\underbrace{\widetilde{W}_{a} D_{a-1}^{\prime}\left(\widetilde{g}_{a-1}+g_{a-1}^{\prime}\right)}_{(\ominus)})
$$

where $\left\|\overrightarrow{\operatorname{err}}_{1}\right\| \leq 2 c_{1} L^{1.5} \omega+O\left(\omega^{2} L^{4} \sqrt{\log m}\right)$. We next look at each term in ( () . For each $a=2,3, \ldots, \ell$, we let

$$
x \stackrel{\text { def }}{=} D_{a-1}^{\prime}\left(\widetilde{g}_{a-1}+g_{a-1}^{\prime}\right)=D_{a-1}^{\prime}\left(\widetilde{W}_{a-1} \widetilde{h}_{a-1}+g_{a-1}^{\prime}\right) .
$$

If we re-scale $x$ by $\frac{1}{\left\|\tilde{h}_{a-1}\right\|}$ (which is a constant in $[0.75,1.5]$ ), we can apply Claim 5.9.2 (with parameter choices in Corollary 5.9.3) on $x$ and this tells us, with probability at least
$1-e^{-\Omega\left(m \omega^{2 / 3} L\right)}:$

$$
\begin{equation*}
\|x\|_{0} \leq O\left(m \omega^{2 / 3} L\right) \quad \text { and } \quad\|x\| \leq O\left(\omega L^{3 / 2}\right) \tag{5.7}
\end{equation*}
$$

Next, each term in $(\Omega)$ contributes to $g_{\ell}^{\prime}$ by

$$
y=\left(\prod_{b=\ell}^{a+1} \widetilde{W}_{b} \widetilde{D}_{b-1}\right) \widetilde{W}_{a}\left(D_{a-1}^{\prime}\left(\widetilde{g}_{a-1}+g_{a-1}^{\prime}\right)\right)
$$

using (5.7) and Claim 5.9 .4 (with $s=O\left(m \omega^{2 / 3} L\right)$ ), we have with probability at least 1 -$e^{-\Omega(s \log m)}$, one can write $y=y_{1}+y_{2}$ for

$$
\left\|y_{1}\right\| \leq O\left(\omega L^{3 / 2} \cdot L^{1 / 2} \omega^{1 / 3} \log m\right) \quad \text { and } \quad\left\|y_{2}\right\|_{\infty} \leq O\left(\omega L^{3 / 2} \cdot \frac{\sqrt{\log m}}{\sqrt{m}}\right)
$$

And therefore by triangle inequality we can write

$$
g_{\ell}^{\prime}=\overrightarrow{\operatorname{err}}_{1}+\overrightarrow{\operatorname{err}}_{2}+\overrightarrow{\operatorname{err}}_{3}
$$

where $\left\|\overrightarrow{e r r}_{2}\right\| \leq O\left(L \cdot \omega L^{3 / 2} \cdot L^{1 / 2} \omega^{1 / 3} \log m\right)=O\left(\omega^{4 / 3} L^{3} \log m\right)$ and $\left\|\overrightarrow{e r r}_{3}\right\|_{\infty} \leq O\left(L \cdot \omega L^{3 / 2}\right.$. $\left.\frac{\sqrt{\log m}}{\sqrt{m}}\right)$. Together with the upper bound on $\overrightarrow{\operatorname{err}}_{1}$, we have

$$
\left\|\overrightarrow{e r r}_{1}+\overrightarrow{\operatorname{err}}_{2}\right\| \leq 2 c_{1} L^{1.5} \omega+O\left(\omega^{2} L^{4} \sqrt{\log m}+\omega^{4 / 3} L^{3} \log m\right)
$$

We emphasize that the above big- $O$ notion can hide polynomial dependency on $c_{1}$. Nevertheless, when $\omega$ is sufficiently small, the above term is at most $4 c_{1} L^{1.5} \omega$. This finishes the proof of Lemma 5.9.1a for layer $\ell$ without blowing up the constant. Finally,

- Lemma 5.9.1b is due to (5.7),
- $g_{\ell}^{\prime}$ part of Lemma 5.9.1c is a simple corollary of Lemma 5.9.1a, and
- $h_{\ell}^{\prime}$ part of Lemma 5.9.1c is due to $h_{\ell}^{\prime}=D_{\ell} g_{\ell}^{\prime}+D_{\ell}^{\prime} g_{\ell}$ together with the bound on $\left\|g_{\ell}^{\prime}\right\|$ and the bound on $D_{\ell}^{\prime} g_{\ell}$ from Lemma 5.9.1b.


### 5.9.1.1 Auxiliary Claim

Claim 5.9.2. Suppose $\delta_{2} \in[0, O(1)]$ and $\delta_{\infty} \in\left[0, \frac{1}{4 \sqrt{m}}\right]$. Suppose $\widetilde{W} \in \mathbb{R}^{m \times m}$ is a random matrix with entries drawn i.i.d. from $\mathcal{N}\left(0, \frac{2}{m}\right)$. With probability at least $1-e^{-\Omega\left(m^{3 / 2} \delta_{\infty}\right)}$, the following holds. Fix any unit vector $\widetilde{h} \in \mathbb{R}^{m}$, and for all $g^{\prime} \in \mathbb{R}^{m}$ that can be written as

$$
g^{\prime}=g_{1}^{\prime}+g_{2}^{\prime} \text { where }\left\|g_{1}^{\prime}\right\| \leq \delta_{2} \text { and }\left\|g_{2}^{\prime}\right\|_{\infty} \leq \delta_{\infty}
$$

Let $D^{\prime} \in \mathbb{R}^{m \times m}$ be the diagonal matrix where $\left(D^{\prime}\right)_{k, k}=\mathbb{1}_{\left(\widetilde{W h}+g^{\prime}\right)_{k} \geq 0}-\mathbb{1}_{(\widetilde{W} \widetilde{h})_{k} \geq 0}, \forall k \in[m]$. Then, letting $x=D^{\prime}\left(\widetilde{W} \widetilde{h}+g^{\prime}\right) \in \mathbb{R}^{m}$, we have

$$
\|x\|_{0} \leq\left\|D^{\prime}\right\|_{0} \leq O\left(m\left(\delta_{2}\right)^{2 / 3}+\delta_{\infty} m^{3 / 2}\right) \quad \text { and } \quad\|x\| \leq O\left(\delta_{2}+\left(\delta_{\infty}\right)^{3 / 2} m^{3 / 4}\right)
$$

Corollary 5.9.3. In particular, if $\omega L^{3 / 2} \leq O(1)$, then with probability at least $1-e^{-\Omega\left(m \omega^{2 / 3} L\right)}$, for every $g^{\prime}=g_{1}^{\prime}+g_{2}^{\prime}$ with $\left\|g_{1}^{\prime}\right\| \leq O\left(\omega L^{3 / 2}\right)$ and $\left\|g_{2}^{\prime}\right\|_{\infty} \leq O\left(\frac{\omega^{2 / 3} L}{m^{1 / 2}}\right)$, it satisfies

$$
\left\|D^{\prime}\right\|_{0} \leq O\left(m \omega^{2 / 3} L\right) \quad \text { and } \quad\|x\| \leq O\left(\omega L^{3 / 2}\right)
$$

Proof of Claim 5.9.2. We first observe $\widetilde{g}=\widetilde{W} \widetilde{h}$ follows from $\mathcal{N}\left(0, \frac{2 I}{m}\right)$ regardless of the choice of $\widetilde{h}$. Therefore, in the remainder of the proof, we just focus on the randomness of $\widetilde{g}$.

We also observe that $\left(D^{\prime}\right)_{j, j}$ is non-zero for some diagonal $j \in[m]$ only if

$$
\begin{equation*}
\left|\left(g_{1}^{\prime}+g_{2}^{\prime}\right)_{j}\right|>\left|(\widetilde{g})_{j}\right| \tag{5.8}
\end{equation*}
$$

Let $\xi \leq \frac{1}{2 \sqrt{m}}$ be a parameter to be chosen later. We shall make sure that $\left\|g_{2}^{\prime}\right\|_{\infty} \leq \xi / 2$.

- We denote by $S_{1} \subseteq[m]$ the index sets where $j$ satisfies $\left|(\widetilde{g})_{j}\right| \leq \xi$. Since we know $(\widetilde{g})_{j} \sim \mathcal{N}(0,2 / m)$, we have $\operatorname{Pr}\left[\left|(\widetilde{g})_{j}\right| \leq \xi\right] \leq O(\xi \sqrt{m})$ for each $j \in[m]$. Using Chernoff bound for all $j \in[m]$, we have with probability at least $1-e^{-\Omega\left(m^{3 / 2} \xi\right)}$,

$$
\left|S_{1}\right|=\left|\left\{i \in[m]:\left|(\widetilde{g})_{j}\right| \leq \xi\right\}\right| \leq O\left(\xi m^{3 / 2}\right) .
$$

Now, for each $j \in S_{1}$ such that $x_{j} \neq 0$, we must have $\left|x_{j}\right|=\left|\left(\widetilde{g}+g_{1}^{\prime}+g_{2}^{\prime}\right)_{j}\right| \leq\left|\left(g_{1}^{\prime}\right)_{j}\right|+2 \xi$ so we can calculate the $\ell_{2}$ norm of $x$ on $S_{1}$ :

$$
\sum_{i \in S_{1}} x_{j}^{2} \leq O\left(\left\|g_{1}^{\prime}\right\|^{2}+\xi^{2}\left|S_{1}\right|\right) \leq O\left(\left\|g_{1}^{\prime}\right\|^{2}+\xi^{3} m^{3 / 2}\right)
$$

- We denote by $S_{2} \subseteq[m] \backslash S_{1}$ the index set of all $j \in[m] \backslash S_{1}$ where $x_{j} \neq 0$. Using (5.8), we have for each $j \in S_{2}$ :

$$
\left|\left(g_{1}^{\prime}\right)_{j}\right| \geq\left|(\widetilde{g})_{j}\right|-\left|\left(g_{2}^{\prime}\right)_{j}\right| \geq \xi-\left\|g_{2}^{\prime}\right\|_{\infty} \geq \xi / 2
$$

This means

$$
\left|S_{2}\right| \leq \frac{4\left\|g_{1}^{\prime}\right\|^{2}}{\xi^{2}}
$$

Now, for each $j \in S_{2}$ where $x_{j} \neq 0$, we know that the signs of $\left(\widetilde{g}+g_{1}^{\prime}+g_{2}^{\prime}\right)_{j}$ and $(\widetilde{g})_{j}$ are opposite. Therefore, we must have

$$
\left|x_{j}\right|=\left|\left(\widetilde{g}+g_{1}^{\prime}+g_{2}^{\prime}\right)_{j}\right| \leq\left|\left(g_{1}^{\prime}+g_{2}^{\prime}\right)_{j}\right| \leq\left|\left(g_{1}^{\prime}\right)_{j}\right|+\xi / 2 \leq 2\left|\left(g_{1}^{\prime}\right)_{j}\right|
$$

and therefore

$$
\sum_{j \in S_{2}} x_{j}^{2} \leq 4 \sum_{j \in S_{2}}\left(g_{1}^{\prime}\right)_{j}^{2} \leq 4\left\|g_{1}^{\prime}\right\|^{2}
$$

From above, we have $\|x\|_{0} \leq\left|S_{1}\right|+\left|S_{2}\right| \leq O\left(\xi m^{3 / 2}+\frac{\left(\delta_{2}\right)^{2}}{\xi^{2}}\right)$ and $\|x\|^{2} \leq O\left(\left(\delta_{2}\right)^{2}+\xi^{3} m^{3 / 2}\right)$. Choosing $\xi=\max \left\{2 \delta_{\infty}, \Theta\left(\frac{\left(\delta_{2}\right)^{2 / 3}}{m^{1 / 2}}\right)\right\}$ for the former, and choosing $\xi=2 \delta_{\infty}$ for the latter, we have the desired result.

Claim 5.9.4. For any $2 \leq a \leq b \leq L$ and any positive integer $s \leq O\left(\frac{m}{L \log m}\right)$, with probability at least $1-e^{-\Omega(s \log m)}$, for all $x \in \mathbb{R}^{m}$ with $\|x\| \leq 1$ and $\|x\|_{0} \leq s$, letting $y=\widetilde{W}_{b} \widetilde{D}_{b-1} \widetilde{W}_{b-1} \cdots \widetilde{D}_{a} \widetilde{W}_{a} x$, we can write $y=y_{1}+y_{2}$ with

$$
\left\|y_{1}\right\| \leq O(\sqrt{s / m} \log m) \quad \text { and } \quad\left\|y_{2}\right\|_{\infty} \leq \frac{2 \sqrt{\log m}}{\sqrt{m}}
$$

Proof of Claim 5.9.4. First of all, fix any $x$, we can let $u=\widetilde{D}_{b-1} \widetilde{W}_{b-1} \cdots \widetilde{D}_{a} \widetilde{W}_{a} x$ and the same proof of Lemma 5.8 .3 implies that with probability at least $1-e^{-\Omega(m / L)}$ we have $\|u\| \leq O(\|x\|)$. We next condition on this event happens.

Let $\beta=\sqrt{\log m} / \sqrt{m}$. If $u$ is fixed and using only the randomness of $W_{b}$, we have $y_{i} \sim \mathcal{N}\left(0, \frac{2\|u\|^{2}}{m}\right)$ so for every $p \geq 1$, by Gaussian tail bound

$$
\operatorname{Pr}\left[\left|y_{i}\right| \geq \beta p\right] \leq e^{-\Omega\left(\beta^{2} p^{2} m /\|x\|^{2}\right)} \leq e^{-\Omega\left(\beta^{2} p^{2} m\right)} .
$$

As long as $\beta^{2} p^{2} m \geq \beta^{2} m \geq \Omega(\log m)$, we know that if $\left|y_{i}\right| \geq \beta p$ occurs for $q / p^{2}$ indices $i$ out of $[m]$, this cannot happen with probability more than

$$
\binom{m}{q / p^{2}} \times\left(e^{-\Omega\left(\beta^{2} p^{2} m\right)}\right)^{q / p^{2}} \leq e^{\frac{q}{p^{2}}\left(O(\log m)-\Omega\left(\beta^{2} p^{2} m\right)\right)} \leq e^{-\Omega\left(\beta^{2} q m\right)}
$$

In other words,

$$
\operatorname{Pr}\left[\left|\left\{i \in[m]:\left|y_{i}\right| \geq \beta p\right\}\right|>q / p^{2}\right] \leq e^{-\Omega\left(\beta^{2} q m\right)}
$$

Finally, by applying union bound over $p=1,2,4,8,16, \ldots$ we have with probability $\geq$ $1-e^{-\Omega\left(\beta^{2} q m\right)} \cdot \log q$,

$$
\begin{equation*}
\sum_{i:\left|y_{i}\right| \geq \beta} y_{i}^{2} \leq \sum_{k=0}^{\lceil\log q\rceil}\left(2^{k+1} \beta\right)^{2}\left|\left\{i \in[m]:\left|y_{i}\right| \geq 2^{k} \beta\right\}\right| \leq \sum_{k=0}^{\lceil\log q\rceil}\left(2^{k+1} \beta\right)^{2} \cdot \frac{q}{2^{2 k}} \leq O\left(q \beta^{2} \log q\right) \tag{5.9}
\end{equation*}
$$

In other words, vector $y$ can be written as $y=y_{1}+y_{2}$ where $\left\|y_{2}\right\|_{\infty} \leq \beta$ and $\left\|y_{1}\right\|^{2} \leq$ $O\left(q \beta^{2} \log q\right)$.

Finally, we want to take $\epsilon$-net over all $s$-sparse inputs $x$. This requires $\beta^{2} q m \geq$ $\Omega(s \log m)$, so we can choose $q=\Theta\left(\frac{s \log m}{m \beta^{2}}\right)=\Theta(s)$.

### 5.9.2 Intermediate Layers

Lemma 5.9.5 (intermediate perturbation). For any integer $s$ with $1 \leq s \leq O\left(\frac{m}{L^{3} \log m}\right)$, with probability at least $1-e^{-\Omega(s \log m)}$ over the randomness of $\vec{W}^{(0)}, A$,

- for every $i \in[n], 1 \leq a \leq b \leq L$,
- for every diagonal matrices $D_{i, 0}^{\prime \prime}, \ldots, D_{i, L}^{\prime \prime} \in[-3,3]^{m \times m}$ with at most $s$ non-zero entries.
- for every perturbation matrices $W_{1}^{\prime}, \ldots, W_{L}^{\prime} \in \mathbb{R}^{m \times m}$ with $\left\|\vec{W}^{\prime}\right\|_{2} \leq \omega \in[0,1]$.
we have
(a) $\left\|\widetilde{W}_{b}\left(\widetilde{D}_{i, b-1}+D_{i, b-1}^{\prime \prime}\right) \cdots\left(\widetilde{D}_{i, a}+D_{i, a}^{\prime \prime}\right) \widetilde{W}_{a}\right\|_{2} \leq O(\sqrt{L})$.
(b) $\left\|\left(\widetilde{W}_{b}+W_{b}^{\prime}\right)\left(\widetilde{D}_{i, b-1}+D_{i, b-1}^{\prime \prime}\right) \cdots\left(\widetilde{D}_{i, a}+D_{i, a}^{\prime \prime}\right)\left(\widetilde{W}_{a}+W_{a}^{\prime}\right)\right\|_{2} \leq O(\sqrt{L})$ if $\omega \leq O\left(\frac{1}{L^{1.5}}\right)$.

Proof. For notational simplicity we ignore subscripts in $i$ in the proofs.
(a) Note that each $D_{\ell}^{\prime \prime}$ can be written as $D_{\ell}^{\prime \prime}=D_{\ell}^{0 / 1} D_{\ell}^{\prime \prime} D_{\ell}^{0 / 1}$, where each $D_{\ell}^{0 / 1}$ is a diagonal matrix satisfying

$$
\left(D_{\ell}^{0 / 1}\right)_{k, k}=\left\{\begin{array}{ll}
1, & \left(D_{\ell}^{\prime \prime}\right)_{k, k}=0 ; \\
0, & \left(D_{\ell}^{\prime \prime}\right)_{k, k}=0 .
\end{array} \quad \text { and } \quad\left\|D_{\ell}^{0 / 1}\right\|_{0} \leq s\right.
$$

In order to bound the spectral norm of $\widetilde{W}_{b}\left(\widetilde{D}_{b-1}+D_{b-1}^{\prime \prime}\right) \widetilde{W}_{b-1} \cdots\left(\widetilde{D}_{a}+D_{a}^{\prime \prime}\right) \widetilde{W}_{a}$, by triangle inequality, we can expend it into $2^{b-a}$ matrices and bound their spectral norms individually. Each such matrix can be written as (ignoring the subscripts)

$$
\begin{equation*}
\left(\widetilde{W} \widetilde{D} \cdots \widetilde{W} D^{0 / 1}\right) D^{\prime \prime}\left(D^{0 / 1} \widetilde{W} \widetilde{D} \cdots \widetilde{W} D^{0 / 1}\right) D^{\prime \prime} \cdots D^{\prime \prime}\left(D^{0 / 1} \widetilde{W} \widetilde{D} \cdots \widetilde{W}\right) \tag{5.10}
\end{equation*}
$$

Therefore, it suffices for us to bound the spectral norm of the following four types of matrices:

- $\widetilde{W} \widetilde{D} \cdots \widetilde{W} D^{0 / 1}$, such matrix has spectral norm at most 2 owing to Lemma 5.8.3b; - $D^{0 / 1} \widetilde{W} \widetilde{D} \cdots \widetilde{W}$, such matrix has spectral norm at most $O(1)$ owing to Lemma 5.8.3c;
- $D^{0 / 1} \widetilde{W} \widetilde{D} \cdots \widetilde{W} D^{0 / 1}$, such matrix has spectral norm at most $\frac{1}{100 L^{1.5}}$ owing to Lemma 5.8.3d and our choice $s \leq O\left(\frac{m}{L^{3} \log m}\right)$;
- $D^{\prime \prime}$, such matrix has spectral norm at most 3 .

Together, we have

$$
\begin{aligned}
& \left\|\widetilde{W}_{b}\left(\widetilde{D}_{b-1}+D_{b-1}^{\prime \prime}\right) \widetilde{W}_{b-1} \cdots\left(\widetilde{D}_{a}+D_{a}^{\prime \prime}\right) \widetilde{W}_{a}\right\| \\
\leq & O(\sqrt{L})+\sum_{j=1}^{b-a}\binom{b-a}{j} \cdot O(1) \cdot\left(\frac{1}{100 L^{1.5}}\right)^{j-1} \cdot 3^{j} \cdot O(1) \leq O(\sqrt{L}) .
\end{aligned}
$$

(b) In order to bound the spectral norm of $\left(\widetilde{W}_{b}+W_{b}^{\prime}\right)\left(\widetilde{D}_{b-1}+D_{b-1}^{\prime \prime}\right) \cdots\left(\widetilde{D}_{a}+D_{a}^{\prime \prime}\right)\left(\widetilde{W}_{a}+W_{a}^{\prime}\right)$, by triangle inequality, we can expend it into $2^{b-a+1}$ matrices in terms of $W^{\prime}$ and bound their spectral norms individually. Each such matrix can be written as (ignoring the subscripts, and denoting $\breve{D}=\widetilde{D}+D^{\prime}$ )

$$
(\widetilde{W} \breve{D} \cdots \widetilde{W} \breve{D}) W^{\prime}(\breve{D} \widetilde{W} \cdots \widetilde{W} \breve{D}) \cdots W^{\prime}(\breve{D} \widetilde{W} \cdots \breve{D} \widetilde{W})
$$

Moreover, from Lemma 5.9.5a, we know the following three types of matrices

- $\widetilde{W} \breve{D} \cdots \widetilde{W} \breve{D}$,
- $\breve{D} \widetilde{W} \cdots \widetilde{W} \breve{D}$, and
- $\breve{D} \widetilde{W} \cdots \breve{D} \widetilde{W}$
all have spectral norm at most $O(\sqrt{L})$. Together, using $\left\|W_{\ell}^{\prime}\right\|_{2} \leq O\left(\frac{1}{L^{1.5}}\right)$, we have

$$
\begin{aligned}
& \left\|\left(\widetilde{W}_{b}+W_{b}^{\prime}\right)\left(\widetilde{D}_{b-1}+D_{b-1}^{\prime \prime}\right)\left(\widetilde{W}_{b-1}+W_{b-1}^{\prime}\right) \cdots\left(\widetilde{D}_{a}+D_{a}^{\prime \prime}\right)\left(\widetilde{W}_{a}+W_{a}^{\prime}\right)\right\| \\
\leq & \sum_{j=0}^{b-a+1}\binom{b-a+1}{j} \cdot(O(\sqrt{L}))^{j+1} \cdot\left(O\left(\frac{1}{L^{1.5}}\right)\right)^{j} \leq O(\sqrt{L}) .
\end{aligned}
$$

### 5.9.3 Backward

Lemma 5.9 .6 (backward perturbation). For any integer $s \in\left[\Omega\left(\frac{d}{\log m}\right), O\left(\frac{m}{L^{3} \log m}\right)\right]$, for $d \leq O\left(\frac{m}{L \log m}\right)$, with probability at least $1-e^{-\Omega(s \log m)}$ over the randomness of $\vec{W}^{(0)}, A, B$,

- for all $i \in[n], a=1,2, \ldots, L+1$,
- for every diagonal matrices $D_{i, 0}^{\prime \prime}, \ldots, D_{i, L}^{\prime \prime} \in[-3,3]^{m \times m}$ with at most $s$ non-zero entries,
- for every perturbation matrices $W_{i, 1}^{\prime}, \ldots, W_{i, L}^{\prime} \in \mathbb{R}^{m \times m}$ with $\left\|\vec{W}^{\prime}\right\|_{2} \leq \omega=O\left(\frac{1}{L^{1.5}}\right)$,
it satisfies $\left\|B\left(\widetilde{D}_{i, L}+D_{i, L}^{\prime \prime}\right)\left(\widetilde{W}_{L}+W_{L}^{\prime}\right) \cdots\left(\widetilde{W}_{a+1}+W_{a+1}^{\prime}\right)\left(\widetilde{D}_{i, a}+D_{i, a}^{\prime \prime}\right)-B \widetilde{D}_{i, L} \widetilde{W}_{L} \cdots \widetilde{W}_{a+1} \widetilde{D}_{i, a}\right\|_{2} \leq$ $O\left(\frac{\sqrt{L^{3} s \log m+\omega^{2} L^{3} m}}{\sqrt{d}}\right)$. Note that if $s=O\left(m \omega^{2 / 3} L\right)$, this upper bound becomes $O\left(\frac{\omega^{1 / 3} L^{2} \sqrt{m \log m}}{\sqrt{d}}\right)$.

Proof. For notational simplicity we ignore subscripts in $i$ in the proofs.
Ignoring the subscripts for cleanness, we have

$$
\begin{aligned}
& \left\|B\left(\widetilde{D}_{i, L}+D_{i, L}^{\prime \prime}\right)\left(\widetilde{W}_{L}+W_{L}^{\prime}\right) \cdots\left(\widetilde{W}_{a+1}+W_{a+1}^{\prime}\right)\left(\widetilde{D}_{i, a}+D_{i, a}^{\prime \prime}\right)-B \widetilde{D}_{i, L} \widetilde{W}_{L} \cdots \widetilde{W}_{a+1} \widetilde{D}_{i, a}\right\|_{2} \\
\leq & \sum_{\ell=a}^{L} \underbrace{\left\|B \widetilde{D}_{i, L} \widetilde{W}_{L} \cdots \widetilde{W}_{\ell+1} D_{\ell}^{0 / 1}\right\|_{2}}_{\text {Lemma 5.8.4a }}\left\|D_{\ell}^{\prime \prime}\right\|_{2} \underbrace{\left\|D_{\ell}^{0 / 1}\left(\widetilde{W}_{\ell}+W_{\ell}^{\prime}\right) \cdots\left(\widetilde{D}_{i, a}+D_{i, a}^{\prime \prime}\right)\right\|_{2}}_{\text {Lemma 5.9.5b }} \\
& +\sum_{\ell=a+1}^{L} \underbrace{\left\|B \widetilde{D}_{i, L} \widetilde{W}_{L} \cdots \widetilde{W}_{\ell+1} \widetilde{D}_{\ell}\right\|_{2}}_{\text {Lemma 5.8.4b }}\left\|W_{\ell}^{\prime}\right\|_{2} \underbrace{\left\|\left(\widetilde{D}_{\ell-1}+D_{\ell-1}^{\prime \prime}\right)\left(\widetilde{W}_{\ell-1}+W_{\ell-1}^{\prime}\right) \cdots\left(\widetilde{D}_{i, a}+D_{i, a}^{\prime \prime}\right)\right\|_{2}}_{\text {Lemma 5.9.5b }} \\
\leq & L \cdot O\left(\frac{\sqrt{s \log m}}{\sqrt{d}}\right) \cdot O(\sqrt{L})+L \cdot O(\sqrt{m / d}) \cdot \omega \cdot O(\sqrt{L})
\end{aligned}
$$

### 5.10 Gradient Bound at Random Initialization

Throughout this section we assume $\vec{W}, A$ and $B$ are randomly generated according to Definition 5.2.2. The diagonal sign matrices $D_{i, \ell}$ are also determined according to this random initialization.

Recall we have defined Back $_{i, \ell} \stackrel{\text { def }}{=} B D_{i, L} W_{L} \cdots D_{i, \ell} W_{\ell} \in \mathbb{R}^{d \times m}$. In this section, we introduce the following notion

Definition 5.10.1. For any vector tuple $\vec{v}=\left(v_{1}, \ldots, v_{n}\right) \in\left(\mathbb{R}^{d}\right)^{n}$ (viewed as a fake loss vector), for each $\ell \in[L]$, we define

$$
\begin{aligned}
& \widehat{\nabla}_{\left[W_{\ell}\right]_{k}}^{\vec{v}} F(\vec{W}) \stackrel{\text { def }}{=} \sum_{i=1}^{n}\left(\operatorname{Back}_{i, \ell+1}^{\top} \mathrm{v}_{i}\right)_{k} \cdot h_{i, \ell-1} \cdot \mathbb{1}_{\left\langle\left[W_{\ell}\right]_{k}, h_{i, \ell-1}\right\rangle \geq 0}, \forall k \in[m] \\
& \widehat{\nabla}_{W_{\ell}}^{\vec{v}} F(\vec{W}) \stackrel{\rightharpoonup}{=} \stackrel{\text { def }}{=} \sum_{i=1}^{n} \widehat{\nabla}_{W_{\ell}}^{\vec{v}} F_{i}(\vec{W}) \quad \text { where } \quad \widehat{\nabla}_{W_{\ell}}^{\vec{v}} F_{i}(\vec{W}) \stackrel{\text { def }}{=} D_{i, \ell}\left(\operatorname{Back}_{i, \ell+1}^{\top} \mathrm{v}_{i}\right) h_{i, \ell-1}^{\top}
\end{aligned}
$$

Remark 5.10.1. It is an easy exercise to check that, if letting $\vec{v}=\left(\mathrm{v}_{1}, \ldots, \mathrm{v}_{n}\right)$ where $\mathrm{v}_{i}=$ $B h_{i, L}-y_{i}^{*}$, then $\widehat{\nabla}_{\left[W_{\ell}\right]_{k}}^{\vec{v}} F(\vec{W})=\nabla_{\left[W_{\ell}\right]_{k}} F(\vec{W})$ and $\widehat{\nabla}_{W_{\ell}}^{\vec{v}} F_{i}(\vec{W})=\nabla_{W_{\ell}} F_{i}(\vec{W})$.

Our main lemma of this section is the following.
Lemma 5.10.1 (gradient bound at random initialization). Fix any $\vec{v} \in\left(\mathbb{R}^{d}\right)^{n}$, with probability at least $1-e^{-\Omega(\delta m / n)}$ over the randomness of $A, \vec{W}, B$, it satisfies for every $\ell \in[L]$ :

$$
\left.\left\|\widehat{\nabla}_{W_{\ell}}^{\vec{v}} F_{i}(\vec{W})\right\|_{F}^{2} \leq O\left(\frac{\left\|\mathrm{v}_{i}\right\|^{2}}{d} \times m\right) \quad \right\rvert\, \widehat{\nabla}_{W_{\ell}}^{\vec{v}_{\ell}} F(\vec{W}) \|_{F}^{2} \leq O\left(\frac{\|\overrightarrow{\mathrm{v}}\|^{2}}{d} \times m n\right)
$$

### 5.10.1 Proof of Lemma 5.10.1: Upper Bound

For each $i \in[n], \ell \in[L]$, we can calculate that

$$
\begin{aligned}
\left\|\widehat{\nabla}_{W_{\ell}}^{\vec{v}} F_{i}(\vec{W})\right\|_{F} & =\left\|D_{i, \ell}\left(\operatorname{Back}_{i, \ell+1}^{\top} \cdot \mathrm{v}_{i}\right) \cdot h_{i, \ell-1}^{\top}\right\|_{F} \\
& =\left\|D_{i, \ell}\left(\operatorname{Back}_{i, \ell+1}^{\top} \cdot \mathrm{v}_{i}\right)\right\|_{2} \cdot\left\|h_{i, \ell-1}\right\|_{2} \\
& \leq\left\|\operatorname{Back}_{i, \ell+1}\right\|_{2} \cdot\left\|\mathrm{v}_{i}\right\|_{2} \cdot\left\|h_{i, \ell-1}\right\|_{2} \\
& \leq\left\|B W_{L} D_{L-1} \cdots D_{i, \ell+1} W_{\ell+1}\right\|_{2} \cdot\left\|\mathrm{v}_{i}\right\|_{2} \cdot\left\|h_{i, \ell-1}\right\|_{2} \\
& \leq O(\sqrt{m / d}) \cdot O(1) \cdot\left\|\mathrm{v}_{i}\right\|_{2} .
\end{aligned}
$$

where inequality (1) uses Lemma 5.8.4b and Lemma 5.8 .1 with high probability. Applying triangle inequality with respect to all $\ell \in[L]$, taking square on both sides, and summing up over all $i \in[n]$ finish the proof.

### 5.10.2 Proof of Lemma 5.10.1: Lower Bound

Let $i^{*}=\operatorname{argmax}_{i \in[n]}\left\{\left\|\mathrm{v}_{i}\right\|\right\}$. Recall

$$
\widehat{\nabla}_{\left[W_{L}\right]_{k}} F(\vec{W})=\sum_{i=1}^{n}\left\langle B_{k}, \mathrm{v}_{i}\right\rangle \cdot h_{i, L-1} \cdot \mathbb{1}_{\left(W_{L} h_{i, L-1}\right)_{k} \geq 0}
$$

Let $\widehat{h} \stackrel{\text { def }}{=} \frac{h_{i^{*}, L-1}}{\left\|h_{i^{*}, L-1}\right\|}$. For analysis purpose, after $\widehat{h}$ is fixed (so after fixing the randomness of $A, W_{1}, \ldots, W_{L-1}$, we redefine $W_{L} \widehat{h}=\sqrt{1-\theta^{2}} \widehat{g}_{1}+\theta \widehat{g}_{2}$ where $\widehat{g}_{1}$ and $\widehat{g}_{2}$ are generated independently from $\mathcal{N}\left(0, \frac{2 I}{m}\right)$. We can do so because the two sides are equal in distribution. In other words, we can set

$$
W_{L}^{\prime} \stackrel{\text { def }}{=} W_{L}\left(I-\widehat{h} \widehat{h}^{\top}\right)-\sqrt{1-\theta^{2}} \widehat{g}_{1} \widehat{h}^{\top} \text { and } W_{L}^{\prime \prime} \stackrel{\text { def }}{=} \theta \widehat{g}_{2} \widehat{h}^{\top},
$$

then we have $W_{L}=W_{L}^{\prime}+W_{L}^{\prime \prime}$. In particular, the randomness of $W_{L}^{\prime}$ and $W_{L}^{\prime \prime}$ are independent. In the remainder of the proof, let us choose $\theta \stackrel{\text { def }}{=} \frac{\delta}{5 n} \leq \frac{1}{5}$.

We first make two technical claims, and the proof of the first one can be found in Section 5.10.2.1.

Claim 5.10.2. We have $\operatorname{Pr}_{W_{L}^{\prime}, W_{L-1}, \ldots, W_{1}, A}\left[\left|N_{2}\right| \geq \frac{\delta}{40 n} m\right] \geq 1-e^{\Omega(\delta m / n)}$
$N_{2} \xlongequal{\text { def }}\left\{k \in[m]:\left(\left|\left(W_{L}^{\prime} h_{i^{*}, L-1}\right)_{k}\right| \leq \frac{\delta}{10 n \sqrt{m}}\right) \bigwedge\left(\forall i \in[n] \backslash\left\{i^{*}\right\}, \quad\left|\left(W_{L}^{\prime} h_{i, L-1}\right)_{k}\right| \geq \frac{\delta}{4 n \sqrt{m}}\right)\right\}$
Claim 5.10.3. Given set $N_{2} \subset[m]$ and $\overrightarrow{\mathrm{v}}$, we have

$$
\operatorname{Pr}_{B_{k}}\left[\left|\left\{k \in N_{2}:\left|\left\langle B_{k}, \mathrm{v}_{i^{*}}\right\rangle\right| \geq \frac{\left\|\mathrm{v}_{i^{*}}\right\|}{\sqrt{d}}\right\}\right| \geq \frac{\left|N_{2}\right|}{2}\right] \geq 1-e^{-\Omega\left(\left|N_{2}\right|\right)}
$$

Proof of Claim 5.10.3. Observe that each $\left\langle B_{k}, \mathrm{v}_{i^{*}}\right\rangle$ follows from $\mathcal{N}\left(0,\left\|\mathrm{v}_{i^{*}}\right\|^{2} / d\right)$, so with probability at least 0.68 it satisfies $\left|\left\langle B_{k}, \mathrm{v}_{i^{*}}\right\rangle\right| \geq \frac{\left\|\mathrm{v}_{i^{*}}\right\|}{\sqrt{d}}$. Using Chernoff bound we have the desired claim.

Combining Claim 5.10.4 and Claim 5.10.3, we can obtain a set $N \subseteq[m]$ satisfying

$$
\begin{gathered}
N \stackrel{\text { def }}{=}\left\{k \in[m]:\left(\left|\left(W_{L}^{\prime} h_{i^{*}, L-1}\right)_{k}\right| \leq \frac{\delta}{10 n \sqrt{m}}\right) \bigwedge\left(\forall i \in[n] \backslash\left\{i^{*}\right\}, \quad\left|\left(W_{L}^{\prime} h_{i, L-1}\right)_{k}\right| \geq \frac{\delta}{4 n \sqrt{m}}\right)\right. \\
\left.\bigwedge\left|\left\langle B_{k}, \mathrm{v}_{i^{*}}\right\rangle\right| \geq \frac{\left\|\mathrm{v}_{i^{*}}\right\|}{\sqrt{d}}\right\}
\end{gathered}
$$

of cardinality $|N| \geq \frac{\delta}{100 n} m$. Let us fix the randomness of $W_{L}^{\prime}$ so that $N$ is fixed. Let $k$ be any index in $N$. We can write

$$
\widehat{\nabla}_{\left[W_{L}\right]_{k}}^{\vec{v}} F(\vec{W})=\sum_{i=1}^{n}\left\langle B_{k}, v_{i}\right\rangle \cdot h_{i, L-1} \cdot \mathbb{1}_{\left(W_{L}^{\prime} h_{i, L-1}\right)_{k}+\left(W_{L}^{\prime \prime} h_{i, L-1}\right)_{k} \geq 0} .
$$

The only remaining source of randomness comes from $W_{L}^{\prime \prime}=\theta \widehat{g}_{2} \widehat{h}^{\top}$.
Recalling that $\theta=\frac{1}{5 n}$ and $\widehat{g}_{2} \sim \mathcal{N}\left(0, \frac{2}{m} I\right)$, so since $\theta\left(\widehat{g}_{2}\right)_{k} \sim \mathcal{N}\left(0, \frac{2 \theta^{2}}{m}\right)$, using numerical values of Gaussian CDF, one can verify that

$$
\operatorname{Pr}_{\widehat{g}_{2}}\left[\left|\theta\left(\widehat{g}_{2}\right)_{k}\right| \in\left(\frac{\delta}{9 n \sqrt{m}}, \frac{\delta}{5 n \sqrt{m}}\right)\right] \geq 0.2 .
$$

Let us denote this event of $\widehat{g}_{2}$ as $\mathfrak{E}_{k}$. Conditioning on $\mathfrak{E}_{k}$ happens, recalling $\left\|h_{i, L-1}\right\| \in$ [0.9, 1.1] from Lemma 5.8.1,

- For every $i \in[n] \backslash\left\{i^{*}\right\}$, we have

$$
\left|\left(W_{L}^{\prime \prime} h_{i, L-1}\right)_{k}\right|=\left|\left(\theta \widehat{g}_{2} \widehat{h}^{\top} h_{i, L-1}\right)_{k}\right| \leq\left|\left(\theta \widehat{g}_{2}\right)_{k}\right| \cdot\left\|h_{i, L-1}\right\|<\frac{\delta}{5 n \sqrt{m}} \cdot 1.1<\left|\left(W_{L}^{\prime} h_{i, L-1}\right)_{k}\right|
$$

and this means $\mathbb{1}_{\left(W_{L} h_{i, L-1}\right)_{k} \geq 0}=\mathbb{1}_{\left(W_{L}^{\prime} h_{i, L-1}\right)_{k} \geq 0}$.

- For $i=i^{*}$, we have
$\left|\left(W_{L}^{\prime \prime} h_{i^{*}, L-1}\right)_{k}\right|=\left|\left(\theta \widehat{g}_{2} \widehat{h}^{\top} h_{i^{*}, L-1}\right)_{k}\right|=\left|\left(\theta \widehat{g}_{2}\right)_{k}\right| \cdot\left\|h_{i^{*}, L-1}\right\|>\frac{\delta}{9 n \sqrt{m}} \cdot 0.9>\left|\left(W_{L}^{\prime} h_{i^{*}, L-1}\right)_{k}\right|$ and this means $\mathbb{1}_{\left(W_{L} h_{i^{*}, L-1}\right)_{k} \geq 0} \neq \mathbb{1}_{\left(W_{L}^{\prime} h_{i^{*}, L-1}\right)_{k} \geq 0}$ with probability exactly $\frac{1}{2}$ - this is because, conditioning on event $\mathfrak{E}_{k}$, the sign of $\left(\theta \widehat{g}_{2}\right)_{k}$ is $\pm 1$ each with half probability.

Recall that for every $k \in N$,

$$
\widehat{\nabla}_{\left[W_{L}\right]_{k}} F(\vec{W})=\underbrace{\left\langle B_{k}, \mathrm{v}_{i^{*}}\right\rangle \cdot h_{i^{*}, L-1} \cdot \mathbb{1}_{\left(W_{L} h_{i^{*}, L-1}\right)_{k} \geq 0}}_{\boldsymbol{*}}+\sum_{i \in[n] \backslash\left\{i^{*}\right\}} \underbrace{\left\langle B_{k}, \mathrm{v}_{i}\right\rangle \cdot h_{i, L-1} \cdot \mathbb{1}_{\left(W_{L} h_{i, L-1}\right)_{k} \geq 0}}_{\boldsymbol{*}}
$$

Now, fix the randomness of $A, B, W_{1}, \ldots, W_{L-1}, W_{L}^{\prime}$ and let $\widehat{g}_{2}$ be the only randomness. Conditioning on $\mathfrak{E}_{k}$, we have that each term in $\boldsymbol{\&}$ is fixed (i.e., independent of $\widehat{g}_{2}$ ) because $\mathbb{1}_{\left(W_{L} h_{i, L-1}\right)_{k} \geq 0}=\mathbb{1}_{\left(W_{L}^{\prime} h_{i, L-1}\right)_{k} \geq 0}$. In contrast, conditioning on $\mathfrak{E}_{k}$, the indicator $\mathbb{1}_{\left(W_{L} h_{i^{*}, L-1}\right)_{k} \geq 0}$ of the term may be 1 or 0 each with half probability. This means,

$$
\underset{\left(\widehat{g}_{2}\right)_{k}}{\operatorname{Pr}}\left[\left\|\widehat{\nabla}_{\left[W_{L}\right]_{k}}^{\overrightarrow{\mathrm{v}}} F(\vec{W})\right\|^{2} \geq\left|\left\langle B_{k}, \mathrm{v}_{i^{*}}\right\rangle\right|^{2} \cdot\left\|h_{i^{*}, L-1}\right\|^{2} \mid k \in N \wedge \mathfrak{E}_{k}\right] \geq \frac{1}{2}
$$

Taking into account the fact that $\left|\left\langle B_{k}, \mathrm{v}_{i^{*}}\right\rangle\right| \geq \frac{\| \mathrm{v}^{*} *}{\sqrt{d}}$ (by definition of $N$ ), the fact that $\left\|h_{i, L-1}\right\| \geq 0.9$, and the fact that $\operatorname{Pr}_{\left(\widehat{g}_{2}\right)_{k}}[\mathfrak{E}] \geq 0.2$, we have

$$
\underset{\left(\widehat{g}_{2}\right)_{k}}{\operatorname{Pr}}\left[\left.\left\|\widehat{\nabla}_{\left[W_{L}\right]_{k}} F(\vec{W})\right\|^{2} \geq 0.8 \frac{\left\|\mathrm{v}_{i^{*}}\right\|^{2}}{d} \right\rvert\, k \in N\right] \geq \frac{1}{10} .
$$

Using the independence of $\left(\widehat{g}_{2}\right)_{k}$ with respect to different $k \in N$, we can apply Chernoff bound and derive:

$$
\underset{\widehat{g}_{2}}{\operatorname{Pr}}\left[\left.\sum_{k \in N}\left\|\widehat{\nabla}_{\left[W_{L}\right]_{k}} F(\vec{W})\right\|^{2} \geq 0.8 \frac{\left\|\mathrm{v}_{i^{*}}\right\|^{2}}{d} \cdot \frac{|N|}{15} \right\rvert\, N\right] \geq 1-e^{-\Omega(|N|)} .
$$

Finally, using and $|N| \geq \frac{\delta}{100 n} m$, we have

$$
\operatorname{Pr}\left[\left\|\widehat{\nabla}_{W_{L}}^{\vec{v}} F(\vec{W})\right\|_{F}^{2} \geq \frac{\left\|\mathrm{v}_{i^{*}}\right\|^{2}}{d} \frac{\delta}{2000 n} m\right] \geq 1-e^{-\Omega(\delta m / n)}
$$

We finish the upper bound proof of Lemma 5.10.1.

### 5.10.2.1 Proof of Claim 5.10.4

Claim 5.10.4. We have $\operatorname{Pr}_{W_{L}^{\prime}, W_{L-1}, \ldots, W_{1}, A}\left[\left|N_{2}\right| \geq \frac{\delta}{40 n} m\right] \geq 1-e^{\Omega(\delta m / n)}$
$N_{2} \stackrel{\text { def }}{=}\left\{k \in[m]:\left(\left|\left(W_{L}^{\prime} h_{i^{*}, L-1}\right)_{k}\right| \leq \frac{\delta}{10 n \sqrt{m}}\right) \bigwedge\left(\forall i \in[n] \backslash\left\{i^{*}\right\}, \quad\left|\left(W_{L}^{\prime} h_{i, L-1}\right)_{k}\right| \geq \frac{\delta}{4 n \sqrt{m}}\right)\right\}$
Proof of Claim 5.10.4. Throughout the proof we assume $W_{L-1}, \ldots, A$ are good enough so that Lemma 5.8.1 holds (for $\epsilon=0.01$ ) and we fix their randomness. Define

$$
N_{1} \stackrel{\text { def }}{=}\left\{k \in[m]:\left|\left(W_{L}^{\prime} h_{i^{*}, L-1}\right)_{k}\right| \leq \frac{\delta}{10 n \sqrt{m}}\right\}
$$

Since $\left\|h_{i^{*}, L-1}\right\|^{2} \leq 1.1$ by Lemma 5.8.1, and since by definition of $W_{L}^{\prime}$ we have $\left(W_{L}^{\prime} h_{i^{*}, L-1}\right)_{k} \sim$ $\mathcal{N}\left(0, \frac{2\left(1-\theta^{2}\right)\left\|h_{i^{*}, L-1}\right\|^{2}}{m}\right.$ ). By standard properties of Gaussian CDF (see Fact 5.10.5), we know $\left|\left(W_{L}^{\prime} h_{i^{*}, L-1}\right)_{k}\right| \leq \frac{\delta}{10 n \sqrt{m}}$ with probability at least $\frac{\delta}{25 n}$ for each $k \in[m]$. By Chernoff bound,

$$
\operatorname{Pr}_{W_{L}^{\prime}}\left[\left|N_{1}\right| \geq \frac{\delta}{30 n} m\right] \geq 1-e^{-\Omega(\delta m / n)}
$$

Next, suppose we fix the randomness of $W_{L}^{\prime} \widehat{h}$. Define

$$
N_{2} \stackrel{\text { def }}{=}\left\{k \in N_{1}: \forall i \in[n] \backslash\left\{i^{*}\right\}, \quad\left|\left(W_{L}^{\prime} h_{i, L-1}\right)_{k}\right| \geq \frac{\delta}{4 n \sqrt{m}}\right\}
$$

For each $k \in N_{1}$ and $i \in[n] \backslash\left\{i^{*}\right\}$, we can write

$$
W_{L}^{\prime} h_{i, L-1}=W_{L}^{\prime} \widehat{h}\left(\widehat{h}^{\top} h_{i, L-1}\right)+W_{L}^{\prime}\left(I-\widehat{h} \widehat{h}^{\top}\right) h_{i, L-1} .
$$

Above, the first term on the right hand side is fixed (because we have fixed the randomness of $\left.W_{L}^{\prime} \widehat{h}\right)$; however, $W_{L}^{\prime}\left(I-\widehat{h} \widehat{h}^{\top}\right) h_{i, L-1}$ is still fresh new random Gaussian. In symbols,

$$
W_{L}^{\prime} h_{i, L-1} \sim \mathcal{N}\left(W_{L}^{\prime} \widehat{h} \widehat{h}^{\top} h_{i, L-1}, \frac{2\left\|\left(I-\widehat{h} \widehat{h}^{\top}\right) h_{i, L-1}\right\|^{2}}{m} I\right) .
$$

According to Lemma 5.8.5, the variance here is at least $\frac{2}{m}\left\|\left(I-\widehat{h} \widehat{h}^{\top}\right) h_{i, L-1}\right\|^{2} \geq \frac{\delta^{2}}{2 m}$. Using standard properties of Gaussian CDF (see Fact 5.10.5), we know $\left|\left(W_{L}^{\prime} h_{i, L-1}\right)_{k}\right| \geq \frac{\delta}{4 n \sqrt{m}}$ with probability at least $1-\frac{1}{8 n}$ for each $k \in[m]$. By union bound, for this $k \in[m]$, with probability at least $\frac{7}{8}$ we know $\left|\left(W_{L}^{\prime} h_{i, L-1}\right)_{k}\right| \geq \frac{\delta}{4 n \sqrt{m}}$ for all $i \in[n] \backslash\left\{i^{*}\right\}$. By Chernoff bound (over all $k \in N_{1}$ ), we conclude that

$$
\underset{W_{L}^{\prime}}{\operatorname{Pr}}\left[\left.\left|N_{2}\right| \geq \frac{3}{4}\left|N_{1}\right| \right\rvert\, N_{1}\right] \geq 1-e^{-\Omega\left(\left|N_{1}\right|\right)}=1-e^{-\Omega(\delta m / n)} .
$$

Combining the two bounds we finish the proof.
Fact 5.10.5. Suppose $x \sim \mathcal{N}\left(0, \sigma^{2}\right)$ is a Gaussian random variable. For any $t \in(0, \sigma)$ we have

$$
\operatorname{Pr}[x \geq t] \in\left[\frac{1}{2}\left(1-\frac{4}{5} \frac{t}{\sigma}\right), \frac{1}{2}\left(1-\frac{2}{3} \frac{t}{\sigma}\right)\right]
$$

Similarly, if $x \sim \mathcal{N}\left(\mu, \sigma^{2}\right)$, for any $t \in(0, \sigma)$, we have

$$
\operatorname{Pr}[|x| \geq t] \in\left[1-\frac{4}{5} \frac{t}{\sigma}, 1-\frac{2}{3} \frac{t}{\sigma}\right] .
$$

### 5.11 Theorem 5.11.1: Gradient Bound at After Perturbation

In this section we prove our main theorem on the gradient upper and lower bounds.
Theorem 5.11.1 (gradient bound, restated). Let $\omega \stackrel{\text { def }}{=} O\left(\frac{\delta^{3 / 2}}{n^{9 / 2} L^{6} \log ^{3} m}\right)$. With probability at least $1-e^{-\Omega\left(m \omega^{2 / 3} L\right)}$ over the randomness of $\vec{W}^{(0)}, A, B$, it satisfies for every $\ell \in[L]$, every $i \in[n]$, and every $\vec{W}$ with $\left\|\vec{W}-\vec{W}^{(0)}\right\|_{2} \leq \omega$,

$$
\left\|\nabla_{W_{\ell}} F_{i}(\vec{W})\right\|_{F}^{2} \leq O\left(\frac{F_{i}(\vec{W})}{d} \times m\right) \quad \nabla_{W_{\ell}} F(\vec{W}) \|_{F}^{2} \leq O\left(\frac{F(\vec{W})}{d} \times m n\right)
$$

Remark 5.11.1. Our Theorem 5.11.1 only gives gradient lower bound on $\left\|\nabla_{W_{L}} F(\vec{W})\right\|_{F}$. In principle, one can derive similar lower bounds on $\left\|\nabla_{W_{\ell}} F(\vec{W})\right\|_{F}$ for all $\ell=1,2, \ldots, L-1$. However, the proof will be significantly more involved. We choose not to derive those bounds at the expense of losing a polynomial factor in $L$ in the final running time. For readers interested in the techniques for obtaining those bounds, we refer to them to the "randomness decomposition" part of [AZLS18].

Proof of Theorem 5.11.1. Again we denote by $\widetilde{D}_{i, \ell}$ and $D_{i, \ell}$ respectively the sign matrix at the initialization $\vec{W}^{(0)}$ and at the current point $\vec{W}$; and by $\widetilde{h}_{i, \ell}$ and $h_{i, \ell}$ respectively the forward vector at $\vec{W}^{(0)}$ and at $\vec{W}$. Let us choose $s=O\left(m \omega^{2 / 3} L\right)$ which bounds the sparsity of $\left\|D_{i, \ell}-\widetilde{D}_{i, \ell}\right\|_{0}$ by Lemma 5.9.1b. Recall

$$
\begin{align*}
& \widehat{\nabla}_{W_{\ell}}^{\vec{v}} F\left(\vec{W}^{(0)}\right)-\widehat{\nabla}_{W_{\ell}}^{\vec{v}} F(\vec{W}) \\
= & \sum_{i=1}^{n}\left(\left(\mathrm{v}_{i}^{\top} B \widetilde{D}_{i, L} \widetilde{W}_{L} \cdots \widetilde{W}_{\ell+1} \widetilde{D}_{i, \ell}\right)^{\top}\left(\widetilde{h}_{i, \ell-1}\right)^{\top}-\left(\mathrm{v}_{i}^{\top} B D_{i, L} W_{L} \cdots W_{\ell+1} D_{i, \ell}\right)^{\top}\left(h_{i, \ell-1}\right)^{\top}\right) \tag{5.11}
\end{align*}
$$

By Lemma 5.9.6, we know that
$\left\|\mathbf{v}_{i}^{\top} B \widetilde{D}_{i, L} \widetilde{W}_{L} \cdots \widetilde{D}_{i, a} \widetilde{W}_{a} \widetilde{D}_{i, a-1}-\mathbf{v}_{i}^{\top} B D_{i, L} W_{L} \cdots D_{i, a} W_{a} D_{i, a-1}\right\| \leq O\left(\omega^{1 / 3} L^{2} \sqrt{m \log m} / \sqrt{d}\right) \cdot\left\|\mathbf{v}_{i}\right\|$

By Lemma 5.8.4b we know

$$
\left\|\mathbf{v}_{i}^{\top} B \widetilde{D}_{i, L} \widetilde{W}_{L} \cdots \widetilde{D}_{i, a} \widetilde{W}_{a} \widetilde{D}_{i, a-1}\right\| \leq O(\sqrt{m / d}) \cdot\left\|\mathbf{v}_{i}\right\|
$$

By Lemma 5.8.1 and Lemma 5.9.1c, we have

$$
\left\|h_{i, \ell-1}\right\| \leq 1.1 \quad \text { and } \quad\left\|h_{i, \ell-1}-\widetilde{h}_{i, \ell-1}\right\| \leq O\left(\omega L^{5 / 2} \sqrt{\log m}\right)
$$

Together, they imply

$$
\begin{align*}
\left\|\widehat{\nabla}_{W_{\ell}}^{\overrightarrow{v_{e}}} F\left(\vec{W}^{(0)}\right)-\widehat{\nabla}_{W_{\ell}}^{\vec{v}} F(\vec{W})\right\|_{F}^{2} & \leq n\|\overrightarrow{\mathrm{v}}\|^{2} \cdot O\left(\omega^{1 / 3} L^{2} \sqrt{m \log m} / \sqrt{d}+\sqrt{m / d} \times \omega L^{5 / 2} \sqrt{\log m}\right)^{2} \\
& \leq n\|\overrightarrow{\mathrm{v}}\|^{2} \cdot O\left(\frac{m \log m}{d} \cdot \omega^{2 / 3} L^{4}\right) \tag{5.12}
\end{align*}
$$

With our parameter assumption on $\omega$, this together with Lemma 5.10.1 implies the same upper and lower bounds at point $\vec{W}=\vec{W}^{(0)}+\vec{W}^{\prime}$ :

$$
\begin{aligned}
\left\|\widehat{\nabla}_{W_{\ell}}^{\vec{v}} F_{i}\left(\vec{W}^{(0)}+\vec{W}^{\prime}\right)\right\|_{F}^{2} \leq O\left(\frac{\left\|\mathrm{v}_{i}\right\|^{2}}{d} \times m\right) & \left\|\widehat{\nabla}_{W_{\ell}}^{\vec{v}} F\left(\vec{W}^{(0)}+\vec{W}^{\prime}\right)\right\|_{F}^{2} \leq O\left(\frac{\|\overrightarrow{\mathrm{v}}\|^{2}}{d} \times m n\right) \\
& \left\|\widehat{\nabla}_{W_{L}}^{\vec{v}_{L}} F\left(\vec{W}^{(0)}+\vec{W}^{\prime}\right)\right\|_{F}^{2} \geq \Omega\left(\frac{\max _{i \in[n]}\left\|\mathrm{v}_{i}\right\|^{2}}{d n / \delta} \times m\right)
\end{aligned}
$$

Finally, taking $\epsilon$-net over all possible vectors $\vec{v}=\left(\mathrm{v}_{1}, \ldots, \mathrm{v}_{n}\right) \in\left(\mathbb{R}^{d}\right)^{n}$, we know that the above bounds hold not only for fixed $\vec{v}$ but for all $\vec{v}$. In particular, we can now plug in the choice of $\mathrm{v}_{i}=$ loss $_{i}=B h_{i, L}-y_{i}^{*}$ and it implies our desired bounds on the true gradients.

### 5.12 Theorem 5.12.1: Objective Semi-Smoothness

The purpose of this section is to prove

Theorem 5.12.1 (objective semi-smoothness, restated). Let $\omega \in\left[\Omega\left(\frac{d^{3 / 2}}{m^{3 / 2} L^{3 / 2} \log ^{3 / 2} m}\right), O\left(\frac{1}{L^{4.5} \log ^{3} m}\right)\right]$ and $\vec{W}^{(0)}, A, B$ be at random initialization. With probability at least $1-e^{-\Omega\left(m \omega^{2 / 3} L\right)}$ over the randomness of $\vec{W}^{(0)}, A, B$, we have for every $\stackrel{\breve{W}}{W} \in\left(\mathbb{R}^{m \times m}\right)^{L}$ with $\left\|\stackrel{\breve{W}}{W}-\vec{W}^{(0)}\right\|_{2} \leq \omega$, and for every $\vec{W}^{\prime} \in\left(\mathbb{R}^{m \times m}\right)^{L}$ with $\left\|\vec{W}^{\prime}\right\|_{2} \leq \omega$, we have

We introduce the following notations before we go to proofs.

Definition 5.12.1. For $i \in[n]$ and $\ell \in[L]$ :

$$
\begin{array}{lll}
\widetilde{g}_{i, 0}=A x_{i} & \breve{g}_{i, 0}=A x_{i} & g_{i, 0}=A x_{i} \\
\widetilde{h}_{i, 0}=\phi\left(A x_{i}\right) & \hbar_{i, 0}=\phi\left(A x_{i}\right) & h_{i, 0}=\phi\left(A x_{i}\right) \\
\widetilde{g}_{i, \ell}=\widetilde{W}_{\ell} \widetilde{h}_{i, \ell-1} & \breve{g}_{i, \ell}=\breve{W}_{\ell} \hbar_{i, \ell-1} & g_{i, \ell}=\left(\breve{W}_{\ell}+W_{\ell}^{\prime}\right) h_{i, \ell-1} \\
\widetilde{h}_{i, \ell}=\phi\left(\widetilde{W}_{\ell} \widetilde{h}_{i, \ell-1}\right) & \hbar_{i, \ell}=\phi\left(\breve{W}_{\ell} \hbar_{i, \ell-1}\right) & h_{i, \ell}=\phi\left(\left(\breve{W}_{\ell}+W_{\ell}^{\prime}\right) h_{i, \ell-1}\right) \\
& \operatorname{loss}_{i}=B \hbar_{i, L}-y_{i}^{*} &
\end{array}
$$

Define diagonal matrices $\widetilde{D}_{i, \ell} \in \mathbb{R}^{m \times m}$ and $\breve{D}_{i, \ell} \in \mathbb{R}^{m \times m}$ respectively by letting

$$
\left(\widetilde{D}_{i, \ell}\right)_{k, k}=\mathbb{1}_{\left(\widetilde{g}_{i, \ell}\right)_{k} \geq 0} \text { and }\left(\breve{D}_{i, \ell}\right)_{k, k}=\mathbb{1}_{\left(\breve{g}_{i, \ell}\right)_{k} \geq 0}, \forall k \in[m] .
$$

The following claim gives rise to a new recursive formula to calculate $h_{i, \ell}-\hbar_{i, \ell}$.

Claim 5.12.2. There exist diagonal matrices $D_{i, \ell}^{\prime \prime} \in \mathbb{R}^{m \times m}$ with entries in $[-1,1]$ such that,

$$
\begin{equation*}
\forall i \in[n], \forall \ell \in[L]: \quad h_{i, \ell}-\hbar_{i, \ell}=\sum_{a=1}^{\ell}\left(\breve{D}_{i, \ell}+D_{i, \ell}^{\prime \prime}\right) \breve{W}_{\ell} \cdots \breve{W}_{a+1}\left(\breve{D}_{i, a}+D_{i, a}^{\prime \prime}\right) W_{a}^{\prime} h_{i, a-1} \tag{5.13}
\end{equation*}
$$

Furthermore, we have $\left\|h_{i, \ell}-\hbar_{i, \ell}\right\| \leq O\left(L^{1.5}\right)\left\|W^{\prime}\right\|_{2},\left\|B h_{i, \ell}-B \hbar_{i, \ell}\right\| \leq O(L \sqrt{m / d})\left\|W^{\prime}\right\|_{2}$ and $\left\|D_{i, \ell}^{\prime \prime}\right\|_{0} \leq O\left(m \omega^{2 / 3} L\right)$.

Proof of Theorem 5.12.1. First of all, since
$\frac{1}{2}\left\|B h_{i, L}-y_{i}^{*}\right\|^{2}=\frac{1}{2}\left\|\operatorname{loss}_{i}+B\left(h_{i, L}-\hbar_{i, L}\right)\right\|^{2}=\frac{1}{2}\left\|\operatorname{loss}_{i}\right\|^{2}+\operatorname{loss}_{i}^{\top} B\left(h_{i, L}-\hbar_{i, L}\right)+\frac{1}{2}\left\|B\left(h_{i, L}-\hbar_{i, L}\right)\right\|^{2}$
we can write

$$
\begin{align*}
& \quad F\left(\stackrel{\breve{W}}{\vec{W}}+\overrightarrow{W^{\prime}}\right)-F(\vec{W})-\left\langle\nabla F(\vec{W}), \vec{W}^{\prime}\right\rangle \\
& \stackrel{\oplus}{=}-\left\langle\nabla F(\stackrel{\breve{W}}{\vec{W}}), \overrightarrow{W^{\prime}}\right\rangle+\frac{1}{2} \sum_{i=1}^{n}\left\|B h_{i, L}-y_{i, L}^{*}\right\|^{2}-\left\|B \hbar_{i, L}-y_{i, L}^{*}\right\|^{2} \\
& \stackrel{(2)}{=}-\left\langle\nabla F(\stackrel{\breve{W}}{\vec{W}}), \overrightarrow{W^{\prime}}\right\rangle+\sum_{i=1}^{n} \breve{\operatorname{loss}}_{i}^{\top} B\left(h_{i, L}-\hbar_{i, L}\right)+\frac{1}{2}\left\|B\left(h_{i, L}-\hbar_{i, L}\right)\right\|^{2} \\
& \stackrel{(3)}{=} \sum_{i=1}^{n} \operatorname{loss}_{i}^{\top} B\left(\left(h_{i, L}-\hbar_{i, L}\right)-\sum_{\ell=1}^{L} \breve{D}_{i, L} \breve{W}_{L} \cdots \breve{W}_{\ell+1} \breve{D}_{i, \ell} W_{\ell}^{\prime} \hbar_{i, \ell-1}\right)+\frac{1}{2}\left\|B\left(h_{i, L}-\hbar_{i, L}\right)\right\|^{2} \\
& \stackrel{\oplus( }{=} \sum_{i=1}^{n} \breve{\operatorname{loss}}_{i}^{\top} B\left(\sum_{\ell=1}^{L}\left(\breve{D}_{i, L}+D_{i, L}^{\prime \prime}\right) \breve{W}_{L} \cdots \breve{W}_{\ell+1}\left(\breve{D}_{i, \ell}+D_{i, \ell}^{\prime \prime}\right) W_{\ell}^{\prime} h_{i, \ell-1}-\breve{D}_{i, L} \breve{W}_{L} \cdots \breve{W}_{\ell+1} \breve{D}_{i, \ell} W_{\ell}^{\prime} \hbar_{i, \ell-1}\right) \\
& \quad+\frac{1}{2} \sum_{i=1}^{n}\left\|B\left(h_{i, L}-\hbar_{i, L}\right)\right\|^{2} \tag{5.15}
\end{align*}
$$

Above, (1) is by the definition of $F(\cdot)$; (2) is by (5.14); (3) is by the definition of $\nabla F(\cdot)$ (see Fact 5.2.3 for an explicit form of the gradient).

We next bound the RHS of (5.15). We first note that by Lemma 5.9.1b, we have $\left\|\breve{D}_{i, \ell}+D_{i, \ell}^{\prime \prime}-\widetilde{D}_{i, \ell}\right\|_{0} \leq s$ and $\left\|\breve{D}_{i, \ell}-\widetilde{D}_{i, \ell}\right\|_{0} \leq s$ for $s=O\left(m \omega^{2 / 3} L\right)$.

We ignore subscripts in $i$ for notational convenience. We first use Claim 5.12.2 to get

$$
\begin{equation*}
\left\|B\left(h_{L}-\hbar_{L}\right)\right\| \leq O(L \sqrt{m / d}) \cdot\left\|\vec{W}^{\prime}\right\|_{2} \tag{5.16}
\end{equation*}
$$

Next we calculate that

$$
\begin{align*}
& \left|\operatorname{loss}_{i}^{\top} B\left(\breve{D}_{L}+D_{L}^{\prime \prime}\right) \breve{W}_{L} \cdots\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right) W_{\ell}^{\prime} h_{\ell-1}-\operatorname{loss}_{i}^{\top} B \breve{D}_{L} \breve{W}_{L} \cdots \breve{D}_{\ell} W_{\ell}^{\prime} h_{\ell-1}\right| \\
\leq & \left\|\breve{\operatorname{loss}}_{i}\right\| \cdot \underbrace{\left\|B\left(\breve{D}_{L}+D_{L}^{\prime \prime}\right) \breve{W}_{L} \cdots \breve{W}_{\ell-1}\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right)-B \breve{D}_{L} \breve{W}_{L} \cdots \breve{W}_{\ell-1} \breve{D}_{\ell}\right\|_{2}}_{\text {Lemma 5.9.6 with } s=O\left(m \omega^{2 / 3} L\right)} \cdot\left\|W_{\ell}^{\prime} h_{\ell-1}\right\| \\
\leq & \left\|\operatorname{loss}_{i}\right\| \cdot O\left(\frac{\sqrt{L^{3} \omega^{2 / 3} L m \log m}}{\sqrt{d}}\right) \cdot O\left(\left\|W_{\ell}^{\prime}\right\|_{2}\right) .
\end{align*}
$$

Finally, we also have

$$
\begin{align*}
& \quad\left|\breve{\operatorname{loss}}_{i}^{\top} B \breve{D}_{L} \breve{W}_{L} \cdots \breve{D}_{\ell} W_{\ell}^{\prime}\left(h_{\ell-1}-\hbar_{\ell-1}\right)\right| \\
& \stackrel{(1)}{\leq}\left\|\operatorname{loss}_{i}\right\| \cdot O\left(\sqrt{m / d}+\frac{\omega^{1 / 3} L^{2} \sqrt{m \log m}}{\sqrt{d}}\right) \cdot\left\|W_{\ell}^{\prime}\right\|_{2} \cdot\left\|h_{\ell}-\hbar_{\ell}\right\|_{2} \\
& \stackrel{(2)}{\leq} O\left(L^{0.5} \sqrt{m / d}\right) \cdot\left\|\operatorname{loss}_{i}\right\|_{2} \cdot L^{1.5}\left\|W_{\ell}^{\prime}\right\|^{2} \tag{5.18}
\end{align*}
$$

where (1) uses Lemma 5.8.4b (and Lemma 5.9.6 for bounding the perturbation) and (2) uses Claim 5.12.2 to bound $\left\|h_{\ell}-\hbar_{\ell}\right\|_{2}$ and our choice of $\omega$.

Putting (5.16), (5.17) and (5.18) back to (5.15), and using triangle inequality, we have the desired result.

### 5.12.1 Proof of Claim 5.12.2

We first present a simple proposition about the ReLU function.

Proposition 5.12.3. Given vectors $a, b \in \mathbb{R}^{m}$ and $D \in \mathbb{R}^{m \times m}$ the diagonal matrix where $D_{k, k}=\mathbb{1}_{a_{k} \geq 0}$. Then, then there exists a diagonal matrix $D^{\prime \prime} \in \mathbb{R}^{m \times m}$ with

- $\left|D_{k, k}+D_{k, k}^{\prime \prime}\right| \leq 1$ and $\left|D_{k, k}^{\prime \prime}\right| \leq 1$ for every $k \in[m]$,
- $D_{k, k}^{\prime \prime} \neq 0$ only when $\mathbb{1}_{a_{k} \geq 0} \neq \mathbb{1}_{b_{k} \geq 0}$, and
- $\phi(a)-\phi(b)=\left(D+D^{\prime \prime}\right)(a-b)$

Proof. We verify coordinate by coordinate for each $k \in[m]$.

- If $a_{k} \geq 0$ and $b_{k} \geq 0$, then $(\phi(a)-\phi(b))_{k}=a_{k}-b_{k}=(D(a-b))_{k}$.
- If $a_{k}<0$ and $b_{k}<0$, then $(\phi(a)-\phi(b))_{k}=0-0=(D(a-b))_{k}$.
- If $a_{k} \geq 0$ and $b_{k}<0$, then $(\phi(a)-\phi(b))_{k}=a_{k}=\left(a_{k}-b_{k}\right)+\frac{b_{k}}{a_{k}-b_{k}}\left(a_{k}-b_{k}\right)=$ $\left(D(a-b)+D^{\prime \prime}(a-b)\right)_{k}$, if we define $\left(D^{\prime \prime}\right)_{k, k}=\frac{b_{k}}{a_{k}-b_{k}} \in[-1,0]$.
- If $a_{k}<0$ and $b_{k} \geq 0$, then $(\phi(a)-\phi(b))_{k}=-b_{k}=0 \cdot\left(a_{k}-b_{k}\right)-\frac{b_{k}}{b_{k}-a_{k}}\left(a_{k}-b_{k}\right)=$ $\left(D(a-b)+D^{\prime \prime}(a-b)\right)_{k}$, if we define $\left(D^{\prime \prime}\right)_{k, k}=\frac{b_{k}}{b_{k}-a_{k}} \in[0,1]$.

Proof of Claim 5.12.2. We ignore the subscript in $i$ for cleanness, and calculate that

$$
\begin{aligned}
h_{\ell}-\hbar_{\ell} & \stackrel{\oplus}{=} \phi\left(\left(\breve{W}_{\ell}+W_{\ell}^{\prime}\right) h_{\ell-1}\right)-\phi\left(\breve{W}_{\ell} \hbar_{\ell-1}\right) \\
& \stackrel{\ominus}{=}\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right)\left(\left(\breve{W}_{\ell}+W_{\ell}^{\prime}\right) h_{\ell-1}-\breve{W}_{\ell} \hbar_{\ell-1}\right) \\
& =\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right) \breve{W}_{\ell}\left(h_{\ell-1}-\hbar_{\ell-1}\right)+\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right) W_{\ell}^{\prime} h_{\ell-1} \\
& \stackrel{\ominus}{=} \sum_{a=1}^{\ell}\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right) \breve{W}_{\ell} \cdots \breve{W}_{a+1}\left(\breve{D}_{a}+D_{a}^{\prime \prime}\right) W_{a}^{\prime} h_{a-1}
\end{aligned}
$$

Above, (1) is by the recursive definition of $h_{\ell}$ and $\hbar_{\ell}$; (2) is by Proposition 5.12 .3 and $D_{\ell}^{\prime \prime}$ is defined according to Proposition 5.12.3; and inequality (3) is by recursively computing $h_{\ell-1}-\hbar_{\ell-1}$. As for the remaining properties:

- We have $\left\|D_{\ell}^{\prime \prime}\right\|_{0} \leq O\left(m \omega^{2 / 3} L\right)$.

This is because, $\left(D_{\ell}^{\prime \prime}\right)_{k, k}$ is non-zero only at the coordinates $k \in[m]$ where the signs of $\breve{g}_{\ell}$ and $g_{\ell}$ are opposite (by Proposition 5.12.3). Such a coordinate $k$ must satisfy either $\left(\widetilde{D}_{\ell}\right)_{k, k} \neq\left(\breve{D}_{\ell}\right)_{k, k}$ or $\left(\widetilde{D}_{\ell}\right)_{k, k} \neq\left(D_{\ell}\right)_{k, k}$, and therefore by Lemma 5.9.1b there are at most $O\left(m \omega^{2 / 3} L\right)$ such coordinates $k$.

- We have $\left\|h_{\ell}-\hbar_{\ell}\right\| \leq O\left(L^{1.5}\right)\left\|\vec{W}^{\prime}\right\|_{2}$.

This is because we have $\left\|\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right) \breve{W}_{\ell} \cdots \breve{W}_{a+1}\left(\breve{D}_{a}+D_{a}^{\prime \prime}\right)\right\|_{2} \leq O(\sqrt{L})$ from Lemma 5.9.5b, we have $\left\|h_{a-1}\right\| \leq O(1)\left(\right.$ by $\left\|\widetilde{h}_{a-1}\right\| \leq O(1)$ from Lemma 5.8.1 and $\left\|\widetilde{h}_{a-1}-h_{a-1}\right\| \leq o(1)$ from Lemma 5.9.1c $)$; and and $\left\|W_{a}^{\prime} h_{a-1}\right\| \leq\left\|W_{a}^{\prime}\right\|_{2}\left\|h_{a-1}\right\| \leq O\left(\left\|\vec{W}^{\prime}\right\|_{2}\right)$.

- We have $\left\|B h_{\ell}-B \hbar_{\ell}\right\| \leq O(L \sqrt{m / d})\left\|\overrightarrow{W^{\prime}}\right\|_{2}$.

This is because we have $\left\|B\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right) \breve{W}_{\ell} \cdots \breve{W}_{a+1}\left(\breve{D}_{a}+D_{a}^{\prime \prime}\right)\right\|_{2} \leq O(\sqrt{m / d})$ from Lemma 5.8.4b (along with perturbation bound Lemma 5.9.6), we have $\left\|h_{a-1}\right\| \leq O(1)$
(by $\left\|\widetilde{h}_{a-1}\right\| \leq O(1)$ from Lemma 5.8.1 and $\left\|\widetilde{h}_{a-1}-h_{a-1}\right\| \leq o(1)$ from Lemma 5.9.1c); and and $\left\|W_{a}^{\prime} h_{a-1}\right\| \leq\left\|W_{a}^{\prime}\right\|_{2}\left\|h_{a-1}\right\| \leq O\left(\left\|\vec{W}^{\prime}\right\|_{2}\right)$.

### 5.13 Theorem 5.13.1: Convergence Rate of GD

Theorem 5.13.1 (gradient descent, restated). For any $\epsilon \in(0,1], \delta \in\left(0, O\left(\frac{1}{L}\right)\right]$. Let $m \geq$ $\widetilde{\Omega}\left(\operatorname{poly}\left(n, L, \delta^{-1}\right) d\right), \eta=\Theta\left(\frac{d \delta}{\operatorname{poly}(n, L) m}\right)$, and $\vec{W}^{(0)}, A, B$ are at random initialization. Then, with probability at least $1-e^{-\Omega\left(\log ^{2} m\right)}$, suppose we start at $\vec{W}^{(0)}$ and for each $t=0,1, \ldots, T-1$,

$$
\vec{W}^{(t+1)}=\vec{W}^{(t)}-\eta \nabla F\left(\vec{W}^{(t)}\right) .
$$

Then, it satisfies

$$
F\left(\vec{W}^{(T)}\right) \leq \epsilon \quad \text { for } \quad T=\Theta\left(\frac{\operatorname{poly}(n, L)}{\delta^{2}} \log \frac{1}{\epsilon}\right)
$$

In other words, the training loss drops to $\epsilon$ in a linear convergence speed.

Proof of Theorem 5.13.1. Using Lemma 5.8.1 we have $\left\|h_{i, L}\right\|_{2} \leq 1.1$ and then using the randomness of $B$, it is easy to show that $\left\|B \widetilde{h}_{i, L}-y_{i}^{*}\right\|^{2} \leq O\left(\log ^{2} m\right)$ with at least $1-e^{-\Omega\left(\log ^{2} m\right)}$ (where $\widetilde{h}_{i, L}$ is defined with respect to the random initialization $\vec{W}^{(0)}$ ), and therefore

$$
F\left(\vec{W}^{(0)}\right) \leq O\left(n \log ^{2} m\right)
$$

Let us assume for every $t=0,1, \ldots, T-1$, the following holds

$$
\begin{equation*}
\left\|\vec{W}^{(t)}-\vec{W}^{(0)}\right\|_{F} \leq \omega \stackrel{\text { def }}{=} O\left(\frac{n^{3} \sqrt{d}}{\delta \sqrt{m}} \log m\right) \tag{5.19}
\end{equation*}
$$

We shall prove the convergence of GD assuming (5.19) holds, so that previous statements such as Theorem 5.12.1 and Theorem 5.11.1 can be applied. At the end of the proof, we shall verify that (5.19) is satisfied.

To make the proof simple, we choose

$$
\begin{aligned}
m & \geq \Omega\left(\frac{n^{24} L^{12} d \log ^{5} m}{\delta^{8}}\right), \\
\eta & =\Theta\left(\frac{d \delta}{n^{4} L^{2} m}\right), \\
T & =\Theta\left(\frac{n^{6} L^{2}}{\delta^{2}} \log \frac{1}{\epsilon}\right)
\end{aligned}
$$

We emphasize that

- Most of the polynomial dependency in $n, L, \delta^{-1}$ come from the non-smoothness of the ReLU activation; if one instead studies smooth activations, their power can be significantly reduced. For instance, for smooth activation functions, one does not need the semi-smoothness Theorem 5.12.1.
- We have not tried to tighten the polynomial dependency on $n, L, \delta^{-1}$. We are aware of many ways to improve the constant in the exponents at the expense of complicating the proofs. Since the main focus of this paper is to derive the first polynomial running time, we do not include such improvements.

Letting $\nabla_{t}=\nabla F\left(\vec{W}^{(t)}\right)$, we calculate that

$$
\begin{align*}
& F\left(\vec{W}^{(t+1)}\right) \\
& \stackrel{(1)}{\leq} F\left(\vec{W}^{(t)}\right)-\eta\left\|\nabla F\left(\vec{W}^{(t)}\right)\right\|_{F}^{2}+\eta \sqrt{n F\left(\vec{W}^{(t)}\right)} \cdot O\left(\frac{\omega^{1 / 3} L^{2} \sqrt{m \log m}}{\sqrt{d}}\right) \cdot\left\|\nabla_{t}\right\|_{2}+O\left(\eta^{2} \frac{n L^{2} m}{d}\right)\left\|\nabla_{t}\right\|_{2}^{2} \\
& \stackrel{(2)}{\leq} F\left(\vec{W}^{(t)}\right)-\eta\left\|\nabla F\left(\vec{W}^{(t)}\right)\right\|_{F}^{2}+O\left(\frac{\eta n L^{2} m \omega^{1 / 3} \sqrt{\log m}}{d}+\frac{\eta^{2} n^{2} L^{2} m^{2}}{d^{2}}\right) \cdot F\left(\vec{W}^{(t)}\right) \\
& \stackrel{\text { ® }}{\leq}\left(1-\Omega\left(\frac{\eta \delta m}{d n^{2}}\right)\right) F\left(\vec{W}^{(t)}\right) . \tag{5.20}
\end{align*}
$$

Above, (1) uses Theorem 5.12.1; (2) uses Theorem 5.11.1 (which gives $\left\|\nabla_{t}\right\|_{2}^{2} \leq \max _{\ell \in[L]}\left\|\nabla_{W_{\ell}} F\left(\vec{W}^{(t)}\right)\right\|_{F}^{2} \leq$ $\left.O\left(\frac{F\left(\vec{W}^{(t)}\right)}{d} \times m n\right)\right)$; (3) use gradient lower bound from Theorem 5.11.1 and our choice of $\eta$. In other words, after $T=\Theta\left(\frac{d n^{2}}{\eta \delta m}\right) \log \frac{n \log m}{\epsilon}$ iterations we have $F\left(\vec{W}^{(T)}\right) \leq \epsilon$.

We need to verify for each $t,\left\|\vec{W}^{(t)}-\vec{W}^{(0)}\right\|_{F}$ is small so that (5.19) holds. By Theorem 5.11.1,

$$
\begin{aligned}
\left\|W_{\ell}^{(t)}-W_{\ell}^{(0)}\right\|_{F} & \leq \sum_{i=0}^{t-1}\left\|\eta \nabla_{W_{\ell}} F\left(\vec{W}^{(i)}\right)\right\|_{F} \leq O(\eta \sqrt{n m / d}) \cdot \sum_{i=0}^{t-1} \sqrt{F\left(\vec{W}^{(i)}\right)} \\
& \leq O(\eta \sqrt{n m / d}) \cdot \Theta\left(\frac{d n^{2}}{\eta \delta m}\right) \cdot O\left(\sqrt{n \log ^{2} m}\right) \leq O\left(\frac{n^{3} \sqrt{d}}{\delta \sqrt{m}} \log m\right)
\end{aligned}
$$

where the last step follows by our choice of $T$.

### 5.14 Theorem 5.14.1: Convergence Rate of SGD

Theorem 5.14.1 (stochastic gradient descent, stated). For any $\epsilon \in(0,1], \delta \in\left(0, O\left(\frac{1}{L}\right)\right]$, $b \in[n]$. Let $m \geq \widetilde{\Omega}\left(\frac{\operatorname{poly}\left(n, L, \delta^{-1}\right) \cdot d}{b}\right), \eta \stackrel{\text { def }}{=} \Theta\left(\frac{b \delta d}{\operatorname{poly}(n, L) m \log ^{2} m}\right)$, and $\vec{W}^{(0)}, A, B$ are at random initialization. Suppose we start at $W^{(0)}$ and for each $t=0,1, \ldots, T-1$,

$$
W^{(t+1)}=W^{(t)}-\eta \cdot \frac{n}{\left|S_{t}\right|} \sum_{i \in S_{t}} \nabla F\left(W^{(t)}\right)
$$

(for a random subset $S_{t} \subseteq[n]$ of fixed cardinality b.)
Then, it satisfies with probability at least $1-e^{-\Omega\left(\log ^{2} m\right)}$ over the randomness of $S_{1}, \ldots, S_{T}$ :

$$
F\left(W^{(T)}\right) \leq \epsilon \quad \text { for all } \quad T=\Theta\left(\frac{\operatorname{poly}(n, L) \log ^{2} m}{b \delta^{2}} \log \frac{n \log m}{\epsilon}\right)
$$

The proof of Theorem 5.14.1 is the same as Theorem 5.13.1 plus the careful use of martingale concentration.

Proof of Theorem 5.14.1. Using similar argument as the proof of Theorem 5.13.1, we have with at least $1-e^{-\Omega\left(\log ^{2} m\right)}$ probability

$$
F\left(\vec{W}^{(0)}\right) \leq O\left(n \log ^{2} m\right)
$$

Let us assume for every $t=0,1, \ldots, T-1$, the following holds

$$
\begin{equation*}
\left\|\vec{W}^{(t)}-\vec{W}^{(0)}\right\|_{F} \leq \omega \stackrel{\text { def }}{=} O\left(\frac{n^{3.5} \sqrt{d}}{\delta \sqrt{b m}} \log m\right) \tag{5.21}
\end{equation*}
$$

We shall prove the convergence of SGD assuming (5.21) holds, so that previous statements such as Theorem 5.12.1 and Theorem 5.11.1 can be applied. At the end of the proof, we shall verify that (5.21) is satisfied throughout the SGD with high probability.

To make the proof simple, we choose

$$
\begin{gathered}
m \geq \Omega\left(\frac{n^{24} L^{12} b d \log ^{5} m}{\delta^{8}}\right), \\
\eta=\Theta\left(\frac{b \delta d}{n^{5} L^{2} m \log ^{2} m}\right), \\
T=\Theta\left(\frac{d n^{2}}{\eta \delta m} \log \frac{n \log m}{\epsilon}\right) \\
=\Theta\left(\frac{n^{7} L^{2} \log ^{2} m}{b \delta^{2}} \log \frac{n \log m}{\epsilon}\right)
\end{gathered}
$$

We emphasize that

- Most of the polynomial dependency in $n, L, \delta^{-1}$ come from the non-smoothness of the ReLU activation; if one instead studies smooth activations, their power can be significantly reduced. For instance, for smooth activation functions, one does not need the semi-smoothness Theorem 5.12.1.
- We have not tried to tighten the polynomial dependency on $n, L, \delta^{-1}$. We are aware of many ways to improve the constant in the exponents at the expense of complicating the proofs. Since the main focus of this paper is to derive the first polynomial running time, we do not include such improvements.

For each $t=0,1, \ldots, T-1$, using the same notation as Theorem 5.13.1, except that
we choose $\nabla_{t}=\frac{n}{\left|S_{t}\right|} \sum_{i \in S_{t}} \nabla F_{i}\left(\vec{W}^{(t)}\right)$. We have $\mathbb{E}_{S_{t}}\left[\nabla_{t}\right]=\nabla F\left(\vec{W}^{(t)}\right)$ and therefore

$$
\begin{align*}
& \underset{S_{t}}{\mathbb{E}}\left[F\left(\vec{W}^{(t+1)}\right)\right] \\
& \stackrel{\otimes}{\leq} F\left(\vec{W}^{(t)}\right)-\eta\left\|\nabla F\left(\vec{W}^{(t)}\right)\right\|_{F}^{2}+\eta \sqrt{n F\left(\vec{W}^{(t)}\right)} \cdot O\left(\frac{\omega^{1 / 3} L^{2} \sqrt{m \log m}}{\sqrt{d}}\right) \cdot \underset{S_{t}}{\mathbb{E}}\left[\left\|\nabla_{t}\right\|_{2}\right] \\
& +O\left(\eta^{2} \frac{n L^{2} m}{d}\right) \underset{S_{t}}{\mathbb{E}}\left[\left\|\nabla_{t}\right\|_{2}^{2}\right] \\
& \stackrel{(2)}{\leq} F\left(\vec{W}^{(t)}\right)-\eta\left\|\nabla_{t}\right\|_{F}^{2}+O\left(\frac{\eta n L^{2} m \omega^{1 / 3} \sqrt{\log m}}{d}+\frac{\eta^{2} n^{2} L^{2} m^{2}}{d^{2}}\right) \cdot F\left(\vec{W}^{(t)}\right) \\
& \stackrel{(3)}{\leq}\left(1-\Omega\left(\frac{\eta \delta m}{d n^{2}}\right)\right) F\left(\vec{W}^{(t)}\right) . \tag{5.22}
\end{align*}
$$

Above, (1) uses Theorem 5.12.1 and $\mathbb{E}_{S_{t}}\left[\nabla_{t}\right]=\nabla F\left(\vec{W}^{(t)}\right)$; (2) uses Theorem 5.11 .1 which give

$$
\begin{aligned}
& \underset{S_{t}}{\mathbb{E}}\left[\left\|\nabla_{t}\right\|_{2}^{2}\right] \leq \frac{n^{2}}{b} \underset{S_{t}}{\mathbb{E}}\left[\sum_{i \in S_{t}} \max _{\ell \in[L]}\left\|\nabla_{W_{\ell}} F_{i}\left(\vec{W}^{(t)}\right)\right\|_{F}^{2}\right] \leq O\left(\frac{n m F\left(\vec{W}^{(t)}\right)}{d}\right) \\
& \underset{S_{t}}{\mathbb{E}}\left[\left\|\nabla_{t}\right\|_{2}\right] \leq\left(\frac{\underset{S_{t}}{\mathbb{E}}}{\mathbb{E}}\left[\left\|\nabla_{t}\right\|_{2}^{2}\right]\right)^{1 / 2} \leq O\left(\left(\frac{n m F\left(\vec{W}^{(t)}\right)}{d}\right)^{1 / 2}\right)
\end{aligned}
$$

(3) use gradient lower bound from Theorem 5.11.1 and our choice of $\eta$.

At the same time, we also have the following absolute value bound:

$$
\begin{align*}
F\left(\vec{W}^{(t+1)}\right) \stackrel{\oplus}{\leq} & F\left(\vec{W}^{(t)}\right)+\eta\left\|\nabla F\left(\vec{W}^{(t)}\right)\right\|_{F} \cdot\left\|\nabla_{t}\right\|_{F} \\
& +\eta \sqrt{n F\left(\vec{W}^{(t)}\right)} \cdot O\left(\frac{\omega^{1 / 3} L^{2} \sqrt{m \log m}}{\sqrt{d}}\right) \cdot\left\|\nabla_{t}\right\|_{2}+O\left(\eta^{2} \frac{n L^{2} m}{d}\right) \cdot\left\|\nabla_{t}\right\|_{2}^{2} \\
\stackrel{(2)}{\leq} & F\left(\vec{W}^{(t)}\right)+\eta \cdot O\left(\sqrt{\frac{L F\left(\vec{W}^{(t)}\right) m n}{d}}\right) \cdot O\left(\sqrt{\frac{n^{2} m L F\left(\vec{W}^{(t)}\right)}{b d}}\right) \\
& +\eta \sqrt{n F\left(\vec{W}^{(t)}\right)} \cdot O\left(\frac{\omega^{1 / 3} L^{2} \sqrt{m \log m}}{\sqrt{d}}\right) \cdot \frac{\sqrt{n^{2} m F\left(\vec{W}^{(t)}\right)}}{\sqrt{b d}} \\
& +O\left(\eta^{2} \frac{n L^{2} m}{d}\right) \cdot \frac{n^{2}}{b} O\left(\frac{m F\left(\vec{W}^{(t)}\right)}{d}\right) \\
\stackrel{(3)}{\leq}(1 & \left.+O\left(\frac{\eta L m n^{1.5}}{\sqrt{b} d}+\frac{\eta n^{1.5} \omega^{1 / 3} L^{2} m \sqrt{\log m}}{\sqrt{b} d}+\frac{\eta^{2} n^{3} L^{2} m^{2}}{d^{2} b}\right)\right) F\left(\vec{W}^{(t)}\right) \cdot \tag{5.23}
\end{align*}
$$

Above, (1) uses Theorem 5.12.1 and Cauchy-Schwarz $\langle A, B\rangle \leq\|A\|_{F}\|B\|_{F}$, and (2) uses Theorem 5.11.1 which give

$$
\begin{aligned}
& \left\|\nabla_{t}\right\|_{2}^{2} \leq \frac{n^{2}}{b}\left[\sum_{i \in S_{t}} \max _{\ell \in[L]}\left\|\nabla_{W_{\ell}} F_{i}\left(\vec{W}^{(t)}\right)\right\|_{F}^{2}\right] \leq \frac{n^{2}}{b} O\left(\frac{m F\left(\vec{W}^{(t)}\right)}{d}\right) \\
& \left\|\nabla_{t}\right\|_{F}^{2} \leq \frac{n^{2}}{b}\left[\sum_{i \in S_{t}} \sum_{\ell=1}^{L}\left\|\nabla_{W_{\ell}} F_{i}\left(\vec{W}^{(t)}\right)\right\|_{F}^{2}\right] \leq \frac{L n^{2}}{b} O\left(\frac{m F\left(\vec{W}^{(t)}\right)}{d}\right)
\end{aligned}
$$

and the derivation from (5.22).
Next, taking logarithm on both sides of (5.22) and (5.23), and using Jensen's inequality $\mathbb{E}[\log X] \leq \log \mathbb{E}[X]$, we have $\mathbb{E}\left[\log F\left(\vec{W}^{(t+1)}\right)\right] \leq \log F\left(\vec{W}^{(t)}\right)-\Omega\left(\frac{\eta \delta m}{d n^{2}}\right) \quad$ and $\quad \log F\left(\vec{W}^{(t+1)}\right) \leq \log F\left(\vec{W}^{(t)}\right)+O\left(\frac{\eta L m n^{1.5}}{\sqrt{b} d}\right)$ By (one-sided) martingale concentration, we have with probability at least $1-e^{-\Omega\left(\log ^{2} m\right)}$, for every $t=1,2, \ldots, T$ :

$$
\log F\left(\vec{W}^{(t)}\right)-\mathbb{E}\left[\log F\left(\vec{W}^{(t)}\right)\right] \leq \sqrt{t} \cdot O\left(\frac{\eta L m n^{1.5}}{\sqrt{b} d}\right) \cdot \log m
$$

This implies for every $t=1,2, \ldots, T$, we have

$$
\begin{aligned}
\log F\left(\vec{W}^{(t)}\right) \leq & \sqrt{t} \cdot O\left(\frac{\eta L m n^{1.5}}{\sqrt{b} d}\right) \cdot \log m+\log F\left(\vec{W}^{(0)}\right)-\Omega\left(\frac{\eta \delta m}{d n^{2}}\right) t \\
\stackrel{\oplus}{=} & \log F\left(\vec{W}^{(0)}\right)-\left(\sqrt{\frac{\eta \delta m}{d n^{2}}} \cdot \Omega(\sqrt{t})-\sqrt{\frac{d n^{2}}{\eta \delta m}} \cdot O\left(\frac{\eta L m n^{1.5}}{\sqrt{b} d} \log m\right)\right)^{2} \\
& +O\left(\frac{\eta L^{2} m n^{5}}{b \delta d} \log ^{2} m\right) \\
& \stackrel{(2)}{\leq} \log F\left(\vec{W}^{(0)}\right)+1-\left(\sqrt{\frac{\eta \delta m}{d n^{2}}} \cdot \Omega(\sqrt{t})-\sqrt{\frac{d n^{2}}{\eta \delta m}} \cdot O\left(\frac{\eta L m n^{1.5}}{\sqrt{b} d} \log m\right)\right)^{2} \\
& \stackrel{\otimes}{\leq} \log F\left(\vec{W}^{(0)}\right)+1-\mathbb{1}\left[t \geq \Theta\left(\frac{L^{2} n^{7}}{b \delta^{2}} \log ^{2} m\right)\right] \cdot \Omega\left(\frac{\eta \delta m}{d n^{2}} t\right) \\
& \stackrel{\oplus}{\leq} \log F\left(\vec{W}^{(0)}\right)+1-\mathbb{1}\left[t \geq \Theta\left(\frac{L^{2} n^{7}}{b \delta^{2}} \log ^{2} m\right)\right] \cdot \Omega\left(\frac{b \delta^{2}}{L^{2} n^{7} \log ^{2} m} t\right) .
\end{aligned}
$$

Above, in (1) we have used $2 a \sqrt{t}-b^{2} t=-(b \sqrt{t}-a / b)^{2}+a^{2} / b^{2}$; in (2) we have used our choice of $\eta$; in (3) we have used $-(a \sqrt{t}-b)^{2} \leq-\mathbb{1}\left[t \geq 2 b^{2} / a^{2}\right] \cdot \frac{a^{2} t}{4}$; and in (4) we have used our choice of $\eta$ again. We can read two things from the above formula:

- If $T \geq \Omega\left(\frac{L^{2} n^{7}}{b \delta^{2}} \log ^{2} m \log \frac{n \log m}{\epsilon}\right)$ then we have

$$
\log F\left(\vec{W}^{(T)}\right) \leq \log O\left(n \log ^{2} m\right)-\Omega\left(\log \frac{n \log ^{2} m}{\epsilon}\right) \leq \log \epsilon
$$

so $F\left(\vec{W}^{(T)}\right) \leq \epsilon$.

- Letting $T_{0}=\Omega\left(\frac{L^{2} n^{7}}{b \delta^{2}} \log ^{2} m\right)$, we have

$$
\sum_{i=0}^{t-1} \sqrt{F\left(\vec{W}^{(i)}\right)} \leq \sqrt{n \log ^{2} m} \cdot 2 T_{0}+\frac{\sqrt{n \log ^{2} m}}{2} \cdot 2 T_{0}+\frac{\sqrt{n \log ^{2} m}}{4} \cdot 2 T_{0}+\cdots \leq O\left(\sqrt{n \log ^{2} m} T_{0}\right)
$$

and therefore one can verify that $\left\|\vec{W}^{(t)}-\vec{W}^{(0)}\right\|_{F}$ is small and (5.21) holds: by Theorem 5.11.1,

$$
\begin{aligned}
\left\|W_{\ell}^{(t)}-W_{\ell}^{(0)}\right\|_{F} & \leq \sum_{i=0}^{t-1}\left\|\eta \frac{n}{\left|S_{t}\right|} \sum_{i \in S_{t}} \nabla_{W_{\ell}} F_{i}\left(\vec{W}^{(t)}\right)\right\|_{F} \leq O\left(\eta \sqrt{\frac{n^{2} m}{b d}}\right) \cdot \sum_{i=0}^{t-1} \sqrt{F\left(\vec{W}^{(i)}\right)} \\
& \leq O\left(\eta \sqrt{\frac{n^{2} m}{b d}}\right) \cdot O\left(T_{0} \sqrt{n} \log m\right) \leq O\left(\frac{n^{3.5} \sqrt{d}}{\delta \sqrt{b m}} \log m\right)
\end{aligned}
$$

### 5.15 Theorem 5.15.1: Equivalence to Neural Tangent Kernel

Recall on input $x \in \mathbb{R}^{\boldsymbol{d}}$, the network output $y(\vec{W} ; x) \stackrel{\text { def }}{=} y=B h_{L} \in \mathbb{R}^{d}$ is a function of the weights $\vec{W}$. The neural tangent kernel (NTK) [JGH18] is usually referred to as the feature space defined by the network gradient at random initialization. In other words, for the $j$-th output dimension,

- the NTK kernel function $K_{j}^{\text {ntk }}(x, \widetilde{x}) \stackrel{\text { def }}{=}\left\langle\nabla y_{j}\left(\vec{W}^{(0)} ; x\right), \nabla y_{j}\left(\vec{W}^{(0)} ; \widetilde{x}\right)\right\rangle$
- the NTK objective $y_{j}^{\text {ntk }}\left(\vec{W}^{\prime} ; x\right) \stackrel{\text { def }}{=}\left\langle\nabla y_{j}\left(\vec{W}^{(0)} ; x\right), \vec{W}^{\prime}\right\rangle$.

We have the following theorem whose proof is subsumed by the proofs of Theorem 5.11.1 and 5.12.1. We prove it here for completeness' sake.

Theorem 5.15.1. Let $\vec{W}^{(0)}, A, B$ be at random initialization. For every fixed unit vector $x \in \mathbb{R}^{\mathbb{d}}$, every (small) parameter $\omega \in\left[\Omega\left(\frac{d^{3 / 2}}{m^{3 / 2} L^{3 / 2} \log ^{3 / 2} m}\right), O\left(\frac{1}{L^{4.5} \log ^{3} m}\right)\right]$, with probability at least $1-e^{-\Omega\left(m \omega^{2 / 3} L\right)}$ over $\vec{W}^{(0)}, A, B$, we have for all $\vec{W}^{\prime} \in\left(\mathbb{R}^{m \times m}\right)^{L}$ with $\left\|\vec{W}^{\prime}\right\|_{2} \leq \omega$, for all $j \in[d]$,
(a) $\left\|\nabla y_{j}\left(\vec{W}^{(0)}+\vec{W}^{\prime} ; x\right)-\nabla y_{j}^{\text {ntk }}\left(\vec{W}^{\prime} ; x\right)\right\|_{F} \leq O\left(\sqrt{\log m} \cdot \omega^{1 / 3} L^{3}\right) \cdot\left\|\nabla y_{j}^{\text {ntk }}\left(\vec{W}^{\prime} ; x\right)\right\|_{F}$; and
(b) $y_{j}\left(\vec{W}^{(0)}+\vec{W}^{\prime} ; x\right)=y_{j}\left(\vec{W}^{(0)} ; x\right)+y_{j}^{\text {ntk }}\left(\vec{W}^{\prime} ; x\right)+O\left(\frac{L^{3} \omega^{4 / 3} \sqrt{m \log m}}{\sqrt{d}}\right)$.
(c) If $x, \widetilde{x} \in \mathbb{R}^{\mathcal{D}}$ are two fixed unit vectors, and $\omega \leq O\left(\frac{1}{L^{9} \log ^{3 / 2} m}\right)$, then

$$
\begin{aligned}
\mid\left\langle\nabla y_{j}\left(\vec{W}^{(0)}+\vec{W}^{\prime} ; x\right), \nabla y_{j}\left(\vec{W}^{(0)}+\right.\right. & \left.\left.\vec{W}^{\prime} ; \widetilde{x}\right)\right\rangle-K_{j}^{\mathrm{ntk}}(x, \widetilde{x}) \mid \\
& \leq O\left(\sqrt{\log m} \cdot \omega^{1 / 3} L^{3}\right) \cdot \sqrt{K_{j}^{\mathrm{ntk}}(x, x) K_{j}^{\mathrm{ntk}}(\widetilde{x}, \widetilde{x})}
\end{aligned}
$$

Proof of Theorem 5.15.1. As before we denote by $\widetilde{D}_{1}, \ldots \widetilde{D}_{L}$ and $\widetilde{h}_{1}, \ldots, \widetilde{h}_{L}$ the diagonal sign matrices and forward vectors determined at random initialization $\vec{W}^{(0)}$ and by $D_{1}, \ldots, D_{L}$ and $h_{1}, \ldots, h_{L}$ those determined at $\vec{W} \stackrel{\text { def }}{=} \vec{W}^{(0)}+\vec{W}^{\prime}$. Recall $\left\|D_{\ell}^{\prime}\right\|_{0}=\left\|D_{\ell}-\widetilde{D}_{\ell}\right\|_{0} \leq s \stackrel{\text { def }}{=}$ $O\left(m \omega^{2 / 3} L\right)$ from Lemma 5.9.1b.
(a) Let $e_{j} \in \mathbb{R}^{d}$ be the $j$-th basis vector and $\ell$ be in $[L]$. We have

$$
\begin{aligned}
& \nabla_{W_{\ell}} y_{j}^{\mathrm{ntk}}\left(\vec{W}^{\prime} ; x\right)-\nabla_{W_{\ell}} y_{j}\left(\vec{W}^{(0)}+\vec{W}^{\prime} ; x\right) \\
= & \left(e_{j}^{\top} B \widetilde{D}_{L} \widetilde{W}_{L} \cdots \widetilde{W}_{\ell+1} \widetilde{D}_{\ell}\right)^{\top}\left(\widetilde{h}_{\ell-1}\right)^{\top}-\left(e_{j}^{\top} B D_{L} W_{L} \cdots W_{\ell+1} D_{\ell}\right)^{\top}\left(h_{\ell-1}\right)^{\top}
\end{aligned}
$$

This difference matrix is precisely (5.11) (by setting $n=1$ and $v=e_{j}$ ). Using the bound (5.12) we have its Frobenius norm is at most $O\left(\sqrt{m \log m / d} \cdot \omega^{1 / 3} L^{2}\right)$. On the other hand, one can calculate for every $k \in[m]$,

$$
\nabla_{\left[W_{L}\right]_{k}} y_{j}^{\text {ntk }}\left(\overrightarrow{W^{\prime}} ; x\right)=\left(B^{\top} e_{j}\right)_{k} \cdot \widetilde{h}_{L-1} \cdot \mathbb{1}_{\left\langle\left[W_{L}\right]_{k}, \widetilde{h}_{L-1}\right\rangle \geq 0} .
$$

We already know $\left\|\widetilde{h}_{L-1}\right\| \geq \Omega(1)$ from Lemma 5.8.1. Now, regardless of the randomness of $\widetilde{h}_{L-1}$, we have $\mathbb{1}_{\left\langle\left[W_{L}\right]_{k}, \widetilde{h}_{L-1}\right\rangle \geq 0}=1$ with exactly half probability; also, regardless of the randomness of $\vec{W}^{(0)}$ and $A$, we have $\left(B^{\top} e_{j}\right)_{k} \sim \mathcal{N}\left(0, \frac{1}{d}\right)$. Therefore, we conclude that with probability at least $1-e^{-\Omega(m)}$ it satisfies $\left\|\nabla_{W_{L}} y_{j}^{\text {ntk }}\left(\overrightarrow{W^{\prime}} ; x\right)\right\|_{F}^{2} \geq \Omega(m / d)$. Putting the two abounds together we finish the proof.
(b) This statement can be derived from (5.15), (5.17) and (5.18). For completeness' sake,
below we provide a direct proof without invoking them. We first calculate that

$$
\begin{align*}
& \quad\left|y_{j}\left(\vec{W}^{(0)}+\vec{W}^{\prime} ; x\right)-e_{j}^{\top} B \widetilde{D}_{L} W_{L} \cdots \widetilde{D}_{1} W_{1} x\right|=\left|\sum_{\ell=1}^{L} e_{j}^{\top} B \widetilde{D}_{L} W_{L} \cdots \widetilde{D}_{\ell+1} W_{\ell+1} D_{\ell}^{\prime} g_{\ell-1}\right| \\
& \leq \sum_{\ell=1}^{L}(\underbrace{\left\|B \widetilde{D}_{L} \widetilde{W}_{L} \cdots \widetilde{D}_{\ell+1} \widetilde{W}_{\ell+1} D_{\ell}^{\prime}\right\|_{2}}_{\text {Lemma 5.8.4a }} \\
& \quad+\underbrace{\left\|B \widetilde{D}_{L} \widetilde{W}_{L} \cdots \widetilde{D}_{\ell+1} \widetilde{W}_{\ell+1}-B \widetilde{D}_{L} W_{L} \cdots \widetilde{D}_{\ell+1} W_{\ell+1}\right\|_{2}}_{\text {Lemma } 5.9 .6}) \cdot \underbrace{\left\|D_{\ell}^{\prime} g_{\ell-1}\right\|}_{\text {Lemma 5.9.1b }} \\
& \leq L \cdot\left(O\left(\frac{\sqrt{s \log m}}{\sqrt{d}}\right)+O\left(\omega L^{1.5} \frac{\sqrt{m}}{\sqrt{d}}\right)\right) \cdot O\left(\omega L^{3 / 2}\right) \leq O\left(\frac{L^{3} \omega^{4 / 3} \sqrt{m \log m}}{\sqrt{d}}\right) \tag{5.24}
\end{align*}
$$

We next calculate that

$$
\begin{align*}
& \left|e_{j}^{\top} B \widetilde{D}_{L} W_{L} \cdots \widetilde{D}_{1} W_{1} x-y_{j}\left(\vec{W}^{(0)}\right)-y^{\mathrm{ntk}}\left(\vec{W}^{\prime} ; x\right)\right| \\
= & \left|\sum_{\ell=1}^{L} e_{j}^{\top} B \widetilde{D}_{L} W_{L} \cdots W_{\ell+1} \widetilde{D}_{\ell} W_{\ell}^{\prime} \widetilde{h}_{\ell-1}-e_{j}^{\top} B \widetilde{D}_{L} \widetilde{W}_{L} \cdots \widetilde{W}_{\ell+1} \widetilde{D}_{\ell} W_{\ell}^{\prime} \widetilde{h}_{\ell-1}\right| \\
\leq & \sum_{\ell=1}^{L} \underbrace{\left\|B\left(\widetilde{D}_{L} W_{L} \cdots W_{\ell+1} \widetilde{D}_{\ell}-\widetilde{D}_{L} \widetilde{W}_{L} \cdots \widetilde{W}_{\ell+1} \widetilde{D}_{\ell}\right)\right\|_{2}}_{\text {Lemma 5.9.6 }} \cdot\left\|W_{\ell}^{\prime}\right\|_{2} \cdot \underbrace{\left\|\widetilde{h}_{\ell-1}\right\|}_{\text {Lemma } 5.8 .1} \\
\leq & L \cdot O\left(\omega L^{1.5} \frac{\sqrt{m}}{\sqrt{d}}\right) \cdot \omega \cdot O(1) \leq O\left(\omega^{2} L^{2.5} \frac{\sqrt{m}}{\sqrt{d}}\right) \tag{5.25}
\end{align*}
$$

Putting (5.24) and (5.25) together finishes the proof.
(c) This is a direct corollary of (a).

### 5.16 Extension to Other Loss Functions

For simplicity, in them main body of this paper we have used the $\ell_{2}$ regression loss. Our results generalize easily to other Lipschitz smooth (but possibly nonconvex) loss functions.

Suppose we are given loss function $f(z ; y)$ that takes as input a neural-network output $z \in \mathbb{R}^{d}$ and a label $y$. Then, our training objective for the $i$-th training sample becomes $F_{i}(W)=f\left(B h_{i, L} ; y_{i}^{*}\right)$. We redefine the loss vector $\operatorname{loss}_{i} \stackrel{\text { def }}{=} \nabla f\left(B h_{i, L} ; y_{i}^{*}\right) \in \mathbb{R}^{d}$ (where the gradient is with respect to $z$ ). Note that if $f(z ; y)=\frac{1}{2}\|z-y\|^{2}$ is the $\ell_{2}$ loss, then this notion coincides with Section 5.2. We assume that $f(z ; y)$ is 1-Lipscthiz (upper) smooth with respect to $z .{ }^{17}$

All the results in Section 5.8, 5.9 and 5.10 remain unchanged. Section 5.11 also remains unchanged, except we need to restate Theorem 5.11 .1 with respect to this new notation:

$$
\begin{array}{ll}
\left\|\nabla_{W_{\ell}} F_{i}(\vec{W})\right\|_{F}^{2} \leq O\left(\frac{\left\|\operatorname{loss}_{i}\right\|^{2}}{d} \times m\right) \quad & \left\|\nabla_{W_{\ell}} F(\vec{W})\right\|_{F}^{2} \leq O\left(\frac{\|\operatorname{loss}\|^{2}}{d} \times m n\right) \\
& \left\|\nabla_{W_{L}} F(\vec{W})\right\|_{F}^{2} \geq \Omega\left(\frac{\max _{i \in[n]}\left\|\operatorname{loss}_{i}\right\|^{2}}{d n / \delta} \times m\right)
\end{array}
$$

Section 5.12 also remains unchanged, except that we need to replace the precise definition of $\ell_{2}$ loss in (5.14) with the semi-smoothness condition:

$$
\begin{align*}
F_{i}(\vec{W})=f\left(B h_{i, L} ; y_{i}^{*}\right) & \leq f\left(B \hbar_{i, L} ; y_{i}^{*}\right)+\left\langle\nabla f\left(B \hbar_{i, L}, y_{i}^{*}\right), B\left(h_{i, L}-\hbar_{i, L}\right)\right\rangle+\frac{1}{2}\left\|B\left(h_{i, L}-\hbar_{i, L}\right)\right\|^{2} \\
& =F_{i}(\stackrel{\breve{W}}{W})+\left\langle\operatorname{loss}_{i}, B\left(h_{i, L}-\hbar_{i, L}\right)\right\rangle+\frac{1}{2}\left\|B\left(h_{i, L}-\hbar_{i, L}\right)\right\|^{2} \tag{5.26}
\end{align*}
$$

[^26]and the rest of the proof remains unchanged.
As for the final convergence theorem of gradient descent, we can replace (5.20) with
\[

$$
\begin{equation*}
F\left(\vec{W}^{(t+1)}\right) \leq F\left(\vec{W}^{(t)}\right)-\Omega\left(\frac{\eta \delta m}{d n^{2}}\right) \cdot\left\|\operatorname{loss}^{(t)}\right\|^{2} \tag{5.27}
\end{equation*}
$$

\]

This means many things:

- If the loss is nonconvex but satisfies the Polyak-Łojasiewicz condition $\|\nabla f(z ; y)\|^{2} \geq$ $\sigma\left(f(z ; y)-f\left(z^{*} ; y\right)\right)$, then in $T=\Omega\left(\frac{d n^{2}}{\eta \delta m \sigma}\right)=O\left(\frac{n^{6} L^{2}}{\delta^{2} \sigma} \log \frac{1}{\epsilon}\right)$ iterations, GD can find a point $\vec{W}^{(T)}$ with $\left\|\operatorname{loss}{ }^{(T)}\right\| \leq \epsilon$. It suffices to choose $m \geq \widetilde{\Omega}\left(\operatorname{poly}\left(n, L, \delta^{-1}\right) \cdot d \sigma^{-2}\right)$ for same reason as before.
- If the loss is nonconvex but bounded (say, $|f(z ; y)| \leq O(1))$, then in $T=O\left(\frac{d n^{2}}{\eta \delta m \epsilon^{2}}\right)=$ $O\left(\frac{n^{6} L^{2}}{\delta^{2} \epsilon^{2}}\right)$ iterations, we can find a point $\vec{W}^{(T)}$ with $\left\|\operatorname{loss}^{(T)}\right\| \leq \epsilon$. If suffices to choose $\left.m \geq \widetilde{\Omega}\left(\operatorname{poly}\left(n, L, \delta^{-1}\right) \cdot d \epsilon^{-1}\right).\right)$
- If the loss is convex and its minimizer has bounded norm, meaning there exists $z^{*}$ so that $f\left(z^{*} ; y\right)=\min _{z} f(z ; y)$ and $\left\|z-z^{*}\right\| \leq D$. Then, by convexity

$$
f(z ; y)-f\left(z^{*} ; y\right) \leq\left\langle\nabla f(z ; y), z-z^{*}\right\rangle \leq D\|\nabla f(z ; y)\|
$$

Putting this into (5.27), we have (here $\vec{W}^{*}=\operatorname{argmin}_{\vec{W}} F_{i}(\vec{W})$ for all $i \in[n]$ )

$$
\begin{aligned}
F\left(\vec{W}^{(t+1)}\right)-F\left(\vec{W}^{*}\right) & \leq F\left(\vec{W}^{(t)}\right)-F\left(\vec{W}^{*}\right)-\Omega\left(\frac{\eta \delta m}{d n^{2} D^{2}}\right) \cdot \sum_{i \in[n]}\left(F_{i}\left(\vec{W}^{(t)}\right)-F_{i}\left(\vec{W}^{*}\right)\right)^{2} \\
& \leq F\left(\vec{W}^{(t)}\right)-F\left(\vec{W}^{*}\right)-\Omega\left(\frac{\eta \delta m}{d n^{3} D^{2}}\right) \cdot\left(F\left(\vec{W}^{(t)}\right)-F\left(\vec{W}^{*}\right)\right)^{2} .
\end{aligned}
$$

This implies (see for instance the classical calculation steps in [Nes04]) that after $T=$ $O\left(\frac{d n^{3} D^{2}}{\eta \delta m \epsilon}\right)=O\left(\frac{n^{7} L^{2} D^{2}}{\delta^{2} \epsilon}\right)$ iterations, we can have $F\left(\vec{W}^{(T)}\right)-F\left(\vec{W}^{*}\right) \leq \epsilon$. The amount of over-parameterization needed is $m \geq \widetilde{\Omega}\left(\operatorname{poly}\left(n, L, \delta^{-1}\right) \cdot d \log \epsilon^{-1}\right)$.

- If the loss is cross entropy $f(z ; y)=\frac{e^{z_{y}}}{\sum_{i=1}^{d} e^{z_{i}}}$ for classification, then $\|\nabla f(z ; y)\|<1 / 4$ implies perfect classification. ${ }^{18}$ Thus, we have $100 \%$ training accuracy in $T=O\left(\frac{n^{6} L^{2}}{\delta^{2}}\right)$ iterations. If suffices to choose $m \geq \widetilde{\Omega}\left(\operatorname{poly}\left(n, L, \delta^{-1}\right) \cdot d\right)$

[^27]
### 5.17 Extension to Convolutional Neural Networks

There are numerous versions of convolutional neural networks (CNNs) that are used in practice. To demonstrate the capability of applying our techniques to such convolutional settings, in this section, we study a simple enough CNN for the $\ell_{2}$ regression task.

A Simple CNN Model We assume that for the input layer (corresponding to $A$ ) and for each hidden layer $\ell=1,2, \ldots, L-1$ (corresponding to $W_{1}, \ldots, W_{L-1}$ ), there are $\mathfrak{d}$ positions each consisting of $m$ channels. (Each position can be thought as a pixel of an image in computer vision tasks.) We assume the last hidden layer $\ell=L$ (corresponding to $W_{L}$ ) and the output layer (corresponding to $B$ ) are fully connected. We assume for each $j \in[\mathfrak{d}]$, there exists a set $Q_{j} \subseteq[\mathfrak{d}]$ of fixed cardinality $q \in[\mathfrak{d}]$ so that the value at position $j$ in any convolutional layer is completely determined by positions $k \in Q_{j}$ of the previous layer.

Assumption 5.17.1. We assume that $\left(Q_{1}, \ldots, Q_{\mathfrak{0}}\right)$ give rise to a $q$-regular bipartite graph: each $Q_{j}$ has exactly $q$ entries and each $k \in[\mathfrak{d}]$ appears in exactly $q$ different sets $Q_{j}$.
(In vision tasks, if $3 \times 3$ kernels are used then $\left|Q_{j}\right|=9$. We ignore the padding issue for simplicity.)

The output of each convolutional layer $\ell=0,1,2, \ldots, L-1$ is represented by a $\mathfrak{d} m$ dimensional vector $h_{\ell}=\left(h_{\ell, 1}, \ldots, h_{\ell, \mathfrak{l}}\right)$ where each $h_{\ell, j} \in \mathbb{R}^{m}, \forall j \in[\mathfrak{d}]$. In the input layer and each $j \in[\mathfrak{d}]$, we assume

$$
h_{0, j}=\phi\left(A_{j} x_{Q_{j}}\right) \in \mathbb{R}^{m}
$$

where $x_{Q_{j}} \in \mathbb{R}^{q}$ denotes the concatenation of $x_{k}$ for all $k \in Q_{j}$ given input $x \in \mathbb{R}^{\mathfrak{d}}$, and $A_{j} \in \mathbb{R}^{m \times q}$ is randomly initialized at $\mathcal{N}\left(0, \frac{2}{\sqrt{q m}}\right)$ per entry. For notational simplicity, we
define matrix $A \in \mathbb{R}^{\boldsymbol{\jmath} m \times \boldsymbol{\mathcal { O }}}$ so that it satisfies $h_{1}=\phi(A x)$. Each row of $A$ has $q$ non-zero entries.

For each layer $\ell=1, \ldots, L-1$ and each $j \in[\mathfrak{d}]$, we assume

$$
h_{\ell, j}=\phi\left(W_{\ell, j} h_{\ell-1, Q_{j}}+\tau \cdot \mathbf{b}_{\ell, j}\right) \in \mathbb{R}^{m}
$$

where $h_{\ell-1, Q_{j}} \in \mathbb{R}^{q m}$ denotes the concatenation of $h_{\ell-1, k}$ for all $k \in Q_{j}$, the weights $W_{\ell, j} \in$ $\mathbb{R}^{m \times(q m)}$ and the bias the $\mathbf{b}_{\ell, j} \in \mathbb{R}^{m}$ are randomly initialized at $\mathcal{N}\left(0, \frac{2}{q m}\right)$ per entry, and $\tau$ is a small parameter (say, $\tau=\frac{\delta^{2}}{100 L}$ ) for bias. For notational simplicity, we define matrix $W_{\ell} \in \mathbb{R}^{\mathfrak{d} m \times \boldsymbol{d} m}$ and vector $\mathbf{b}_{\ell} \in \mathbb{R}^{\boldsymbol{d} m}$ so that it satisfies $h_{\ell}=\phi\left(W_{\ell} h_{\ell-1}+\tau \mathbf{b}_{\ell}\right)$, and define vector $g_{\ell} \xlongequal{\text { def }} W_{\ell} h_{\ell-1}+\tau \mathbf{b}_{\ell} \in \mathbb{R}^{\boldsymbol{\jmath} m}$. Note that each row of $W_{\ell}$ has $q m$ non-zero entries.

We assume the last layer $W_{L}$ and the output layer $B$ are simply fully connected (say without bias). That is, each entry of $W_{L} \in \mathbb{R}^{\boldsymbol{\mathcal { D } m \times \delta} m}$ is from $\mathcal{N}\left(0, \frac{2}{q m}\right)$, and of $B \in \mathbb{R}^{d \times \mathrm{o} m}$ is from $\mathcal{N}\left(0, \frac{1}{d}\right)$.

We denote by $h_{i, \ell}$ the value of $h_{\ell}$ when the input vector is $x_{i}$, and define $g_{i, \ell}, D_{i, \ell}$ in the same way as before.

### 5.17.1 Changes in the Proofs

If one is willing to loose polynomial factors in $L$ and $\mathfrak{d}$ in the final complexity, then changes to each of the lemmas of this paper is very little. ${ }^{19}$

Changes to Section 5.8 The first main result is Lemma 5.8.1: $\left\|h_{i, \ell}\right\|$ is in $[1-\epsilon, 1+\epsilon]$ with high probability. In the CNN case, for every $j \in[\mathfrak{d}]$, recalling that $h_{i, \ell, j}=\phi\left(W_{\ell, j} h_{i, \ell-1, Q_{j}}+\right.$ $\tau \mathbf{b}_{\ell}$. Applying Fact 5.8.2, we have that $\frac{q m\left\|h_{i, \ell, j}\right\|^{2}}{2\left(\left\|h_{i, \ell-1, Q_{j}}\right\|^{2}+\tau^{2}\right)}$ is distributed as $\chi_{\omega}^{2}$ distribution with $\omega \sim \mathcal{B}\left(m, \frac{1}{2}\right)$. Due to the concentration of $\chi^{2}$ distribution and the concentration of binomial distribution, $\left\|h_{i, \ell, j}\right\|^{2}$ is extremely close to $\frac{\left\|h_{i, \ell-1, Q_{j}}\right\|^{2}+\tau^{2}}{q}$ (a careful argument of this can be found in the proof of Lemma 5.8.1). Summing this up over all $j \in[\mathfrak{d}]$, and using Assumption 5.17.1, we have $\left\|h_{i, \ell, j}\right\|^{2}$ is concentrated at $\left\|h_{i, \ell-1}\right\|^{2}+\frac{\tau^{2} \mathrm{~J}}{q}$. Applying induction, we have $\left\|h_{i, \ell}\right\|$ is in $[1-\epsilon, 1+\epsilon]$ with probability at least $1-e^{-\Omega\left(m \epsilon^{2} / L^{2}\right)}$, as long as $\tau^{2} \leq \frac{\epsilon q}{100 L} .{ }^{20}$

The changes to Lemma 5.8.3 and Lemma 5.8.4 are the same as above, but we loose some polynomial factors in $L$ (because we are not careful in the argument above). For instance, the intermediate bound in Lemma 5.8.3a becomes $\left\|W_{b} D_{i, b-1} W_{b-1} \cdots D_{i, a} W_{a}\right\|_{2} \leq$ $O(L)$.

As for the $\delta$-separateness Lemma 5.8.5, we need to redefine the notion of $\delta$-separateness between $h_{i, \ell}$ and $h_{j, \ell}$ :

$$
\begin{equation*}
\sum_{k \in[\mathfrak{p}]}\left\|\left(I-\frac{h_{i, \ell, k} h_{i, \ell, k}^{\top}}{\left\|h_{i, \ell, k}\right\|^{2}}\right) h_{j, \ell, k}\right\|^{2} \geq \Omega\left(\delta^{2}\right) \tag{5.28}
\end{equation*}
$$

[^28]Then, denoting by $\widehat{h}_{k}=h_{i, \ell-1, k} /\left\|h_{i, \ell-1, k}\right\|$, we have

$$
h_{j, \ell, k}=\phi\left(W_{\ell, k} h_{j, \ell-1, Q_{j}}+\tau \mathbf{b}_{\ell, k}\right)=\phi\left(\vec{g}_{1}+\left(\sum_{z \in Q_{k}}\left\|\left(I-\widehat{h}_{z} \hat{h}_{z}^{\top}\right) h_{j, \ell-1, z}\right\|^{2}\right)^{1 / 2} \vec{g}_{2}\right)
$$

where $\vec{g}_{2} \sim \mathcal{N}\left(0, \frac{2}{q m} I\right)$ is independent of the randomness of $h_{i, \ell, k}$ once $A, W_{1}, \ldots, W_{\ell-1}$ are fixed. One can use this to replace (5.6) and the rest of the proof follows.

Changes to Section 5.9 The first main result is Lemma 5.9.1, and we discuss necessary changes here to make it work for CNN. The first change in the proof is to replace $2 c_{1} L^{1.5}$ with $2 c_{1} L^{2}$ due to the above additional factor from Lemma 5.8.3a. Next, call that the proof of Lemma 5.9.1 relied on Claim 5.9.2 and Claim 5.9.4:

- For Claim 5.9.2, we can replace the definition of $x$ with $x=D^{\prime}\left(\widetilde{W} \widetilde{h}+\tau \mathbf{b}+g^{\prime}\right)$ for $\mathbf{b} \in \mathcal{N}\left(0, \frac{2}{q m} I\right)$. This time, instead of using the randomness of $\widetilde{W}$ like in the old proof (because $\widetilde{W}$ is no longer a full matrix), we use the randomness of $\tau \mathbf{b}$. The new statement becomes

$$
\|x\|_{0} \leq O\left(\frac{\mathfrak{d} m}{\tau^{2 / 3}}\left\|g_{1}^{\prime}\right\|^{2 / 3}+\frac{1}{\tau}\left\|g_{2}^{\prime}\right\|_{\infty}(\mathfrak{d} m)^{3 / 2}\right) \quad \text { and } \quad\|x\| \leq O\left(\left\|g_{1}^{\prime}\right\|+\frac{1}{\sqrt{\tau}}\left\|g_{2}^{\prime}\right\|_{\infty}^{3 / 2}(\mathfrak{d} m)^{3 / 4}\right)
$$

and its proof is by re-scaling $x$ by $\frac{1}{\tau}$ and then applying the old proof (with dimension $m$ replaced with $\mathfrak{d} m)$.

- For Claim 5.9.4, it becomes $\left\|y_{1}\right\| \leq O(\sqrt{q s / m} \log m) \quad$ and $\quad\left\|y_{2}\right\|_{\infty} \leq \frac{2 \sqrt{\log m}}{\sqrt{q m}}$.

After making all of these changes, we loose at most some polynomial factors in $L$ and $\mathfrak{d}$ for the new statement of Lemma 5.9.1:
(a) $\left\|D_{i, \ell}^{\prime}\right\|_{0} \leq m \omega^{2 / 3} \operatorname{poly}(L, \mathfrak{d})$ and $\left\|D_{i, \ell}^{\prime} g_{i, \ell}\right\| \leq \omega \operatorname{poly}(L, \mathfrak{d})$.
(b) $\left\|g_{i, \ell}^{\prime}\right\|,\left\|h_{i, \ell}^{\prime}\right\| \leq \omega \operatorname{poly}(L, \mathfrak{d}) \sqrt{\log m}$.

Finally, the statements of Lemma 5.9.5 and Lemma 5.9.6 only loose polynomial factors in $L$ and $\mathfrak{d}$.

Changes to Section 5.10 The norm upper bound part is trivial to modify so we only focus on the gradient norm lower bound. Since we have assumed $W_{L}$ to be fully connected, the gradient on $W_{L}$ is the same as before:

$$
\widehat{\nabla}_{\left[W_{L}\right]_{k}}^{\hat{\sigma}^{2}} F(\vec{W})=\sum_{i=1}^{n}\left\langle B_{k}, v_{i}\right\rangle \cdot h_{i, L-1} \cdot \mathbb{1}_{\left(W_{L} h_{i, L-1}\right)_{k} \geq 0}
$$

Since we still have $\delta$-separateness (5.28), one can verify for $\ell=L-1$,

$$
\left\|h_{i, \ell}-h_{j, \ell}\right\|^{2}=\sum_{k \in[0]}\left\|h_{i, \ell, k}-h_{j, \ell, k}\right\|^{2} \geq \sum_{k \in[0]}\left\|\left(I-\frac{h_{i, \ell, k} h_{i, \ell, k}^{\top}}{\left\|h_{i, \ell, k}\right\|^{2}}\right) h_{j, \ell, k}\right\|^{2} \geq \Omega\left(\delta^{2}\right)
$$

Since $\left\|h_{i, \ell}\right\| \approx 1$ and $\left\|h_{j, \ell}\right\| \approx 1$, this gives back the old definition of $\delta$-separateness: $\left(I-h_{i, \ell} h_{i, \ell}^{\top} /\left\|h_{i, \ell}\right\|^{2}\right) h_{j, \ell}$ has norm at least $\Omega(\delta)$. Therefore, the entire rest of Section 5.10 follows as before.

Final Theorem Since Section 5.11 and 5.12 rely on previous sections, they do not need to be changed (besides some polynomial factor blowup in $L$ and $\mathfrak{d}$ ). Our final theorem becomes

Theorem 5.17.2 (CNN). Let $m \geq \widetilde{\Omega}\left(\operatorname{poly}\left(n, L, \mathfrak{d}, \delta^{-1}\right) \cdot d\right)$. For the convolutional neural network defined in this section, with probability at least $1-e^{-\Omega\left(\log ^{2} m\right)}$ over the random initialization, GD and SGD respectively need at most $T=\frac{\operatorname{poly}(n, L, \mathbf{0})}{\delta^{2}} \log \frac{1}{\epsilon}$ and $T=$ $\frac{\operatorname{poly}(n, L, \mathfrak{)}) \cdot \log ^{2} m}{\delta^{2}} \log \frac{1}{\epsilon}$ iterations to find a point $F(\vec{W}) \leq \epsilon$.

### 5.18 Extension to Residual Neural Networks

Again as we have discussed in Section 5.18, there are numerous versions of residual neural networks that are used in practice. To demonstrate the capability of applying our techniques to residual settings, in this section, we study a simple enough residual network for the $\ell_{2}$ regression task (without convolutional layers).

A Simple Residual Model We consider an input layer $h_{0}=\phi(A x), L-1$ residual layers $h_{\ell}=\phi\left(h_{\ell-1}+\tau W_{\ell} h_{\ell-1}\right)$ for $\ell=1,2, \ldots, L-1$, a fully-connected layer $h_{L}=\phi\left(W_{L} h_{L-1}\right)$ and an output layer $y=B h_{L}$. We assume that $h_{0}, \ldots, h_{L} \in \mathbb{R}^{m}$ and the entries of $W_{\ell} \in \mathbb{R}^{m \times m}$ are from $\mathcal{N}\left(0, \frac{2}{m}\right)$ as before. We choose $\tau=\frac{1}{\Omega(L \log m)}$ which is similar as previous work [ZLSD18].

We denote by $g_{0}=A x, g_{\ell}=h_{\ell-1}+\tau W_{\ell} h_{\ell-1}$ for $\ell=1,2, \ldots, L-1$ and $g_{L}=W_{L} h_{L-1}$. For analysis, we use $h_{i, \ell}$ and $g_{i, \ell}$ to denote the value of $h_{\ell}$ when the input vector is $x_{i}$, and $D_{i, \ell}$ the diagonal sign matrix so that $\left[D_{i, \ell}\right]_{k, k}=\mathbb{1}_{\left(g_{i, \ell}\right)_{k} \geq 0}$.

### 5.18.1 Changes in the Proofs

Conceptually, we need to replace all the occurrences of $W_{\ell}$ with $\left(I+\tau W_{\ell}\right)$ for $\ell=$ $1,2, \ldots, L-1$. Many of the proofs in the residual setting becomes much simpler when residual links are present. The main property we shall use is that the spectral norm

$$
\begin{equation*}
\left\|\left(I+\tau W_{a}\right) D_{i, a+1} \cdots D_{i, b}\left(I+\tau W_{b}\right)\right\|_{2} \leq 1.01 \tag{5.29}
\end{equation*}
$$

for any $L-1 \geq a \geq b \geq 1$ with our choice of $\tau$.

Changes to Section 5.8 For Lemma 5.8.1, ignoring subscripts in $i$ for simplicity, we can combine the old proof with (5.29) to derive that $\left\|h_{\ell}\right\| \leq 1.02$ for every $i$ and $\ell$. We also have $\left\|h_{\ell}\right\| \geq \frac{1}{\sqrt{20}}$ by the following argument.

- Fact 5.8.2 says each coordinate of $h_{0}$ follows i.i.d. from a distribution which is 0 with half probability, and $\left|\mathcal{N}\left(0, \frac{2}{m}\right)\right|$ with half probability. Therefore, with high probability, at least $m / 4$ of the coordinates $k \in[m]$ will satisfy $\left|\left(h_{0}\right)_{k}\right| \geq \frac{0.6}{\sqrt{m}}$. Denote this set as $M_{0} \subseteq[m]$.
- In the following layer $\ell=1,\left(h_{\ell}\right)_{k} \geq\left(h_{\ell-1}\right)_{k}-\tau\left|\left(W_{\ell} h_{\ell-1}\right)_{k}\right|$. Since $W_{\ell} h_{\ell-1} \sim \mathcal{N}\left(0, \frac{2\left\|h_{\ell-1}\right\|^{2}}{m} I\right)$ and $\left\|h_{\ell-1}\right\| \leq 1.02$, we know with high probability, at least $1-\frac{1}{10 L}$ fraction of the coordinates in $M_{0}$ will satisfy $\left|\left(W_{\ell} h_{\ell-1}\right)_{k}\right| \leq O\left(\frac{\log L}{\sqrt{m}}\right)$. Therefore, for each of these $\left(1-\frac{1}{10 L}\right)\left|M_{0}\right|$ coordinates, we have $\left(h_{\ell}\right)_{k} \geq\left(h_{\ell-1}\right)_{k}-\frac{1}{10 L}$ by our choice of $\tau$. Denote this set as $M_{1} \subseteq M_{0}$, then we have $\left(h_{\ell}\right)_{k} \geq \frac{0.6}{\sqrt{m}}-\frac{1}{10 L \sqrt{m}}$ for each $k \in M_{1}$.
- Continuing this argument for $\ell=2,3, \ldots, L-1$, we know that every time we move from $M_{\ell-1}$ to $M_{\ell}$, its size shrinks by a factor $1-\frac{1}{10 L}$, and the magnitude of $\left(h_{\ell}\right)_{k}$ for
$k \in M_{\ell}$ decreases by $\frac{1}{10 L \sqrt{m}}$. Putting this together, we know $\left\|h_{\ell}\right\|^{2} \geq\left(\frac{0.6}{\sqrt{m}}-\frac{1}{10 \sqrt{m}}\right)^{2}$. $\left(1-\frac{1}{10 L}\right)^{L} \cdot \frac{m}{4} \geq \frac{1}{20}$ for all $\ell=1,2, \ldots, L-1$. The proof of the last layer $h_{L}$ is the same as the old proof.

Lemma 5.8.3 is not needed anymore because of (5.29). Lemma 5.8.4 becomes trivial to prove using (5.29): for instance for Lemma 5.8.4a, we have $\| D_{i, L} W_{L} D_{i, L-1}\left(I+\tau W_{L-1}\right) \cdots D_{i, a}(I+$ $\left.\tau W_{a}\right) u \| \leq O(\|u\|)$ and thus $\left\|B D_{i, L} W_{L} D_{i, L-1}\left(I+\tau W_{L-1}\right) \cdots D_{i, a}\left(I+\tau W_{a}\right) u\right\| \leq O\left(\frac{\sqrt{s} \log m}{\sqrt{d}}\right)\|u\|$ for all $s$-sparse vectors $u$.

Lemma 5.8.5 needs the following changes in the same spirit as our changes to Lemma 5.8.1. With probability at least $1-e^{-\Omega\left(\log ^{2} m\right)}$ it satisfies $\left\|W_{\ell} h_{i, \ell}\right\|_{\infty} \leq O\left(\frac{\log m}{\sqrt{m}}\right)$ for all $i \in[n]$ and $\ell \in L$. In the following proof we condition on this event happens. ${ }^{21}$ Consider $i, j \in[n]$ with $i \neq j$.

- In the input layer, since $\left\|x_{i}-x_{j}\right\| \geq \delta$, the same Claim 5.8.6 shows that, with high probability, there are at least $\frac{3}{4} m$ coordinates $k \in[m]$ with $\left|\left(h_{i, 0}-h_{j, 0}\right)_{k}\right| \geq \frac{\delta}{10 \sqrt{m}}$. At the same time, at least $\frac{3}{4} m$ coordinates $k \in[m]$ will satisfy $\left(h_{i, 0}\right)_{k} \geq \frac{1}{10 \sqrt{m}}$ and $\left(h_{j, 0}\right)_{k} \geq \frac{1}{10 \sqrt{m}}$. Denote $M_{0} \subseteq[m]$ as the set of coordinates $k$ satisfying both properties. We have $\left|M_{0}\right| \geq \frac{m}{2}$ and $\sum_{k \in M_{0}}\left|\left(h_{i, 0}-h_{j, 0}\right)_{k}\right| \geq \frac{\delta}{20} \sqrt{m}$.
- In the following layer $\ell=1$, we have

$$
\left(h_{i, \ell}-h_{j, \ell}\right)_{k}=\phi\left(\left(h_{i, \ell-1}\right)_{k}+\tau\left(W_{\ell} h_{i, \ell-1}\right)_{k}\right)-\phi\left(\left(h_{j, \ell-1}\right)_{k}+\tau\left(W_{\ell} h_{j, \ell-1}\right)_{k}\right)
$$

[^29]Using $\left\|W_{\ell} h_{i, \ell}\right\|_{\infty} \leq O\left(\frac{\log m}{\sqrt{m}}\right)$ and our choice of $\tau$, we know for every $k \in M_{0}$, it satisfies $\left(h_{i, \ell}\right)_{k} \geq \frac{1}{10 \sqrt{m}}-\frac{1}{100 L \sqrt{m}}$ and $\left(h_{j, \ell}\right)_{k} \geq \frac{1}{10 \sqrt{m}}-\frac{1}{100 L \sqrt{m}}$. Therefore, the ReLU activation becomes identity for such coordinates $k \in M_{0}$ and

$$
\Delta_{k} \stackrel{\text { def }}{=}\left(h_{i, \ell}-h_{j, \ell}\right)_{k}=\left(h_{i, \ell-1}-h_{j, \ell-1}\right)_{k}+\tau\left(W_{\ell}\left(h_{i, \ell-1}-h_{j, \ell-1}\right)\right)_{k} .
$$

Let $s_{k}=1$ if $\left(h_{i, \ell-1}-h_{j, \ell-1}\right)_{k} \geq 0$ and $s_{k}=-1$ otherwise. Then,

$$
\sum_{k \in M_{0}}\left|\Delta_{k}\right| \geq \sum_{k \in M_{0}} s_{k} \cdot \Delta_{k}=\sum_{k \in M_{0}}\left|\left(h_{i, \ell-1}-h_{j, \ell-1}\right)_{k}\right|+\tau \cdot s_{k}\left(W_{\ell}\left(h_{i, \ell-1}-h_{j, \ell-1}\right)\right)_{k}
$$

Note that when $h_{i, \ell-1}$ and $h_{j, \ell-1}$ are fixed, the values $s_{k}\left(W_{\ell}\left(h_{i, \ell-1}-h_{j, \ell-1}\right)\right)_{k}$ are independent Gaussian with mean zero. This means, with probability at least $1-e^{-\Omega\left(\log ^{2} m\right)}$, the summation $\sum_{k \in M_{0}} s_{k}\left(W_{\ell}\left(h_{i, \ell-1}-h_{j, \ell-1}\right)\right)_{k}$ is at most $O(\log m)$ in absolute value. Putting this into the above equation, we have

$$
\sum_{k \in M_{0}}\left|\Delta_{k}\right| \geq \sum_{k \in M_{0}}\left|\left(h_{i, \ell-1}-h_{j, \ell-1}\right)_{k}\right|-O(\tau \log m) \geq \frac{\delta}{20} \sqrt{m}-O(\tau \log m)
$$

- Continuing this process for $\ell=2,3, \ldots, L-1$, we can conclude that $\sum_{k \in M_{0}} \mid\left(h_{i, L-1}-\right.$ $\left.h_{j, L-1}\right)_{k} \left\lvert\, \geq \frac{\delta}{30} \sqrt{m}\right.$ and therefore $\left\|h_{i, L-1}-h_{j, L-1}\right\| \geq \Omega\left(\delta^{2}\right)$. This is the same statement as before that we shall need for the downstream application of Lemma 5.8.5.

Changes to Section 5.9 Lemma 5.9 .1 becomes easy to prove with all the $L$ factors disappear for the following reason. Fixing $i$ and ignoring the subscript in $i$, we have for $\ell=1,2, \ldots, L-1$ :

$$
\begin{aligned}
h_{\ell}^{\prime} & =D_{\ell}^{\prime \prime}\left(\left(I+\tau W_{\ell}+\tau W_{\ell}^{\prime}\right) h_{\ell-1}-\left(I+\tau W_{\ell}\right) \widetilde{h}_{\ell-1}\right) \\
& =D_{\ell}^{\prime \prime}\left(\left(I+\tau W_{\ell}\right) h_{\ell-1}^{\prime}+\tau W_{\ell}^{\prime} h_{\ell-1}\right)
\end{aligned}
$$

For some diagonal matrix $D_{\ell}^{\prime \prime} \in \mathbb{R}^{m \times m}$ with diagonal entries in $[-1,1]$ (see Proposition 5.12.3). By simple spectral norm of matrices bound we have

$$
\left\|h_{\ell}^{\prime}\right\| \leq\left(1+\tau\left\|W_{\ell}\right\|_{2}+\tau\left\|W_{\ell}^{\prime}\right\|_{2}\right)\left\|h_{\ell-1}^{\prime}\right\|+\tau\left\|W_{\ell}^{\prime}\right\|_{2}\left\|\widetilde{h}_{\ell-1}\right\| \leq\left(1+\frac{1}{10 L}\right)\left\|h_{\ell-1}^{\prime}\right\|+O(\tau \omega) \leq \cdots \leq O(\tau \omega)
$$

This implies $\left\|h_{\ell}^{\prime}\right\|,\left\|g_{\ell}^{\prime}\right\| \leq O(\tau \omega)$ for all $\ell \in[L-1]$, and combining with the old proof we have $\left\|h_{L}^{\prime}\right\|,\left\|g_{L}^{\prime}\right\| \leq O(\omega)$.

As for the sparsity $\left\|D_{\ell}^{\prime}\right\|_{0}$, because $\widetilde{g}_{\ell}=\widetilde{h}_{\ell-1}+\tau \widetilde{W}_{\ell} \widetilde{h}_{\ell-1} \sim \mathcal{N}\left(\widetilde{h}_{\ell-1}, \frac{2 \tau^{2}\left\|\widetilde{h}_{\ell-1}\right\|^{2}}{m}\right)$ and $\left\|g_{\ell}^{\prime}\right\| \leq O(\tau \omega)$, applying essentially the same Claim 5.9.2, we have $\left\|D_{\ell}^{\prime}\right\|_{0} \leq O\left(m \omega^{2 / 3}\right)$ for every $\ell=1,2, \ldots, L-1$. One can similarly argue that $\left\|D_{L}^{\prime}\right\|_{0} \leq O\left(m \omega^{2 / 3}\right)$.

Next, Lemma 5.9.5 and Lemma 5.9.6 become trivial to prove (recall we have to change $\widetilde{W}_{\ell}$ with $I+\tau \widetilde{W}_{\ell}$ for $\ell<L$ ) and the $L$ factor also gets improved.

Changes to Section 5.10 The proofs of this section require only notational changes.

Final Theorem Since Section 5.11 and 5.12 rely on previous sections, they do not need to be changed (besides improving polynomial factors in $L$ ). Our final theorem becomes

Theorem 5.18.1 (ResNet). Let $m \geq \widetilde{\Omega}\left(\operatorname{poly}\left(n, L, \delta^{-1}\right) \cdot d\right)$. For the residual neural network defined in this section, with probability at least $1-e^{-\Omega\left(\log ^{2} m\right)}$ over the random initialization, GD needs at most $T=O\left(\frac{n^{6} L^{2}}{\delta^{2}} \log \frac{1}{\epsilon}\right)$ iterations and $S G D$ needs at most $T=$ $O\left(\frac{n^{7} L^{2} \log ^{2} m}{b \delta^{2}} \log \frac{1}{\epsilon}\right)$ iterations to find a point $F(\vec{W}) \leq \epsilon$.






Figure 5.2: ResNet-32 architecture [Yan18] landscape on CIFAR10 vs CIFAR100.





CIFAR-100 dataset, resnet 110 architecture





Figure 5.3: ResNet-110 architecture [Yan18] landscape on CIFAR10 vs CIFAR100.




CIFAR-10 dataset,
vgg19bn architecture









Figure 5.4: VGG19 architecture (with BN) [Yan18] landscape on CIFAR10 vs CIFAR100.


Figure 5.5: ResNet-32 architecture [Yan18] landscape on CIFAR10 vs CIFAR10 (20\% noise), means we have randomly perturbed $20 \%$ of the true labels in the training set.


Figure 5.6: ResNet-110 architecture [Yan18] landscape on CIFAR10 vs CIFAR10 (20\% noise), means we have randomly perturbed $20 \%$ of the true labels in the training set.


Figure 5.7: VGG19 architecture (with BN) [Yan18] landscape on CIFAR10 vs CIFAR10 ( $20 \%$ noise), means we have randomly perturbed $20 \%$ of the true labels in the training set.

## Chapter 6

## Convergence result of Recurrent Neural Network

In this chapter, we focus on recurrent neural networks (RNNs) which are multi-layer networks widely used in natural language processing. They are harder to analyze than feedforward neural networks, because the same recurrent unit is repeatedly applied across the entire time horizon of length $L$, which is analogous to feedforward networks of depth $L$. We show when the number of neurons is sufficiently large, meaning polynomial in the training data size and in $L$, then SGD is capable of minimizing the regression loss in the linear convergence rate. This gives theoretical evidence of how RNNs can memorize data.

More importantly, in this paper we build general toolkits to analyze multi-layer networks with ReLU activations. For instance, we prove why ReLU activations can prevent exponential gradient explosion or vanishing, and build a perturbation theory to analyze first-order approximation of multi-layer networks.

### 6.1 Introduction

Neural networks have been one of the most powerful tools in machine learning over the past a few decades [KSH12, GMH13, $\left.\mathrm{LHP}^{+} 15, \mathrm{AAA}^{+} 16, \mathrm{HZRS}^{2} 6, \mathrm{SHM}^{+} 16, \mathrm{SSS}^{+} 17\right]$. The multi-layer structure of neural network gives it supreme power in expressibility and learning performance. However, it raises complexity concerns: the training objective is generally non-convex and non-smooth. In practice, local-search algorithms such as stochastic gradient descent (SGD) are capable of finding global optima, at least on the training data $\left[\mathrm{ZBH}^{+} 17\right.$, GVS15]. How SGD avoids local minima for such objectives remains an open theoretical question since [GVS15].

In recent years, there have been a number of theoretical results aiming at a better understanding of this phenomenon. Many of them focus on two-layer (thus one-hiddenlayer) neural networks and assume that the inputs are random Gaussian or sufficiently close to Gaussian [BG17, Son17, Tia17, LY17, DLT ${ }^{+}$18, GLM17, PRSZ18, ZSJ ${ }^{+}$17, ZSD17]. Some study deep neural networks but assuming the activation function is linear [HM17, ACGH18, BHL18]. Some study the convex task of training essentially only the last layer of the network [Dal17]. On the technique side, some of these results try to understand the gradient dynamics [Son17, Tia17, BG17, LY17, Dal17, BGMSS18, DLT ${ }^{+}$18, PRSZ18, ZSJ ${ }^{+} 17$, ZSD17, VW18], while others focus on the geometry properties of the training objective [GLM17, SS18, FB17, ZL17, HM17].

More recently, [SS18] provided evidence that, even when inputs are standard Gaussians, two-layer neural networks can indeed have spurious local minima, and suggested that over-parameterization (i.e., increasing the number of neurons) may be the key in avoiding spurious local minima. [LL18] showed that, for two-layer networks with the cross-entropy
loss, in the over-parametrization regime, gradient descent (GD) is capable of finding nearlyglobal optimal solutions on the training data. This result was later extended to the $\ell_{2}$ loss by [DZPS19].

In this paper, we show GD and SGD are capable of training multi-layer neural networks (with ReLU activation) to global minima on any non-degenerate training data set. Furthermore, the running time is polynomial in the number of layers and the number of data points. Since there are many different types of multi-layer networks (convolutional, feedforward, recurrent, etc.), in this present paper, we focus on recurrent neural networks (RNN) as our choice of multi-layer networks, and feedforward networks are only its "special case" (see for instance a follow-up work [AZLS19]).

Recurrent Neural Networks Among different architectures of neural networks, one of the least theoretically-understood structure is the recurrent one [Elm90]. A recurrent neural network recurrently applies the same network unit to a sequence of input tokens, such as a sequence of words in a language sentence. RNN is particularly useful when there are longterm, non-linear interactions between input tokens in the same sequence. These networks are widely used in practice for natural language processing, language generation, machine translation, speech recognition, video and music processing, and many other tasks [MKB ${ }^{+} 10$, $M^{\prime}{ }^{+}$11, SSN12, KB13, SSB14, SVL14, CVMBB14, CGCB14]. On the theory side, while there are some attempts to show that an RNN is more expressive than a feedforward neural network [KNO18], when and how an RNN can be efficiently learned has little theoretical explanation.

In practice, RNN is usually trained by simple local-search algorithms such as SGD.

However, unlike shallow networks, the training process of RNN often runs into the trouble of vanishing or exploding gradient [SMH11]. That is, the value of the gradient becomes exponentially small or large in the time horizon, even when the training objective is still constant. ${ }^{1}$ In practice, one of the popular ways to resolve this is by the long short term memory (LSTM) structure [HS97]. However, one can also use rectified linear units (ReLUs) as activation functions to avoid vanishing or exploding gradient [SBS $\left.{ }^{+} 17\right]$. In fact, one of the earliest adoptions of ReLUs was on applications of RNNs for this purpose twenty years ago [Hah98, SA96].

Since the RNN structure was proposed, a large number of variations have been designed over past decades, including LSTM [HS97], Bidirectional RNN [SP97], Bidirectional LSTM [GS05], gate recurrent unit [CVMG ${ }^{+}$14], statistical recurrent unit [OPS17], Fourier recurrent unit [ZLSD18]. For a more detailed survey, we refer the readers to [SBS $\left.{ }^{+} 17\right]$.

[^30]
### 6.1.1 Our Question

In this paper, we study the following general question

- Can ReLU provably stabilize the training process and avoid vanishing/exploding gradient?
- Can RNN be trained close to zero training error efficiently under mild assumptions?

Remark 6.1.1. When there is no activation function, RNN is known as linear dynamical system. Hardt, Ma and Recht [HMR18] first proved the convergence of finding global minima for such linear dynamical systems. Followups in this line of research include [HSZ17, HLS ${ }^{+}$18].

Motivations One may also want to study whether RNN can be trained close to zero test error. However, unlike feedforward networks, the training error, or the ability to memorize examples, may actually be desirable for RNN. After all, many tasks involving RNN are related to memories, and certain RNN units are even referred to memory cells. Since RNN applies the same network unit to all input tokens in a sequence, the following question can possibly of its own interest:

- How does RNN learn mappings (say from token 3 to token 7) without destroying others?

Another motivation is the following. An RNN can be viewed as a space constraint, differentiable Turing machine, except that the input is only allowed to be read in a fixed order. It was shown in [SS91] that all Turing machines can be simulated by fully-connected recurrent networks built of neurons with non-linear activations. In practice, RNN is also used as a tool to build neural Turing machines [GWD14], equipped with a grand goal of
automatically learning an algorithm based on the observation of the inputs and outputs. To this extent, we believe the task of understanding the trainability as a first step towards understanding RNN can be meaningful on its own.

Our Result To present the simplest result, we focus on the classical Elman network with ReLU activation:

$$
\begin{array}{rr}
h_{\ell}=\phi\left(W \cdot h_{\ell-1}+A x_{\ell}\right) \in \mathbb{R}^{m} & \text { where } W \in \mathbb{R}^{m \times m}, A \in \mathbb{R}^{m \times d_{x}} \\
y_{\ell}=B \cdot h_{\ell} \in \mathbb{R}^{d} & \text { where } B \in \mathbb{R}^{d \times m}
\end{array}
$$

We denote by $\phi$ the ReLU activation function: $\phi(x)=\max (x, 0)$. We note that (fullyconnected) feedforward networks are only "special cases" to this by replacing $W$ with $W_{\ell}$ for each layer. ${ }^{2}$

We consider a regression task where each sequence of inputs consists of vectors $x_{1}, \ldots, x_{L} \in \mathbb{R}^{d_{x}}$ and we perform least-square regression with respect to $y_{1}^{*}, \ldots, y_{L}^{*} \in \mathbb{R}^{d}$. We assume there are $n$ training sequences, each of length $L$. We assume the training sequences are $\delta$-separable (say vectors $x_{1}$ are different by relative distance $\delta>0$ for every pairs of training sequences). Our main theorem can be stated as follows

Theorem 6.1.1. If the number of neurons $m \geq \operatorname{poly}\left(n, d, L, \delta^{-1}, \log \epsilon^{-1}\right)$ is polynomially large, we can find weight matrices $W, A, B$ where the $R N N$ gives $\epsilon$ training error

- if gradient descent (GD) is applied for $T=\Omega\left(\frac{\operatorname{poly}(n, d, L)}{\delta^{2}} \log \frac{1}{\epsilon}\right)$ iterations, starting from random Gaussian initializations; or

[^31]- if (mini-batch or regular) stochastic gradient descent (SGD) is applied for

$$
T=\Omega\left(\frac{\operatorname{poly}(n, d, L)}{\delta^{2}} \log \frac{1}{\epsilon}\right)
$$

iterations, starting from random Gaussian initializations. ${ }^{3}$
(To present the simplest possible result, we have not tried to tighten the polynomial dependency with respect to $n, d$ and $L$. We only tightened the dependency with respect to $\delta$ and $\epsilon$.)

Our Contribution We summarize our contributions as follows.

- We believe this is the first proof of convergence of GD/SGD for training the hidden layers of recurrent neural networks (or even for any multi-layer networks of more than two layers) when activation functions are present. ${ }^{4}$
- Our results provide arguably the first theoretical evidence towards the empirical finding of [GVS15] on multi-layer networks, regarding the ability of SGD to avoid (spurious) local minima. Our theorem does not exclude the existence of bad local minima

[^32]- We build new technical toolkits to analyze multi-layer networks with ReLU activation, which have now found many applications [AZLS19, ALL18c, AL19a, CG19]. For instance, combining this paper with new techniques, one can derive guarantees on testing error for RNN in the PAC-learning language [AL19a].

Extension: Deep RNN Elman RNN is also referred to as three-layer RNN, and one may also study the convergence of RNNs with more hidden layers. This is referred to as deep RNN [SBS $\left.{ }^{+} 17\right]$. Our theorem also applies to deep RNNs (by combining this paper together with [AZLS19]).

Extension: Loss functions For simplicity, in this paper we have adopted the $\ell_{2}$ regression loss. Our results generalize to other Lipschitz smooth (but possibly nonconvex) loss functions, by combining with the techniques of [AZLS19].

### 6.1.2 Other Related Works

Another relevant work is [BGMSS18] where the authors studied over-paramterization in the case of two-layer neural network under a linear-separable assumption.

Instead of using randomly initialized weights like this paper, there is a line of work proposing algorithms using weights generated from some "tensor initialization" process [ABGM14, SA15, JSA15, Tia17, ZSJ ${ }^{+}$17].

There is huge literature on using the mean-field theory to study neural networks [MMN18, YS17, DFS16, XBSD ${ }^{+} 18$, LBN $^{+} 17$, YS18, CPS18, PW17, PB17, PSG17, PLR ${ }^{+} 16$, SGGSD17]. At a high level, they study the network dynamics at random initialization when the number of hidden neurons grow to infinity, and use such initialization theory to predict performance after training. However, they do not provide theoretical convergence rate for the training process (at least when the number of neurons is finite).

### 6.2 Notations and Preliminaries

We denote by $\|\cdot\|_{2}$ (or sometimes $\|\cdot\|$ ) the Euclidean norm of vectors, and by $\|\cdot\|_{2}$ the spectral norm of matrices. We denote by $\|\cdot\|_{\infty}$ the infinite norm of vectors, $\|\cdot\|_{0}$ the sparsity of vectors or diagonal matrices, and $\|\cdot\|_{F}$ the Frobenius norm of matrices. Given matrix $W$, we denote by $W_{k}$ or $w_{k}$ the $k$-th row vector of $W$. We denote the row $\ell_{p}$ norm for $W \in \mathbb{R}^{m \times d}$ as

$$
\begin{equation*}
\|W\|_{2, p} \stackrel{\text { def }}{=}\left(\sum_{i \in[m]}\left\|w_{i}\right\|_{2}^{p}\right)^{1 / p} \tag{6.1}
\end{equation*}
$$

By definition, $\|W\|_{2,2}=\|W\|_{F}$.
We use $\mathcal{N}(\mu, \sigma)$ to denote Gaussian distribution with mean $\mu$ and variance $\sigma$; or $\mathcal{N}(\mu, \Sigma)$ to denote Gaussian vector with mean $\mu$ and covariance $\Sigma$. We use $\mathbb{1}_{\text {event }}$ to denote the indicator function of whether event is true. We denote by $e_{k}$ the $k$-th standard basis vector. We use $\phi(\cdot)$ to denote the $\operatorname{ReLU}$ function, namely $\phi(x)=\max \{x, 0\}=\mathbb{1}_{x \geq 0} \cdot x$. Given univariate function $f: \mathbb{R} \rightarrow \mathbb{R}$, we also use $f$ to denote the same function over vectors: $f(x)=\left(f\left(x_{1}\right), \ldots, f\left(x_{m}\right)\right)$ if $x \in \mathbb{R}^{m}$.

Given vectors $v_{1}, \ldots, v_{n} \in \mathbb{R}^{m}$, we define $U=\mathrm{GS}\left(v_{1}, \ldots, v_{n}\right)$ as their Gram-Schmidt orthonormalization. Namely, $U=\left[\widehat{v}_{1}, \ldots, \widehat{v}_{n}\right] \in \mathbb{R}^{m \times n}$ where

$$
\widehat{v}_{1}=\frac{v_{1}}{\left\|v_{1}\right\|} \quad \text { and } \quad \text { for } i \geq 2: \quad \widehat{v}_{i}=\frac{\prod_{j=1}^{i-1}\left(I-\widehat{v}_{j} \widehat{v}_{j}^{\top}\right) v_{i}}{\left\|\prod_{j=1}^{i-1}\left(I-\widehat{v}_{\gamma_{v}} \widehat{v}_{j}^{\top}\right) v_{i}\right\|} .
$$

Note that in the occasion that $\prod_{j=1}^{i-1}\left(I-\widehat{v}_{j} \widehat{v}_{j}^{\top}\right) v_{i}$ is the zero vector, we let $\widehat{v}_{i}$ be an arbitrary unit vector that is orthogonal to $\widehat{v}_{1}, \ldots, \widehat{v}_{i-1}$.

### 6.2.1 Elman Recurrent Neural Network

We assume $n$ training inputs are given: $\left(x_{i, 1}, x_{i, 2}, \ldots, x_{i, L}\right) \in\left(\mathbb{R}^{d_{x}}\right)^{L}$ for each input $i \in[n]$. We assume $n$ training labels are given: $\left(y_{i, 1}^{*}, y_{i, 2}^{*}, \ldots, y_{i, L}^{*}\right) \in\left(\mathbb{R}^{d}\right)^{L}$ for each input $i \in[n]$. Without loss of generality, we assume $\left\|x_{i, \ell}\right\| \leq 1$ for every $i \in[n]$ and $\ell \in[L]$. Also without loss of generality, we assume $\left\|x_{i, 1}\right\|=1$ and its last coordinate $\left[x_{i, 1}\right]_{d_{x}}=\frac{1}{\sqrt{2}}$ for every $i \in[n] .{ }^{5}$

We make the following assumption on the input data:

Assumption 6.2.1. $\left\|x_{i, 1}-x_{j, 1}\right\| \geq \delta$ for some parameter $\delta \in(0,1]$ and every pair of $i \neq j \in[n]$.

Given weight matrices $W \in \mathbb{R}^{m \times m}, A \in \mathbb{R}^{m \times d_{x}}, B \in \mathbb{R}^{d \times m}$, we introduce the following notations to describe the evaluation of RNN on the input sequences. For each $i \in[n]$ and $j \in[L]:$

$$
\begin{array}{ll}
h_{i, 0}=0 \in \mathbb{R}^{m} & g_{i, \ell}=W \cdot h_{i, \ell-1}+A x_{i, \ell} \in \mathbb{R}^{m} \\
y_{i, \ell}=B \cdot h_{i, \ell} \in \mathbb{R}^{d} & h_{i, \ell}=\phi\left(W \cdot h_{i, \ell-1}+A x_{i, \ell}\right) \in \mathbb{R}^{m}
\end{array}
$$

A very important notion that this entire paper relies on is the following:
Definition 6.2.1. For each $i \in[n]$ and $\ell \in[L]$, let $D_{i, \ell} \in \mathbb{R}^{m \times m}$ be the diagonal matrix where

$$
\left(D_{i, \ell}\right)_{k, k}=\mathbf{1}_{\left(W \cdot h_{i, \ell-1}+A x_{i, \ell}\right)_{k} \geq 0}=\mathbf{1}_{\left(g_{i, \ell}\right)_{k} \geq 0}
$$

As a result, we can write $h_{i, \ell}=D_{i, \ell} W h_{i, \ell-1}$.

[^33]We consider the following random initialization distributions for $W, A$ and $B$.

Definition 6.2.2. We say that $W, A, B$ are at random initialization, if the entries of $W$ and $A$ are i.i.d. generated from $\mathcal{N}\left(0, \frac{2}{m}\right)$, and the entries of $B_{i, j}$ are i.i.d. generated from $\mathcal{N}\left(0, \frac{1}{d}\right)$.

Throughout this paper, for notational simplicity, we refer to index $\ell$ as the $\ell$-th layer of RNN, and $h_{i, \ell}, x_{i, \ell}, y_{i, \ell}$ respectively as the hidden neurons, input, output on the $\ell$-th layer. We acknowledge that in certain literatures, one may regard Elman network as a three-layer RNN.

Assumption 6.2.2. We assume $m \geq \operatorname{poly}\left(n, d, L, \frac{1}{\delta}, \log \frac{1}{\epsilon}\right)$ for some sufficiently large polynomial.

Without loss of generality, we assume $\delta \leq \frac{1}{C L^{2} \log ^{3} m}$ for some sufficiently large constant $C$ (if this is not satisfied one can decrease $\delta$ ). Throughout the paper except the detailed appendix, we use $\widetilde{O}, \widetilde{\Omega}$ and $\widetilde{\Theta}$ notions to hide polylogarithmic dependency in $m$. To simplify notations, we denote by

$$
\rho \stackrel{\text { def }}{=} n L d \log m \quad \text { and } \quad \varrho \stackrel{\text { def }}{=} n L d \delta^{-1} \log (m / \epsilon)
$$

### 6.2.2 Objective and Gradient

For simplicity, we only optimize over the weight matrix $W \in \mathbb{R}^{m \times m}$ and let $A$ and $B$ be at random initialization. As a result, our $\ell_{2}$-regression objective is a function over $W$ : ${ }^{6}$

$$
f(W) \stackrel{\text { def }}{=} \sum_{i=1}^{n} f_{i}(W) \quad \text { and } \quad f_{i}(W) \stackrel{\text { def }}{=} \frac{1}{2} \sum_{\ell=2}^{L}\left\|\operatorname{loss}_{i, \ell}\right\|_{2}^{2} \quad \text { where } \quad \operatorname{loss}_{i, \ell} \stackrel{\text { def }}{=} B h_{i, \ell}-y_{i, \ell}^{*} .
$$

Using chain rule, one can write down a closed form of the (sub-)gradient:

Fact 6.2.3. For $k \in[m]$, the gradient with respect to $W_{k}$ (denoted by $\nabla_{k}$ ) and the full gradient are

$$
\begin{aligned}
\nabla_{k} f(W) & =\sum_{i=1}^{n} \sum_{a=2}^{L} \sum_{\ell=1}^{a-1}\left(\operatorname{Back}_{i, \ell+1 \rightarrow a}^{\top} \cdot \operatorname{loss}_{i, a}\right)_{k} \cdot h_{i, \ell} \cdot \mathbf{1}_{\left\langle W_{k}, h_{i, \ell}\right\rangle+\left\langle A_{k}, x_{i, \ell+1}\right\rangle \geq 0} \\
\nabla f(W) & =\sum_{i=1}^{n} \sum_{a=2}^{L} \sum_{\ell=1}^{a-1} D_{i, \ell+1}\left(\operatorname{Back}_{i, \ell+1 \rightarrow a}^{\top} \cdot \operatorname{loss}_{i, a}\right) \cdot h_{i, \ell}^{\top}
\end{aligned}
$$

where for every $i \in[n], \ell \in[L]$, and $a=\ell+1, \ell+2, \ldots, L$ :

$$
\text { Back }_{i, \ell \rightarrow \ell} \stackrel{\text { def }}{=} B \in \mathbb{R}^{d \times m} \quad \text { and } \quad \text { Back }_{i, \ell \rightarrow a} \stackrel{\text { def }}{=} B D_{i, a} W \cdots D_{i, \ell+1} W \in \mathbb{R}^{d \times m}
$$

[^34]
### 6.3 Our Results

Our main results can be formally stated as follows.

Theorem 6.3.1 (GD). Suppose $\eta=\widetilde{\Theta}\left(\frac{\delta}{m} \operatorname{poly}(n, d, L)\right)$ and $m \geq \operatorname{poly}\left(n, d, L, \delta^{-1}, \log \epsilon^{-1}\right)$.
Let $W^{(0)}, A, B$ be at random initialization. With high probability over the randomness of $W^{(0)}, A, B$, if we apply gradient descent for $T$ steps $W^{(t+1)}=W^{(t)}-\eta \nabla f\left(W^{(t)}\right)$, then it satisfies

$$
f\left(W^{(T)}\right) \leq \epsilon \quad \text { for } \quad T=\widetilde{\Omega}\left(\frac{\operatorname{poly}(n, d, L)}{\delta^{2}} \log \frac{1}{\epsilon}\right) .
$$

Theorem 6.3.2 (SGD). Suppose $\eta=\widetilde{\Theta}\left(\frac{\delta}{m} \operatorname{poly}(n, d, L)\right)$ and $m \geq \operatorname{poly}\left(n, d, L, \delta^{-1}, \log \epsilon^{-1}\right)$. Let $W^{(0)}, A, B$ be at random initialization. If we apply stochastic gradient descent for $T$ steps $W^{(t+1)}=W^{(t)}-\eta \nabla f_{i}\left(W^{(t)}\right)$ for a random index $i \in[n]$ per step, then with high probability (over $W^{(0)}, A, B$ and the randomness of $S G D$ ), it satisfies

$$
f\left(W^{(T)}\right) \leq \epsilon \quad \text { for } \quad T=\widetilde{\Omega}\left(\frac{\operatorname{poly}(n, d, L)}{\delta^{2}} \log \frac{1}{\epsilon}\right)
$$

In both cases, we essentially have linear convergence rates. ${ }^{7}$ Notably, our results show that the dependency of the number of layers $L$, is polynomial. Thus, even when RNN is applied to sequences of long input data, it does not suffer from exponential gradient explosion or vanishing (e.g., $2^{\Omega(L)}$ or $\left.2^{-\Omega(L)}\right)$ through the entire training process.

Main Technical Theorems Our main Theorem 6.17.1 and Theorem 6.18.1 are in fact natural consequences of the following two technical theorems. They both talk about the

[^35]first-order behavior of RNNs when the weight matrix $W$ is sufficiently close to some random initialization.

The first theorem is similar to the classical Polyak-Łojasiewicz condition [Pol63, Loj63], and says that $\|\nabla f(W)\|_{F}^{2}$ is at least as large as the objective value.

Theorem 6.3.3. With high probability over random initialization $\widetilde{W}, A, B$, it satisfies

$$
\begin{aligned}
& \forall W \in \mathbb{R}^{m \times m} \text { with }\|W-\widetilde{W}\|_{2} \leq \frac{\operatorname{poly}(\varrho)}{\sqrt{m}}: \\
& \|\nabla f(W)\|_{F}^{2} \geq \frac{\delta}{\operatorname{poly}(\rho)} \times m \times f(W) \\
& \|\nabla f(W)\|_{F}^{2},\left\|\nabla f_{i}(W)\right\|_{F}^{2} \leq \operatorname{poly}(\rho) \times m \times f(W) .
\end{aligned}
$$

The second theorem shows a special "semi-smoothness" property of the objective.

Theorem 6.3.4. With high probability over random initialization $\widetilde{W}, A, B$, it satisfies for every $\breve{W} \in \mathbb{R}^{m \times m}$ with $\|\breve{W}-\widetilde{W}\| \leq \frac{\text { poly }(e)}{\sqrt{m}}$, and for every $W^{\prime} \in \mathbb{R}^{m \times m}$ with $\left\|W^{\prime}\right\| \leq \frac{\tau_{0}}{\sqrt{m}}$,
$f\left(\breve{W}+W^{\prime}\right) \leq f(\breve{W})+\left\langle\nabla f(\breve{W}), W^{\prime}\right\rangle+\operatorname{poly}(\varrho) m^{1 / 3} \cdot \sqrt{f(W)} \cdot\left\|W^{\prime}\right\|_{2}+\operatorname{poly}(\rho) m\left\|W^{\prime}\right\|_{2}^{2}$.

At a high level, the convergence of GD and SGD are careful applications of the two technical theorems above: indeed, Theorem 6.15.2 shows that as long as the objective value is high, the gradient is large; and Theorem 6.16 .1 shows that if one moves in the (negative) gradient direction, then the objective value can be sufficiently decreased. These two technical theorems together ensure that GD/SGD does not hit any saddle point or (bad) local minima along its training trajectory. This was practically observed by [GVS15] and a theoretical justification was open since then.

An Open Question We did not try to tighten the polynomial dependencies of ( $n, d, L$ ) in the proofs. When $m$ is sufficiently large, we make use of the randomness at initialization to argue that, for all the points within a certain radius from initialization, for instance Theorem 6.15.2 holds. In practice, however, the SGD can create additional randomness as time goes; also, in practice, it suffices for those points on the SGD trajectory to satisfy Theorem 6.15.2. Unfortunately, such randomness can - in principle - be correlated with the SGD trajectory, so we do not know how to use that in the proofs. Analyzing such correlated randomness is certainly beyond the scope of this paper, but can possibly explain why in practice, the size of $m$ needed is not that large.

### 6.3.1 Conclusion

Overall, we provide the first proof of convergence of GD/SGD for non-linear neural networks that have more two layers. We show with overparameterization GD/SGD can avoid hitting any (bad) local minima along its training trajectory. This was practically observed by [GVS15] and a theoretical justification was open since then. We present our result using recurrent neural networks (as opposed to the simpler feedforward networks [AZLS19]) in this very first paper, because memorization in RNN could be of independent interest. Also, our result proves that RNN can learn mappings from different input tokens to different output tokens simultaneously using the same recurrent unit.

Last but not least, we build new tools to analyze multi-layer networks with ReLU activations that could facilitate many new research on deep learning. For instance, our techniques in Section 6.5 provide a general theory for why ReLU activations avoid exponential exploding (see e.g. (6.2), (6.5)) or exponential vanishing (see e.g. (6.2), (6.4)); and our techniques in Section 6.6 give a general theory for the stability of multi-layer networks against adversarial weight perturbations, which is at the heart of showing the semi-smoothness Theorem 6.16.1, and used by all the follow-up works [AZLS19, ALL18c, AL19a, CG19].

### 6.4 Proof Sketch

The main difficulty of this paper is to prove Theorem 6.15.2 and 6.16.1, and we shall sketch the proof ideas in Section 6.5 through 6.8. In such high-level discussions, we shall put our emphasize on

- how to avoid exponential blow up in $L$, and
- how to deal with the issue of randomness dependence across layers.

We genuinely hope that this high-level sketch can (1) give readers a clear overview of the proof without the necessity of going to the appendix, and (2) appreciate our proof and understand why it is necessarily long. ${ }^{8}$

[^36]
### 6.5 Basic Properties at Random Initialization

In this section we derive basic properties of the RNN when the weight matrices $W, A, B$ are all at random initialization. The corresponding precise statements and proofs are in Section 6.11.

The first one says that the forward propagation neither explodes or vanishes, that is,

$$
\begin{equation*}
\frac{1}{2} \leq\left\|h_{i, \ell}\right\|_{2},\left\|g_{i, \ell}\right\|_{2} \leq O(L) \tag{6.2}
\end{equation*}
$$

Intuitively, (6.2) very reasonable. Since the weight matrix $W$ is randomly initialized with entries i.i.d. from $\mathcal{N}\left(0, \frac{2}{m}\right)$, the norm $\|W z\|_{2}$ is around $\sqrt{2}$ for any fixed vector $z$. Equipped with ReLU activation, it "shuts down" roughly half of the coordinates of $W z$ and reduces the norm $\|\phi(W z)\|$ to one. Since in each layer $\ell$, there is an additional unit-norm signal $x_{i, \ell}$ coming in, we should expect the final norm of hidden neurons to be at most $O(L)$.

Unfortunately, the above argument cannot be directly applied since the weight matrix $W$ is reused for $L$ times so there is no fresh new randomness across layers. Let us explain how we deal with this issue carefully, because it is at the heart of all of our proofs in this paper. Recall, each time $W$ is applied to some vector $h_{i, \ell}$, it only uses "one column of randomness" of $W$. Mathematically, letting $U_{\ell} \in \mathbb{R}^{m \times n \ell}$ denote the column orthonormal matrix using Gram-Schmidt

$$
U_{\ell} \stackrel{\text { def }}{=} \mathrm{GS}\left(h_{1,1}, \ldots, h_{n, 1}, h_{1,2}, \ldots, h_{n, 2}, \ldots, h_{1, \ell}, \ldots, h_{n, \ell}\right),
$$

we have $W h_{i, \ell}=W U_{\ell-1} U_{\ell-1}^{\top} h_{i, \ell}+W\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}$.

- The second term $W\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}$ has new randomness independent of the previous layers. ${ }^{9}$
- The first term $W U_{\ell-1} U_{\ell-1}^{\top} h_{i, \ell}$ relies on the randomness of $W$ in the directions of $h_{i, a}$ for $a<\ell$ of the previous layers. We cannot rely on the randomness of this term, because when applying inductive argument till layer $\ell$, the randomness of $W U_{\ell-1}$ is already used.

Fortunately, $W U_{\ell-1} \in \mathbb{R}^{m \times n(\ell-1)}$ is a rectangular matrix with $m \gg n(\ell-1)$ (thanks to overparameterization!) so one can bound its spectral norm by roughly $\sqrt{2}$. This ensures that no matter how $h_{i, \ell}$ behaves (even arbitrarily correlated with $W U_{\ell-1}$ ), the norm of the first term cannot be too large. It is crucial here that $W U_{\ell-1}$ is a rectangular matrix, because for a square random matrix such as $W$, its spectral norm is 2 and using that, the forward propagation bound will exponentially blow up.

This summarizes the main idea for proving $\left\|h_{i, \ell}\right\| \leq O(L)$ in (6.2); the lower bound $\frac{1}{2}$ is similar. Our next property says in each layer, the amount of "fresh new randomness" is non-negligible:

$$
\begin{equation*}
\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}\right\|_{2} \geq \widetilde{\Omega}\left(\frac{1}{L^{2}}\right) \tag{6.3}
\end{equation*}
$$

This relies on a more involved inductive argument than (6.2). At high level, one needs to show that in each layer, the amount of "fresh new randomness" reduces only by a factor at most $1-\frac{1}{10 L}$.

[^37]Using (6.2) and (6.3), we obtain the following property about the data separability:

$$
\begin{equation*}
\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{i, \ell+1} \text { and }\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{j, \ell+1} \text { are }(\delta / 2) \text {-separable, } \forall i, j \in[n] \text { with } i \neq j \tag{6.4}
\end{equation*}
$$

Here, we say two vectors $x$ and $y$ are $\delta$-separable if $\left\|\left(I-y y^{\top} /\|y\|_{2}^{2}\right) x\right\| \geq \delta$ and vice versa. Property (6.4) shows that the separability information (say on input token 1) does not diminish by more than a polynomial factor even if the information is propagated for $L$ layers.

We prove (6.4) by induction. In the first layer $\ell=1$ we have $h_{i, 1}$ and $h_{j, 1}$ are $\delta$-separable which is a consequence of Assumption 6.2.1. If having fresh new randomness, given two $\delta$ separable vectors $x, y$, one can show that $\phi(W x)$ and $\phi(W y)$ are also $\delta\left(1-o\left(\frac{1}{L}\right)\right)$ separable. Again, in RNN, we do not have fresh new randomness, so we rely on (6.3) to give us reasonably large fresh new randomness. Applying a careful induction helps us to derive that (6.4) holds for all layers. ${ }^{10}$

Intermediate Layers and Backward Propagation Training neural network is not only about forward propagation. We also have to bound intermediate layers and backward propagation.

The first two results we derive are the following. For every $\ell_{1} \geq \ell_{2}$ and diagonal

[^38]matrices $D^{\prime}$ of sparsity $s \in\left[\rho^{2}, m^{0.49}\right]$ :
\[

$$
\begin{align*}
\left\|W D_{i, \ell_{1}} \cdots D_{i, \ell_{2}} W\right\|_{2} & \leq O\left(L^{3}\right)  \tag{6.5}\\
\left\|D^{\prime} W D_{i, \ell_{1}} \cdots D_{i, \ell_{2}} W D^{\prime}\right\|_{2} & \leq \widetilde{O}(\sqrt{s} / \sqrt{m}) \tag{6.6}
\end{align*}
$$
\]

Intuitively, one cannot use spectral bound argument to derive (6.5) or (6.6): the spectral norm of $W$ is 2 , and even if ReLU activations cancel half of its mass, the spectral norm $\|D W\|_{2}$ remains to be $\sqrt{2}$. When stacked together, this grows exponential in $L$.

Instead, we use an analogous argument to (6.2) to show that, for each fixed vector $z$, the norm of $\left\|W D_{i, \ell_{1}} \cdots D_{i, \ell_{2}} W z\right\|_{2}$ is at most $O(1)$ with extremely high probability $1-$ $e^{-\Omega\left(m / L^{2}\right)}$. By standard $\epsilon$-net argument, $\left\|W D_{i, \ell_{1}} \cdots D_{i, \ell_{2}} W z\right\|_{2}$ is at most $O(1)$ for all $\frac{m}{L^{3}}$ sparse vectors $z$. Finally, for a possible dense vector $z$, we can divide it into $L^{3}$ chunks each of sparsity $\frac{m}{L^{3}}$. Finally, we apply the upper bound for $L^{3}$ times. This proves (6.5). One can use similar argument to prove (6.6).

Remark 6.5.1. We did not try to tighten the polynomial factor here in $L$. We conjecture that proving an $O(1)$ bound may be possible, but that question itself may be a sufficiently interesting random matrix theory problem on its own.

The next result is for back propagation. For every $\ell_{1} \geq \ell_{2}$ and diagonal matrices $D^{\prime}$ of sparsity $s \in\left[\rho^{2}, m^{0.49}\right]$ :

$$
\begin{equation*}
\left\|B D_{i, \ell_{1}} \cdots D_{i, \ell_{2}} W D^{\prime}\right\|_{2} \leq \widetilde{O}(\sqrt{s}) \tag{6.7}
\end{equation*}
$$

Its proof is in the same spirit as (6.6), with the only difference being the spectral norm of $B$ is around $\sqrt{m / d}$ as opposed to $O(1)$.

### 6.6 Stability After Adversarial Perturbation

In this section we study the behavior of RNN after adversarial perturbation. The corresponding precise statements and proofs are in Section 6.12.

Letting $\widetilde{W}, A, B$ be at random initialization, we consider some matrix $W=\widetilde{W}+W^{\prime}$ for $\left\|W^{\prime}\right\|_{2} \leq \frac{\operatorname{poly}(\rho)}{\sqrt{m}}$. Here, $W^{\prime}$ may depend on the randomness of $\widetilde{W}, A$ and $B$, so we say it can be adversarially chosen. The results of this section will later be applied essentially twice:

- Once for those updates generated by GD or SGD, where $W^{\prime}$ is how much the algorithm has moved away from the random initialization.
- The other time (see Section 6.7.3) for a technique that we call "randomness decomposition" where we decompose the true random initialization $W$ into $W=\widetilde{W}+W^{\prime}$, where $\widetilde{W}$ is a "fake" random initialization but identically distributed as $W$. Such technique at least traces back to smooth analysis [ST04].

To illustrate our high-level idea, from this section on (so in Section 6.6, 6.7 and 6.8) we ignore the polynomial dependency in $\varrho$ and hide it in the big-O notion.

We denote by $\widetilde{D}_{i, \ell}, \widetilde{g}_{i, \ell}, \widetilde{h}_{i, \ell}$ respectively the values of $D_{i, \ell}, g_{i, \ell}$ and $h_{i, \ell}$ determined by $\widetilde{W}$ and $A$ at random initialization; and by $D_{i, \ell}=\widetilde{D}_{i, \ell}+D_{i, \ell}^{\prime}, g_{i, \ell}=\widetilde{g}_{i, \ell}+g_{i, \ell}^{\prime}$ and $h_{i, \ell}=\widetilde{h}_{i, \ell}+h_{i, \ell}^{\prime}$ respectively those determined by $W=\widetilde{W}+W^{\prime}$ after the adversarial perturbation.

Forward Stability Our first, and most technical result is the following:

$$
\begin{equation*}
\left\|g_{i, \ell}^{\prime}\right\|_{2},\left\|h_{i, \ell}^{\prime}\right\|_{2} \leq O\left(m^{-1 / 2}\right), \quad\left\|D_{i, \ell}^{\prime}\right\|_{0} \leq O\left(m^{2 / 3}\right) \quad \text { and } \quad\left\|D_{i, \ell}^{\prime} g_{i, \ell}\right\|_{2} \leq O\left(m^{-1 / 2}\right) \tag{6.8}
\end{equation*}
$$

Intuitively, one may hope to prove (6.8) by induction, because we have (ignoring subscripts in $i$ )

$$
g_{\ell^{\prime}}^{\prime}=\underbrace{W^{\prime} D_{\ell^{\prime}-1} g_{\ell^{\prime}-1}}_{(1)}+\underbrace{\widetilde{W} D_{\ell^{\prime}-1}^{\prime} g_{\ell^{\prime}-1}}_{(2}+\underbrace{\widetilde{W} \widetilde{D}_{\ell^{\prime}-1} g_{\ell^{\prime}-1}^{\prime}}_{(3}
$$

The main issue here is that, the spectral norm of $\widetilde{W} \widetilde{D}_{\ell^{\prime}-1}$ in (3) is greater than 1, so we cannot apply naive induction due to exponential blow up in $L$. Neither can we apply techniques from Section 6.5, because the changes such as $g_{\ell^{\prime}-1}$ can be adversarial.

In our actual proof of (6.8), instead of applying induction on (3), we recursively expand (3) by the above formula. This results in a total of $L$ terms of (1) type and $L$ terms of (2) type. The main difficulty is to bound a term of (2) type, that is:

$$
\left\|\widetilde{W} \widetilde{D}_{\ell_{1}} \cdots \widetilde{D}_{\ell_{2}+1} \widetilde{W} D_{\ell_{2}}^{\prime} g_{\ell_{2}}\right\|_{2}
$$

Our argument consists of two conceptual steps.

1. Suppose $g_{\ell_{2}}=\widetilde{g}_{\ell_{2}}+g_{\ell_{2}}^{\prime}=\widetilde{g}_{\ell_{2}}+g_{\ell_{2}, 1}^{\prime}+g_{\ell_{2}, 2}^{\prime}$ where $\left\|g_{\ell_{2}, 1}^{\prime}\right\|_{2} \leq m^{-1 / 2}$ and $\left\|g_{\ell_{2}, 2}^{\prime}\right\|_{\infty} \leq m^{-1}$, then we argue that $\left\|D_{\ell_{2}}^{\prime} g_{\ell_{2}}\right\|_{2} \leq O\left(m^{-1 / 2}\right)$ and $\left\|D_{\ell_{2}}^{\prime} g_{\ell_{2}}\right\|_{0} \leq O\left(m^{2 / 3}\right)$.
2. Suppose $x \in \mathbb{R}^{m}$ with $\|x\|_{2} \leq m^{-1 / 2}$ and $\|x\|_{0} \leq m^{2 / 3}$, then we show that $y=$ $\widetilde{W} \widetilde{D}_{\ell_{1}} \cdots \widetilde{D}_{\ell_{2}+1} \widetilde{W} x$ can be written as $y=y_{1}+y_{2}$ with $\left\|y_{1}\right\|_{2} \leq O\left(m^{-2 / 3}\right)$ and $\left\|y_{2}\right\|_{\infty} \leq$ $O\left(m^{-1}\right)$.

The two steps above enable us to perform induction without exponential blow up. Indeed, they together enable us to go through the following logic chain:

$$
\left.\begin{array}{l}
\|\cdot\|_{2} \leq m^{-1 / 2} \text { and }\|\cdot\|_{\infty} \leq m^{-1} \stackrel{(1)}{\Longrightarrow} \\
\|\cdot\|_{2} \leq m^{-2 / 3} \text { and }\|\cdot\|_{\infty} \leq m^{-1} \\
\Longleftarrow
\end{array}\right\}\|\cdot\|_{2} \leq m^{-1 / 2} \text { and }\|\cdot\|_{0} \leq m^{2 / 3}
$$

Since there is a gap between $m^{-1 / 2}$ and $m^{-2 / 3}$, we can make sure that all blow-up factors are absorbed into this gap, using the property that $m$ is polynomially large. This enables us to perform induction to prove (6.8) without exponential blow-up.

Intermediate Layers and Backward Stability Using (6.8), and especially using the sparsity $\left\|D^{\prime}\right\|_{0} \leq m^{2 / 3}$ from (6.8), one can apply the results in Section 6.5 to derive the following stability bounds for intermediate layers and backward propagation:

$$
\begin{gather*}
\left\|D_{i, \ell_{1}} W \cdots D_{i, \ell_{2}} W-\widetilde{D}_{i, \ell_{1}} \widetilde{W} \cdots \widetilde{D}_{i, \ell_{2}} \widetilde{W}\right\|_{2} \leq O\left(L^{7}\right)  \tag{6.9}\\
\left\|B D_{i, \ell_{1}} W \cdots D_{i, \ell_{2}} W-B \widetilde{D}_{i, \ell_{1}} \widetilde{W} \cdots \widetilde{D}_{i, \ell_{2}} \widetilde{W}\right\|_{2} \leq O\left(m^{1 / 3}\right) . \tag{6.10}
\end{gather*}
$$

Special Rank-1 Perturbation For technical reasons, we also need two bounds in the special case of $W^{\prime}=y z^{\top}$ for some unit vector $z$ and sparse $y$ with $\|y\|_{0} \leq \operatorname{poly}(\varrho)$. We prove that, for this type of rank-one adversarial perturbation, it satisfies for every $k \in[m]$ :

$$
\begin{align*}
\left|\left(\left(\widetilde{W}+W^{\prime}\right) h_{i, \ell}^{\prime}\right)_{k}\right| & \leq O\left(m^{-2 / 3}\right)  \tag{6.11}\\
\left\|B D_{i, \ell_{1}} W \cdots D_{i, \ell_{2}} W e_{k}-B \widetilde{D}_{i, \ell_{1}} \widetilde{W} \cdots \widetilde{D}_{i, \ell_{2}} \widetilde{W} e_{k}\right\|_{2} & \leq O\left(m^{-1 / 6}\right) \tag{6.12}
\end{align*}
$$

### 6.7 Proof Sketch of Theorem 6.15.2: Polyak-Łojasiewicz Condition

The upper bound in Theorem 6.15.2 is easy to prove (based on Section 6.5 and 6.6 ), but the lower bound (a.k.a. the Polyak-Łojasiewicz condition) is the most technically involved result to prove in this paper. We introduce the notion of "fake gradient". Given fixed vectors $\left\{\operatorname{loss}_{i, a}\right\}_{i \in[n], a \in\{2, \ldots, L\}}$, we define

$$
\begin{equation*}
\widehat{\nabla}_{k} f(W) \stackrel{\text { def }}{=} \sum_{i=1}^{n} \sum_{\ell=1}^{L-1}\left(u_{i, \ell}\right)_{k} \cdot h_{i, \ell} \cdot \mathbf{1}_{\left(g_{i, \ell+1}\right)_{k} \geq 0} \tag{6.13}
\end{equation*}
$$

where $u_{i, \ell} \stackrel{\text { def }}{=} \sum_{a=\ell+1}^{L} \operatorname{Back}_{i, \ell+1 \rightarrow a}^{\top} \cdot \operatorname{loss}_{i, a}$. Note that if $\operatorname{loss}_{i, a}=B h_{i, a}-y_{i, a}^{*}$ is the true loss vector, then $\widehat{\nabla}_{k} f(W)$ will be identical to $\nabla_{k} f(W)$ by Fact 6.2.3. Our main technical theorem is the following:

Theorem 6.7.1. For every fixed vectors $\left\{\operatorname{loss}_{i, a}\right\}_{i \in[n], a \in\{2, \ldots, L\}}$, if $W, A, B$ are at random initialization, then with high probability

$$
\|\widehat{\nabla} f(W)\|_{F}^{2} \geq \widetilde{\Omega}\left(\frac{\delta m}{\operatorname{poly}(\rho)}\right) \times \max _{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|^{2}\right\}
$$

There are only two conceptually simple steps from Theorem 6.14.2 to Theorem 6.15.2 (see Section 6.15).

- First, one can use the stability lemmas in Section 6.6 to show that, the fake gradient $\left\|\widehat{\nabla} f\left(W+W^{\prime}\right)\right\|_{F}$ after adversarial perturbation $W^{\prime}\left(\right.$ with $\left.\left\|W^{\prime}\right\|_{2} \leq \frac{1}{\sqrt{m}}\right)$ is also large.
- Second, one can apply $\epsilon$-net and union bound to turn "fixed loss" into "for all loss". This allows us to turn the lower bound on the fake gradient into a lower bound on the true gradient $\left\|\nabla_{k} f\left(W+W^{\prime}\right)\right\|_{F}$.

Therefore, in the rest of this section, we only sketch the ideas behind proving Theorem 6.14.2.
Let $\left(i^{*}, \ell^{*}\right)=\operatorname{argmax}_{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|_{2}\right\}$ be the sample and layer corresponding to the largest loss. Recall $W, A, B$ are at random initialization.

### 6.7.1 Indicator and Backward Coordinate Bounds

There are three factors in the notion of fake gradient (6.13): $\left(u_{i, \ell}\right)_{k} \in \mathbb{R}$ the backward coordinate, $h_{i, \ell} \in \mathbb{R}^{m}$ the forward vector, and $\mathbf{1}_{\left(g_{i, \ell+1}\right)_{k} \geq 0} \in\{0,1\}$ the indicator coordinate. We already know very well how the forward vector $h_{i, \ell}$ behaves from the previous sections. Let us provide bounds on the other two factors at the random initialization. (Details in Section 6.13.)

Our "backward coordinate bound" controls the value of $\left(u_{i, \ell}\right)_{k}$ : at random initialization,

$$
\begin{equation*}
\left|\left(u_{i^{*}, \ell^{*}}\right)_{k}\right| \geq \frac{\left\|\operatorname{loss}_{i^{*}, e^{*}}\right\|_{2}}{\rho} \quad \text { for at least } 1-o(1) \text { fraction of } k \in[m] \tag{6.14}
\end{equation*}
$$

The main idea behind proving (6.14) is to use the randomness of $B$. For a fixed $k \in[m]$, it is in fact not hard to show that $\left|\left(u_{i^{*}, \ell^{*}}\right)_{k}\right|$ is large with high probability. Unfortunately, the randomness of $B$ are shared for different coordinates $k$. We need to also bound the correlation between pairs of coordinates $k_{1}, k_{2} \in[m]$, and resort to MiDiarmid inequality to provide a high concentration bound with respect to all the coordinates.

Our indicator coordinate bound controls the value $\left(g_{i, \ell+1}\right)_{k}$ inside the indicator functions $\mathbf{1}_{\left(g_{i, \ell+1}\right)_{k} \geq 0}$. It says, letting $\beta_{+} \stackrel{\text { def }}{=} \frac{\delta}{\rho^{2}}$ and $\beta_{-} \stackrel{\text { def }}{=} \frac{\delta}{\rho^{10}}$, then at random initialization, for at least $\frac{\delta}{\operatorname{poly}(\rho)}$ fraction of the coordinates $k \in[m]$,

$$
\left|\left(g_{i^{*}, \ell^{*}+1}\right)_{k}\right| \leq \frac{\beta_{-}}{\sqrt{m}} \quad \text { and } \quad\left|\left(g_{i, \ell+1}\right)_{k}\right| \geq \frac{\beta_{+}}{\sqrt{m}} \quad \begin{align*}
& \text { for } i \neq i^{*} \text { and } \ell=\ell^{*}, \text { or }  \tag{6.15}\\
& \text { for } i \in[n] \text { and } \ell>\ell^{*}
\end{align*}
$$

This should be quite intuitive to prove, in the following two steps.

- First, there are $\frac{\delta}{\operatorname{poly}(\rho)} m$ coordinates $k \in[m]$ with $\left|\left(g_{i^{*}, e^{*}+1}\right)_{k}\right| \leq \frac{\beta_{-}}{\sqrt{m}}$.

To show this, we write $\left(g_{i^{*}, \ell^{*}+1}\right)_{k}=\left(W U_{\ell^{*}} U_{\ell^{*}} h_{i^{*}, \ell^{*}}+A x_{i^{*}, \ell^{*}+1}\right)_{k}$, and prove that for every $z \in \mathbb{R}^{n \ell^{*}}$ with bounded norms (by $\epsilon$-net), there are at least $\frac{\delta}{\operatorname{poly}(\rho)} m$ coordinates $k \in[m]$ with $\left|\left(W U_{\ell^{*}} z+A x_{i^{*}, \ell^{*}+1}\right)_{k}\right| \leq \frac{\beta_{-}}{\sqrt{m}}$. This is possible using the independence between $W U_{\ell^{*}}$ and $A$.

- Then, conditioning on the first event happens, we look at $\left|\left(g_{i, \ell+1}\right)_{k}\right|$ for (1) each $i \neq i^{*}$ and $\ell=\ell^{*}$, or (2) each $i \in[n]$ and $\ell>\ell^{*}$. In both cases, even though the value of $g_{i^{*}, \ell^{*}+1}$ is fixed, we still have sufficient fresh new randomness (by invoking (6.4) for case (1) and (6.3) for case (2)). Such additional randomness can make sure that, with high probability (over the fresh new randomness), the value of $\left(g_{i, \ell+1}\right)_{k}$ is larger than $\frac{\beta_{+}}{\sqrt{m}}$.


### 6.7.2 Thought Experiment: Adding Small Rank-One Perturbation

We now focus on a fixed coordinate $k \in[m]$ satisfying (6.15) and (6.14) in the Section 6.7.1, and denote by $v_{i^{*}, \ell^{*}} \xlongequal{\text { def }} \frac{\left(I-U_{\ell^{*}-1} U_{\ell^{*}-1}^{\top}\right) h_{i^{*}, \ell^{*}}}{\|\left(I-U_{\ell^{*}-1} U_{\ell^{*}-1}^{\top} h_{i^{*}}, \ell^{*} \|\right.}$.

For analysis purpose, imagine that we apply a small random perturbation $W_{k}^{\prime}$ to the already-randomly initialized matrix $W$, in the rank-one direction $e_{k} v_{i^{*}, \ell^{*}}^{\top}$. Namely, we set $W_{k}^{\prime}=g \cdot e_{k} v_{i^{*}, \ell^{*}}^{\top}$ where $g \sim \mathcal{N}\left(0, \theta^{2}\right)$ and $\theta$ is a parameter satisfying $\frac{\beta_{-}}{\sqrt{m}} \ll \theta \ll \frac{\beta_{+}}{\sqrt{m}}$. Using the fact that $k$ satisfies (6.15), one can show that ${ }^{11}$
(a) $\mathbf{1}_{\left(g_{i, \ell+1}\right)_{k} \geq 0}$ stays the same with respect to perturbation $W_{k}^{\prime}$, except for $i=i^{*}$ and $\ell=\ell^{*}$; and
(b) $\mathbf{1}_{\left(g_{i^{*}, \ell^{*}+1}\right)_{k} \geq 0}$ can be 0 or 1 each with at least constant probability over $W_{k}^{\prime}$.

At the same time, using the fact that $k$ satisfies (6.14), one can show that
(c) $h_{i, \ell}$ and $\left(u_{i, \ell}\right)_{k}$ does not change by much (owing to Section 6.6); and
(d) $\left|\left(u_{i^{*}, \ell^{*}+1}\right)_{k}\right| \geq \frac{\left\|\operatorname{los}_{i^{*}, \ell^{*}}\right\|_{2}}{\rho}$ is large (owing to (6.14)).

## ${ }^{11}$ Specifically,

- For every $\ell<\ell^{*}$, we have $\left(g_{i, \ell+1}\right)_{k}$ is unchanged with respect to $W_{k}^{\prime}$. This is because $g_{i, \ell+1}$ only depends on the randomness of $W U_{\ell}$, but $v_{i^{*}, \ell^{*}}$ is orthogonal to the columns of $U_{\ell}$.
- For $i=i^{*}$ and $i=\ell^{*}$, we have $\left(g_{i, \ell+1}\right)_{k}$ will change sign with constant probability with respect to $W_{k}^{\prime}$. This is because $\left\|W_{k}^{\prime}\right\|$ is above $\frac{\beta_{-}}{\sqrt{m}}-$ the original $\left(g_{i, \ell+1}\right)_{k}$ before perturbation.
- For any other $i$ and $\ell$, we have $\left(g_{i, \ell+1}\right)_{k}$ will not change sign with high probability with respect to $W_{k}^{\prime}$. This is because $\left\|W_{k}^{\prime}\right\|$ is below $\frac{\beta_{+}}{\sqrt{m}}$ - the original $\left(g_{i, \ell+1}\right)_{k}$ before perturbation.

Putting (a), (b), (c), and (d) together, we know for such specially chosen $k$, at least with constant probability over the random perturbation of $W_{k}^{\prime}$,

$$
\begin{equation*}
\left\|\widehat{\nabla}_{k} f\left(W+W_{k}^{\prime}\right)\right\|_{F} \geq \Omega\left(\frac{\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}}{\rho}\right) . \tag{6.16}
\end{equation*}
$$

### 6.7.3 Real Proof: Randomness Decomposition and McDiarmid's Inequality

There are only two main differences between (6.16) and our desired Theorem 6.14.2. First, (6.16) gives a gradient lower bound at $W+W_{k}^{\prime}$, while in Theorem 6.14.2 we need a gradient lower bound at random initialization $W$. Second, (6.16) gives a lower bound on $\widehat{\nabla}_{k} f(\cdot)$ with constant probability for a small fraction of good coordinates $k$, but in Theorem 6.14.2 we need a lower bound for the entire $\widehat{\nabla} f(\cdot)$.

Randomness Decomposition To fix the first issue, we resort to a randomness decomposition technique at least tracing back to the smooth analysis of [ST04]:

Proposition 6.7.2. Given small constant $\theta \in(0,1)$ and m-dimensional random $g \sim \mathcal{N}\left(0, \frac{1}{m} I\right)$, we can rewrite $g=g_{1}+g_{3}$ where $g_{1}$ follows from $\mathcal{N}\left(0, \frac{1}{m} I\right)$ and $g_{3}$ is very close to $\mathcal{N}\left(0, \frac{\theta^{2}}{m} I\right)$. (Note that there is no contradiction here because $g_{1}$ and $g_{3}$ shall be correlated.)

Proof. Let $g_{1}, g_{2} \in \mathbb{R}^{m}$ be two independent random vectors sampled from $\mathcal{N}\left(0, \frac{1}{m} I\right)$. We can couple them and make sure $g=\sqrt{1-\theta^{2}} g_{1}+\theta g_{2}$. We now choose $g_{3}=\theta g_{2}-\left(1-\sqrt{1-\theta^{2}}\right) g_{1}$. Since $\left(1-\sqrt{1-\theta^{2}}\right) \leq O\left(\theta^{2}\right)$, when $\theta$ is sufficiently small, we know that $g_{3}$ is close to being generated from distribution $\mathcal{N}\left(0, \frac{\theta^{2}}{m} I\right)$.

Using Proposition 6.7.2, for each good coordinate $k$, instead of "adding" perturbation $W_{k}^{\prime}$ to $W$, we can instead decompose $W$ into $W=W_{0}+W_{k}^{\prime}$, where $W_{0}$ is distributed in the same way as $W$. In other words, $W_{0}$ is also at random initialization. If this idea is carefully implemented, one can immediately turn (6.16) into

$$
\begin{equation*}
\left\|\widehat{\nabla}_{k} f(W)\right\|_{F}=\left\|\widehat{\nabla}_{k} f\left(W_{0}+W_{k}^{\prime}\right)\right\|_{F} \geq \Omega\left(\frac{\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}}{\rho}\right) \tag{6.17}
\end{equation*}
$$

Extended McDiarmid's Inequality To fix the second issue, one may wish to consider all the indices $k \in[m]$ satisfying (6.15) and (6.14). Since there are at least $\frac{\delta}{\text { poly }(\rho)}$ fraction of such coordinates, if all of them satisfied (6.17), then we would have already proved Theorem 6.14.2. Unfortunately, neither can we apply Chernoff bound (because the events with different $k \in$ [ $m$ ] are correlated), nor can we apply union bound (because the event occurs only with constant probability).

Our technique is to resort to (an extended probabilistic variant of) McDiarmid's inequality (see Section 6.10.6) in a very non-trivial way to boost the confidence.

Give any fixed subset $N$ of cardinality $|N|=(\operatorname{poly}(\rho) / \delta)^{2}$, one can show that there again exists $\frac{\delta}{\text { poly }(\rho)}$ fraction of coordinate $k \in N$ satisfying (6.15) and (6.14). ${ }^{12}$ Now, instead of decomposing $W=W_{0}+W_{k}^{\prime}$, we decompose it as $W=W_{1}+W_{N}^{\prime}$ for $W_{N}^{\prime}=u_{N} v_{i^{*}, \ell^{*}}^{\top}$ where $u_{N}$ is only supported on coordinates $k \in N$. In other words, we simultaneously perturb in the directions of $W_{k}^{\prime}$ for all $k \in N$. Since this perturbation is small enough -i.e., $\left\|W_{N}^{\prime}\right\|_{2} \leq \frac{\operatorname{poly}(\varrho)}{\sqrt{m}}$ - one can show that (6.17) remains true, that is, for a large fraction of $k \in N$, with at least constant probability over $W_{N}^{\prime}$ :

$$
\begin{equation*}
\left\|\widehat{\nabla}_{k} f\left(W_{1}+W_{N}^{\prime}\right)\right\|_{F} \geq \Omega\left(\frac{\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}}{\rho}\right) \tag{6.18}
\end{equation*}
$$

In order to apply McDiarmid's, we next need to bound on the difference between $\| \widehat{\nabla}_{k} f\left(W_{1}+\right.$ $\left.W_{N}^{\prime}\right) \|_{F}$ and $\left\|\widehat{\nabla}_{k} f\left(W_{1}+W_{N}^{\prime}+W_{j}^{\prime \prime}\right)\right\|_{F}$ for an arbitrary (but small) perturbation $W_{j}^{\prime \prime}$ (in the direction of $\left.e_{j} v_{i^{*}, \ell^{*}}^{\top}\right)$. We show that,

$$
\begin{equation*}
\sum_{k \in[N]}\left|\left\|\widehat{\nabla}_{k}\left(W_{1}+W_{N}^{\prime}\right)\right\|_{2}^{2}-\left\|\widehat{\nabla}_{k}\left(W_{1}+W_{N}^{\prime}+W_{j}^{\prime \prime}\right)\right\|_{2}^{2}\right| \leq O\left(\rho^{8}+\frac{|N|}{m^{1 / 6}}\right) \tag{6.19}
\end{equation*}
$$

[^39]In other words, although there are $|N|$ difference terms, their total summation only grows in rate $\frac{|N|}{m^{1 / 6}}$ according to (6.19). After applying a variant of McDiarmid's inequality, we derive that with high probability over $W_{N}^{\prime}$, it satisfies

$$
\begin{equation*}
\sum_{k \in N}\left\|\widehat{\nabla}_{k} f(W)\right\|_{F}^{2}=\sum_{k \in N}\left\|\widehat{\nabla}_{k} f\left(W_{1}+W_{N}^{\prime}\right)\right\|_{F}^{2} \geq \Omega\left(\frac{\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}^{2}}{\operatorname{poly}(\rho)}|N|\right) \tag{6.20}
\end{equation*}
$$

Finally, by sampling sufficiently many random sets $N$ to cover the entire space [ $m$ ], we can show

$$
\|\widehat{\nabla} f(W)\|_{F}^{2}=\sum_{k \in[m]}\left\|\widehat{\nabla}_{k} f(W)\right\|_{F}^{2} \geq \Omega\left(\frac{\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}^{2}}{\operatorname{poly}(\rho)} m\right)
$$

and therefore Theorem 6.14.2 holds.

### 6.8 Proof Sketch of Theorem 6.16.1: Objective Semi-Smoothness

The objective semi-smoothness Theorem 6.16.1 turns out to be much simpler to prove than Theorem 6.15.2. It only relies on Section 6.5 and 6.6, and does not need randomness decomposition or McDiarmid's inequality. (Details in Section 6.16.)

Recall that in Theorem 6.16.1, $\widetilde{W}, A, B$ are at random initialization. $\breve{W}$ is an adversarially chosen matrix with $\|\breve{W}-\widetilde{W}\| \leq \frac{\operatorname{poly}(\rho)}{\sqrt{m}}$, and $W^{\prime}$ is some other adversarial perturbation on top of $\breve{W}$, satisfying $\left\|W^{\prime}\right\| \leq \frac{\tau_{0}}{\sqrt{m}}$. We denote by

- $\widetilde{D}_{i, \ell}, \widetilde{g}_{i, \ell}, \widetilde{h}_{i, \ell}$ respectively the values of $D_{i, \ell}, g_{i, \ell}, h_{i, \ell}$ determined by weight matrix $\widetilde{W}$;
- $\breve{D}_{i, \ell}, \breve{g}_{i, \ell}, \hbar_{i, \ell}, \operatorname{loss}_{i, \ell}$ respectively those of $D_{i, \ell}, g_{i, \ell}, h_{i, \ell}$ and $\operatorname{loss}_{i, \ell}$ at weight matrix $\breve{W}$; and
- $D_{i, \ell}, g_{i, \ell}, h_{i, \ell}$ respectively the values of $D_{i, \ell}, g_{i, \ell}, h_{i, \ell}$ at weight matrix $W=\breve{W}+W^{\prime}$.

Our main tool is to derive the following strong formula for $h_{i, \ell}-\hbar_{i, \ell}$ : there exist diagonal matrices $D_{i, \ell}^{\prime \prime} \in \mathbb{R}^{m \times m}$ with entries in $[-1,1]$ and sparsity $\left\|D_{i, \ell}^{\prime \prime}\right\|_{0} \leq O\left(m^{2 / 3}\right)$ such that,

$$
\begin{equation*}
h_{i, \ell}-\hbar_{i, \ell}=\sum_{a=1}^{\ell-1}\left(\breve{D}_{i, \ell}+D_{i, \ell}^{\prime \prime}\right) \breve{W} \cdots \breve{W}\left(\breve{D}_{i, a+1}+D_{i, a+1}^{\prime \prime}\right) W^{\prime} h_{i, a} \tag{6.21}
\end{equation*}
$$

In particular, (6.21) implies $\left\|h_{i, \ell}-\hbar_{i, \ell}\right\| \leq O\left(L^{9}\right)\left\|W^{\prime}\right\|_{2}$ after careful linear-algebraic manipulations (esp. using (6.6)). The main take-away message from (6.21) is that, this difference $h_{i, \ell}-\hbar_{i, \ell}$ is proportional to the norm of the perturbation, $\left\|W^{\prime}\right\|_{2}$, no matter how small it is. In contrast, in (6.2), we only derived a weak upper bound of the form $\left\|h_{i, \ell}-\widetilde{h}_{i, \ell}\right\| \leq O\left(m^{-1 / 2}\right)$. Nevertheless, the proof of (6.21) relies on (6.2), so we are not duplicating proofs.

Finally, we carefully derive that

$$
\begin{align*}
& f\left(\breve{W}+W^{\prime}\right)-f(W)-\left\langle\nabla f(W), W^{\prime}\right\rangle \\
= & \sum_{i=1}^{n} \sum_{\ell=2}^{L} \breve{\operatorname{loss}}_{i, \ell}^{\top} B\left(\left(h_{i, \ell}-\hbar_{i, \ell}\right)-\sum_{a=1}^{\ell-1} \breve{D}_{i, \ell} \breve{W} \cdots \breve{W} \breve{D}_{i, a+1} W^{\prime} \hbar_{i, a}\right)+\frac{1}{2}\left\|B\left(h_{i, \ell}-\hbar_{i, \ell}\right)\right\|^{2} \tag{6.22}
\end{align*}
$$

and plug (6.21) into (6.22) to derive our final Theorem 6.16.1.

### 6.9 Roadmap of Later Proof Details

- Section 6.10 recalls some old lemmas and derives some new lemmas in probability theory.
- Section 6.11 serves for Section 6.5, the basic properties at random initialization.
- Section 6.12 serves for Section 6.6, the stability after adversarial perturbation.
- Section 6.13, 6.14 and 6.15 together serve for Section 6.7 and prove Theorem 6.15.2, the Polyak-Łojasiewicz condition and gradient upper bound. In particular:
- Section 6.13 serves for Section 6.7.1 (the indicator and backward coordinate bounds).
- Section 6.14 serves for Section 6.7 .3 (the randomness decomposition and McDiarmid's inequality) and proves Theorem 6.14.2.
- Section 6.15 shows how to go from Theorem 6.14.2 to Theorem 6.15.2.
- Section 6.16 serves for Section 6.8, the proof of Theorem 6.16.1, the objective semismoothness.
- Section 6.17 gives the final proof for Theorem 6.17.1, the GD convergence theorem.
- Section 6.18 gives the final proof for Theorem 6.18.1, the SGD convergence theorem.

Parameters We also summarize a few parameters we shall use in the proofs.

- In Definition 6.13.1, we shall introduce two parameters

$$
\beta_{+} \stackrel{\text { def }}{=} \frac{\delta}{\rho^{2}} \text { and } \beta_{-} \xlongequal{\xlongequal[\text { def }]{=}} \frac{\delta}{\rho^{10}}
$$

to control the thresholds of indicator functions (recall Section 6.7.1).

- In Definition 6.14.2, we shall introduce parameter

$$
\theta \in\left[\rho^{4} \cdot \beta_{-}, \rho^{-3} \cdot \beta_{+}\right]
$$

to describe how much randomness we want to decompose out of $W$ (recall Section 6.7.2).

- In (6.60) of Section 6.14.4, we shall choose

$$
N=\frac{\rho^{22}}{\beta_{-}^{2}}
$$

which controls the size of the set $N$ where we apply McDiarmid's inequality (recall Section 6.7.3).

### 6.10 Preliminaries on Probability Theory

The goal of this section is to present a list of probability tools.

- In Section 6.10.1, we recall how to swap randomness.
- In Section 6.10.2, we recall concentration bounds for the chi-square distribution.
- In Section 6.10.3, we proved a concentration bound of sum of squares of ReLU of Gaussians.
- In Section 6.10.4 and Section 6.10.5, we show some properties for random Gaussian vectors.
- In Section 6.10.6, we recall the classical McDiarmid's inequality and then prove a general version of it.


### 6.10.1 Swapping Randomness

Fact 6.10.1 (probability splitting). If $f(X, Y)$ holds with probability at least $1-\epsilon$, then

- with probability at least $1-\sqrt{\epsilon}$ (over randomness of $X$ ), the following event holds,
$-f(X, Y)$ holds with probability at least $1-\sqrt{\epsilon}$ (over randomness of $Y$ ).
In other words,

$$
\operatorname{Pr}_{X}\left[\operatorname{Pr}_{Y}[f(X, Y)] \geq 1-\sqrt{\epsilon} \geq 1-\sqrt{\epsilon}\right.
$$

Fact 6.10.2 (swapping probability and expectation). If $\operatorname{Pr}_{X, Y}[f(X, Y) \geq a] \geq \epsilon$, then

$$
\operatorname{Pr}_{X}[\underset{Y}{\mathbb{E}}[f(X, Y)] \geq a \epsilon / 2] \geq \epsilon / 2 .
$$

Proof. We prove it by making a contradiction,
Suppose

$$
\operatorname{Pr}_{X}[\underset{Y}{\mathbb{E}}[f(X, Y)] \leq a \epsilon / 2] \geq 1-\epsilon / 2,
$$

which implies that

$$
\operatorname{Pr}_{X, Y}[f(X, Y) \geq a \mid X] \leq \epsilon .
$$

### 6.10.2 Concentration of Chi-Square Distribution

Lemma 6.10.3 (Lemma 1 on page 1325 of [LM00]). Let $X \sim X_{k}^{2}$ be a chi-squared distributed random variable with $k$ degrees of freedom. Each one has zero mean and $\sigma^{2}$ variance. Then

$$
\begin{aligned}
& \operatorname{Pr}\left[X-k \sigma^{2} \geq(2 \sqrt{k t}+2 t) \sigma^{2}\right] \leq \exp (-t) \\
& \operatorname{Pr}\left[k \sigma^{2}-X \geq 2 \sqrt{k t} \sigma^{2}\right] \leq \exp (-t)
\end{aligned}
$$

One straightforward application is
Lemma 6.10.4. Let $x_{1}, x_{2}, \cdots, x_{n}$ denote i.i.d. samples from $\mathcal{N}\left(0, \sigma^{2}\right)$. For any $b \geq 1$, we have

$$
\operatorname{Pr}\left[\left|\|x\|_{2}^{2}-n \sigma^{2}\right| \geq \frac{n}{b} \sigma^{2}\right] \leq 2 \exp \left(-n /\left(8 b^{2}\right)\right)
$$

Proof. We choose $t=k /\left(8 b^{2}\right)$ in Lemma 6.10.3,

$$
\operatorname{Pr}\left[\left|\|x\|_{2}^{2}-n \sigma^{2}\right| \geq\left(\frac{2 n}{\sqrt{8} b}+\frac{2 n}{8 b^{2}}\right) \sigma^{2}\right] \leq 2 \exp \left(-n /\left(8 b^{2}\right)\right)
$$

Since $\frac{2 n}{\sqrt{8} b}+\frac{2 n}{8 b^{2}} \leq \frac{2 n}{\sqrt{8} b}+\frac{2 n}{8 b} \leq n / b$. Thus,

$$
\operatorname{Pr}\left[\left|\|x\|_{2}^{2}-n \sigma^{2}\right| \geq \frac{n}{b} \sigma^{2}\right] \leq 2 \exp \left(-n /\left(8 b^{2}\right)\right)
$$

which completes the proof.

Lemma 6.10.5. Let $x_{1}, x_{2}, \cdots, x_{m}$ denote i.i.d. samples from $\mathcal{N}(0,1)$, and $y_{i}=\max \left\{x_{i}^{2}-\right.$ $\log m, 0\}$. We have

$$
\operatorname{Pr}\left[\sum_{i=1}^{m} y_{i} \geq 2 \sqrt{m}\right] \leq e^{-\Omega(\sqrt{m})}
$$

Proof. First of all, letting $g=\sqrt{\log m}$,

$$
\mathbb{E}\left[y_{i}\right]=\int_{g}^{\infty} \frac{\exp \left(-\frac{x^{2}}{2}\right)(x-g)^{2}}{\sqrt{2 \pi}} d x=\frac{1}{2}\left(g^{2}+1\right) \operatorname{erfc}\left(\frac{g}{\sqrt{2}}\right)-\frac{e^{-\frac{g^{2}}{2}} g}{\sqrt{2 \pi}} \leq \frac{1}{e^{\frac{g^{2}}{2}}}=\frac{1}{\sqrt{m}} .
$$

On the other hand, each random variable $y_{i}$ is $O(1)$-subgaussian. By subgaussian concentration,

$$
\operatorname{Pr}\left[\sum_{i=1}^{m} y_{i} \geq 2 \sqrt{m}\right] \leq e^{-\Omega(\sqrt{m})}
$$

### 6.10.3 Concentration of Sum of Squares of ReLU of Gaussians

Lemma 6.10.6 (Upper bound). Given $n$ i.i.d. Gaussian random variables $x_{1}, x_{2}, \cdots, x_{n} \sim$ $\mathcal{N}\left(0, \sigma^{2}\right)$, we have

$$
\operatorname{Pr}\left[\left(\sum_{i=1}^{n} \max \left(x_{i}, 0\right)^{2}\right)^{1 / 2}<(1+\epsilon) \sqrt{n / 2} \sigma\right] \geq 1-\exp \left(-\epsilon^{2} n / 100\right)
$$

Proof. Using Chernoff bound, we know that with probability $1-\exp \left(-\epsilon^{2} n / 6\right), \sum_{i=1}^{n} \max \left(x_{i}, 0\right)^{2}$ is a at most degree- $(1+\epsilon) \frac{n}{2}$ Chi-square random variable. Let us say this is the first event.

Using Lemma 6.10.3, we have

$$
\begin{array}{rlr} 
& \operatorname{Pr}\left[X \geq k \sigma^{2}+(2 \sqrt{k t}+2 t) \sigma^{2}\right] \leq \exp (-t) \\
\Longrightarrow & \operatorname{Pr}\left[X \geq k \sigma^{2}+\left(2 \epsilon k+2 \epsilon^{2} k\right) \sigma^{2}\right] \leq \exp \left(-\epsilon^{2} k\right) \\
\Longrightarrow & \operatorname{Pr}\left[X \geq k(1+4 \epsilon) \sigma^{2}\right] \leq \exp \left(-\epsilon^{2} k\right) \\
\Longrightarrow & \operatorname{Pr}\left[X \geq(1+\epsilon)(1+4 \epsilon) \frac{n}{2} \sigma^{2}\right] \leq \exp \left(-\epsilon^{2}(1+\epsilon) n / 2\right) \quad \text { by choosing } t=\epsilon^{2} k \\
\hline
\end{array} \quad \text { by } k=(1+\epsilon) n / 2
$$

Thus, we have with probability at least $1-\exp \left(-\epsilon^{2} n / 2\right)$,

$$
X \leq(1+4 \epsilon)^{2} \frac{n}{2} \sigma^{2}
$$

Let the above event denote the second event.
By taking the union bound of two events, we have with probability $1-2 \exp \left(-\epsilon^{2} n / 6\right)$

$$
\left(\sum_{i=1}^{n} \max \left(x_{i}, 0\right)^{2}\right)^{1 / 2} \leq X \leq(1+4 \epsilon) \frac{n}{2} \sigma
$$

Then rescaling the $\epsilon$, we get the desired result.

Lemma 6.10.7 (Lower bound). Given $n$ i.i.d. Gaussian random variables $x_{1}, x_{2}, \cdots, x_{n} \sim$ $\mathcal{N}\left(0, \sigma^{2}\right)$, we have

$$
\operatorname{Pr}\left[\left(\sum_{i=1}^{n} \max \left(x_{i}, 0\right)^{2}\right)^{1 / 2}>(1-\epsilon) \sqrt{n / 2} \sigma\right] \geq 1-\exp \left(\epsilon^{2} n / 100\right)
$$

Proof. Using Chernoff bound, we know that with probability $1-\exp \left(-\epsilon^{2} n / 6\right), \sum_{i=1}^{n} \max \left(x_{i}, 0\right)^{2}$ is a at most degree- $(1-\epsilon) \frac{n}{2}$ Chi-square random variable. Let us say this is the first event.

Using Lemma 6.10.3, we have

$$
\begin{array}{rlr} 
& \operatorname{Pr}\left[X \leq k \sigma^{2}-2 \sqrt{k t} \sigma^{2}\right] \leq \exp (-t) \\
\Longrightarrow & \operatorname{Pr}\left[X \leq k \sigma^{2}-2 \epsilon k \sigma^{2}\right] \leq \exp \left(-\epsilon^{2} k\right) \\
\Longrightarrow & \operatorname{Pr}\left[X \leq k \sigma^{2}(1-2 \epsilon)\right] \leq \exp \left(-\epsilon^{2} k\right) \\
\Longrightarrow & \operatorname{Pr}\left[X \leq(1-\epsilon)(1-2 \epsilon) \frac{n}{2} \sigma^{2}\right] \leq \exp \left(-\epsilon^{2}(1-\epsilon) n / 2\right) \quad \text { by choosing } t=\epsilon^{2} k \\
\hline
\end{array}
$$

Thus, we have with probability at least $1-\exp \left(-\epsilon^{2} n / 4\right)$,

$$
X \geq(1-2 \epsilon)^{2} \frac{n}{2} \sigma^{2}
$$

Let the above event denote the second event.
By taking the union bound of two events, we have with probability at least 1 $2 \exp \left(-\epsilon^{2} n / 6\right)$,

$$
\left(\sum_{i=1}^{n} \max \left(x_{i}, 0\right)^{2}\right)^{1 / 2} \geq X \geq(1-2 \epsilon) \frac{n}{2} \sigma .
$$

Then rescaling the $\epsilon$, we get the desired result.

Combining Lemma 6.10.7 and Lemma 6.10.6, we have

Lemma 6.10.8 (Two sides bound). Given $n$ i.i.d. Gaussian random variables $x_{1}, x_{2}, \cdots, x_{n} \sim$ $\mathcal{N}\left(0, \sigma^{2}\right)$, let $\phi(a)=\max (a, 0)^{2}$. We have

$$
\operatorname{Pr}_{x}\left[\|\phi(x)\|_{2} \in((1-\epsilon) \sqrt{n / 2} \sigma,(1+\epsilon) \sqrt{n / 2} \sigma)\right] \geq 1-2 \exp \left(-\epsilon^{2} n / 100\right)
$$

Corollary 6.10.9 (Two sides bound for single matrix). Let $x \in \mathbb{R}^{n}$ denote a fixed vector. Given a random Gaussian matrix $A \in \mathbb{R}^{m \times n}$ where each entry is i.i.d. sampled from $\mathcal{N}\left(0,2 \sigma^{2} / m\right)$.

$$
\underset{A}{\operatorname{Pr}}\left[\|\phi(A x)\|_{2} \in\left((1-\epsilon)\|x\|_{2} \sigma,(1+\epsilon)\|x\|_{2} \sigma\right)\right] \geq 1-2 \exp \left(-\epsilon^{2} m / 100\right) .
$$

Proof. For each $i \in[m]$, let $y_{i}=(A x)_{i}$. Then $y_{i} \sim \mathcal{N}\left(0, \widetilde{\sigma}^{2}\right)$, where $\widetilde{\sigma}^{2}=2 \sigma^{2} / m \cdot\|x\|_{2}^{2}$. Using Corollary 6.10.8, we have

$$
\operatorname{Pr}_{y}\left[\|\phi(y)\|_{2} \in((1-\epsilon) \sqrt{m / 2} \widetilde{\sigma},(1+\epsilon) \sqrt{m / 2} \widetilde{\sigma})\right] \geq 1-2 \exp (-\epsilon m / 100)
$$

which implies

$$
\underset{y}{\operatorname{Pr}}\left[\|\phi(y)\|_{2} \in\left((1-\epsilon)\|x\|_{2} \sigma,(1+\epsilon)\|x\|_{2} \sigma\right)\right] \geq 1-2 \exp \left(-\epsilon^{2} m / 100\right)
$$

Since $y=A x$, thus we complete the proof.

Corollary 6.10.10 (Two sides bound for multiple matrices). Let $x_{1}, x_{2}, \cdots, x_{k}$ denote $k$ fixed vectors where $x_{i} \in \mathbb{R}^{n_{i}}$. Let $x=\left[\begin{array}{lll}x_{1}^{\top} & x_{2}^{\top} \cdots & x_{k}^{\top}\end{array}\right]^{\top} \in \mathbb{R}^{n}$ where $n=\sum_{i=1}^{k} n_{i}$. Let $A_{1}, A_{2}, \cdots A_{k}$ denote $k$ independent random Gaussian matrices where each entry of $A_{i} \in \mathbb{R}^{m \times n_{i}}$ is i.i.d. sampled from $\mathcal{N}\left(0,2 \sigma_{i}^{2} / m\right)$ for each $i \in[k]$. Let $A=\left[A_{1} A_{2} \cdots A_{k}\right]$. We have

$$
\operatorname{Pr}_{A}\left[\|\phi(A x)\|_{2} \in\left((1-\epsilon)\left(\sum_{i=1}^{k}\left\|x_{i}\right\|_{2}^{2} \sigma_{i}^{2}\right)^{1 / 2},(1+\epsilon)\left(\sum_{i=1}^{k}\left\|x_{i}\right\|_{2}^{2} \sigma_{i}^{2}\right)^{1 / 2}\right)\right] \geq 1-2 \exp \left(-\epsilon^{2} m / 100\right)
$$

Proof. It is similar as Corollary 6.10.9.

Fact 6.10.11 (see e.g. [AZLS19]). Let $h, q \in \mathbb{R}^{p}$ be fixed vectors and $h \neq 0, W \in \mathbb{R}^{m \times p}$ be random matrix with i.i.d. entries $W_{i, j} \sim \mathcal{N}\left(0, \frac{2}{m}\right)$, and vector $v \in \mathbb{R}^{m}$ defined as $v_{i}=$ $\mathbf{1}_{\left\langle W_{i}, h+q\right\rangle \geq 0}\left\langle W_{i}, h\right\rangle$. Then,

- $\left|v_{i}\right|$ follows i.i.d. from the following distribution: with half probability $\left|v_{i}\right|=0$, and with the other half probability $\left|v_{i}\right|$ follows from folded Gaussian distributions $\left|\mathcal{N}\left(0, \frac{2\|h\|^{2}}{m}\right)\right|$.
- $\frac{m\|v\|^{2}}{2\|h\|^{2}}$ is in distribution identical to $\chi_{\omega}^{2}$ (chi-square distribution of order $\omega$ ) where $\omega$ follows from binomial distribution $\mathcal{B}(m, 1 / 2)$ ( $m$ trials each with success rate $1 / 2$ ).


### 6.10.4 Gaussian Vector Percentile: Center

Fact 6.10.12. Suppose $x \sim \mathcal{N}\left(0, \sigma^{2}\right)$ is a Gaussian random variable. For any $t \in(0, \sigma]$ we have

$$
\operatorname{Pr}[x \geq t] \in\left[\frac{1}{2}\left(1-\frac{4}{5} \frac{t}{\sigma}\right), \frac{1}{2}\left(1-\frac{2}{3} \frac{t}{\sigma}\right)\right] .
$$

Similarly, if $x \sim \mathcal{N}\left(\mu, \sigma^{2}\right)$, for any $t \in(0, \sigma]$, we have

$$
\operatorname{Pr}[|x| \geq t] \in\left[1-\frac{4}{5} \frac{t}{\sigma}, 1-\frac{2}{3} \frac{t}{\sigma}\right]
$$

Lemma 6.10.13. Let $x \sim \mathcal{N}\left(0, \sigma^{2} I\right)$. For any $\alpha \in(0,1 / 2)$, we have with probability at least $1-\exp \left(-\alpha^{2} m / 100\right)$,

- there exists at least $\frac{1}{2}(1-\alpha)$ fraction of $i$ such that $x_{i} \geq 5 \alpha \sigma / 16$, and
- there exists at least $\frac{1}{2}(1-\alpha)$ fraction of $i$ such that $x_{i} \leq-5 \alpha \sigma / 16$.

Proof. Let $c_{1}=4 / 5$. For each $i \in[m]$, we define random variable $y_{i}$ as

$$
y_{i}= \begin{cases}1, & \text { if } x_{i} \geq \alpha \sigma /\left(4 c_{1}\right) \\ 0, & \text { otherwise }\end{cases}
$$

Let $p=\operatorname{Pr}\left[y_{i}=1\right]$. Using Fact 6.10 .12 with $t=\alpha \sigma /\left(4 c_{1}\right)$, we know that $p \geq \frac{1}{2}(1-\alpha / 4)$. Letting $Y=\sum_{i=1}^{m} y_{i}$, we have $\mu=\mathbb{E}[Y]=m p \geq \frac{1}{2}(1-\alpha / 4) m$. Using Chernoff bound, we have

$$
\operatorname{Pr}[Y \leq(1-\delta) \mu] \leq \exp \left(-\delta^{2} \mu / 2\right)
$$

Choosing $\delta=\frac{1}{4} \alpha$, we have

$$
\operatorname{Pr}[Y \leq(1-\alpha / 4) \mu] \leq \exp \left(-\alpha^{2} \mu / 32\right)
$$

Further we have

$$
\begin{array}{rlrl}
\operatorname{Pr}\left[Y \leq \frac{1}{2}(1-\alpha) m\right] & \leq \operatorname{Pr}\left[Y \leq(1-\alpha / 4) \frac{1}{2}(1-\alpha / 4) m\right] & & \text { by }(1-\alpha) \leq(1-\alpha / 4)^{2} \\
& \leq \operatorname{Pr}[Y \leq(1-\alpha / 4) \mu] & & \text { by } \mu \geq \frac{1}{2}(1-\alpha / 4) m \\
& =\operatorname{Pr}[Y \leq(1-\delta) \mu] & & \\
& \leq \exp \left(-\alpha^{2} \mu / 32\right) & & \\
& \leq \exp \left(-\alpha^{2} \frac{1}{2}(1-\alpha / 4) m / 32\right) & \text { by } \mu \geq \frac{1}{2}(1-\alpha / 4) m \\
& =\exp \left(-\alpha^{2}(1-\alpha / 4) m / 64\right) . &
\end{array}
$$

We provide the definition of $(\alpha, \sigma)$-good. Note that this definition will be used often in the later proof.

Definition 6.10.1 $\left((\alpha, \sigma)\right.$-good). Given $w \in \mathbb{R}^{m}$, we say $w$ is $(\alpha, \sigma)$-good if the following two conditions holding:

- there are at least $\frac{1}{2}(1-\alpha)$ fraction coordinates satisfy that $w_{i} \geq \alpha \sigma$; and
- there are at least $\frac{1}{2}(1-\alpha)$ fraction coordinates satisfy that $w_{i} \leq-\alpha \sigma$.

Lemma 6.10.13 gives the following immediate corollary:

Corollary 6.10.14 (random Gaussian is ( $\alpha, \sigma / 4$ )-good). Let $x \sim \mathcal{N}\left(0, \sigma^{2} I\right)$. For any $\alpha \in$ $(0,1 / 2)$, we have with probability at least $1-\exp \left(-\alpha^{2} m / 100\right)$ that $x$ is ( $\left.\alpha, \sigma / 4\right)$-good.

Corollary 6.10.15. Let $x_{1}, x_{2}, \cdots, x_{k}$ be $k$ fixed vectors where $x_{i} \in \mathbb{R}^{n_{i}}$, and $A_{1}, A_{2}, \cdots, A_{k}$ be $k$ independent random Gaussian matrices where each entry of $A_{i} \in \mathbb{R}^{m \times n_{i}}$ is i.i.d. sampled from $\mathcal{N}\left(0, \sigma_{i}^{2}\right)$ for each $i \in[k]$. Denote by $x=\left(x_{1}, \ldots, x_{k}\right) \in \mathbb{R}^{n}$ for $n=\sum_{i=1}^{k} n_{i}, A=$
$\left[A_{1}, A_{2}, \cdots, A_{k}\right] \in \mathbb{R}^{m \times n}$, and $\sigma=\left(\sum_{i=1}^{k} \sigma_{i}^{2}\left\|x_{i}\right\|_{2}^{2}\right)^{1 / 2}$. For any fixed parameter $\alpha \in(0,1 / 2)$, we have

$$
A x=A_{1} x_{1}+\cdots A_{k} x_{k} \text { is }(\alpha, \sigma / 4) \text {-good with probability at least } 1-\exp \left(-\alpha^{2} m / 100\right)
$$

Proof. It is clear that $A x$ follows from a Gaussian distribution $\mathcal{N}\left(0,\left(\sum_{i=1}^{k} \sigma_{i}^{2}\left\|x_{i}\right\|_{2}^{2}\right) I\right)$, so we can directly apply Corollary 6.10.14.

### 6.10.5 Gaussian Vector Percentile: Tail

Lemma 6.10.16. Suppose $W \in \mathbb{R}^{m \times n}$ is a random matrix with entries drawn i.i.d. from $\mathcal{N}\left(0, \frac{2}{m}\right)$. Given $d \leq s \leq m$, with probability at least $1-e^{-\Omega\left(s \log ^{2} m\right)}$, for all $x \in \mathbb{R}^{n}$, letting $y=W x$, we can write $y=y_{1}+y_{2}$ with

$$
\left\|y_{1}\right\| \leq \frac{\sqrt{s} \log ^{2} m}{\sqrt{m}} \cdot\|x\| \quad \text { and } \quad\left\|y_{2}\right\|_{\infty} \leq \frac{\log m}{\sqrt{m}} \cdot\|x\| .
$$

Proof of Lemma 6.10.16. Without loss of generality we only prove the result for $\|x\|=1$.
Fixing any such $x$ and letting $\beta=\frac{\log m}{2 \sqrt{m}}$, we have $y_{i} \sim \mathcal{N}\left(0, \frac{2}{m}\right)$ so for every $p \geq 1$, by Gaussian tail bound

$$
\operatorname{Pr}\left[\left|y_{i}\right| \geq \beta p\right] \leq e^{-\Omega\left(\beta^{2} p^{2} m\right)} .
$$

Since $\beta^{2} p^{2} m \geq \beta^{2} m \gg \Omega(\log m)$, we know that if $\left|y_{i}\right| \geq \beta p$ occurs for $q / p^{2}$ indices $i$ out of [ $m$ ], this cannot happen with probability more than

$$
\binom{m}{q / p^{2}} \times\left(e^{-\Omega\left(\beta^{2} p^{2} m\right)}\right)^{q / p^{2}} \leq e^{\frac{q}{p^{2}}\left(O(\log m)-\Omega\left(\beta^{2} p^{2} m\right)\right)} \leq e^{-\Omega\left(\beta^{2} q m\right)}
$$

In other words,

$$
\operatorname{Pr}\left[\left|\left\{i \in[m]:\left|y_{i}\right| \geq \beta p\right\}\right|>q / p^{2}\right] \leq e^{-\Omega\left(\beta^{2} q m\right)}
$$

Finally, by applying union bound over $p=1,2,4,8,16, \ldots$ we have with probability $\geq$ $1-e^{-\Omega\left(\beta^{2} q m\right)} \cdot \log q$,

$$
\begin{equation*}
\sum_{i:\left|y_{i}\right| \geq \beta} y_{i}^{2} \leq \sum_{k=0}^{\lceil\log q\rceil}\left(2^{k+1} \beta\right)^{2}\left|\left\{i \in[m]:\left|y_{i}\right| \geq 2^{k} \beta\right\}\right| \leq \sum_{k=0}^{\lceil\log q\rceil}\left(2^{k+1} \beta\right)^{2} \cdot \frac{q}{2^{2 k}} \leq 4 q \beta^{2} \log q \tag{6.23}
\end{equation*}
$$

In other words, vector $y$ can be written as $y=y_{1}+y_{2}$ where $\left\|y_{2}\right\|_{\infty} \leq \beta$ and $\left\|y_{1}\right\|^{2} \leq$ $4 q \beta^{2} \log q$.

At this point, we can choose $q=\frac{s \log ^{2} m}{m \beta^{2}}=4 s$ so the above event happens with probability at least $1-e^{-\Omega\left(\beta^{2} q m\right)} \geq 1-e^{-\Omega\left(s \log ^{2} m\right)}$. Finally, applying standard $\epsilon$-net argument over all unit vectors $x \in \mathbb{R}^{n}$, we have for each such $x \in \mathbb{R}^{n}$, we can decompose $y=W x$ into $y=y_{1}+y_{2}$ where

$$
\left\|y_{2}\right\|_{\infty} \leq 2 \beta=\frac{\log m}{\sqrt{m}} \text { and }\left\|y_{1}\right\|^{2} \leq \frac{8 s \log ^{3} m}{m}
$$

### 6.10.6 McDiarmid's Inequality and An Extension

We state the standard McDiarmid's inequality,

Lemma 6.10.17 (McDiarmid's inequality). Consider independent random variables $x_{1}, \cdots, x_{n} \in$ $X$ and a mapping $f: X^{n} \rightarrow \mathbb{R}$. If for all $i \in[n]$ and for all $y_{1}, \cdots, y_{n}, y_{i}^{\prime} \in X$, the function $f$ satisfies

$$
\left|f\left(y_{1}, \cdots, y_{i-1}, y_{i}, y_{i+1}, \cdots, y_{n}\right)-f\left(y_{1}, \cdots, y_{i-1}, y_{i}^{\prime}, y_{i+1}, \cdots, y_{n}\right)\right| \leq c_{i}
$$

Then

$$
\begin{array}{r}
\operatorname{Pr}\left[f\left(x_{1}, \cdots, x_{n}\right)-\mathbb{E} f \geq t\right] \geq \exp \left(\frac{-2 t^{2}}{\sum_{i=1}^{n} c_{i}^{2}}\right), \\
\operatorname{Pr}\left[f\left(x_{1}, \cdots, x_{n}\right)-\mathbb{E} f \leq-t\right] \geq \exp \left(\frac{2 t^{2}}{\sum_{i=1}^{n} c_{i}^{2}}\right) .
\end{array}
$$

We prove a more general version of McDiarmid's inequality,

Lemma 6.10.18 (McDiarmid extension). Let $w_{1}, \ldots, w_{N}$ be independent random variables and
$f:\left(w_{1}, \ldots, w_{N}\right) \mapsto[0,1]$. Suppose it satisfies:

- $\mathbb{E}_{w_{1}, \ldots, w_{N}}\left[f\left(w_{1}, \ldots, w_{N}\right)\right] \geq \mu$, and
- With probability at least $1-p$ over $w_{1}, \ldots, w_{N}$, it satisfies

$$
\begin{equation*}
\forall k \in[N], \forall w_{k}^{\prime \prime}:\left|f\left(w_{-k}, w_{k}\right)-f\left(w_{-k}, w_{k}^{\prime \prime}\right)\right| \leq c \tag{6.24}
\end{equation*}
$$

Then, $\operatorname{Pr}\left[f\left(w_{1}, \ldots, w_{N}\right) \geq \mu / 2\right] \geq 1-N^{2} \sqrt{p}-e^{\Omega\left(\frac{-\mu^{2}}{N\left(c^{2}+p\right)}\right)}$.

Proof of Lemma 6.10.18. For each $t \in[N]$, we have with probability at least $1-\sqrt{p}$ over $w_{1}, \ldots, w_{t}$, it satisfies

$$
\operatorname{Pr}_{w_{t+1}, \ldots, w_{N}}\left[\forall k \in[N], \forall w_{k}^{\prime \prime}:\left|f\left(w_{-k}, w_{k}\right)-f\left(w_{-k}, w_{k}^{\prime \prime}\right)\right| \leq c\right] \geq 1-\sqrt{p} .
$$

Define those $\left(w_{1}, \ldots, w_{t}\right)$ satisfying the above event to be $K_{\leq t}$.
Define random variable $X_{t}$ (which depends only on $w_{1}, \ldots, w_{t}$ ) as

$$
X_{t}:=\underset{w_{>t}}{\mathbb{E}}\left[f(\vec{w}) \mid w_{\leq t}\right] \mathbb{1}_{\left(w_{\leq 1}, \ldots, w_{\leq t}\right) \in K_{\leq 1} \times \cdots \times K_{\leq t}}+N\left(1-\mathbb{1}_{\left(w_{\leq 1}, \ldots, w_{\leq t}\right) \in K_{\leq 1} \times \cdots \times K_{\leq t}}\right)
$$

For every $t$ and fixed $w_{1}, \ldots, w_{t-1}$.

- If $\left(w_{\leq 1}, \ldots, w_{<t}\right) \notin K_{\leq 1} \times \cdots \times K_{<t}$, then $X_{t}=X_{t-1}=N$.
- If $\left(w_{\leq 1}, \ldots, w_{<t}\right) \in K_{\leq 1} \times \cdots \times K_{<t}$,
- If $w_{\leq t} \notin K_{\leq t}$, then $X_{t}-X_{t-1}=N-\cdots \geq 0$.
- If $w_{\leq t} \in K_{\leq t}$, then

$$
X_{t}-X_{t-1}=\underset{w>t}{\mathbb{E}}\left[f\left(w_{<t}, w_{t}, w_{>t}\right) \mid w_{\leq t}\right]-\underset{w \geq t}{\mathbb{E}}\left[f\left(w_{<t}, w_{t}, w_{>t}\right) \mid w_{<t}\right]
$$

Recall from our assumption that, with probability at least $1-\sqrt{p}$ over $w_{t}$ and $w_{>t}$, it satisfies

$$
\forall w_{t}^{\prime \prime}:\left|f\left(w_{<t}, w_{t}^{\prime \prime}, w_{>t}\right)-f\left(w_{<t}, w_{t}, w_{>t}\right)\right| \leq c
$$

Taking expectation over $w_{t}$ and $w_{>t}$, we have

$$
\forall w_{t}^{\prime \prime}: \underset{w_{>t}}{\mathbb{E}}\left[f\left(w_{<t}, w_{t}^{\prime \prime}, w_{>t}\right)\right]-\underset{w_{\geq t}}{\mathbb{E}}\left[f\left(w_{<t}, w_{t}, w_{>t}\right)\right] \geq-(c+\sqrt{p})
$$

This precisely means $X_{t}-X_{t-1} \geq c+\sqrt{p}$.

In sum, we have just shown that $X_{t}-X_{t-1} \geq-(c+\sqrt{p})$ always holds. By applying martingale concentration (with one-sided bound),

$$
\operatorname{Pr}\left[X_{N}-X_{0} \leq-t\right] \leq \exp \left(\frac{-t^{2}}{N(c+\sqrt{p})^{2}}\right)
$$

Notice that $X_{0}=\mu$ so if we choose $t=\mu / 2$, we have

$$
\operatorname{Pr}\left[X_{N} \geq \mu / 2\right] \geq 1-\exp \left(\frac{-\mu^{2}}{N(c+\sqrt{p})^{2}}\right)
$$

Recalling

$$
X_{N}:=f(\vec{w}) \mathbb{1}_{\left(w_{\leq 1}, \ldots, w_{\leq t}\right) \in K_{\leq 1} \times \cdots \times K_{\leq t}}+N\left(1-\mathbb{1}_{\left.\left(w_{\leq 1}, \ldots, w_{\leq t}\right) \in K_{\leq 1} \times \cdots \times K_{\leq t}\right)}\right)
$$

and we have $X_{N}=f\left(w_{1}, \ldots, w_{N}\right)$ with probability at least $1-N \sqrt{p}$ (and $X_{N}=N$ with the remaining probabilities). Together, we have the desired theorem.

### 6.11 Basic Properties at Random Initialization

Recall that the recursive update equation of RNN can be described as follows

$$
\begin{array}{ll}
h_{i, 0}=0 & \forall i \in[n] \\
h_{i, \ell}=\phi\left(W \cdot h_{i, \ell-1}+A x_{i, \ell}\right) & \forall i \in[n], \forall \ell \in[L] \\
y_{i, \ell}=B \cdot h_{i, \ell} & \\
\forall i \in[n], \forall \ell \in[L]
\end{array}
$$

Throughout this section, we assume that matrices $W \in \mathbb{R}^{m \times m}, A \in \mathbb{R}^{m \times d_{x}}, B \in \mathbb{R}^{d \times m}$ are at their random initialization position: each entry of $W$ and $A$ is sampled i.i.d. from $\mathcal{N}\left(0, \frac{2}{m}\right)$ and each entry of $B$ is sampled i.i.d. from $\mathcal{N}\left(0, \frac{1}{d}\right)$. We recall

Definition 6.11.1. For each $i \in[n]$ and $\ell \in[L]$, let $D_{i, \ell} \in \mathbb{R}^{m \times m}$ be the diagonal matrix where

$$
\left(D_{i, \ell}\right)_{k, k}=\mathbf{1}_{\left(W \cdot h_{i, \ell-1}+A x_{i, \ell}\right)_{k} \geq 0}=\mathbf{1}_{\left(g_{i, \ell}\right)_{k} \geq 0} .
$$

As a result, we can write $h_{i, \ell}=D_{i, \ell} W h_{i, \ell-1}$.

We introduce two notations that shall repeatedly appear in our proofs.

Definition 6.11.2 $\left(U_{\ell}\right)$. Let $U_{\ell} \in \mathbb{R}^{m \times n \ell}$ denote the column orthonormal matrix using Gram-Schmidt

$$
U_{\ell} \stackrel{\text { def }}{=} \mathrm{GS}\left(h_{1,1}, \ldots, h_{n, 1}, h_{1,2}, \ldots, h_{n, 2}, \ldots, h_{1, \ell}, \ldots, h_{n, \ell}\right)
$$

Definition 6.11.3 $\left(v_{i, \ell}\right)$. For each $i \in[n], \ell \in[L]$, we define vector $v_{i, \ell} \in \mathbb{R}^{m}$ as

$$
v_{i, \ell}=\frac{\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}}{\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}\right\|_{2}}
$$

## Roadmap

- Section 6.11.1 proves that the forward propagation $\left\|h_{i, \ell}\right\|_{2}$ neither vanishes nor explodes.
- Section 6.11.2 gives a lower bound on the projected forward propagation $\left\|\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{i, \ell+1}\right\|_{2}$.
- Section 6.11.3 proves that for two data points $i, j \in[n]$, the projected forward propagation $\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{i, \ell+1}$ and $\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{j, \ell+1}$ are separable from either other.
- Section 6.11.4 and Section 6.11.5 prove that the consecutive intermediate layers, in terms of spectral norm, do not explode (for full and sparse vectors respectively).
- Section 6.11.6 proves that the backward propagation does not explode.


### 6.11.1 Forward Propagation

Our first result of this section is on showing upper and lower bounds on the forward propagation.

Lemma 6.11 .1 (c.f. (6.2)). With probability at least $1-\exp \left(-\Omega\left(m / L^{2}\right)\right)$ over the random initialization $W, A$ (see Definition 6.2.2), it satisfies

$$
\begin{array}{r}
\forall i \in[n], \quad \forall \ell \in\{0,1, \ldots, L-1\} \quad: \quad(1-1 /(4 L))^{\ell} \leq\left\|h_{i, \ell+1}\right\|_{2} \leq 2 \ell+4 \\
\left\|g_{i, \ell+1}\right\|_{2} \leq 4 \ell+8
\end{array}
$$

We prove Lemma 6.11 .1 by induction on $\ell \in[L]$. We only prove the $h_{i, \ell+1}$ part and the $g_{i, \ell+1}$ part is completely analogous.

For the base case $\ell=0$, we have $\left\|x_{i, 1}\right\|_{2}=1$ and $h_{i, 1}=\phi\left(A x_{i, 1}\right)$. Using Corollary 6.10 .10 we have $1-1 / 4 L \leq\left\|h_{i, 1}\right\|_{2} \leq 1+1 / 4 L$ with probability at least $1-\exp \left(-\Omega\left(-m / L^{2}\right)\right)$. We proceed the proof for the case of $\ell \geq 1$, assuming that Lemma 6.11.1 already holds for $0,1, \ldots, \ell-1$. We first show

Claim 6.11.2. With probability at least $1-\exp \left(-\Omega\left(m / L^{2}\right)\right)$ over $W$ and $A$,

$$
\forall i \in[n]:\left\|h_{i, \ell+1}\right\|_{2} \leq(1+1 / L)\left(\left\|h_{i, \ell}\right\|_{2}+1\right)
$$

Note if Claim 6.11.2 holds for all $\ell=0,1, \ldots, \ell-1$, then we must have $\left\|h_{i, \ell}\right\|_{2} \leq$ $(1+1 / L)^{\ell}\left\|h_{i, 0}\right\|_{2}+\sum_{i=1}^{\ell}(1+1 / L)^{i}=(1+1 / L)^{\ell}+\sum_{i=1}^{\ell}(1+1 / L)^{i} \leq 2 \ell+4$.

Proof of Claim 6.11.2. Recall the definition of $h_{i, \ell} \in \mathbb{R}^{m}$, we have

$$
\left\|h_{i, \ell+1}\right\|_{2}=\left\|\phi\left(W \cdot h_{i, \ell}+A x_{i, \ell+1}\right)\right\|_{2}
$$

We can rewrite vector $W h_{i, \ell} \in \mathbb{R}^{m}$ as follows

$$
\begin{aligned}
W h_{i, \ell} & =W U_{\ell-1} U_{\ell-1}^{\top} h_{i, \ell}+W\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell} \\
& =W U_{\ell-1} U_{\ell-1}^{\top} h_{i, \ell}+W \frac{\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}}{\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}\right\|_{2}} \cdot\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}\right\|_{2} \\
& =W U_{\ell-1} U_{\ell-1}^{\top} h_{i, \ell}+W v_{i, \ell} \cdot\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}\right\|,
\end{aligned}
$$

where the last step follows by definition of $v_{i, \ell}$ (See Definition 6.11.3). Define $z_{1} \in \mathbb{R}^{n(\ell-1)}, z_{2} \in$ $\mathbb{R}, z_{3} \in \mathbb{R}^{d}$ as follows

$$
\begin{equation*}
z_{1}=U_{\ell-1}^{\top} h_{i, \ell}, \quad z_{2}=\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}\right\|_{2}, \quad z_{3}=x_{i, \ell+1} . \tag{6.25}
\end{equation*}
$$

Then

$$
\begin{align*}
\left\|z_{1}\right\|_{2}^{2}+z_{2}^{2}+\left\|z_{3}\right\|_{2}^{2} & =\left\|U_{\ell-1}^{\top} h_{i, \ell}\right\|_{2}^{2}+\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}\right\|_{2}^{2}+\left\|x_{i, \ell+1}\right\|_{2}^{2} \\
& =\left\|h_{i, \ell}\right\|_{2}^{2}+\left\|x_{i, \ell+1}\right\|_{2}^{2} \leq\left\|h_{i, \ell}\right\|_{2}^{2}+1 . \tag{6.26}
\end{align*}
$$

We can thus rewrite

$$
\begin{aligned}
W h_{i, \ell}+A x_{i, \ell+1} & =W U_{\ell-1} z_{1}+W v_{i, \ell} z_{2}+A z_{3} \\
& =\left[\begin{array}{lll}
M_{1} & M_{2} & M_{3}
\end{array}\right] \cdot\left[\begin{array}{l}
z_{1} \\
z_{2} \\
z_{3}
\end{array}\right] \\
& =M \cdot z
\end{aligned}
$$

where the last step follows by defining $M \in \mathbb{R}^{m \times(n(\ell-1)+1+d)}$ as follows,

$$
M=\left[\begin{array}{lll}
M_{1} & M_{2} & M_{3}
\end{array}\right]=\left[\begin{array}{lll}
W U_{\ell-1} & W v_{i, \ell} & A \tag{6.27}
\end{array}\right]
$$

We stress here that the entries of $M_{1}, M_{2}, M_{3}$ are i.i.d. from $\mathcal{N}\left(0, \frac{2}{m}\right) .{ }^{13}$ We have $\left\|h_{i, \ell+1}\right\|=$ $\left\|\phi\left(W h_{i, \ell}+A x_{i, \ell+1}\right)\right\|=\|\phi(M \cdot z)\|$.

Next, applying Corollary 6.10.10, we know that if $z_{1}, z_{2}, z_{3}$ are fixed (instead of defined as in (6.25)), then, letting Choosing $\epsilon=1 / 2 L$, we have

$$
\operatorname{Pr}_{M}\left[\left\|h_{i, \ell+1}\right\| \leq(1+\epsilon) \cdot \sqrt{\left\|z_{1}\right\|_{2}^{2}+z_{2}^{2}+\left\|z_{3}\right\|_{2}^{2}}\right] \geq 1-\exp \left(-\Omega\left(\epsilon^{2} m\right)\right)
$$

To move from fixed choices of $z_{1}$ and $z_{2}$ to all choices of $z_{1}$ and $z_{2}$, we perform a standard $\epsilon$-net argument. Since the dimension of $z_{1}$ and $z_{2}$ are respectively $n(\ell-1)$ and 1 , the size of $\epsilon$-net for $z_{1}$ and $z_{2}$ is at most $e^{O(n L \log L)}$. Thus, with probability at least $1-e^{O(n L \log L+\log n)}$. $\exp \left(-\Omega\left(\epsilon^{2} m\right)\right) \geq 1-\exp \left(-\Omega\left(\epsilon^{2} m\right)\right)$ we have: for all $z_{1} \in \mathbb{R}^{n(\ell-1)}$ and $z_{2} \in \mathbb{R}$ satisfying $\left\|z_{1}\right\| \leq 2 L+4$ and $0 \leq z_{2} \leq 2 L+4$, and fixed $z_{3}=x_{i, \ell+1}$,

$$
\left\|h_{i, \ell+1}\right\| \leq(1+2 \epsilon) \cdot \sqrt{\left\|z_{1}\right\|_{2}^{2}+z_{2}^{2}+\left\|z_{3}\right\|_{2}^{2}}
$$

In particular, since we have "for all" quantifies on $z_{1}$ and $z_{2}$ above, we can substitute the choice of $z_{1}$ and $z_{2}$ in (6.25) (which may depend on the randomness of $W$ and $A$ ). This, together with (6.26), gives

$$
\operatorname{Pr}\left[\left\|h_{\ell+1}\right\|_{2} \leq(1+\epsilon) \cdot \sqrt{\left\|h_{i, \ell}\right\|_{2}^{2}+1}\right] \geq 1-\exp \left(-\Omega\left(\epsilon^{2} m\right)\right)
$$

Similarly, we can prove a lower bound

[^40]Claim 6.11.3. With probability at least $1-\exp \left(-\Omega\left(m / L^{2}\right)\right)$ over $W$ and $A$,

$$
\left\|h_{i, \ell+1}\right\|_{2} \geq\left(1-\frac{1}{4 L}\right)\left\|h_{i, \ell}\right\|_{2} .
$$

Proof. We can define $z_{1}, z_{2}, z_{3}$ in the same way as (6.25). This time, we show a lower bound

$$
\begin{equation*}
\left\|z_{1}\right\|_{2}^{2}+z_{2}^{2}+\left\|z_{3}\right\|_{2}^{2} \geq\left\|z_{1}\right\|_{2}^{2}+z_{2}^{2}=\left\|h_{i, \ell}\right\|_{2}^{2} \tag{6.28}
\end{equation*}
$$

Applying Corollary 6.10.10, we know if $z_{1}, z_{2}, z_{3}$ are fixed (instead of defined as in (6.25)), then choosing $\epsilon=1 / 8 L$,

$$
\operatorname{Pr}\left[\left\|h_{i, \ell+1}\right\| \geq(1-\epsilon) \cdot \sqrt{\left\|z_{1}\right\|_{2}^{2}+z_{2}^{2}+\left\|z_{3}\right\|_{2}^{2}}\right] \geq 1-\exp \left(-\Omega\left(\epsilon^{2} m\right)\right) . .
$$

Again, after applying $\epsilon$-net, we know with probability at least $1-e^{O(n L \log L+\log n)} \cdot \exp \left(-\Omega\left(\epsilon^{2} m\right)\right) \geq$ $1-\exp \left(-\Omega\left(\epsilon^{2} m\right)\right)$, for all $z_{1} \in \mathbb{R}^{n(\ell-1)}$ and $z_{2} \in \mathbb{R}$ satisfying $\left\|z_{1}\right\| \leq 2 L+4$ and $0 \leq z_{2} \leq$ $2 L+4$, and fixed $z_{3}=x_{i, \ell+1}$,

$$
\left\|h_{i, \ell+1}\right\|_{2} \geq(1-\epsilon) \cdot \sqrt{\left\|z_{1}\right\|_{2}^{2}+z_{2}^{2}+\left\|z_{3}\right\|_{2}^{2}}
$$

Substituting the choice of $z_{1}, z_{2}, z_{3}$ in (6.25), and the lower bound (6.28), we have

$$
\operatorname{Pr}\left[\left\|h_{i, \ell+1}\right\| \geq(1-\epsilon) \cdot\left\|h_{i, \ell}\right\|_{2}\right] \geq 1-\exp \left(-\Omega\left(\epsilon^{2} m\right)\right)
$$

Choosing $\epsilon=1 /(4 L)$ gives the desired statement.

Finally, recursively applying Claim 6.11.2 and Claim 6.11.3 for all $\ell=0,1, \ldots, L-1$, we have with probability at least $1-L^{2} \exp \left(-\Omega\left(m / L^{2}\right)\right) \geq 1-\exp \left(-\Omega\left(m / L^{2}\right)\right)$, we have

$$
\left\|h_{i, \ell}\right\|_{2} \geq\left(1-\frac{1}{4 L}\right)^{\ell}, \quad\left\|h_{i, \ell}\right\|_{2} \leq 2 \ell+4
$$

This finishes the proof of Lemma 6.11.1.

### 6.11.2 Forward Correlation

This subsection proves the following lemma which, as discussed in Section 6.5, bounds how much "fresh new randomness" is left after propagating to layer $\ell$.

Lemma 6.11.4 (c.f. (6.3)). With probability at least $1-e^{-\Omega(\sqrt{m})}$ over the random initialization $W, A$ in Definition 6.2.2, letting $U_{\ell}$ be defined in Definition 6.11.2, we have

$$
\forall i \in[n], \forall \ell \in\{0,1, \cdots, L-1\} \quad: \quad\left\|\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{i, \ell+1}\right\|_{2} \geq \frac{1}{2 \cdot 10^{6} L^{2} \log ^{3} m}
$$

To prove Lemma 6.11.4, we inductively (with the increasing order of $\ell$ ) show for each $i \in[n]$, for each $\ell \in\{0,1, \cdots, L-1\}$, we have

$$
\begin{equation*}
\left\|\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{i, \ell+1}\right\|_{2} \geq \xi_{\ell} \stackrel{\text { def }}{=} \frac{1}{10^{6} L^{2} \log ^{3} m}\left(1-2 \alpha-\frac{1}{4 L}\right)^{\ell} \quad \text { where } \quad \alpha \stackrel{\text { def }}{=} \frac{1}{2 \cdot 10^{4} L^{2} \log ^{2} m} \tag{6.29}
\end{equation*}
$$

We first show (6.29) in the base case $\ell=0$. Since $U_{0}$ is an empty matrix, we have $\left(I-U_{0} U_{0}^{\top}\right) h_{i, 1}=h_{i, 1}$ so according to Lemma 6.11.1 we have $\left\|h_{i, 1}\right\|_{2} \geq(1-1 / 4 L)$.

The remainder of the proof assumes (6.29) already holds for $\ell-1$. We can write

$$
h_{i, \ell+1}=\phi\left(W h_{i, \ell}+A x_{i, \ell+1}\right)=\phi\left(M_{1} z_{1}+M_{2} z_{2}+M_{3} z_{3}\right) .
$$

where $M_{1} \in \mathbb{R}^{m \times n(\ell-1)}, M_{2} \in \mathbb{R}^{m \times 1}, M_{3} \in \mathbb{R}^{m \times d}$ and $z_{1} \in \mathbb{R}^{n(\ell-1)}, z_{2} \in \mathbb{R}, z_{3} \in \mathbb{R}^{d}$ are defined as

$$
\begin{align*}
M_{1} & =W U_{\ell-1} & M_{2} & =W v_{i, \ell}
\end{align*} M_{3}=A
$$

in the same way as (6.25) and (6.27) as in the proof of Lemma 6.11.1. We again have the entries of $M_{1}, M_{2}, M_{3}$ are i.i.d. from $\mathcal{N}\left(0, \frac{2}{m}\right)$ (recall Footnote 13). For the quantity $M_{2} z_{2}$, we can further decompose it as follows

$$
\begin{equation*}
M z_{2}=\nu \cdot\left(z_{2}-c_{5} \alpha\right)_{+}+\nu^{\prime} z_{2}^{\prime} \tag{6.31}
\end{equation*}
$$

where $c_{5}=\frac{1}{16 \log m}$ is some fixed parameter, $\nu$ and $\nu^{\prime}$ denote two vectors that are independently generated from $\mathcal{N}\left(0, \frac{2 I}{m}\right)$, and

$$
\left(z_{2}-c_{5} \alpha\right)_{+}=\left\{\begin{array}{ll}
0, & \text { if } z_{2}<c_{5} \alpha ; \\
\sqrt{z_{2}^{2}-c_{5}^{2} \alpha^{2}}, & \text { if } z_{2} \geq c_{5} \alpha
\end{array} \quad z_{2}^{\prime}= \begin{cases}z_{2}, & \text { if } z_{2}<c_{5} \alpha \\
c_{5} \alpha, & \text { if } z_{2} \geq c_{5} \alpha\end{cases}\right.
$$

It is clear that the two sides of (6.31) are identical in distribution (because $M_{2} \sim \mathcal{N}\left(0, \frac{2 I}{m}\right)$ and $\left.\left(z_{2}-c_{5} \alpha\right)_{+}^{2}+\left(z_{2}^{\prime}\right)^{2}=z_{2}^{2}\right)$.

Next, suppose $z_{1}$ and $z_{2}$ are fixed (instead of depending on the randomness of $W$ and A) and satisfies ${ }^{14}$

$$
\begin{equation*}
\frac{1}{\sqrt{2}} \leq\left\|z_{1}\right\| \leq 2 L+6 \text { and }\left|z_{2}\right| \leq 2 L+6 \tag{6.32}
\end{equation*}
$$

We can apply Corollary 6.10 .15 to obtain the following statement: with probability at least $1-\exp \left(-\Omega\left(\alpha^{2} m\right)\right)$,

$$
w \stackrel{\text { def }}{=} M_{1} z_{1}+\nu\left(z_{2}-c_{5} \alpha\right)_{+}+M_{3} z_{3}
$$

is $(\alpha, \sigma / 4)$-good where $\sigma=\left(\frac{2}{m}\left\|z_{1}\right\|_{2}^{2}+\frac{2}{m}\left(\left(z_{2}-c_{5} \alpha\right)_{+}\right)^{2}+\frac{2}{m}\left\|z_{3}\right\|_{2}^{2}\right)^{1 / 2}$. Using $\left\|z_{1}\right\|_{2}^{2} \geq 1 / 2$, we can lower bound $\sigma^{2}$ as

$$
\sigma^{2} \geq \frac{2}{m}\left(\left\|z_{1}\right\|_{2}^{2}+z_{2}^{2}-c_{5}^{2} \alpha^{2}\right) \geq \frac{2}{m}\left(\frac{1}{2}-c_{5}^{2} \alpha^{2}\right) \geq \frac{1}{2 m}
$$

[^41]In other words, $w$ is $\left(\alpha, \frac{1}{4 \sqrt{2} \sqrt{m}}\right)$-good. Next, applying standard $\epsilon$-net argument, we have with probability at least $1-\exp \left(-\Omega\left(\alpha^{2} m\right)\right)$, for all vectors $z_{1} \in \mathbb{R}^{n(\ell-1)}$ and $z_{2} \in \mathbb{R}$ satisfying (6.32), it satisfies $w$ is $\left(\alpha, \frac{1}{8 \sqrt{m}}\right)$-good. This allows us to plug in the random choice of $z_{1}$ and $z_{2}$ in (6.30).

We next apply Lemma 6.11 .5 with

$$
w=M_{1} z_{1}+\nu\left(z_{2}-c_{5} \alpha\right)_{+}+M_{3} z_{3}, \quad r=z_{2}^{\prime}, \quad v=\nu, \quad U=U_{\ell} .
$$

(We can do so because the randomness of $v$ is independent of the randomness of $U_{\ell}$ and $w$.) Lemma 6.11 .5 tells us that, with probability at least $1-\exp (-\Omega(\sqrt{m}))$ over the randomness of $v$,

$$
\left\|\left(I-U_{\ell} U_{\ell}^{\top}\right) \cdot \phi(w+r v)\right\| \geq r(1-2 \alpha)-\frac{\alpha^{1.5}}{4}
$$

By induction hypothesis, we know

$$
r=z_{2}^{\prime} \geq \min \left\{z_{2}, c_{5} \alpha\right\} \geq \min \left\{\xi_{\ell-1}, c_{5} \alpha\right\} \geq \xi_{\ell-1}
$$

where the last step follows by $\xi_{\ell-1} \leq c_{5} \alpha$. Thus, we have
$\left\|\left(I-U_{\ell} U_{\ell}^{\top}\right) \cdot \phi(w+r v)\right\| \geq r(1-2 \alpha)-\frac{\alpha^{1.5}}{4} \geq \xi_{\ell-1}(1-2 \alpha)-\frac{\alpha^{1.5}}{4} \geq \xi_{\ell-1}\left(1-2 \alpha-\frac{1}{4 L}\right)=\xi_{\ell}$, where the last inequality uses $\alpha^{1.5} \leq \frac{\xi_{\ell-1}}{L}$.

### 6.11.2.1 Tools

Lemma 6.11.5. Suppose $\alpha \in\left[\frac{1}{100 L^{4}}, 1 / 2\right), m \geq 4 \widetilde{d} / \alpha$ and $r \in\left(0, \frac{\alpha}{16 \log m}\right]$. Suppose $w$ is a fixed vector that is $\left(\alpha, \frac{1}{8 \sqrt{m}}\right)$-good (see Definition 6.10.1), and $U \in \mathbb{R}^{m \times \widetilde{d}}$ is a fixed column
orthonormal matrix. Then, if $v \sim \mathcal{N}\left(0, \frac{2}{m} I\right)$, with probability at least $1-\exp (-\Omega(\sqrt{m}))$, it satisfies

$$
\left\|\left(I-U U^{\top}\right) \cdot \phi(w+r v)\right\| \geq r(1-2 \alpha)-\frac{\alpha^{1.5}}{4}
$$

Proof. We split $w \in \mathbb{R}^{m}$ into three pieces $w=\left(w_{1}, w_{2}, w_{3}\right) \in \mathbb{R}^{m}$ with disjoint support such that:

- $w_{1}$ corresponds to the coordinates that are $\geq \alpha / 8 \sqrt{m}$,
- $w_{2}$ is the remaining, which corresponds to the coordinates that are within $(-\alpha / 8 \sqrt{m}, \alpha / 8 \sqrt{m})$.
- $w_{3}$ corresponds to the coordinates that are $\leq-\alpha / 8 \sqrt{m}$, and

We write $v=\left(v_{1}, v_{2}, v_{3}\right)$ according to the same partition. We consider the following three cases.

- For each index $k$ in the first block, we have $\left(\phi\left(w_{1}+r v_{1}\right)\right)_{k} \neq\left(w_{1}+r v_{1}\right)_{k}$ only if $\left(r v_{1}\right)_{k} \leq-\alpha / 8 \sqrt{m}$. However, if this happens, we have

$$
0 \leq\left(\phi\left(w_{1}+r v_{1}\right)\right)_{k}-\left(w_{1}+r v_{1}\right)_{k} \leq \max \left\{\left(r v_{1}\right)_{k}-\alpha / 8 \sqrt{m}, 0\right\}
$$

Applying Lemma 6.10.5, we know with probability at least $1-e^{-\Omega(\sqrt{m})}$,

$$
\delta_{1} \stackrel{\text { def }}{=} \phi\left(w_{1}+r v_{1}\right)-\left(w_{1}+r v_{1}\right) \quad \text { satisfies } \quad\left\|\delta_{1}\right\|^{2} \leq 2 \sqrt{m} \frac{2 r^{2}}{m} \leq 4 \alpha^{2} / \sqrt{m}
$$

- Similarly, for each index $k$ in the third block, we have $\left(\phi\left(w_{1}+r v_{1}\right)\right)_{k} \neq 0$ only if $\left(r v_{1}\right)_{k} \geq \alpha / 8 \sqrt{m}$. Therefore, we can similarly derive that

$$
\delta_{3} \stackrel{\text { def }}{=} \phi\left(w_{3}+r v_{3}\right) \quad \text { satisfies } \quad\left\|\delta_{3}\right\|^{2} \leq 4 \alpha^{2} / \sqrt{m} .
$$

- For the second block, we claim that

$$
\begin{equation*}
\delta_{2} \stackrel{\text { def }}{=} \phi\left(w_{2}+r v_{2}\right) \quad \text { satisfies } \quad\left\|\delta_{2}\right\|_{2} \leq \alpha^{3 / 2} / 4 \tag{6.33}
\end{equation*}
$$

To prove (6.33), we use triangle inequality,

$$
\left\|\phi\left(w_{2}+r v_{2}\right)\right\|_{2} \leq\left\|\phi\left(w_{2}\right)\right\|_{2}+\left\|\phi\left(r v_{2}\right)\right\|_{2}
$$

Since $\left\|\phi\left(w_{2}\right)\right\|_{\infty} \leq \alpha / 8 \sqrt{m}$ and the size of support of $\phi\left(w_{2}\right)$ is at most $\alpha m$, we have

$$
\left\|\phi\left(w_{2}\right)\right\|_{2} \leq\left((\alpha / 8 \sqrt{m})^{2} \alpha m\right)^{1 / 2} \leq \frac{\alpha^{3 / 2}}{8}
$$

Since $v \sim \mathcal{N}\left(0, \frac{2}{m} I\right)$, and since the size of support of $v_{2}$ is at most $\alpha m$, we have $\left\|v_{2}\right\| \leq 2 \sqrt{\alpha}$ with probability at least $1-e^{-\Omega(\alpha m)}$ (due to chi-square distribution concentration). Thus

$$
\left\|\phi\left(r v_{2}\right)\right\|_{2}=r \cdot\left\|\phi\left(v_{2}\right)\right\|_{2} \leq \frac{\alpha}{16} \cdot 2 \sqrt{\alpha}=\frac{\alpha^{3 / 2}}{8} .
$$

Together, by triangle inequality we have $\left\|\phi\left(w_{2}+r v_{2}\right)\right\| \leq \frac{\alpha^{3 / 2}}{4}$. This finishes the proof of (6.33).

Denoting by $\delta=\left(\delta_{1}, \delta_{2}, \delta_{3}\right)$, we have

$$
\phi(w+r v)=w_{1}+r v_{1}+\delta
$$

Taking the norm on both sides,

$$
\begin{aligned}
\left\|\left(I-U U^{\top}\right) \phi(w+r v)\right\| & =\left\|\left(I-U U^{\top}\right)\left(w_{1}+r v_{1}+\delta\right)\right\| \\
& \geq\left\|\left(I-U U^{\top}\right)\left(w_{1}+r v_{1}\right)\right\|-\left\|\left(I-U U^{\top}\right) \delta\right\| \\
& \left.\geq\left\|\left(I-U U^{\top}\right)\left(w_{1}+r v_{1}\right)\right\|-\|\delta\| \quad \quad \quad \text { (by triangle inequality) }\left\|\left(I-U U^{\top}\right) \delta\right\| \leq\|\delta\|\right) \\
& \geq\left\|\left(I-U U^{\top}\right)\left(w_{1}+r v_{1}\right)\right\|-\frac{\alpha^{3 / 2}}{4}-\frac{4 \alpha}{m^{1 / 4}} \\
& \geq\left\|\left(I-U U^{\top}\right)\left(w_{1}+r v_{1}\right)\right\|-\frac{\alpha^{3 / 2}}{2} \\
& \geq\left\|\left(I-U U^{\top}\right) w_{1}+r v_{1}\right\|-r\left\|U^{\top} v_{1}\right\|-\frac{\alpha^{3 / 2}}{2} .
\end{aligned}
$$

Now, since $U \in \mathbb{R}^{m \times \widetilde{d}}$, it is not hard to show that $\left\|U^{\top} v_{1}\right\|_{2}^{2} \leq \frac{2 \widetilde{d}}{m}$ with probability at least $1-\exp (-\Omega(m))$. Using $\frac{2 \widetilde{d}}{m} \leq \frac{\alpha}{2}$ (owing to our assumption $m \geq 4 \widetilde{d} / \alpha$ ), we have $r\left\|U^{\top} v_{1}\right\| \leq \frac{r \alpha}{2}$.

On the other hand, the random vector $z=\left(I-U U^{\top}\right) w_{1}+r v_{1}$ follows from distribution $\mathcal{N}\left(\mu, \frac{2 r^{2}}{m}\right)$ for some fixed vector $\mu=\left(I-U U^{\top}\right) w_{1}$ and has at least $\frac{m}{2}(1-\alpha)$ dimensions. By chi-square concentration, we have $\|z\| \geq r(1-3 \alpha / 2)$ with probability at least $1-e^{-\Omega\left(\alpha^{2} m\right)}$. Putting these together, we have

$$
\left\|\left(I-U U^{\top}\right) \phi(w+r v)\right\| \geq r(1-3 \alpha / 2)-\frac{r \alpha}{2}-\frac{\alpha^{3 / 2}}{2}=r(1-2 \alpha)-\frac{\alpha^{3 / 2}}{2}
$$

### 6.11.3 Forward $\delta$-Separateness

We first give the definition of $\delta$-Separable,

Definition 6.11.4 ( $\delta$-separable vectors). For any two vectors $x, y$, we say $x$ and $y$ are $\delta$-separable if

$$
\left\|\left(I-\frac{y y^{\top}}{\|y\|_{2}^{2}}\right) x\right\| \geq \delta \quad \text { and } \quad\left\|\left(I-\frac{x x^{\top}}{\|x\|_{2}^{2}}\right) y\right\| \geq \delta \ldots
$$

We say a finite set $X$ is $\delta$-separable if for any two vectors $x, y \in X, x$ and $y$ are $\delta$-separable.

The goal of this subsection is to prove the $\delta$-separateness over all layers $\ell$.

Lemma 6.11 .6 (c.f. (6.4)). Let $\left\{x_{i, 1}\right\}_{i \in[n]}$ be $\delta$-separable with $\delta \leq \frac{1}{10^{6} L^{2} \log ^{3} m}$. With probability at least $1-e^{-\Omega(\sqrt{m})}$, for all $\ell=0,1, \ldots, L-1$, for all $i, j \in[n]$ with $i \neq j$, we have

$$
\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{i, \ell+1} \text { and }\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{j, \ell+1} \text { are } \frac{\delta}{2} \text {-separable. }
$$

(Note that since we have assumed $\left\|x_{i, 1}-x_{j, 1}\right\| \geq \delta$ in Assumption 6.2.1 and assumed without loss of generality that $\left(x_{i, 1}\right)_{d_{x}}=\frac{1}{\sqrt{2}}$, it automatically satisfies that $\left\{x_{i, 1}\right\}_{i \in[n]}$ is $O(\delta)$ separable.)

To prove Lemma 6.11.6, we inductively (with the increasing order of $\ell$ ) show for each $i \in[n]$, for each $\ell \in\{0,1, \cdots, L-1\}$, we have

$$
\begin{equation*}
\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{i, \ell+1} \text { and }\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{j, \ell+1} \text { are } \delta_{\ell} \text { separable for } \delta_{\ell} \stackrel{\text { def }}{=} \delta\left(1-2 \alpha-\frac{1}{4 L}\right)^{\ell} \tag{6.34}
\end{equation*}
$$

where $\alpha \stackrel{\text { def }}{=} 16 \delta \log m$.

We skip the base case and only prove (6.34) for $\ell \geq 1$ by assuming (6.34) already holds for $\ell-1 .{ }^{15}$

Claim 6.11.7. For $i \neq j$, if $\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}$ and $\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{j, \ell}$ are $\delta_{\ell-1-\text {-separable, }}$ then letting $U=\left[U_{\ell}, \widehat{h}\right]$ where $\widehat{h}=\frac{\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{j, \ell+1}}{\left\|\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{j, \ell+1}\right\|_{2}}$, we have

$$
\left\|\left(I-U U^{\top}\right) h_{i, \ell+1}\right\| \geq \delta_{\ell}=\delta_{\ell-1}(1-3 \alpha)
$$

holds with probability at least $1-\exp (-\Omega(\sqrt{m}))$.

Since it is easy to verify that

$$
\left\|\left(I-\widehat{h} \widehat{h}^{\top}\right)\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{i, \ell+1}\right\|=\left\|\left(I-U U^{\top}\right) h_{i, \ell+1}\right\|_{2},
$$

Claim 6.11.7 immediately implies (6.34) for layer $\ell$, and thus finishes the proof of Lemma 6.11.6. Therefore, we only need to prove Claim 6.11 .7 below.

Proof of Claim 6.11.7. Let $x=\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{j, \ell}$ and $y=\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}$. If $x$ and $y$ are $\delta_{\ell-1}$-separable, we split $y$ into two parts where $y_{1}$ is parallel to $x$ and $y_{2}$ is orthogonal to $x$

$$
y=y_{1}+y_{2}, \quad y_{1}=\frac{\langle x, y\rangle x}{\|x\|_{2}^{2}}, \quad y_{2}=\left(I-x x^{\top} /\|x\|_{2}^{2}\right) y
$$

This also implies that the randomness in $y_{1}$ is independent of the randomness in $y_{2}$.
It is easy to see that $\left\|y_{1}\right\|_{2}=\frac{\langle x, y\rangle}{\|x\|_{2}}$, so we can rewrite $W y$ as follows

$$
W y=W y_{1}+W y_{2}=W \frac{\langle x, y\rangle x}{\|x\|_{2}^{2}}+W y_{2}=\left(\left\|y_{1}\right\|_{2} /\|x\|_{2}\right) W x+W y_{2} .
$$

[^42]Therefore,

$$
\begin{aligned}
W h_{i, \ell}+A x_{i, \ell+1} & =W U_{\ell-1} U_{\ell-1}^{\top} h_{i, \ell}+W\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}+A x_{i, \ell+1} \\
& =W U_{\ell-1} U_{\ell-1}^{\top} h_{i, \ell}+W y+A x_{i, \ell+1} \quad \quad \quad \quad \text { by definition of } y \text { ) } \\
& =W U_{\ell-1} U_{\ell-1}^{\top} h_{i, \ell}+\left(\left\|y_{1}\right\|_{2} /\|x\|_{2}\right) W x+A x_{i, \ell+1}+W y_{2} \\
& =M_{1} z_{1}+M_{2} z_{2}+M_{3} z_{3}+M_{4} z_{4}
\end{aligned} \quad \text { (by rewritting } W y \text { ) }
$$

where

$$
\begin{align*}
M_{1} & =W U_{\ell-1} & M_{2} & =W \frac{x}{\|x\|_{2}}
\end{align*} \begin{array}{lrl} 
& =A & M_{4}
\end{array}=W \frac{y_{2}}{\left\|y_{2}\right\|_{2}}
$$

and we know the entries of $M_{1}, M_{2}, M_{3}, M_{4}$ are i.i.d. from $\mathcal{N}\left(0, \frac{2}{m}\right)$, owing to a similar treatment as Footnote 13. For the vector $M_{4} z_{4}$, we further rewrite it as

$$
M_{4} z_{4}=\nu \cdot\left(z_{4}-c_{5} \alpha\right)_{+}+\nu^{\prime} z_{4}^{\prime}
$$

where $c_{5}=\frac{1}{16 \log m}$ is a fixed parameter, and $\nu$ and $\nu^{\prime}$ denote two vectors that are independently generated from the same distribution $\mathcal{N}\left(0, \frac{2}{m} I\right)$ as vector $M_{4} \in \mathbb{R}^{m}$, and

$$
\left(z_{4}-c_{5} \alpha\right)_{+}=\left\{\begin{array}{ll}
0, & \text { if } z_{4}<c_{5} \alpha ; \\
\sqrt{z_{4}^{2}-c_{5}^{2} \alpha^{2}}, & \text { if } z_{4} \geq c_{5} \alpha
\end{array} \quad z_{4}^{\prime}= \begin{cases}z_{4}, & \text { if } z_{4}<c_{5} \alpha \\
c_{5} \alpha, & \text { if } z_{4} \geq c_{5} \alpha\end{cases}\right.
$$

Together, we can write

$$
\begin{aligned}
\left(I-U U^{\top}\right) h_{i, \ell+1} & =\left(I-U U^{\top}\right) \phi\left(W h_{i, \ell}+A x_{i, \ell+1}\right) \\
& =\left(I-U U^{\top}\right) \phi\left(M_{1} z_{1}+M_{2} z_{2}+M_{3} z_{3}+M_{4} z_{4}\right) \\
& =\left(I-U U^{\top}\right) \phi\left(M_{1} z_{1}+M_{2} z_{2}+M_{3} z_{3}+\nu \cdot\left(z_{4}-c_{5} \alpha\right)_{+}+\nu^{\prime} z_{4}^{\prime}\right) \\
& =\left(I-U U^{\top}\right) \phi\left(w+\nu^{\prime} z_{4}^{\prime}\right)
\end{aligned}
$$

where the last step follows by defining

$$
w=M_{1} z_{1}+M_{2} z_{2}+M_{3} z_{3}+\nu \cdot\left(z_{4}-c_{5} \alpha\right)_{+}
$$

Our plan is to first use the randomness in $w$ to argue that $w$ is $(\alpha, \gamma / 4 \sqrt{m})$-good. Then we conditioned $w$ is good, and prove that the norm of $\left(I-U U^{\top}\right) \phi\left(w+\nu^{\prime} z_{4}^{\prime}\right)$ is lower bounded (using (6.11.5)).

Now, suppose that $z_{1} \in \mathbb{R}^{n(\ell-1)}, z_{2} \in \mathbb{R}$, and $z_{4} \in \mathbb{R}$ are fixed (instead of computed based on the randomness of $W$ and $A$ ), and satisfies ${ }^{16}$

$$
\begin{equation*}
\frac{1}{\sqrt{2}} \leq\left\|z_{1}\right\| \leq 2 L+6 \text { and }\left|z_{2}\right|,\left|z_{4}\right| \leq 2 L+6 \tag{6.36}
\end{equation*}
$$

we can use Corollary 6.10 .15 to obtain the following statement: $w$ is $(\alpha, \sigma / 4)$-good with probability at least $1-\exp \left(-\Omega\left(\alpha^{2} m\right)\right)$ where $\sigma=\left(\frac{2}{m}\left\|z_{1}\right\|_{2}^{2}+\frac{2}{m} z_{2}^{2}+\frac{2}{m}\left\|z_{3}\right\|_{2}^{2}+\frac{2}{m}\left(z_{4}^{2}-c_{5}^{2} \alpha^{2}\right)\right)^{1 / 2}$. Using $\left\|z_{1}\right\|_{2}^{2} \geq 1 / 2$, we can lower bound $\sigma^{2}$ as

$$
\sigma^{2} \geq \frac{2}{m}\left(\left\|z_{1}\right\|_{2}^{2}+z_{2}^{2}+z_{4}^{2}-c_{5}^{2} \alpha^{2}\right) \geq \frac{2}{m}\left(\frac{1}{2}-c_{5}^{2} \alpha^{2}\right)
$$

In other words, $w$ is $\left(\alpha, \frac{1}{4 \sqrt{2} \sqrt{m}}\right)$-good. Next, applying standard $\epsilon$-net argument, we have with probability at least $1-\exp \left(-\Omega\left(\alpha^{2} m\right)\right.$ ), for all $z_{1}, z_{2}, z_{4}$ satisfying (6.36), it satisfies $w$ is $\left(\alpha, \frac{1}{8 \sqrt{m}}\right)$-good. This allows us to plug in the random choice of $z_{1}, z_{2}, z_{4}$ in (6.35).

We apply Lemma 6.11 .5 with the following setting

$$
U=\left[U_{\ell}, \frac{\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{j, \ell+1}}{\left\|\left(I-U_{\ell} U_{\ell}^{\top}\right) h_{j, \ell+1}\right\|_{2}}\right], \quad w=w, \quad v=\nu^{\prime}, \quad r=z_{4}^{\prime} \in\left[0, c_{5} \alpha\right]
$$

[^43]where $w$ is $(\alpha, \gamma / 8 \sqrt{m})$-good. (We can do so because the randomness of $v$ is independent of the randomness of $U$ and $w$.) Lemma 6.11 .5 tells us that, with probability at least $1-\exp (-\Omega(\sqrt{m}))$ over the randomness of $v$,
$$
\left\|\left(I-U U^{\top}\right) \cdot \phi(w+r v)\right\| \geq r(1-2 \alpha)-\frac{\alpha^{1.5}}{4}
$$

By induction hypothesis, we know

$$
r=z_{4}^{\prime} \geq \min \left\{z_{4}, c_{5} \alpha\right\} \geq \min \left\{\delta_{\ell-1}, c_{5} \alpha\right\} \geq \delta_{\ell-1}
$$

where the last step follows by $\delta_{\ell-1} \leq c_{5} \alpha$. Thus, we have
$\left\|\left(I-U U^{\top}\right) \cdot \phi(w+r v)\right\| \geq r(1-2 \alpha)-\frac{\alpha^{1.5}}{4} \geq \delta_{\ell-1}(1-2 \alpha)-\frac{\alpha^{1.5}}{4} \geq \delta_{\ell-1}\left(1-2 \alpha-\frac{1}{4 L}\right)=\delta_{\ell}$,
where the last inequality uses $\alpha^{1.5} \leq \frac{\delta_{\ell-1}}{L}$.

### 6.11.4 Intermediate Layers: Spectral Norm

The following lemma bounds the spectral norm of (consecutive) intermediate layers.

Lemma 6.11 .8 (c.f. (6.5)). With probability at least $1-\exp \left(-\Omega\left(m / L^{2}\right)\right)$, we have for all $L \geq \ell_{2} \geq \ell_{1} \geq 0$ and $i \in[n]$

$$
\left\|\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} W\right\|_{2} \leq O\left(L^{3}\right)
$$

We start with an important claim whose proof is almost identical to Lemma 6.11.1.

Claim 6.11.9. Given $\ell>b \geq 1$, given $i \in[n]$, and given $z_{b-1} \in \mathbb{R}^{m}$ a fixed vector, letting

$$
z_{\ell}=D_{i, \ell} W D_{i, \ell-1} \cdots D_{i, b} W z_{b-1}
$$

we have with probability at least $1-\exp \left(-\Omega\left(m / L^{2}\right)\right)$, it satisfies $\left\|z_{\ell}\right\|_{2} \leq(1+1 / L)^{\ell-b+1}\left\|z_{b-1}\right\|_{2}$.

Proof. Let $U$ denote the following column orthonormal matrix using Gram-Schmidt (its first $n(\ell-1)$ columns coincide with $\left.U_{\ell-1}\right)$ :

$$
U \stackrel{\text { def }}{=} \mathrm{GS}\left(h_{1,1}, \ldots, h_{n, 1}, h_{1,2}, \ldots, h_{n, 2}, \ldots, h_{1, \ell-1}, \ldots, h_{n, \ell-1}, z_{b-1}, \cdots z_{\ell-1}\right)
$$

We can rewrite $\left\|z_{\ell}\right\|_{2}$ as follows:

$$
\begin{aligned}
\left\|z_{\ell}\right\|_{2} & =\left\|D_{i, \ell} W z_{\ell-1}\right\|_{2} \\
& =\left\|\mathbf{1}_{W h_{i, \ell-1}+A x_{i, \ell} \geq 0} \cdot W z_{\ell-1}\right\|_{2} \\
& =\left\|\mathbf{1}_{W U U^{\top} h_{i, \ell-1}+A x_{i, \ell} \geq 0} \cdot W U U^{\top} z_{\ell-1}\right\|_{2} \\
& =\left\|\mathbf{1}_{M y+A x \geq 0} M z\right\|_{2} .
\end{aligned}
$$

where in the last step we have defined $M=W U, y=U^{\top} h_{i, \ell-1}, x=x_{i, \ell}$ and $z=U^{\top} z_{\ell-1}$. We stress here that the entries of $M$ and $A$ are i.i.d. from $\mathcal{N}\left(0, \frac{2}{m}\right) .{ }^{17}$

Now, Claim 6.11.10 tells us for fixed $x \in \mathbb{R}^{d}$ and fixed $y, z \in \mathbb{R}^{n \ell+\ell-b}$, we have with probability $1-\exp \left(-\Omega\left(m / L^{2}\right)\right.$ ) (over the randomness of $M, A$ ),

$$
\left\|\mathbf{1}_{M y+A x \geq 0} M z\right\|_{2} \leq\|z\|_{2}(1+1 / 2 L) .
$$

After taking $\epsilon$-net over all possible $y, z$, we have that for fixed $x \in \mathbb{R}^{d}$ but all $y, z \in \mathbb{R}^{n \ell+\ell-b}$ :

$$
\left\|\mathbf{1}_{M y+A x \geq 0} M z\right\|_{2} \leq\|z\|_{2}(1+1 / L) .
$$

We can thus plug in the choice $y=U^{\top} h_{i, \ell-1}$ and $z=U^{\top} z_{\ell-1}$ (both of which may depend on the randomness of $W$ and $A$ ). Using $\|z\|_{2} \leq\left\|z_{\ell-1}\right\|_{2}$, we have

$$
\left\|z_{\ell}\right\|_{2} \leq\left\|z_{\ell-1}\right\|_{2}(1+1 / L) .
$$

Finally, taking union bound over all possible $\ell$ and applying induction, we have

$$
\left\|z_{\ell}\right\|_{2} \leq(1+1 / L)^{\ell-b+1}\left\|z_{b-1}\right\|_{2}
$$

Now, to prove the spectral norm bound in Lemma 6.11.8, we need to go from "for each $z_{b-1}$ (see Claim 6.11.9)" to "for all $z_{b-1}$ ". Since $z_{b-1}$ has $m$ dimensions, we cannot afford taking union bound over all possible $z_{b-1}$ (or its $\epsilon$-net).

To bypass this issue, we partition the coordinates of $z_{b-1}$ into $L^{3}$ trunk (each of length $\left.m / L^{3}\right)$. We write $z_{b-1}=\sum_{j=1}^{L^{3}}\left(z_{b-1}\right)_{j}$ where each $\left(z_{b-1}\right)_{j} \in \mathbb{R}^{m}$ denotes a vector that only

[^44]has non-zero entries on $m / L^{3}$ coordinates. For each $j \in\left[L^{3}\right]$, we have with probability $1-\exp \left(-m / L^{2}\right)$,
$$
\left\|\left(z_{\ell}\right)_{j}\right\|_{2} \leq(1+1 / L)^{\ell-b} \cdot\left\|\left(z_{b-1}\right)_{j}\right\|_{2} \leq 2\left\|\left(z_{b-1}\right)_{j}\right\|_{2}
$$

By applying an $\epsilon$-net argument over all such possible (but sparse) $\left(z_{b-1}\right)_{j}$, we have with probability at least $1-2^{O\left(m / L^{3}\right)} \exp \left(-\Omega\left(m / L^{2}\right)\right) \geq 1-\exp \left(-\Omega\left(m / L^{2}\right)\right)$, the above equation holds for all possible $\left(z_{b-1}\right)_{j}$.

Next, taking a union bound over all $j \in\left[L^{3}\right]$, we have with probability at least $1-\exp \left(-\Omega\left(m / L^{2}\right)\right):$

$$
\forall z_{b-1} \in \mathbb{R}^{m}:\left\|z_{\ell}\right\|_{2} \leq O\left(L^{3}\right)\left\|z_{b-1}\right\|_{2}
$$

Taking a union bound over all $\ell, b$, we complete the proof of Lemma 6.11.8.

### 6.11.4.1 Tools

Claim 6.11.10. For fixed $x \in \mathbb{R}^{d}, y, z \in \mathbb{R}^{k}$. Let $M \in \mathbb{R}^{m \times k}$ and $A \in \mathbb{R}^{m \times d}$ denote random Gaussian matrices where each entry is i.i.d. sampled from $\mathcal{N}(0,2 / m)$. We have with probability at least $1-\exp \left(-\Omega\left(m / L^{2}\right)\right)$

$$
\left\|\mathbf{1}_{M y+A x \geq 0} \cdot M z\right\|_{2} \leq\|z\|_{2}(1+1 / 2 L)
$$

Proof. Without loss of generality we assume $\|z\|_{2}=1$. We can rewrite $M y$ as follows

$$
\begin{aligned}
M y & =M\left(z z^{\top}\right) y+M\left(I-z z^{\top}\right) y \\
& =M z \cdot z^{\top} y+\frac{M\left(I-z z^{\top}\right) y}{\left\|\left(I-z z^{\top}\right) y\right\|_{2}} \cdot\left\|\left(I-z z^{\top}\right) y\right\|_{2} \\
& =M_{1} z_{1}+M_{2} z_{2},
\end{aligned}
$$

where $M_{1}, M_{2} \in \mathbb{R}^{m}$ and $z_{1}, z_{2} \in \mathbb{R}$ are defined as follows

$$
\begin{aligned}
M_{1} & =M z, & M_{2}=\frac{M\left(I-z z^{\top}\right) y}{\left\|\left(I-z z^{\top}\right) y\right\|_{2}} \\
z_{1} & =z^{\top} y, & z_{2}=\left\|\left(I-z z^{\top}\right) y\right\|_{2}
\end{aligned}
$$

It is easy to see that $M_{1}$ is independent of $M_{2}$. We can rewrite

$$
\left\|\mathbf{1}_{M y+A x \geq 0} \cdot M z\right\|_{2}=\left\|\mathbf{1}_{M_{1} z_{1}+M_{2} z_{2}+A x \geq 0} \cdot M_{1}\right\|_{2}
$$

Using Fact 6.10.11 together with concentration bounds (for binomial distribution and for chi-square distribution), we have with probability at least $1-\exp \left(-\Omega\left(m / L^{2}\right)\right)$,

$$
\left\|\mathbf{1}_{M_{1} z_{1}+M_{2} z_{2}+A x \geq 0} \cdot M_{1}\right\|_{2} \leq\|z\|_{2}(1+1 / 2 L) .
$$

Thus, we complete the proof.

### 6.11.5 Intermediate Layers: Sparse Spectral Norm

This section proves two results corresponding to the spectral norm of intermediate layers with respect to sparse vectors. We first show Lemma 6.11.11 and our Corollary 6.11.12 and 6.11.13 shall be direct applications of Lemma 6.11.11.

Lemma 6.11.11. For every $k \in[m]$ and $t \geq 2$, with probability at least

$$
1-n L^{2} e^{O(k \log (m))}\left(e^{-\Omega\left(m / L^{2}\right)}+e^{-\Omega\left(n L t^{2}\right)}\right)
$$

it satisfies, for all $i \in[n]$, for all $L>\ell_{2} \geq \ell_{1} \geq 1$, for all $k$-sparse vectors $z, y \in \mathbb{R}^{m}$

$$
\left|y^{\top} W\left(\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} W\right) z\right| \leq \frac{5 t \sqrt{n L}}{m^{1 / 2}} \cdot\|y\|_{2} \cdot\|z\|_{2}
$$

Proof. Without loss of generality we assume $\|y\|=\|z\|=1$. Fixing $\ell_{2} \geq \ell_{1}$, fixing $i$, and fixing $z_{\ell_{1}-1}=z$, we have according to Claim 6.11.9, letting $z_{\ell_{2}}=D_{i, \ell_{2}} W D_{i, \ell_{2}-1} W \cdots D_{i, \ell_{1}} z_{\ell_{1}-1}$, then with probability at least $1-e^{-\Omega\left(m / L^{2}\right)}$

$$
\left\|z_{\ell_{2}}\right\|_{2} \leq(1+1 / L)^{\ell_{2}-\ell_{1}-1}\left\|z_{\ell_{1}-1}\right\|_{2} \leq 2 .
$$

Applying $\epsilon$-net over all $k$-sparse vectors $z$, we have with probability at least $1-e^{O(k \log m)} e^{-\Omega\left(m / L^{2}\right)}$, it satisfies $\left\|z_{\ell_{2}}\right\|_{2} \leq 3$ for all $k$-sparse $z$ :

Similar to the proof of Lemma 6.11.8, we let $U$ denote the following column orthonormal matrix using Gram-Schmidt:

$$
U \stackrel{\text { def }}{=} \mathrm{GS}\left(h_{1,1}, \ldots, h_{n, 1}, h_{1,2}, \ldots, h_{n, 2}, \ldots, h_{1, \ell_{2}-1}, \ldots, h_{n, \ell_{2}-1}, z_{\ell_{1}-1}, \cdots z_{\ell_{2}}\right)
$$

and we have

$$
\left|y^{\top} W\left(\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} W\right) z\right|=\left|y^{\top} W U U^{\top} z_{\ell_{2}}\right| \leq\left\|y^{\top} W U\right\| \cdot\left\|U^{\top} z_{\ell_{2}}\right\| \leq 3\left\|y^{\top} W U\right\| .
$$

Next, observe the entries of $W U \in \mathbb{R}^{m \times n\left(\ell_{2}-1\right)+\left(\ell_{2}-\ell_{1}+2\right)}$ are i.i.d. drawn from $\mathcal{N}\left(0, \frac{2}{m}\right)$ (following a similar argument as Footnote 17). Therefore, if $y$ is a fixed vector, then $y^{\top} W U$ is in distribution identical to a Gaussian vector $\mathcal{N}\left(0, \frac{2}{m} I\right)$ of $n\left(\ell_{2}-1\right)+\left(\ell_{2}-\ell_{1}+2\right) \leq 2 n L$ dimensions. By chi-square distribution tail bound (see Lemma 6.10.3), we have for $t \geq 2$ :

$$
\operatorname{Pr}\left[\left\|y^{\top} W U\right\|^{2} \geq \frac{n L t^{2}}{m}\right] \leq e^{-\Omega\left(n L t^{2}\right)}
$$

Applying $\epsilon$-net over all $k$-sparse vectors $y$, we have with probability at least $1-e^{O(k \log m)} e^{-\Omega\left(n L t^{2}\right)}$, it satisfies $\left\|y^{\top} W U\right\|^{2} \leq \frac{2 n L t^{2}}{m}$ for all $k$-sparse vectors $y$.

Conditioning on both events happen, we have

$$
\left|y^{\top} W\left(\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} W\right) z\right| \leq 3 \frac{\sqrt{2 n L} t}{\sqrt{m}} .
$$

Taking union bound over all possible $i, \ell_{1}, \ell_{2}$ we finish the proof.

Choosing $k=s^{2} m^{2 / 3}$ and $t=\frac{s m^{1 / 3} \log m}{5 \sqrt{n L}}$ in Lemma 6.11.11, we have
Corollary 6.11.12 (c.f. (6.6)). Let $s \in\left[m^{-1 / 4}, m^{1 / 6}\right]$ be a fixed real. With probability at least $1-e^{-\Omega\left(s^{2} m^{2 / 3} \log ^{2} m\right)}$, we have for all $i \in[n]$, for all $L>\ell_{2} \geq \ell_{1} \geq 1$, for all $s^{2} m^{2 / 3}$-sparse vectors $y, z \in \mathbb{R}^{m}$,

$$
\left|z^{\top} W\left(\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} W\right) y\right| \leq \frac{s \log m}{m^{1 / 6}} \cdot\|z\|_{2} \cdot\|y\|_{2}
$$

Choosing $k=1$ and $t=\frac{\sqrt{n L} d \log m}{5}$ in Lemma 6.11.11, we have
Corollary 6.11.13 (c.f. (6.6)). Let $\rho=n L d \log m$. With probability at least $1-\exp \left(-\Omega\left(\rho^{2}\right)\right)$, it satisfies for all $i \in[n]$, for all $L>\ell_{2} \geq \ell_{1} \geq 1$, and for all 1 -sparse vectors $y, z \in \mathbb{R}^{m}$,

$$
\left|z^{\top} W\left(\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} W\right) y\right| \leq \frac{\rho}{m^{1 / 2}} \cdot\|z\|_{2} \cdot\|y\|_{2}
$$

### 6.11.6 Backward Propagation

This section proves upper bound on the backward propagation against sparse vectors. We first show Lemma 6.11 .14 and our Corollary 6.11 .15 and 6.11 .16 shall be immediate corollaries.

Lemma 6.11.14. For any $k \in[m], t \geq 2$ and any $a \in \mathbb{R}^{d}$, with probability at least

$$
1-n L^{2} e^{O(k \log m)}\left(e^{-\Omega\left(m / L^{2}\right)}+e^{-\Omega\left(t^{2}\right)}\right)
$$

over the randomness of $W, A, B$, we have for all $i \in[n]$, for all $L \geq \ell_{2} \geq \ell_{1} \geq 1$, and for all $k$-sparse $y \in \mathbb{R}^{m}$,

$$
\left|a^{\top} B\left(\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} W\right) y\right| \leq \frac{t}{\sqrt{d}} \cdot\|a\|_{2} \cdot\|y\|_{2}
$$

Proof of Lemma 6.11.14. Using Claim 6.11.9, we know that for fixed $z_{\ell_{1}-1}=y$, letting

$$
z_{\ell_{2}}=D_{i, \ell_{2}} W D_{i, \ell_{2}-1} W \cdots D_{i, \ell_{1}} W z_{\ell_{1}-1}
$$

with probability at least $1-\exp \left(-\Omega\left(m / L^{2}\right)\right)$ we have

$$
\left\|z_{\ell_{2}}\right\|_{2} \leq(1+1 / L)^{\ell_{2}-\ell_{1}+1}\left\|z_{\ell_{1}-1}\right\|_{2} \leq 3\left\|z_{\ell_{1}-1}\right\|_{2}
$$

Next, fixing vectors $a$ and $z_{\ell_{2}}$ and letting $B$ be the only source of randomness, we know $a^{\top} B z_{\ell_{2}}$ follows from $\mathcal{N}\left(0,\|a\|_{2}^{2}\left\|z_{\ell_{2}}\right\|_{2}^{2} / d\right)$. Then with probability at least $1-\exp \left(-\Omega\left(t^{2}\right)\right)$ over $B$,

$$
\left|a^{\top} B z_{\ell_{2}}\right| \leq t \cdot\|a\|_{2}\left\|z_{\ell_{2}}\right\|_{2} / \sqrt{d}
$$

Taking a union of the above two events, we get with probability $1-\exp \left(-\Omega\left(m / L^{2}\right)\right)-$ $\exp \left(-\Omega\left(t^{2}\right)\right)$,

$$
\left|a^{\top} B z_{\ell_{2}}\right| \leq \frac{3 t}{\sqrt{d}} \cdot\|a\|_{2} \cdot\|y\|_{2}
$$

At this point, we apply $\epsilon$-net argument for all $k$-sparse vectors $y \in \mathbb{R}^{m}$ (the size of which is at most $\left.e^{O(k \log (m / k))}\right)$. Taking a union bound over all such vectors in the $\epsilon$-net, we have with probability at least

$$
1-\exp (O(d+k \log (m / k)))\left(\exp \left(-m / L^{2}\right)+\exp \left(-\Omega\left(t^{2}\right)\right)\right),
$$

for all $k$-sparse vector $y \in \mathbb{R}^{m}$,

$$
\left|a^{\top} B z_{\ell_{2}}\right| \leq \frac{3 t}{\sqrt{d}} \cdot\|a\|_{2} \cdot\|y\|_{2}
$$

Finally, we also take a union bound over all $\ell_{2}, \ell_{1}$ and $i$, and there are at most $O\left(n L^{2}\right)$ choices.

Using Lemma 6.11 .14 with $k=s^{2} m^{2 / 3}$ and $t=s m^{1 / 3} \log m$, and taking union bound over all $a \in \mathbb{R}^{d}$, give

Corollary 6.11.15 (c.f. (6.7)). Let $s \in\left[m^{-1 / 4}, m^{1 / 6}\right]$ be a fixed real. With probability at least $1-\exp \left(-\Omega\left(s^{2} m^{2 / 3} \log ^{2} m\right)\right)$, we have for all $i \in[n]$, for all $L \geq \ell_{2} \geq \ell_{1} \geq 1$, for all $\left(s^{2} \cdot m^{2 / 3}\right)$-sparse $y \in \mathbb{R}^{m}$, for all $a \in \mathbb{R}^{d}$,

$$
\left|a^{\top} B\left(\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} W\right) y\right| \leq\left(s m^{1 / 3} \log m\right) \cdot\|a\|_{2} \cdot\|y\|_{2}
$$

Using Lemma 6.11 .14 with $k=1$ and $t=n L d \log m$, and taking union bound over all $a \in \mathbb{R}^{d}$, give

Corollary 6.11.16 (c.f. (6.7)). Let $\rho=n L d \log m$. With probability at least $1-\exp \left(-\Omega\left(\rho^{2}\right)\right)$, we have for all $i \in[n]$, for all $L \geq \ell_{2} \geq \ell_{1} \geq 1$, for all 1 -sparse vector $y \in \mathbb{R}^{m}$, and for all $a \in \mathbb{R}^{d}$,

$$
\left|a^{\top} B\left(\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} W\right) y\right| \leq \rho \cdot\|a\|_{2} \cdot\|y\|_{2} .
$$

### 6.12 Stability After Adversarial Perturbation

Throughout this section, we consider some random initialization $\widetilde{W}, A, B$, and some adversarially chosen perturbation $W^{\prime} \in \mathbb{R}^{m \times m}$ which may depend on the randomness of $\widetilde{W}, A, B$. We introduce the following notations in this section

## Definition 6.12.1.

$$
\begin{array}{lll}
\widetilde{g}_{i, 0}=\widetilde{h}_{i, 0}=0 & g_{i, 0}=h_{i, 0}=0 & \text { for } i \in[n] \\
\widetilde{g}_{i, \ell}=\widetilde{W} \widetilde{h}_{i, \ell-1}+A x_{i, \ell} & g_{i, \ell}=\left(\widetilde{W}+W^{\prime}\right) h_{i, \ell-1}+A x_{i, \ell} & \\
\text { for } i \in[n] \text { and } \ell \in[L] \\
\widetilde{h}_{i, \ell}=\phi\left(\widetilde{W} \widetilde{h}_{i, \ell-1}+A x_{i, \ell}\right) & h_{i, \ell}=\phi\left(\left(\widetilde{W}+W^{\prime}\right) h_{i, \ell-1}+A x_{i, \ell}\right) & \text { for } i \in[n] \text { and } \ell \in[L] \\
h_{i, \ell}^{\prime}=h_{i, \ell}-\widetilde{h}_{i, \ell} & g_{i, \ell}^{\prime}=g_{i, \ell}-\widetilde{g}_{i, \ell} & \text { for } i \in[n] \text { and } \ell \in[L]
\end{array}
$$

Define diagonal matrices $\widetilde{D}_{i, \ell}$ and $D_{i, \ell}$ by letting

$$
\left(\widetilde{D}_{i, \ell}\right)_{k, k}=\mathbb{1}_{\left(\widetilde{g}_{i, \ell}\right)_{k} \geq 0} \text { and }\left(D_{i, \ell}\right)_{k, k}=\mathbb{1}_{\left(g_{i, \ell}\right)_{k} \geq 0}
$$

Accordingly, we let diagonal matrix $D_{i, \ell}^{\prime}=D_{i, \ell}-\widetilde{D}_{i, \ell}$.

## Roadmap

- Section 6.12 .1 proves the stability at forward when the perturbation matrix $W^{\prime}$ has small spectral norm. It gives bounds on $\left\|g_{i, \ell}^{\prime}\right\|_{2},\left\|h_{i, \ell}^{\prime}\right\|_{2},\left\|D_{i, \ell}^{\prime}\right\|_{0}$, and $\left\|D_{i, \ell}^{\prime} g_{i, \ell}\right\|_{2}$.
- Section 6.12 .2 proves the stability of intermediate layers. It gives bound on $\left\|\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} \widetilde{W}\right\|_{2}$ and $\left\|\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} W\right\|_{2}$.
- Section 6.12 .3 analyzes the stability for backward. It bounds the difference $\| a^{\top} B\left(\prod_{\ell=\ell_{2}}^{\ell_{1}} \widetilde{D}_{i, \ell} \widetilde{W}-\right.$ $\left.\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} \widetilde{W}\right) \|_{2}$.
- Section 6.12 .4 considers a special type of rank-one perturbation matrix $W^{\prime}$, and provides stability bounds on the forward and backward propagation.

As discussed in Section 6.6, the results of Section 6.12.1, 6.12 .2 and 6.12 .3 shall be used twice, once for the final training updates (see Section 6.16), and once for the randomness decomposition (see Section 6.14). In contrast, the results of Section 6.12 .4 shall only be used once in Section 6.14.

### 6.12.1 Forward

The goal of this section is to prove Lemma 6.12.1,

Lemma 6.12.1 (forward stability, c.f. (6.8)). Letting $\varrho=n L d \delta^{-1} \log (m / \epsilon)$, for any $\tau_{0} \in$ $\left[\varrho^{-100}, \varrho^{100}\right]$, with probability at least $1-e^{-\Omega\left(L^{6} \tau_{0}^{4 / 3} m^{1 / 3}\right)}$ over the randomness of $\widetilde{W}, A, B$, for every $W^{\prime} \in \mathbb{R}^{m \times m}$ with $\left\|W^{\prime}\right\|_{2} \leq \frac{\tau_{0}}{\sqrt{m}}$, for every $i \in[n]$ and $\ell \in[L]$, we have
(a) $\left\|g_{i, \ell}^{\prime}\right\|_{2},\left\|h_{i, \ell}^{\prime}\right\|_{2} \leq O\left(L^{6} \tau_{0}\right) / m^{1 / 2}$,
(b) $\left\|D_{i, \ell}^{\prime}\right\|_{0} \leq O\left(L^{10 / 3} \tau_{0}^{2 / 3}\right) \cdot m^{2 / 3}$, and
(c) $\left\|D_{i, \ell}^{\prime} g_{i, \ell}\right\|_{2} \leq O\left(L^{5} \tau_{0}\right) / m^{1 / 2}$.

Proof of Lemma 6.12.1. Suppose $C>1$ is a large enough constant so that the hidden constant in Lemma 6.11.8 can be $C$. We inductively prove that one can write $g_{i, \ell}^{\prime}=g_{i, \ell, 1}^{\prime}+g_{i, \ell, 2}^{\prime}$ where

$$
\begin{array}{lllll}
\text { I: } & \left\|g_{i, \ell, 1}^{\prime}\right\|_{2} \leq \tau_{1} \cdot \frac{1}{m^{1 / 2}} & \text { II: } \quad\left\|g_{i, \ell, 2}^{\prime}\right\|_{\infty} \leq \tau_{2} \cdot \frac{1}{m} \quad \text { III: } \quad\left\|g_{i, \ell}^{\prime}\right\|_{2} \leq \tau_{3} \cdot \frac{1}{m^{1 / 2}} \\
\text { IV: } & \left\|D_{i, \ell}^{\prime}\right\|_{0} \leq \tau_{4} \cdot m^{2 / 3} & \text { V: } & \left\|D_{i, \ell}^{\prime} g_{i, \ell}\right\|_{2} \leq \tau_{5} \cdot \frac{1}{m^{1 / 2}} & \text { VI: } \quad\left\|h_{i, \ell}^{\prime}\right\|_{2} \leq\left(\tau_{3}+\tau_{5}\right) \frac{1}{m^{1 / 2}} \tag{6.37}
\end{array}
$$

Above, we choose parameters

$$
\tau_{1}=5 C L^{4}(L+2) \tau_{0} \quad \tau_{2}=4 L \tau_{5} \log m \quad \tau_{3}=\tau_{1}+\tau_{2} \quad \tau_{4}=10\left(\tau_{1}\right)^{2 / 3} \quad \tau_{5}=3 \tau_{1}
$$

We emphasize that all these parameters are polynomial in $\varrho$ so negligible when comparing to $m$.

Throughout the proof, we focus on some fixed $i \in[n]$ without loss of generality, and one can always take a union bound at the end. We drop the subscript $i$ for notational
simplicity. In order to prove (6.37), we first assume that it holds for all $1,2, \ldots, \ell-1$. In particular, we assume for all $\ell^{\prime} \leq \ell-1$,

$$
\left\|g_{\ell^{\prime}, 1}^{\prime}\right\|_{2} \leq \tau_{1} \cdot \frac{1}{m^{1 / 2}} \quad \text { and } \quad\left\|g_{\ell^{\prime}, 2}^{\prime}\right\|_{\infty} \leq \tau_{2} \cdot \frac{1}{m}
$$

A useful observation is $g_{\ell^{\prime}}^{\prime}$ can be split into three terms

$$
\begin{equation*}
g_{\ell^{\prime}}^{\prime}=\underbrace{W^{\prime} D_{\ell^{\prime}-1} g_{\ell^{\prime}-1}}_{z_{\ell^{\prime}-1,1}}+\underbrace{\widetilde{W} D_{\ell^{\prime}-1}^{\prime} g_{\ell^{\prime}-1}}_{z_{\ell^{\prime}-1,2}}+\underbrace{\widetilde{W} \widetilde{D}_{\ell^{\prime}-1} g_{\ell^{\prime}-1}^{\prime}}_{z_{\ell^{\prime}-1,3}} . \tag{6.38}
\end{equation*}
$$

After recursively applying (6.38), we can write

$$
\begin{equation*}
g_{\ell}^{\prime}=g_{\ell, 1}^{\prime}=\sum_{\ell_{a}=1}^{\ell-1}\left(\widetilde{W} \widetilde{D}_{\ell-1} \cdots \widetilde{W} \widetilde{D}_{\ell-\ell_{a}+1}\right) z_{\ell-\ell_{a}, 1}+\left(\widetilde{W} \widetilde{D}_{\ell-1} \cdots \widetilde{W} \widetilde{D}_{\ell-\ell_{a}+1}\right) z_{\ell-\ell_{a}, 2} \tag{6.39}
\end{equation*}
$$

Applying Claim 6.12.2, with probability at least $1-e^{-\Omega\left(\tau_{1}^{4 / 3} m^{1 / 3}\right)}$, we have for all $\ell^{\prime} \leq \ell-1$,

$$
\begin{equation*}
\left\|z_{\ell^{\prime}, 1}\right\|_{2} \leq \frac{\tau_{0}}{\sqrt{m}}\left(4 \ell^{\prime}+8+\frac{\tau_{1}+\tau_{2}}{\sqrt{m}}\right) \tag{6.40}
\end{equation*}
$$

Applying Claim 6.12.5, we have with probability at least $1-e^{-\Omega\left(\tau_{1}^{4 / 3} m^{1 / 3}\right)}$, one can write

$$
\left(\widetilde{W} \widetilde{D}_{\ell-1} \cdots \widetilde{W} \widetilde{D}_{\ell-\ell_{a}+1}\right) z_{\ell-\ell_{a}, 2}=z_{\ell-\ell_{a}, 2^{\sharp}}+z_{\ell-\ell_{a}, 2^{b}}
$$

where

$$
\begin{equation*}
\left\|z_{\ell-\ell_{a}, 2^{\sharp}}\right\|_{2} \leq \frac{3 \tau_{5} \sqrt{\tau_{4}} \log ^{2} m}{m^{2 / 3}}, \quad\left\|z_{\ell-\ell_{a}, 2^{\natural}}\right\|_{\infty} \leq \frac{4 \tau_{5} \log m}{m} . \tag{6.41}
\end{equation*}
$$

As a result, we can define $g_{\ell}^{\prime}=g_{\ell, 1}^{\prime}+g_{\ell, 2}^{\prime}$ for

$$
g_{\ell, 1}^{\prime}=\sum_{\ell_{a}=1}^{\ell-1}\left(\widetilde{W} \widetilde{D}_{\ell-1} \cdots \widetilde{W} \widetilde{D}_{\ell-\ell_{a}+1}\right) z_{\ell-\ell_{a}, 1}+z_{\ell-\ell_{a}, 2^{\sharp}} \quad \text { and } \quad g_{\ell, 2}^{\prime}=\sum_{\ell_{a}=1}^{\ell-1} z_{\ell-\ell_{a}, 2^{b}}
$$

We first bound $\left\|g_{\ell, 1}^{\prime}\right\|_{2}$,

$$
\begin{aligned}
\left\|g_{\ell, 1}^{\prime}\right\|_{2} & \leq \sum_{\ell_{a}=1}^{\ell-1}\left\|\widetilde{W} \widetilde{D}_{i, \ell-1} \cdots \widetilde{W} \widetilde{D}_{i, \ell-\ell_{a}+1}\right\|_{2} \cdot\left\|z_{\ell-\ell_{a}, 1}\right\|_{2}+\left\|z_{\ell-\ell_{a}, 2^{\sharp}}\right\|_{2} \\
& \stackrel{(1)}{\leq} \sum_{\ell_{a}=1}^{\ell}\left(C L^{3}\left\|z_{\ell-\ell_{a}, 1}\right\|_{2}+\left\|z_{\ell-\ell_{a}, 2^{\sharp}}\right\|_{2}\right) \\
& \stackrel{(2)}{\leq} C L^{4} \frac{\tau_{0}}{\sqrt{m}}\left(4 L+8+\frac{\tau_{1}+\tau_{2}}{\sqrt{m}}\right)+L \frac{3 \tau_{5} \sqrt{\tau_{4}} \log ^{2} m}{m^{2 / 3}} \\
& \leq 5 C L^{4} \cdot \frac{\tau_{0}(L+2)}{\sqrt{m}} \stackrel{B}{\leq} \tau_{1} \frac{1}{\sqrt{m}} .
\end{aligned}
$$

Above, inequality (1) follows from Lemma 6.11.8, inequality (2) follows from (6.40) and (6.41), and (3) follows from our choice of $\tau_{1}=5 C L^{4}(L+2) \tau_{0}$. We next bound $\left\|g_{\ell, 2}^{\prime}\right\|_{\infty}$,

$$
\left\|g_{\ell, 2}^{\prime}\right\|_{\infty} \leq \sum_{\ell_{a}=1}^{\ell-1}\left\|z_{\ell-\ell_{a}, 2^{b}}\right\|_{\infty} \stackrel{(1}{\leq} L \cdot \frac{4 \tau_{5} \log m}{m} \stackrel{(2}{\leq} \tau_{2} \frac{1}{m}
$$

where inequality (1) is due to (6.41), and (2) follows from our choice of $\tau_{2}=4 L \tau_{5} \log \mathrm{~m}$. Thus, we have showed I and II of (6.37):

$$
\left\|g_{\ell, 1}^{\prime}\right\|_{2} \leq \tau_{1} \cdot \frac{1}{m^{1 / 2}} \quad \text { and } \quad\left\|g_{\ell, 2}^{\prime}\right\|_{\infty} \leq \tau_{2} \cdot \frac{1}{m}
$$

They together further imply

$$
\left\|g_{\ell}^{\prime}\right\|_{2} \leq\left(\tau_{1}+\tau_{2}\right) \frac{1}{m^{1 / 2}}=\tau_{3} \frac{1}{m^{1 / 2}}
$$

so III of (6.37) holds. IV and V of (6.37) are implied by Claim 6.12.3, and VI is implied because

$$
\left\|h_{\ell}^{\prime}\right\|_{2}=\left\|\phi\left(g_{\ell}\right)-\phi\left(\widetilde{g}_{\ell}\right)\right\|_{2} \leq\left\|g_{\ell}-\widetilde{g}_{\ell}\right\|_{2} \leq\left\|g_{\ell}^{\prime}\right\|_{2}
$$

### 6.12.1.1 Tools

Claim 6.12.2. Suppose $g_{\ell-1}=\widetilde{g}_{\ell-1}+g_{\ell-1}^{\prime}=\widetilde{g}_{\ell-1}+g_{\ell-1,1}^{\prime}+g_{\ell-1,2}^{\prime}$ where

$$
\left\|g_{\ell-1,1}^{\prime}\right\|_{2} \leq \frac{\tau_{1}}{\sqrt{m}} \quad \text { and } \quad\left\|g_{\ell-1,2}^{\prime}\right\|_{\infty} \leq \frac{\tau_{2}}{m} .
$$

Then, we have with probability at least $1-e^{-\Omega\left(m / L^{2}\right)}$

$$
\left\|W^{\prime} D_{\ell-1} g_{\ell-1}\right\|_{2} \leq \frac{\tau_{0}}{\sqrt{m}}\left(4 \ell+8+\frac{\tau_{1}+\tau_{2}}{\sqrt{m}}\right) .
$$

Proof of Claim 6.12.2. Using triangle inequality, we can calculate

$$
\left\|W^{\prime} D_{\ell-1} g_{\ell-1}\right\|_{2} \leq\left\|W^{\prime}\right\|_{2} \cdot\left\|D_{\ell-1}\right\|_{2} \cdot\left(\left\|\widetilde{g}_{\ell-1}\right\|_{2}+\left\|g_{\ell-1,1}^{\prime}\right\|_{2}+\left\|g_{\ell-1,2}^{\prime}\right\|_{2}\right)
$$

Using Lemma 6.11.1, we have $\left\|\widetilde{g}_{\ell-1}\right\| \leq 4 \ell+8$. Using elementary calculation, we have

$$
\left\|g_{\ell-1,1}^{\prime}\right\|_{2}+\left\|g_{\ell-1,2}^{\prime}\right\|_{2} \leq \frac{\tau_{1}}{\sqrt{m}}+\sqrt{m} \frac{\tau_{2}}{m} \leq \frac{\tau_{1}+\tau_{2}}{\sqrt{m}}
$$

Together we finish the proof.
Claim 6.12.3. With probability at least $1-e^{-\Omega\left(\tau_{1}^{4 / 3} m^{1 / 3}\right)}$ the following holds. Whenever

$$
\left\|g_{\ell-1,1}^{\prime}\right\|_{2} \leq \tau_{1} \cdot \frac{1}{m^{1 / 2}}, \quad\left\|g_{\ell-1,2}^{\prime}\right\|_{\infty} \leq \tau_{2} \cdot \frac{1}{m}
$$

then letting $\tau_{4}=10\left(\tau_{1}\right)^{2 / 3}$ and $\tau_{5}=3 \tau_{1}$, we have

$$
\left\|D_{\ell-1}^{\prime} g_{\ell-1}\right\|_{2} \leq \tau_{5} \cdot \frac{1}{m^{1 / 2}}, \quad\left\|D_{\ell-1}^{\prime}\right\|_{0} \leq \tau_{4} \cdot m^{2 / 3}
$$

Proof of Claim 6.12.3. We choose parameters $\xi=\frac{\left(\tau_{1}\right)^{2 / 3}}{10^{1 / 3} m^{5 / 6}}$ and $\alpha=10 \xi \sqrt{m}$ in the proof. We have $\left\|g_{\ell-1,2}^{\prime}\right\|_{\infty} \leq \xi$.

First of all, using similar (but simpler) proof as Lemma 6.11.4, one can show with probability at least $1-\exp \left(-\Omega\left(\alpha^{2} m\right)\right)$, the vector $\widetilde{g}_{\ell-1}$ is $\left(\alpha, \frac{1}{5 \sqrt{m}}\right)$-good (recall Definition 6.10.1). ${ }^{18}$ This implies that $\widetilde{g}_{\ell-1}$ has at most $\alpha m=10 \xi m^{3 / 2}$ coordinates $j$ satisfying $\left|\left(\widetilde{g}_{\ell-1}\right)_{j}\right| \leq \frac{\alpha}{5 \sqrt{m}}=2 \xi$.

For each $j \in[m]$, if it satisfies $\left(D_{\ell-1}^{\prime}\right)_{j, j} \neq 0$, then the sign of $\widetilde{g}_{\ell-1}$ and $\widetilde{g}_{\ell-1}+g_{\ell-1,1}^{\prime}+$ $g_{\ell-1,2}^{\prime}$ must differ on coordinate $j$. As a result:

$$
\left|\left(g_{\ell-1,1}^{\prime}+g_{\ell-1,2}^{\prime}\right)_{j}\right|>\left|\left(\widetilde{g}_{\ell-1}\right)_{j}\right|
$$

Define $y=D_{\ell-1}^{\prime} g_{\ell-1}$. There are two possibilities for such $j$ with $\left(D_{\ell-1}^{\prime}\right)_{j, j} \neq 0$.

- Case 1: $\left|\left(\widetilde{g}_{\ell-1}\right)_{j}\right| \leq 2 \xi$. Let such coordinates be $S_{1} \subset[m]$, and we have $\left|S_{1}\right| \leq 10 \xi m^{3 / 2}$ using the above argument.

Next, for each such $j \in S_{1}$, we must have $\left|y_{j}\right|=\left|\left(\widetilde{g}_{\ell-1}+g_{\ell-1,1}^{\prime}+g_{\ell-1,2}^{\prime}\right)_{j}\right| \leq \mid\left(g_{\ell-1,1}^{\prime}+\right.$

[^45]where $M_{1} \in \mathbb{R}^{m \times n(\ell-2)}, M_{3} \in \mathbb{R}^{m \times d}$ and $z_{1} \in \mathbb{R}^{n(\ell-1)}, z_{3} \in \mathbb{R}^{d}$ are defined as
\[

$$
\begin{align*}
M_{1} & =\widetilde{W} U_{\ell-2} & M_{2} & =A \\
z_{1} & =U_{\ell-2}^{\top} \widetilde{h}_{\ell-2} & z_{2} & =x_{\ell-1} . \tag{6.42}
\end{align*}
$$
\]

The entries of $M_{1}$ and $M_{2}$ are i.i.d. from $\mathcal{N}\left(0, \frac{2}{m}\right)$ (recall Footnote 13). Suppose $z_{1}$ and $z_{2}$ are fixed (instead of depending on the randomness of $\widetilde{W}$ and $A$ ) and satisfies $\frac{1}{\sqrt{2}} \leq\left\|z_{1}\right\| \leq 2 L+6$. We can apply Corollary 6.10.15 to obtain the following statement: with probability at least $1-\exp \left(-\Omega\left(\alpha^{2} m\right)\right), \widetilde{g}_{\ell-1}$ is $(\alpha, \sigma / 4)$-good where $\sigma^{2}=\frac{2}{m}\left\|z_{1}\right\|_{2}^{2}+\frac{2}{m}\left\|z_{3}\right\|_{2}^{2} \geq \frac{1}{m}$. In other words, $w$ is $\left(\alpha, \frac{1}{4 \sqrt{m}}\right)$-good. Applying standard $\epsilon$-net argument, we have with probability at least $1-\exp \left(-\Omega\left(\alpha^{2} m\right)\right)$, for all vectors $z_{1} \in \mathbb{R}^{n(\ell-2)}$ satisfying $\frac{1}{\sqrt{2}} \leq\left\|z_{1}\right\| \leq 2 L+6$, it satisfies $\widetilde{g}_{\ell-1}$ is $\left(\alpha, \frac{1}{5 \sqrt{m}}\right)$-good. This allows us to plug in the random choice of $z_{1}$ in (6.42), to conclude that $\widetilde{g}_{\ell-1}$ is $\left(\alpha, \frac{1}{5 \sqrt{m}}\right)$-good.
$\left.g_{\ell-1,2}^{\prime}\right)_{j}\left|\leq\left|\left(g_{\ell-1,1}^{\prime}\right)_{j}\right|+\xi\right.$ so we can calculate the $\ell_{2}$ norm of $y$ on $S_{1}$ :

$$
\sum_{i \in S_{1}} y_{j}^{2} \leq 2\left\|g_{\ell-1,1}^{\prime}\right\|^{2}+2 \xi^{2}\left|S_{1}\right| \leq \frac{2 \tau_{1}^{2}}{m}+20 \xi^{3} m^{3 / 2}=\frac{4 \tau_{1}^{2}}{m}
$$

- Case 2: $\left|\left(\widetilde{g}_{\ell-1}\right)_{j}\right|>2 \xi$. Let such coordinates be $S_{2} \subset[m] \backslash S_{1}$. In this case we must have $\left|\left(g_{\ell-1,1}^{\prime}\right)_{j}\right| \geq\left|\left(\widetilde{g}_{\ell-1}\right)_{j}\right|-\left|\left(g_{\ell-1,2}^{\prime}\right)_{j}\right|>2 \xi-\xi=\xi$. Therefore, $\left|S_{2}\right| \leq \frac{\left\|g_{\ell-1,1}^{\prime}\right\|_{2}^{2}}{\xi^{2}} \leq \frac{\tau_{1}^{2}}{m \xi^{2}}$. Next, for each $j \in S_{2}$, we must have

$$
\left|y_{j}\right|=\left|\left(\widetilde{g}_{\ell-1}+g_{\ell-1,1}^{\prime}+g_{\ell-1,2}^{\prime}\right)_{j}\right| \leq\left|\left(g_{\ell-1,1}^{\prime}+g_{\ell-1,2}^{\prime}\right)_{j}\right| \leq\left|\left(g_{\ell-1,1}^{\prime}\right)_{j}\right|+\xi / 2 \leq \frac{5}{4}\left|\left(g_{\ell-1,1}^{\prime}\right)_{j}\right|
$$

and therefore

$$
\sum_{j \in S_{2}} y_{j}^{2} \leq 2 \sum_{j \in S_{2}}\left(g_{\ell-1,1}^{\prime}\right)_{j}^{2} \leq \frac{2 \tau_{1}^{2}}{m}
$$

In sum, we conclude that

$$
\left\|D_{\ell-1}^{\prime}\right\|_{0} \leq\left|S_{1}\right|+\left|S_{2}\right| \leq 10 \xi m^{3 / 2}+\frac{\tau_{1}^{2}}{m \xi^{2}}<10\left(\tau_{1}\right)^{2 / 3} m^{2 / 3}=\tau_{4} \cdot m^{2 / 3}
$$

and

$$
\|y\|_{2} \leq \frac{\sqrt{4 \tau_{1}^{2}+2 \tau_{1}^{2}}}{\sqrt{m}} \leq \frac{3 \tau_{1}}{\sqrt{m}}=\tau_{5} \cdot \frac{1}{m^{1 / 2}} .
$$

Claim 6.12.4. With probability at least $1-e^{-\Omega\left(\tau_{4} m^{2 / 3} \log ^{2} m\right)}$, for all $L \geq \ell+1 \geq b \geq 1$, for all

$$
x \in \mathbb{R}^{m} \quad \text { with } \quad\|x\|_{2} \leq \frac{\tau_{5}}{\sqrt{m}} \quad \text { and } \quad\|x\|_{0} \leq \tau_{4} \cdot m^{2 / 3}
$$

we have that the vector $y=\widetilde{W} \widetilde{D}_{\ell} \widetilde{W} \cdots \widetilde{D}_{b} \widetilde{W} x$ can be written as

$$
y=y_{1}+y_{2} \quad \text { where } \quad\left\|y_{1}\right\|_{2} \leq \frac{3 \tau_{5} \sqrt{\tau_{4}} \log ^{2} m}{m^{2 / 3}} \quad \text { and } \quad\left\|y_{2}\right\|_{\infty} \leq \frac{4 \tau_{5} \log m}{m}
$$

Proof of Claim 6.12.4. Let $s=\tau_{4} \cdot m^{2 / 3}$ for notational simplicity. Let $z_{b-1}=x$ and assume for now that $x$ is a fixed vector. Letting $z_{\ell}=\widetilde{D}_{\ell} \widetilde{W} \cdots \widetilde{D}_{b} z_{b-1}$, we have according to Claim 6.11.9, with probability at least $1-\exp \left(-\Omega\left(m / L^{2}\right)\right)$,

$$
\left\|z_{\ell}\right\|_{2} \leq(1+1 / L)^{\ell-b-1}\left\|z_{b-1}\right\|_{2} \leq 3\left\|z_{b-1}\right\|_{2} .
$$

Similar to the proof of Lemma 6.11.8, we let $U$ denote the following column orthonormal matrix using Gram-Schmidt:

$$
U \stackrel{\text { def }}{=} \mathrm{GS}\left(h_{1,1}, \ldots, h_{n, 1}, h_{1,2}, \ldots, h_{n, 2}, \ldots, h_{1, \ell-1}, \ldots, h_{n, \ell-1}, z_{b-1}, \cdots z_{\ell}\right),
$$

and we have

$$
\widetilde{W} \widetilde{D}_{\ell} \widetilde{W} \cdots \widetilde{D}_{b} \widetilde{W} x=\widetilde{W} U U^{\top} z_{\ell_{2}}
$$

Observe the entries of $\widetilde{W} U \in \mathbb{R}^{m \times n(\ell-1)+(\ell-b+2)}$ are i.i.d. drawn from $\mathcal{N}\left(0, \frac{2}{m}\right)$ (following a similar argument as Footnote 17). Therefore, according to Lemma 6.10.16, we have can write $y=\widetilde{W} \widetilde{D}_{\ell} \widetilde{W} \cdots \widetilde{D}_{b} \widetilde{W} x$ as $y=y_{1}+y_{2}$ with

$$
\left\|y_{1}\right\| \leq \frac{\sqrt{s} \log ^{2} m}{\sqrt{m}} \cdot\left\|U^{\top} z_{\ell_{2}}\right\| \quad \text { and } \quad\left\|y_{2}\right\|_{\infty} \leq \frac{\log m}{\sqrt{m}} \cdot\left\|U^{\top} z_{\ell_{2}}\right\|
$$

Plugging in $\left\|U^{\top} z_{\ell}\right\|_{2} \leq\left\|z_{\ell}\right\|_{2} \leq 3\left\|z_{b-1}\right\|_{2}=3\|x\|_{2} \leq \frac{3 \tau_{5}}{\sqrt{m}}$, we have

$$
\left\|y_{1}\right\| \leq \frac{3 \tau_{5} \sqrt{s} \log ^{2} m}{m} \quad \text { and } \quad\left\|y_{2}\right\|_{\infty} \leq \frac{3 \tau_{5} \log m}{m}
$$

Finally, taking $\epsilon$-net over all $s$-sparse vectors $x$, we have the desired result.

Combining Claim 6.12.3 and Claim 6.12.4, we have

Claim 6.12.5. With probability at least $1-e^{-\Omega\left(\tau_{1}^{4 / 3} m^{1 / 3}\right)}$, whenever

$$
\left\|g_{\ell-1,1}^{\prime}\right\|_{2} \leq \frac{\tau_{1}}{\sqrt{m}}, \quad\left\|g_{\ell-1,2}^{\prime}\right\|_{\infty} \leq \frac{\tau_{2}}{m}
$$

we have the vector $y=\widetilde{W} \widetilde{D}_{\ell} \widetilde{W} \cdots \widetilde{D}_{b} \widetilde{W} D_{\ell-1}^{\prime} g_{\ell-1}$ can be written as

$$
y=y_{1}+y_{2} \quad \text { where } \quad\left\|y_{1}\right\|_{2} \leq \frac{3 \tau_{5} \sqrt{\tau_{4}} \log ^{2} m}{m^{2 / 3}} \quad \text { and } \quad\left\|y_{2}\right\|_{\infty} \leq \frac{4 \tau_{5} \log m}{m}
$$

### 6.12.2 Intermediate Layers

The goal of this subsection is to prove Lemma 6.12.6.

Lemma 6.12.6 (intermediate stability, c.f. (6.9)). Letting $\varrho=n L d \log m$, for any $\tau_{0} \in$ $\left[\varrho^{-100}, \varrho^{100}\right]$, with probability at least $1-e^{-\Omega\left(L^{6} \tau_{0}^{4 / 3} m^{1 / 3}\right)}$ over the randomness of $\widetilde{W}, A, B$, for every $W^{\prime} \in \mathbb{R}^{m \times m}$ with $\left\|W^{\prime}\right\|_{2} \leq \frac{\tau_{0}}{\sqrt{m}}$, for every $i \in[n]$, for every $L \geq \ell_{2} \geq \ell_{1} \geq 1$, the following holds

- $\left\|\prod_{\ell=\ell_{2}}^{\ell_{1}}\left(\widetilde{D}_{i, \ell}+D_{i, \ell}^{\prime}\right) \widetilde{W}\right\|_{2} \leq O\left(L^{7}\right)$.
- $\left\|\prod_{\ell=\ell_{2}}^{\ell_{1}}\left(\widetilde{D}_{i, \ell}+D_{i, \ell}^{\prime}\right)\left(\widetilde{W}+W^{\prime}\right)\right\|_{2} \leq O\left(L^{7}\right)$.

We show Claim 6.12.7 and then use it to prove Corollary 6.12.7.

Claim 6.12.7. Let $s \in\left[m^{-1 / 4}, m^{1 / 8}\right]$ be any fixed real. With probability at least $1-$ $e^{-\Omega\left(s^{2} m^{2 / 3} \log ^{2} m\right)}$, it satisfies that for every $i \in[n]$, for every $L \geq \ell_{2} \geq \ell_{1} \geq 1$, if $D_{i, \ell}^{\prime} \in \mathbb{R}^{m \times m}$ is an arbitrary diagonal matrix with entries in $\{-1,0,1\}$ of sparsity $\left\|D_{i, \ell}^{\prime}\right\|_{0} \leq s^{2} m^{2 / 3}$ (for each $\left.\ell=\ell_{1}, \ldots, \ell_{2}\right)$, then

$$
\left\|\prod_{\ell=\ell_{2}}^{\ell_{1}}\left(\widetilde{D}_{i, \ell}+D_{i, \ell}^{\prime}\right) \widetilde{W}\right\|_{2} \leq O\left(L^{7}\right) .
$$

Proof of Claim 6.12.7. We define set $\mathcal{C}$ as follows

$$
\mathcal{C}=\left\{\prod_{\ell=\ell_{2}}^{\ell_{1}}\left(D_{i, \ell}^{\prime}\right)^{c_{\ell}}\left(\widetilde{D}_{i, \ell}\right)^{1-c_{\ell}} \widetilde{W} \mid c_{\ell} \in\{0,1\}, \forall \ell \in\left[\ell_{1}, \ell_{2}\right]\right\} .
$$

For $\ell$, we define set $\mathcal{C}_{\ell}$ as follows

$$
\mathcal{C}_{\ell}=\left\{C \in \mathcal{C} \mid D^{\prime} \text { appears } \ell \text { times in } C\right\}
$$

From Lemma 6.11.8, we can bound the $\ell=0$ term

$$
\left\|\sum_{C \in \mathcal{C}_{\ell}} C\right\|_{2} \leq O\left(L^{3}\right)
$$

For each $C \in \mathcal{C}_{\ell}$ with $\ell>0$, we have

$$
\|C\|_{2} \leq O\left(L^{3}\right) \cdot\left(\frac{s \log m}{m^{1 / 6}}\right)^{\ell-1} \cdot O\left(L^{3}\right)
$$

where we apply Lemma 6.11 .8 twice (one on the left, one on the right) and apply Corollary $6.11 .12 \ell-1$ times. The probability of the above event is at least $1-e^{\Omega\left(s^{2} \cdot m^{2 / 3} \log ^{2} m\right)}$. Therefore, for each $\ell \in\left[\ell_{2}-\ell_{1}\right]$, we can bound

$$
\left\|\sum_{C \in \mathrm{C}_{\ell}} C\right\|_{2} \leq O\left(L^{6}\right)\binom{\ell_{2}-\ell_{1}}{\ell}\left(\frac{s \log m}{m^{1 / 6}}\right)^{\ell-1}
$$

which is at most $O\left(L^{7}\right)$ when $\ell=1$ and $o\left(L^{7}\right)$ for $\ell>1$ by our parameter choices. Putting it altogether,

$$
\text { LHS }=\left\|\sum_{\ell=0}^{\ell_{2}-\ell_{1}} \sum_{C \in \mathrm{C}_{\ell}} C\right\|_{2} \leq \sum_{\ell=0}^{\ell_{2}-\ell_{1}}\left\|\sum_{C \in \mathrm{C}_{\ell}} C\right\|_{2} \leq O\left(L^{3}+L^{7}\right)
$$

Proof of Lemma 6.12.6. Applying Lemma 6.12.1c, we know with probability $\geq 1-e^{-\Omega\left(L^{6} \tau_{0}^{4 / 3} m^{1 / 3}\right)}$ it satisfies $\left\|D_{i, \ell}^{\prime}\right\|_{0} \leq s^{2} m^{2 / 3}$ for $s^{2}=O\left(L^{10 / 3} \tau_{0}^{2 / 3}\right)$. Therefore, applying Claim 6.12.7 we have

$$
\left\|\prod_{\ell=\ell_{2}}^{\ell_{1}}\left(\widetilde{D}_{i, \ell}+D_{i, \ell}^{\prime}\right) \widetilde{W}\right\|_{2} \leq O\left(L^{7}\right) .
$$

On the other hand, since $\left\|W^{\prime}\right\| \leq \frac{\tau_{0}}{\sqrt{m}}$, we also have

$$
\left\|\prod_{\ell=\ell_{2}}^{\ell_{1}}\left(\widetilde{D}_{i, \ell}+D_{i, \ell}^{\prime}\right)\left(\widetilde{W}+W^{\prime}\right)\right\|_{2} \leq O\left(L^{7}\right)
$$

by expanding out the $2^{\ell_{2}-\ell_{1}+1}$ terms with respect to $W^{\prime}$.

### 6.12.3 Backward

We state the main result in this section.

Lemma 6.12.8 (backward stability, c.f. (6.10)). Letting $\varrho=n L d \log m$, for any $\tau_{0} \in$ $\left[\varrho^{-100}, \varrho^{100}\right]$, with probability at least $1-e^{-\Omega\left(L^{6} \tau_{0}^{4 / 3} m^{1 / 3}\right)}$ over the randomness of $\widetilde{W}, A, B$, for every $W^{\prime} \in \mathbb{R}^{m \times m}$ with $\left\|W^{\prime}\right\|_{2} \leq \frac{\tau_{0}}{\sqrt{m}}$, for every $1 \leq \ell_{1} \leq \ell_{2} \leq L$, for every $i \in[n]$, the followings hold
(a) $\left\|a^{\top} B \prod_{\ell=\ell_{2}}^{\ell_{1}} \widetilde{D}_{i, \ell} \widetilde{W}-a^{\top} B \prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} \widetilde{W}\right\|_{2} \leq O\left(\tau_{0}^{1 / 3} L^{6} \log m \cdot m^{1 / 3}\right) \cdot\|a\|_{2}$.
(b) $\left\|a^{\top} B \prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} \widetilde{W}-a^{\top} B \prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} W\right\|_{2} \leq O\left(\tau_{0} L^{15}\right) \cdot\|a\|_{2}$.
(c) $\left\|a^{\top} B \prod_{\ell=\ell_{2}}^{\ell_{1}} \widetilde{D}_{i, \ell} \widetilde{W}-a^{\top} B \prod_{\ell=\ell_{2}}^{\ell_{1}} \widetilde{D}_{i, \ell} W\right\|_{2} \leq O\left(\tau_{0} L^{7}\right) \cdot\|a\|_{2}$.
(d) $\left\|a^{\top} B \prod_{\ell=\ell_{2}}^{\ell_{1}} \widetilde{D}_{i, \ell} W-a^{\top} B \prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} W\right\|_{2} \leq O\left(\tau_{0}^{1 / 3} L^{6} \log m \cdot m^{1 / 3}\right) \cdot\|a\|_{2}$.

We first use Lemma 6.12 .1 to derive that, with probability $\geq 1-e^{-\Omega\left(L^{6} \tau_{0}^{4 / 3} m^{1 / 3}\right)}$, $\left\|D_{i, \ell}^{\prime}\right\|_{0} \leq s^{2} m^{2 / 3}$, for parameter $s=O\left(L^{5 / 3} \tau_{0}^{1 / 3}\right)$. We prove the five statements separately.

Proof of Lemma 6.12.8a. Without loss of generality we assume $\|a\|=1$. We define set $\mathcal{C}$ to be

$$
\mathcal{C}=\left\{\prod_{\ell=\ell_{2}}^{\ell_{1}}\left(D_{i, \ell}^{\prime}\right)^{c_{\ell}}\left(\widetilde{D}_{i, \ell}\right)^{1-c_{\ell}} \widetilde{W} \mid c_{\ell_{1}}, \ldots, c_{\ell_{2}} \in\{0,1\}\right\}
$$

For each $\ell \in\left[\ell_{2}-\ell_{1}\right]$, we define set $\mathcal{C}_{\ell}$ to be

$$
\mathcal{C}_{\ell}=\left\{C \in \mathcal{C} \mid D^{\prime} \text { appears } \ell \text { times in } C\right\}
$$

Ignoring subscripts for the ease of presentation, by Corollary 6.11.15, we know

$$
\left\|a^{\top} B \widetilde{D} \widetilde{W} \cdots \widetilde{D} \widetilde{W} D^{\prime}\right\|_{2} \leq(s \log m) m^{1 / 3}
$$

By Corollary 6.11.12, we know

$$
\left\|D^{\prime} \widetilde{W} \widetilde{D} \cdots \widetilde{D} \widetilde{W} D^{\prime}\right\|_{2} \leq(s \log m) m^{-1 / 6}
$$

Thus, for each $\ell$ and $C \in \mathcal{C}_{\ell}$,

$$
\begin{aligned}
\left\|a^{\top} B C\right\|_{2} & \leq\left\|a^{\top} B \widetilde{D} \widetilde{W} \cdots \widetilde{D} \widetilde{W} D^{\prime}\right\|_{2} \cdot\left\|D^{\prime} \widetilde{W} \widetilde{D} \cdots \widetilde{D} \widetilde{W} D^{\prime}\right\|_{2}^{\ell-1} \cdot\left\|D^{\prime} \widetilde{W} \widetilde{D} \cdots \widetilde{D} \widetilde{W}\right\|_{2} \\
& \leq\left(s \cdot m^{1 / 3} \log m\right) \times\left(s \cdot m^{-1 / 6} \log m\right)^{\ell-1} \times O\left(L^{3}\right)
\end{aligned}
$$

Finally, the LHS of the Lemma 6.12 .8 a becomes

$$
\begin{aligned}
\text { LHS } & =\left\|\sum_{\ell=1}^{\ell_{2}-\ell_{1}} \sum_{C \in \mathfrak{C}_{\ell}} a^{\top} B C\right\|_{2} \leq \sum_{\ell=1}^{\ell_{2}-\ell_{1}} \sum_{C \in \mathfrak{C}_{\ell}}\left\|a^{\top} B C\right\|_{2} \\
& \leq \sum_{\ell=1}^{\ell_{2}-\ell_{1}}\binom{\ell_{2}-\ell_{1}}{\ell}\left(s \cdot m^{1 / 3} \log m\right) \times\left(s \cdot m^{-1 / 6} \log m\right)^{\ell-1} \times O\left(L^{3}\right) \leq O\left(L^{4} \cdot s \log m \cdot m^{1 / 3}\right) .
\end{aligned}
$$

Proof of Lemma 6.12.8b. Again we assume $\|a\|=1$ without loss of generality. We define set $\mathcal{C}$ to be

$$
\mathcal{C}=\left\{\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell}\left(W^{\prime}\right)^{c_{\ell}} \widetilde{W}^{1-c_{\ell}} \mid c_{\ell_{1}}, \ldots, c_{\ell_{2}} \in\{0,1\}\right\}
$$

For each $\ell \in\left[\ell_{2}-\ell_{1}\right]$, we define set $\mathcal{C}_{\ell}$

$$
\mathcal{C}_{\ell}=\left\{C \in \mathcal{C} \mid W^{\prime} \text { appears } \ell \text { times in } C\right\}
$$

Ignoring subscripts for the ease of presentation, we have for each $C \in \mathcal{C}_{\ell}$,

$$
\begin{align*}
\left\|a^{\top} B C\right\|_{2} & \leq\left\|a^{\top} B\right\|_{2} \cdot\|\widetilde{W} D \cdots \widetilde{W} D\|_{2} \cdot\left\|W^{\prime}\right\|_{2} \cdot\left(\|D \widetilde{W} \cdots \widetilde{W} D\|_{2}\left\|W^{\prime}\right\|_{2}\right)^{\ell-1} \cdot\|D \widetilde{W} \cdots D \widetilde{W}\|_{2}  \tag{6.43}\\
& \leq O\left(\sqrt{m} L^{7}\right) \cdot\left(\frac{\tau_{0}}{\sqrt{m}}\right) \cdot\left(\frac{O\left(\tau_{0} L^{7}\right)}{\sqrt{m}}\right)^{\ell-1} \cdot O\left(L^{7}\right) \tag{6.44}
\end{align*}
$$

Above, we have used $\left\|W^{\prime}\right\|_{2} \leq \frac{\tau_{0}}{\sqrt{m}},\|\widetilde{W} D \cdots \widetilde{W} D\|_{2} \leq O\left(L^{7}\right)$ from Lemma 6.12.6, and $\|B\|_{2} \leq O(\sqrt{m})$ with high probability. Finally, by triangle inequality, the LHS of Lemma 6.12 .8 b becomes

$$
\begin{aligned}
\text { LHS } & =\left\|\sum_{\ell=1}^{\ell_{2}-\ell_{1}} \sum_{C \in \mathfrak{C}_{\ell}} a^{\top} B C\right\|_{2} \leq \sum_{\ell=1}^{\ell_{2}-\ell_{1}} \sum_{C \in \mathfrak{C}_{\ell}}\left\|a^{\top} B C\right\|_{2} \\
& \leq \sum_{\ell=1}^{\ell_{2}-\ell_{1}}\binom{\ell_{2}-\ell_{1}}{\ell} O\left(\sqrt{m} L^{7}\right) \cdot\left(\frac{\tau_{0}}{\sqrt{m}}\right) \cdot\left(\frac{O\left(\tau_{0} L^{7}\right)}{\sqrt{m}}\right)^{\ell-1} \cdot O\left(L^{7}\right) \\
& \leq O\left(\tau_{0} L^{15}\right)
\end{aligned}
$$

Proof of Lemma 6.12.8c. It is similar to the proof of Lemma 6.12.8b.

Proof of Lemma 6.12.8d. It is similar to the proof of Lemma 6.12.8a.

### 6.12.4 Special Rank-One Perturbation

We prove two lemmas regarding the special rank-one perturbation with respect to a coordinate $k \in[m]$. The first one talks about forward propagation.

Lemma 6.12 .9 (c.f. (6.11)). Let $\rho=n L d \log m, \varrho=n L d \delta^{-1} \log (m / \epsilon), \tau_{0} \in\left[\varrho^{-100}, \varrho^{50}\right]$, $N \in\left[1, \varrho^{100}\right]$. If $\widetilde{W}, A$ are at random initialization, then with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$, for any rank-one perturbation $W^{\prime}=y z^{\top}$ with $\|z\|=1,\|y\|_{0}=N,\|y\|_{\infty} \leq \frac{\tau_{0}}{\sqrt{m}}$, it satisfies

$$
\forall i \in[n], \forall \ell \in[L], \forall k \in[m]: \quad\left|\left(\left(\widetilde{W}+W^{\prime}\right) h_{i, \ell}^{\prime}\right)_{k}\right| \leq O\left(\frac{\rho^{8} N^{2 / 3} \tau_{0}^{5 / 6}}{m^{2 / 3}}\right)
$$

Proof. Without loss of generality we only prove the statement for fixed $i, \ell, k$ because one can take union bound at the end. We can rewrite $W h_{i, \ell}^{\prime}$ as follows

$$
\begin{aligned}
\left(\widetilde{W}+W^{\prime}\right) h_{i, \ell}^{\prime} & =W^{\prime} h_{i, \ell}^{\prime}+\widetilde{W} h_{i, \ell}^{\prime} \\
& =W^{\prime} h_{i, \ell}^{\prime}+\widetilde{W}\left(h_{i, \ell}-\widetilde{h}_{i, \ell}\right) \\
& =W^{\prime} h_{i, \ell}^{\prime}+\widetilde{W}\left(\left(\widetilde{D}_{i, \ell}+D_{i, \ell}^{\prime}\right) g_{i, \ell}-\widetilde{D}_{i, \ell} \widetilde{g}_{i, \ell}\right) \\
& =W^{\prime} h_{i, \ell}^{\prime}+\widetilde{W} \widetilde{D}_{i, \ell} g_{i, \ell}^{\prime}+\widetilde{W} D_{i, \ell}^{\prime} g_{i, \ell} \\
& \left.=W^{\prime} h_{i, \ell}^{\prime}+\widetilde{W} \widetilde{D}_{i, \ell}\left(\widetilde{W}+W^{\prime}\right) h_{i, \ell-1}-\widetilde{W} \widetilde{h}_{i, \ell-1}\right)+\widetilde{W} D_{i, \ell}^{\prime} g_{i, \ell} \\
& \left.=W^{\prime} h_{i, \ell}^{\prime}+\widetilde{W} \widetilde{D}_{i, \ell} \widetilde{W} h_{i, \ell-1}^{\prime}+W^{\prime} h_{i, \ell-1}\right)+\widetilde{W} D_{i, \ell}^{\prime} g_{i, \ell} \\
& =W^{\prime} h_{i, \ell}^{\prime}+\widetilde{W} \widetilde{D}_{i, \ell}^{\widetilde{W} h_{i, \ell-1}^{\prime}}+\widetilde{W} \widetilde{D}_{i, \ell} W^{\prime} h_{i, \ell-1}+\widetilde{W} D_{i, \ell}^{\prime} g_{i, \ell}
\end{aligned}
$$

After recursively calculating as above, using $h_{i, 0}^{\prime}=0$, and applying triangle inequality, we have
$\left|\left(\left(\widetilde{W}+W^{\prime}\right) h_{i, \ell}^{\prime}\right)_{k}\right| \leq\left|\left(W^{\prime} h_{i, \ell}^{\prime}\right)_{k}\right|+\sum_{a=0}^{\ell}\left|\left(\widetilde{W} \widetilde{D}_{i, \ell} \cdots \widetilde{W} \widetilde{D}_{i, a} W^{\prime} h_{i, a}\right)_{k}\right|+\sum_{a=0}^{\ell}\left|\left(\widetilde{W} \widetilde{D}_{i, \ell} \cdots \widetilde{W} \widetilde{D}_{i, a+1} \widetilde{W} D_{i, a}^{\prime} g_{i, a}\right)_{k}\right|$

We bound the three types of terms on the RHS of (6.45) as follows.

- First, we can bound

$$
\left|\left(W^{\prime} h_{i, \ell}^{\prime}\right)_{k}\right|=\left|e_{k}^{\top} y z^{\top} h_{i, \ell}^{\prime}\right|=\left|e_{k}^{\top} y\right| \cdot\left|z^{\top} h_{i, \ell}^{\prime}\right| \leq\|y\|_{\infty} \cdot\left\|h_{i, \ell}^{\prime}\right\|_{2} \leq \frac{\tau_{0}}{\sqrt{m}} \cdot\left(\frac{O\left(L^{6} \tau_{0} \sqrt{N}\right)}{\sqrt{m}}\right)
$$

where the last step follows by $\|y\|_{\infty} \leq \frac{\tau_{0}}{\sqrt{m}},\left\|W^{\prime}\right\|_{2} \leq \frac{\sqrt{N} \tau_{0}}{\sqrt{m}}$ and Lemma 6.12.1a.

- Second, we can bound

$$
\begin{aligned}
\left|\left(\widetilde{W} \widetilde{D}_{i, \ell} \cdots \widetilde{W} \widetilde{D}_{i, a} W^{\prime} h_{i, a}\right)_{k}\right| & =\left|e_{k}^{\top} \widetilde{W} \widetilde{D}_{i, \ell} \cdots \widetilde{W} \widetilde{D}_{i, a} y z^{\top} h_{i, a}\right| \\
& =\left|e_{k}^{\top} \widetilde{W} \widetilde{D}_{i, \ell} \cdots \widetilde{W} \widetilde{D}_{i, a} y\right| \cdot\left|z^{\top} h_{i, a}\right| \\
& \stackrel{\otimes}{ \pm}\left|e_{k}^{\top} \widetilde{W} \widetilde{D}_{i, \ell} \cdots \widetilde{W} \widetilde{D}_{i, a} y\right| \cdot\left\|h_{i, a}\right\|_{2} \\
& \leq\left|e_{k}^{\top} \widetilde{W} \widetilde{D}_{i, \ell} \cdots \widetilde{W} \widetilde{D}_{i, a} y\right| \cdot O(L) \\
& \leq N \cdot\|y\|_{\infty} \max _{j}\left|e_{k}^{\top} \widetilde{W} \widetilde{D}_{i, \ell} \cdots \widetilde{W} \widetilde{D}_{i, a} e_{j}\right| \cdot O(L) \\
& \stackrel{(3}{\leq} N \cdot\|y\|_{\infty} \cdot \frac{\rho}{\sqrt{m}} \cdot O(L) \\
& \stackrel{\oplus}{\leq} N \cdot\left(\tau_{0} \frac{1}{\sqrt{m}}\right) \cdot \frac{\rho}{\sqrt{m}} \cdot O(L) \leq O\left(\frac{N \tau_{0} \rho L}{m}\right) .
\end{aligned}
$$

where (1) follows by $\|z\|_{2} \leq 1$, (2) follows by Lemma 6.11.1, (3) follows by Corollary 6.11.13, and (4) follows by $\|y\|_{\infty} \leq \tau_{0} \frac{1}{\sqrt{m}}$.

- Third, we can bound,

$$
\begin{aligned}
\left|\left(\widetilde{W} \widetilde{D}_{i, \ell} \cdots \widetilde{W} \widetilde{D}_{i, a+1} \widetilde{W} D_{i, a}^{\prime} g_{i, a}\right)_{k}\right| & \leq\left\|e_{k} \widetilde{W} \widetilde{D}_{i, \ell} \cdots \widetilde{W} \widetilde{D}_{i, a+1} \widetilde{W} D_{i, a}^{\prime}\right\|_{2} \cdot\left\|D_{i, a}^{\prime} g_{i, a}\right\|_{2} \\
& \leq O\left(\frac{L^{5 / 3} \tau_{0}^{1 / 3} N^{1 / 6} \log m}{m^{1 / 6}}\right) \cdot O\left(\frac{L^{5} \tau_{0}^{1 / 2} \sqrt{N}}{\sqrt{m}}\right)
\end{aligned}
$$

where the last step follows by Corollary 6.11 .12 (which relies on Lemma 6.12.1b to give $\left.s=O\left(L^{5 / 3} \tau_{0}^{1 / 3} N^{1 / 6}\right)\right)$ together with Lemma 6.12.1c.

Putting the three bounds into (6.45) (where the term one is the dominating term), we have

$$
\left|\left(\left(\widetilde{W}+W^{\prime}\right) h_{i, \ell}^{\prime}\right)_{k}\right| \leq O\left(\frac{\rho^{8} N^{2 / 3} \tau_{0}^{5 / 6}}{m^{2 / 3}}\right)
$$

The next one talks about backward propagation.
Lemma 6.12.10 (c.f. (6.12)). Let $\rho=n L d \log m, \varrho=n L d \delta^{-1} \log (m / \epsilon), \tau_{0} \in\left[\varrho^{-100}, \varrho^{50}\right]$, $N \in\left[1, \varrho^{100}\right]$. If $\widetilde{W}, A, B$ are at random initialization, and given some perturbation $W^{\prime}=y z^{\top}$ with $\|z\|=1,\|y\|_{0}=N,\|y\|_{\infty} \leq \frac{\tau_{0}}{\sqrt{m}}$ where $W^{\prime}$ can only depend on the randomness of $\widetilde{W}$ and A (but not on B), then with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$, for all $i \in[n]$, all $1 \leq \ell_{1} \leq \ell_{2} \leq L$, all $k \in[m]:$
(a) $\left|a^{\top} B\left(\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} \widetilde{W}\right) e_{k}-a^{\top} B\left(\prod_{\ell=\ell_{2}}^{\ell_{1}} \widetilde{D}_{i, \ell} \widetilde{W}\right) e_{k}\right| \leq\|a\|_{2} \cdot O\left(\frac{\rho^{3} \tau_{0}^{1 / 3} N^{1 / 6}}{m^{1 / 6}}\right)$,
(b) $\left|a^{\top} B\left(\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} W\right) e_{k}-a^{\top} B\left(\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} \widetilde{W}\right) e_{k}\right| \leq\|a\|_{2} \cdot O\left(\frac{\rho^{12} N^{5 / 6} \tau_{0}^{5 / 3}}{m^{1 / 3}}\right)$.

Proof. We prove for a fixed $i \in[n]$ and drop the subscript in $i$ for notational simplicity. Without loss of generality we assume $\|a\|=1$.
(a) We define set $\mathcal{C}$ to be

$$
\mathcal{C}=\left\{\prod_{\ell=\ell_{2}}^{\ell_{1}}\left(D_{\ell}^{\prime}\right)^{c_{\ell}}\left(\widetilde{D}_{\ell}\right)^{1-c_{\ell}} \widetilde{W} \mid c_{\ell} \in\{0,1\}, \forall \ell \in\left[\ell_{1}, \ell_{2}\right]\right\} .
$$

For each $\ell \in\left[\ell_{2}-\ell_{1}\right]$, we define set $\mathfrak{C}_{\ell}$ to be

$$
\mathcal{C}_{\ell}=\left\{C \in \mathcal{C} \mid D^{\prime} \text { appears } \ell \text { times in } C\right\} .
$$

Ignoring the subscripts for the ease of presentation, by Lemma 6.11.8, we know

$$
\left\|\widetilde{D} \widetilde{W} \cdots \widetilde{D} \widetilde{W} \cdot D^{\prime}\right\|_{2} \leq\|\widetilde{D} \widetilde{W} \cdots \widetilde{D} \widetilde{W}\|_{2} \leq O\left(L^{3}\right)
$$

By Corollary 6.11.12 (which relies on Lemma 6.12 .1 b to give $s=O\left(L^{5 / 3} \tau_{0}^{1 / 3} N^{1 / 6}\right)$ ), we know

$$
\begin{array}{r}
\left\|D^{\prime} \widetilde{W} \widetilde{D} \cdots \widetilde{D} \widetilde{W} D^{\prime}\right\|_{2} \leq O\left(\frac{L^{5 / 3} \tau_{0}^{1 / 3} N^{1 / 6}}{m^{1 / 6}} \log m\right) \\
\forall k \in[m]:\left\|D^{\prime} \widetilde{W} \widetilde{D} \cdots \widetilde{D} \widetilde{W} e_{k}\right\|_{2} \leq O\left(\frac{L^{5 / 3} \tau_{0}^{1 / 3} N^{1 / 6}}{m^{1 / 6}} \log m\right)
\end{array}
$$

Thus, for each $k \in[m]$ and $C \in \mathcal{C}_{\ell}$, we can bound it

$$
\left\|C e_{k}\right\| \leq L^{3} \times\left(O\left(\frac{L^{5 / 3} \tau_{0}^{1 / 3} N^{1 / 6}}{m^{1 / 6}} \log m\right)\right)^{\ell}
$$

As a consequence, letting

$$
v=\left(\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{i, \ell} \widetilde{W}\right) e_{k}-\left(\prod_{\ell=\ell_{2}}^{\ell_{1}} \widetilde{D}_{i, \ell} \widetilde{W}\right) e_{k}
$$

we have

$$
\begin{aligned}
\|v\| & =\left\|\sum_{\ell=1}^{\ell_{2}-\ell_{1}} \sum_{C \in \mathrm{C}_{\ell}} C e_{k}\right\| \\
& \leq \sum_{\ell=1}^{\ell_{2}-\ell_{1}} \sum_{C \in \mathrm{e}_{\ell}}\left\|C e_{k}\right\| \leq \sum_{\ell=1}^{\ell_{2}-\ell_{1}}\binom{L}{\ell} \cdot\left(O\left(\frac{L^{5 / 3} \tau_{0}^{1 / 3} N^{1 / 6}}{m^{1 / 6}} \log m\right)\right)^{\ell} \leq O\left(\frac{\rho^{3} \tau_{0}^{1 / 3} N^{1 / 6}}{m^{1 / 6}}\right)
\end{aligned}
$$

Finally, we use the randomness of $B$ (recall that $v$ does not depend on $B$ ), we conclude that with probability at least $1-e^{-\Omega(m)}$, it satisfies

$$
\mathrm{LHS}=\left\|a^{\top} B v\right\| \leq O\left(\frac{\rho^{3} \tau_{0}^{1 / 3} N^{1 / 6}}{m^{1 / 6}}\right)
$$

(b) This time, we define set $\mathcal{C}$ to be

$$
\mathcal{C}=\left\{\prod_{\ell=\ell_{2}}^{\ell_{1}} D_{\ell}\left(W^{\prime}\right)^{c_{\ell}} \widetilde{W}^{1-c_{\ell}} \mid c_{\ell} \in\{0,1\}, \forall \ell \in\left[\ell_{1}, \ell_{2}\right]\right\}
$$

For each $\ell \in\left[\ell_{2}-\ell_{1}\right]$, we define set $\mathcal{C}_{\ell}$

$$
\mathcal{C}_{\ell}=\left\{C \in \mathcal{C} \mid W^{\prime} \text { appears } \ell \text { times in } C\right\} .
$$

One can carefully derive that ${ }^{19}$

$$
\left|a^{\top} B D \widetilde{W} \cdots \widetilde{W} D y\right| \leq O\left(\frac{\rho^{4} N^{5 / 6} \tau_{0}^{5 / 3}}{m^{1 / 3}}\right)
$$

By Lemma 6.12 .6 we have

$$
\begin{aligned}
&\left|z^{\top} D \widetilde{W} \cdots \widetilde{W} D y\right| \leq\|z\|_{2} \cdot O\left(L^{7}\right) \cdot\|y\|_{2} \leq O\left(\frac{L^{7} \sqrt{N} \tau_{0}}{\sqrt{m}}\right) \\
&\left|z^{\top} D \widetilde{W} \cdots \widetilde{W} D e_{k}\right| \leq\|z\|_{2} \cdot O\left(L^{7}\right) \cdot\left\|e_{k}\right\|_{2} \leq O\left(L^{7}\right)
\end{aligned}
$$

Therefore, for each $C \in \mathcal{C}_{\ell}$ we can upper bound $\left\|a^{\top} B C\right\|_{2}$ as follows

$$
\left|a^{\top} B C e_{k}\right| \leq O\left(\frac{\rho^{4} N^{5 / 6} \tau_{0}^{5 / 3}}{m^{1 / 3}}\right) \cdot\left(O\left(\frac{L^{7} \sqrt{N} \tau_{0}}{\sqrt{m}}\right)\right)^{\ell-1} \cdot O\left(L^{7}\right)
$$

[^46] we have with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$
$$
\left|a^{\top} B \widetilde{D} \widetilde{W} \cdots \widetilde{W} \widetilde{D} y\right| \leq\|a\|_{2} \cdot \sqrt{|N|} \log m \cdot\|y\|_{2} \leq\|y\|_{\infty}\|y\|_{0} \cdot \rho \leq \frac{\tau_{0}|N|}{\sqrt{m}}
$$

Using binomial expansion again, we have

$$
\begin{aligned}
\left|a^{\top} B \widetilde{D} \widetilde{W} \cdots \widetilde{W} \widetilde{D} y-a^{\top} B D \widetilde{W} \cdots \widetilde{W} D y\right| & \leq \sum_{\ell=1}^{L}\binom{L}{\ell}\left\|a^{\top} B \widetilde{D} \widetilde{W} \cdots \widetilde{W} D^{\prime}\right\| \cdot\left\|D^{\prime} \widetilde{W} \cdots \widetilde{W} D^{\prime}\right\|_{2}^{\ell-1} \cdot\left\|D^{\prime} \widetilde{W} \cdots \widetilde{W} \widetilde{D} y\right\| \\
& \leq \sum_{\ell=1}^{L}\binom{L}{\ell}\left(s m^{1 / 3} \log m\right) \cdot\left(\frac{s \log m}{m^{1 / 6}}\right)^{\ell-1} \cdot\left(\frac{s \log m}{m^{1 / 6}}\right) \cdot\|y\| \\
& \leq O(L)\left(s m^{1 / 3} \log m\right)\left(\frac{s \log m}{m^{1 / 6}}\right) \cdot \frac{\sqrt{N} \tau_{0}}{\sqrt{m}}=O\left(\frac{\rho^{4} N^{5 / 6} \tau_{0}^{5 / 3}}{m^{1 / 3}}\right)
\end{aligned}
$$

Here, we have used Corollary 6.11 .15 and Corollary 6.11 .12 with parameter $s=O\left(L^{5 / 3} \tau_{0}^{1 / 3} N^{1 / 6}\right)$ chosen from Lemma 6.12.1b, as well as $\|y\| \leq \frac{\sqrt{N} \tau_{0}}{\sqrt{m}}$.
and therefore

$$
\begin{aligned}
\text { LHS } & =\left|\sum_{\ell=1}^{\ell_{2}-\ell_{1}} \sum_{C \in \mathcal{C}_{\ell}} a^{\top} B C e_{k}\right| \\
& \leq \sum_{\ell=1}^{\ell_{2}-\ell_{1}}\binom{\ell_{2}-\ell_{1}}{\ell} \cdot O\left(\frac{\rho^{4} N^{5 / 6} \tau_{0}^{5 / 3}}{m^{1 / 3}}\right) \cdot\left(O\left(\frac{L^{7} \sqrt{N} \tau_{0}}{\sqrt{m}}\right)\right)^{\ell-1} \cdot O\left(L^{7}\right) \leq O\left(\frac{\rho^{12} N^{5 / 6} \tau_{0}^{5 / 3}}{m^{1 / 3}}\right) .
\end{aligned}
$$

### 6.13 Indicator and Backward Coordinate Bounds

In this section, let $W \in \mathbb{R}^{m \times m}, A \in \mathbb{R}^{m \times d_{x}}$ and $B \in \mathbb{R}^{d \times m}$ denote three random matrices where each entry of $W$ and $A$ is i.i.d. sampled from $\mathcal{N}\left(0, \frac{2}{m}\right)$ and each entry of $B$ is i.i.d. sampled from $\mathcal{N}\left(0, \frac{1}{d}\right)$. We let $\left\{\operatorname{loss}_{i, \ell} \in \mathbb{R}^{d}: i \in[n], \ell \in[L] \backslash\{1\}\right\}$ be arbitrary fixed vectors, and let $i^{*}, \ell^{*}=\arg \max _{i, \ell}\left\|\operatorname{loss}_{i, \ell}\right\|_{2}$.

- In Section 6.13.1, we prove that there are sufficiently many coordinates $k \in[m]$ such that the quantity $\left|\left(g_{i, \ell+1}\right)_{k}\right|=\left|\left(W h_{i, \ell}+A x_{i, \ell+1}\right)_{k}\right|$ is small for $(i, \ell)=\left(i^{*}, \ell^{*}\right)$ but large for $(i, \ell) \in[n] \times\left\{\ell^{*}, \ell^{*}+1, L\right\} \backslash\left(i^{*}, \ell^{*}\right)$.
- In Section 6.13.2, we prove that there are sufficiently many coordinates $k \in[m]$ such that $\left|\sum_{a}\left(\operatorname{Back}_{i^{*}, \ell^{*} \rightarrow a}^{\top} \cdot \operatorname{loss}_{i^{*}, a}\right)_{k}\right|$ is large.


### 6.13.1 Indicator Coordinate Bound

For this section, we fix some choice of $i^{*} \in[n]$ and $\ell^{*} \in[L] \backslash\{1\}$, and define two parameters

## Definition 6.13.1.

$$
\beta_{+} \stackrel{\text { def }}{=} \frac{\delta}{\rho^{2}} \quad \text { and } \quad \beta_{-} \stackrel{\text { def }}{=} \frac{\delta}{\rho^{10}}
$$

Lemma 6.13 .1 (indicator coordinate bound, c.f. (6.15)). Suppose $W$ and $A$ follow from random initialization. Given fixed set $N_{1} \subseteq[m]$ and define

$$
N_{4} \stackrel{\text { def }}{=}\left\{k \in N_{1} \left\lvert\,\left[\begin{array}{ll}
\left|\left(W h_{i, \ell}+A x_{i, \ell+1}\right)_{k}\right| \leq \beta_{-} / \sqrt{m}, & \text { if } i=i^{*}, \ell=\ell^{*} ; \\
\left|\left(W h_{i, \ell}+A x_{i, \ell+1}\right)_{k}\right| \geq \beta_{+} / \sqrt{m}, & \text { if } i \neq i^{*} \text { and } \ell=\ell^{*} ; \\
\left|\left(W h_{i, \ell}+A x_{i, \ell+1}\right)_{k}\right| \geq \beta_{+} / \sqrt{m}, & \text { if } i \in[n] \text { and } \ell>\ell^{*} .
\end{array}\right\}\right.\right.
$$

Then, as long as $\left|N_{1}\right| \geq n^{2} L^{2} / \beta_{-}$, we have

$$
\operatorname{Pr}_{W, A}\left[\left|N_{4}\right| \geq \frac{\beta_{-}\left|N_{1}\right|}{64 L}\right] \geq 1-e^{-\Omega\left(\beta_{-}\left|N_{1}\right| / L\right)} .
$$

Lemma 6.13.2. Suppose $W$ and $A$ follow from random initialization. Given fixed set $N_{1} \subseteq$ [ $m$ ] and define

$$
N_{2} \stackrel{\text { def }}{=}\left\{k \in N_{1}| |\left\langle W_{k}, h_{i^{*}, \ell^{*}}\right\rangle+\left\langle A_{k}, x_{i^{*}, \ell^{*}+1}\right\rangle \left\lvert\, \leq \frac{\beta_{-}}{\sqrt{m}}\right.\right\} .
$$

Then, we have

$$
\operatorname{Pr}_{W, A}\left[\left|N_{2}\right| \geq \frac{\beta_{-}\left|N_{1}\right|}{16 L}\right] \geq 1-e^{-\Omega\left(\beta_{-}\left|N_{1}\right| / L\right)} .
$$

Proof. Let $i=i^{*}$ and $\ell=\ell^{*}$. Similar to the proof of Lemma 6.11.1, we can write

$$
y=W h_{i, \ell}+A x_{i, \ell+1}=W U_{\ell} U_{\ell}^{\top} h_{i, \ell}+A x_{i, \ell+1}=M_{1} z_{1}+M_{2} z_{2}
$$

where $M_{1} \in \mathbb{R}^{m \times n \ell}, M_{2}=A, z_{1}=U_{\ell}^{\top} h_{i, \ell}$ and $z_{2}=x_{i, \ell+1}$. Again, the entries of $M_{1}$ and $M_{2}$ are i.i.d. from $\mathcal{N}\left(0, \frac{2}{m}\right)$.

Now, suppose $z_{1}$ is fixed (instead of depending on the randomness of $W$ and $A$ ), and it satisfies $\left\|z_{1}\right\| \leq 6 L$. We have each entry of $y$ is is distributed as $\mathcal{N}\left(0, \frac{2\left\|z_{1}\right\|^{2}+2\left\|z_{2}\right\|^{2}}{m}\right)$. By property of Gaussian (see Fact 6.10.12), we have for each $y \in[m]$,

$$
\operatorname{Pr}\left[\left|y_{k}\right| \leq \frac{0.9 \beta_{-}}{\sqrt{m}}\right] \geq \frac{2}{3} \frac{0.9 \beta_{-}}{\sqrt{2\left\|z_{1}\right\|^{2}+2\left\|z_{2}\right\|^{2}}} \geq \frac{\beta_{-}}{15 L}
$$

Since we have $\left|N_{1}\right|$ independent random variables, applying a Chernoff bound, we have

$$
\operatorname{Pr}\left[\left|\left\{k \in N_{1}:\left|y_{k}\right| \leq \frac{0.9 \beta_{-}}{\sqrt{m}}\right\}\right| \geq \frac{\beta_{-}}{16 L}\left|N_{1}\right|\right] \geq 1-\exp \left(-\Omega\left(\beta_{-}\left|N_{1}\right| / L\right)\right)
$$

Finally, by taking a union bound with respect to the $\epsilon$-net over all possible choices of $z_{1} \in \mathbb{R}^{n \ell}$, we have that

$$
\operatorname{Pr}\left[\left|\left\{k \in N_{1}:\left|y_{k}\right| \leq \frac{\beta_{-}}{\sqrt{m}}\right\}\right| \geq \frac{\beta_{-}}{16 L}\left|N_{1}\right|\right] \geq 1-\exp \left(-\Omega\left(\beta_{-}\left|N_{1}\right| / L\right)\right)
$$

for all $z_{1}$ and a fixed $z_{2}$. Plugging in the random choice of $z_{1}=U_{\ell}^{\top} h_{i, \ell}$ (and we have $\left\|z_{1}\right\| \leq 6 L$ by Lemma 6.11.1), we have the desired result.

Observe that the set $N_{2}$ produced by Lemma 6.13.2 only depends on the randomness of $A, W U_{\ell^{*}-1}$ and $W h_{i^{*}, \ell^{*}}$. In other words, for any unit vector $z \in \mathbb{R}^{m}$ that is orthogonal to the columns of $U_{\ell^{*}-1}$ and orthogonal to $h_{i^{*}, \ell^{*}}$, we have that $W z$ is independent of $N_{2}$.

We next proceed to refine $N_{2}$ :

Lemma 6.13.3 $\left(i \neq i^{*}, \ell=\ell^{*}\right)$. Suppose $W$ and $A$ follow from random initialization, and $N_{2} \subseteq[m]$ is a given set (that may only depend on the randomness of $A, W U_{\ell^{*}-1}$ and $W h_{i^{*}, \ell^{*}}$ ).

Define

$$
N_{3}=\left\{k \in N_{2}| |\left\langle W_{k}, h_{i, \ell^{*}}\right\rangle+\left\langle A_{k}, x_{i, \ell^{*}+1}\right\rangle \left\lvert\, \geq \frac{\beta_{-}}{\sqrt{m}}\right., \forall i \in[n] \backslash i^{*}\right\} .
$$

Then

$$
\operatorname{Pr}\left[\left.\left|N_{3}\right| \geq \frac{1}{2}\left|N_{2}\right| \right\rvert\, A, W U_{\ell^{*}-1}, W h_{i^{*}, \ell^{*}}\right] \geq 1-e^{-\Omega\left(\left|N_{2}\right| / n\right)}
$$

Proof. First fix some $i \in[n] \backslash i^{*}$ and $\ell=\ell^{*}$. Define set $N_{3, i}$,

$$
N_{3, i} \stackrel{\text { def }}{=}\left\{k \in N_{2}| |\left\langle W_{k}, h_{i, \ell}\right\rangle+\left\langle A_{k}, x_{i, \ell+1}\right\rangle \left\lvert\, \geq \frac{\beta_{-}}{\sqrt{m}}\right.\right\} .
$$

Let $v \stackrel{\text { def }}{=} \frac{\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i^{*}, \ell}}{\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i^{*}, \ell}\right\|}$, and let column orthonormal matrix

$$
U \stackrel{\text { def }}{=} \mathrm{GS}\left(U_{\ell-1}, h_{i^{*}, \ell}\right)=\left[U_{\ell-1}, v\right] .
$$

We have

$$
\left\langle W_{k}, h_{i, \ell}\right\rangle+\left\langle A_{k}, x_{i, \ell+1}\right\rangle=\left(W U U h_{i, \ell}+A x_{i, \ell+1}+W\left(I-U U^{\top}\right) h_{i, \ell}\right)_{k} .
$$

The rest of the proof is focusing on having the $W U$ and $A$ being fixed (so the first two vectors $W U U h_{i, \ell}+A x_{i, \ell+1}$ are fixed), and letting $W\left(I-U U^{\top}\right) h_{i, \ell}$ be the only random variable. We have

$$
W\left(I-U U^{\top}\right) h_{i, \ell} \sim \mathcal{N}\left(0, \frac{2\left\|\left(I-U U^{\top}\right) h_{i, \ell}\right\|_{2}^{2} \cdot I}{m}\right) .
$$

and therefore for a fixed vector $\mu$ it satisfies
$W h_{i, \ell}+A x_{i, \ell+1}=W U U h_{i, \ell}+A x_{i, \ell+1}+W\left(I-U U^{\top}\right) h_{i, \ell} \sim \mathcal{N}\left(\mu, \frac{2\left\|\left(I-U U^{\top}\right) h_{i, \ell}\right\|_{2}^{2} \cdot I}{m}\right)$.

According to Lemma 6.11.6, with probability at least $1-\exp (-\Omega(\sqrt{m}))$, it satisfies $\|(I-$ $\left.U U^{\top}\right) h_{i, \ell} \|_{2} \geq \frac{\delta}{2}$. By property of Gaussian distribution (see Fact 6.10.12), we have for each $k \in N_{2}:$

$$
\operatorname{Pr}\left[\left.\left|\left(W h_{i, \ell}+A x_{i, \ell+1}\right)_{k}\right| \geq \frac{\beta_{-}}{\sqrt{m}} \right\rvert\, W U, A\right] \geq 1-\frac{\beta_{-}}{\delta} \geq 1-\frac{1}{4 n} .
$$

Applying Chernoff bound, with probability at least $1-e^{-\Omega\left(\left|N_{2}\right| / n\right)}$, we have $\left|N_{3, i}\right| \geq(1-$ $\left.\frac{1}{2 n}\right)\left|N_{2}\right|$. Applying union bound over all possible $i \in[n] \backslash\left\{i^{*}\right\}$, we have $N_{3}=\bigcap_{i \neq i^{*}} N_{3, i}$ has cardinality at least $\frac{\left|N_{2}\right|}{2}$.

Lemma 6.13.4 $\left(\ell>\ell^{*}\right)$. Suppose $W$ and $A$ follow from random initialization, and $N_{3} \subseteq[m]$ is a given set (that may only depend on the randomness of $A, W U_{\ell^{*}}$ ). Define

$$
N_{4}=\left\{k \in N_{3}| |\left\langle W_{k}, h_{i, \ell}\right\rangle+\left\langle A_{k}, x_{i, \ell+1}\right\rangle \left\lvert\, \geq \frac{\beta_{-}}{\sqrt{m}}\right., \forall i \in[n], \ell>\ell^{*}\right\}
$$

Then

$$
\operatorname{Pr}\left[\left.\left|N_{4}\right| \geq \frac{1}{2}\left|N_{3}\right| \right\rvert\, A, W U_{\ell^{*}}\right] \geq 1-e^{-\Omega\left(\left|N_{3}\right| / n L\right)} .
$$

Proof. Letting $N_{4, \ell^{*}} \stackrel{\text { def }}{=} N_{3}$. For any $i \in[n]$ and $\ell=\ell^{*}+1, \ell^{*}+2, \ldots, L$, we define $N_{4, i, \ell}$ as

$$
N_{4, i, \ell} \stackrel{\text { def }}{=}\left\{k \in N_{4, \ell-1}| |\left\langle W_{k}, h_{i, \ell}\right\rangle+\left\langle A_{k}, x_{i, \ell+1}\right\rangle \left\lvert\, \geq \frac{\beta_{-}}{\sqrt{m}}\right.\right\} \quad \text { and } \quad N_{4, \ell} \stackrel{\text { def }}{=} \bigcap_{i \in[n]} N_{4, i, \ell} .
$$

We rewrite

$$
\left\langle W_{k}, h_{i, \ell}\right\rangle+\left\langle A_{k}, x_{i, \ell+1}\right\rangle=\left(W U_{\ell-1} U_{\ell-1}^{\top} h_{i, \ell}+A x_{i, \ell+1}+W\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}\right)_{k} .
$$

The rest of the proof is focusing on having the $W U_{\ell-1}$ and $A$ being fixed (so the first two vectors $W U U h_{i, \ell}+A x_{i, \ell+1}$ are fixed and $N_{4, \ell-1}$ is fixed), and letting $W\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}$ be
the only random variable. We have

$$
W\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell} \sim \mathcal{N}\left(0, \frac{2\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}\right\|_{2}^{2} \cdot I}{m}\right) .
$$

and therefore for a fixed vector $\mu$ it satisfies
$W h_{i, \ell}+A x_{i, \ell+1}=W U_{\ell-1} U_{\ell-1} h_{i, \ell}+A x_{i, \ell+1}+W\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell} \sim \mathcal{N}\left(\mu, \frac{2\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}\right\|_{2}^{2} \cdot I}{m}\right)$.
According to Lemma 6.11.4, with probability at least $1-\exp (-\Omega(\sqrt{m}))$, it satisfies $\left\|\left(I-U U^{\top}\right) h_{i, \ell}\right\|_{2} \geq \frac{\delta}{2}$. By property of Gaussian distribution (see Fact 6.10.12), we have for each $k \in N_{4, \ell-1}$ :

$$
\operatorname{Pr}\left[\left.\left|\left(\widetilde{W} h_{i, \ell}+A x_{i, \ell+1}\right)_{k}\right| \geq \frac{\beta_{-}}{\sqrt{m}} \right\rvert\, W U, A\right] \geq 1-\frac{\beta_{-}}{\delta} \geq 1-\frac{1}{4 n L} .
$$

Applying Chernoff bound, we have

$$
\operatorname{Pr}\left[\left.\left|N_{4, i, \ell}\right| \geq\left(1-\frac{1}{3 n L}\right)\left|N_{4, \ell-1}\right| \right\rvert\, W U_{\ell-1}, A\right] \geq 1-e^{-\Omega\left(\left|N_{4, \ell-1}\right| / n L\right)} .
$$

Applying union bound over all possible $i \in[n]$, we have $N_{4, \ell}=\bigcap_{i \in[n]} N_{4, i, \ell}$ has cardinality at least $\left(1-\frac{1}{3 L}\right)\left|N_{4, \ell-1}\right|$. After telescoping we have $\left|N_{4}\right|=\left|N_{4, L}\right| \geq \frac{1}{2}\left|N_{3}\right|$.

### 6.13.2 Backward Coordinate Bound

Recall the following notions

Definition 6.13.2. Given matrices $W, A, B$. For each $i$, for each $\ell$, we define

$$
\text { Back }_{i \ell \rightarrow \ell}=B \in \mathbb{R}^{d \times m}
$$

For each $i$, for each $a \geq \ell+1$, we define

$$
\operatorname{Back}_{i, \ell \rightarrow a}=B D_{i, a} W \cdots D_{i, \ell+1} W \in \mathbb{R}^{d \times m}
$$

The goal of this section is to prove Lemma 6.13.5.

Lemma 6.13.5 (backward coordinate bound, c.f. (6.14)). Let $\rho=n d L \log m$ and $\varrho=$ $n L d \delta^{-1} \log (m / \epsilon)$. Let $\left\{\operatorname{loss}_{i, \ell} \in \mathbb{R}^{d}: i \in[n], \ell \in[L] \backslash\{1\}\right\}$ be fixed vectors and $i^{*}, \ell^{*}=$ $\arg \max _{i, \ell}\left\|\operatorname{loss}_{i, \ell}\right\|_{2}$. If $W \in \mathbb{R}^{m \times m}, A \in \mathbb{R}^{m \times d_{x}}, B \in \mathbb{R}^{d \times m}$ are random, and $N_{4} \subseteq[m]$ is a set with $\left|N_{4}\right| \in\left[\rho^{4}, \varrho^{100}\right]\left(N_{4}\right.$ can depend on the randomness of $W$ and $A$, but not depend on B). Define

$$
N_{5} \stackrel{\text { def }}{=}\left\{k \in N_{4}| | \sum_{a=\ell^{*}}^{L}\left(\operatorname{Back}_{i^{*}, \ell^{*} \rightarrow a}^{\top} \cdot \operatorname{loss}_{i^{*}, a}\right)_{k} \left\lvert\, \geq \frac{\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}}{6 \sqrt{d} n L}\right.\right\} .
$$

Then with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$, we have

$$
\left|N_{5}\right| \geq\left(1-\frac{1}{2 n L}\right)\left|N_{4}\right| .
$$

Proof. For notational simplicity, in the proof we use $\ell$ to denote $\ell^{*}$, use $N$ to denote $N_{4}$, and drop the subscript $i^{*}$. We denote $B \in \mathbb{R}^{d \times m}$ as $\left[b_{1}, b_{2}, \cdots, b_{m}\right]$ where each $b_{i} \in \mathbb{R}^{d}$. We denote by $C_{a, \ell+1}=D_{a} W \cdots D_{\ell+1} W$.

We define a function $v_{k}\left(b_{1}, b_{2}, \cdots, b_{m}\right) \in \mathbb{R}$ as follows

$$
v_{k}\left(b_{1}, b_{2}, \cdots, b_{m}\right) \stackrel{\text { def }}{=} \sum_{a=\ell}^{L}\left(\operatorname{Back}_{\ell \rightarrow a}^{\top} \cdot \operatorname{loss}_{i, a}\right)_{k} .
$$

We can rewrite

$$
\begin{aligned}
v_{k}\left(b_{1}, b_{2}, \cdots, b_{m}\right) & =\left(B^{\top} \cdot \operatorname{loss}_{i, \ell}+\sum_{a=\ell+1}^{L} \operatorname{Back}_{\ell \rightarrow a}^{\top} \cdot \operatorname{loss}_{i, a}\right)_{k} \\
& =\left(B^{\top} \cdot \operatorname{loss}_{i, \ell}+\sum_{a=\ell+1}^{L}\left(C_{a, \ell+1}\right)^{\top} B^{\top} \cdot \operatorname{loss}_{i, a}\right)_{k} \\
& =\left\langle b_{k}, \operatorname{loss}_{i, \ell}\right\rangle+\sum_{a=\ell+1}^{L}\left\langle\sum_{j=1}^{m}\left(C_{a, \ell+1}\right)_{k, j} \cdot b_{j}, \operatorname{loss}_{i, a}\right\rangle \\
& =\left\langle b_{k}, \operatorname{loss}_{i, \ell}+\sum_{a=\ell+1}^{L}\left(C_{a, \ell+1}\right)_{k, j} \cdot \operatorname{loss}_{i, a}\right\rangle+\sum_{j \in[m] \backslash k}\left\langle b_{j}, \sum_{a=\ell+1}^{L}\left(C_{a, \ell+1}\right)_{k, j} \cdot \operatorname{loss}_{i, a}\right\rangle
\end{aligned}
$$

where we have used $\left(C_{a, \ell+1}\right)_{k, j}$ to denote the $(k, j)$ entry of matrix $C_{a, \ell+1}$.
Choose parameter $\mathfrak{L}_{h}=\frac{3 \sqrt{d} n L}{\left\|\operatorname{loss}_{i, \ell}\right\|_{2}}$ and we can define function $h: \mathbb{R} \rightarrow[0,1]$ by

$$
h(t) \stackrel{\text { def }}{=} \begin{cases}|t| \cdot \mathfrak{L}_{h}, & \text { if }|t| \leq 1 / \mathfrak{L}_{h} \\ 1, & \text { otherwise }\end{cases}
$$

We have $h(t)$ is $\mathfrak{L}_{h}$-Lipschitz continuous.
We define two probabilistic events $E_{1}, E_{2}$.

- Event $E_{1}$ depends on the randomness of $W$ and $A$ :

$$
E_{1} \stackrel{\text { def }}{=}\left\{\left|z^{\top} C_{a, \ell+1} y\right| \leq \frac{\rho}{\sqrt{m}}\|z\|\|y\| \text { for all } a, \ell \text { and for all } 1 \text {-sparse vectors } y, z\right\}
$$

Corollary 6.11 .13 says $\operatorname{Pr}\left[E_{1}\right] \geq 1-e^{-\Omega\left(\rho^{2}\right)}$.

- Event $E_{2}$ depends on the randomness of $B$ :

$$
E_{2} \stackrel{\text { def }}{=}\left\{\left|B_{i, j}\right| \leq \frac{\rho}{\sqrt{d}} \text { for all } i, j\right\}
$$

By Gaussian tail bound, we have $\operatorname{Pr}\left[E_{2}\right] \geq 1-e^{-\Omega\left(\rho^{2}\right)}$.

In the rest of the proof, we assume $W$ and $A$ are fixed and satisfy $E_{1}$. We let $B$ be the only source of randomness but we condition on $B$ satisfies $E_{2}$.

For simplicity we use subscript $-j$ to denote $[m] \backslash j$, and subscript $-N$ to denote $[m] \backslash N$. For instance, we shall write $b=\left(b_{-k}, b_{k}\right)=\left(b_{N}, b_{-N}\right)$.

Step 1 Fixing $b_{-N}$ and letting $b_{N}$ be the only randomness, we claim for each $k \in N$ :

$$
\underset{b_{N}}{\mathbb{E}}\left[h\left(v_{k}\right)\right] \geq \operatorname{Pr}_{b_{N}}\left[\left|v_{k}\right| \geq \frac{1}{\mathfrak{L}_{h}}\right] \geq 1-\frac{1}{4 n L} .
$$

We prove this inequality as follows. We can rewrite

$$
\begin{aligned}
v_{k}= & \left\langle b_{k}, \operatorname{loss}_{i, \ell}+\sum_{a=\ell+1}^{L}\left(C_{a, \ell+1}\right)_{k, k} \cdot \operatorname{loss}_{i, a}\right\rangle+\sum_{j \in N \backslash\{k\}}\left\langle b_{j}, \sum_{a=\ell+1}^{L}\left(C_{a, \ell+1}\right)_{k, j} \cdot \operatorname{loss}_{i, a}\right\rangle \\
& +\sum_{j \in[m\rfloor \backslash N}\left\langle b_{j}, \sum_{a=\ell+1}^{L}\left(C_{a, \ell+1}\right)_{k, j} \cdot \operatorname{loss}_{i, a}\right\rangle
\end{aligned}
$$

and $v_{k}$ is distributed as Gaussian random variable $\mathcal{N}\left(\mu, \sigma^{2}\right)$, and $\mu$ and $\sigma^{2}$ are defined as follows

$$
\begin{aligned}
\mu & =\sum_{j \in[m] \backslash N}\left\langle b_{j}, \sum_{a=\ell+1}^{L}\left(C_{a, \ell+1}\right)_{k, j} \cdot \operatorname{loss}_{i, a}\right\rangle \\
\sigma^{2} & =\frac{1}{d}\left\|\operatorname{loss}_{i, \ell}+\sum_{a=\ell+1}^{L}\left(C_{a, \ell+1}\right)_{k, k} \cdot \operatorname{loss}_{i, a}\right\|_{2}^{2}+\frac{1}{d} \sum_{j \in N \backslash\{k\}}\left\|\sum_{a=\ell+1}^{L}\left(C_{a, \ell+1}\right)_{k, j} \cdot \operatorname{loss}_{i, a}\right\|_{2}^{2}
\end{aligned}
$$

Meanwhile, we calculate that

$$
\begin{align*}
\left\|\sum_{a=\ell+1}^{L}\left(C_{a, \ell+1}\right)_{k, j} \cdot \operatorname{loss}_{i, a}\right\|_{2} & \stackrel{(1)}{\leq} \sum_{a=\ell+1}^{L}\left|\left(C_{a, \ell+1}\right)_{k, j}\right| \cdot\left\|\operatorname{loss}_{i, a}\right\|_{2} \\
& \stackrel{(2)}{\leq} \sum_{a=\ell+1}^{L}\left|\left(C_{a, \ell+1}\right)_{k, j}\right| \cdot\left\|\operatorname{loss}_{i, \ell}\right\|_{2} \stackrel{(3}{\leq} \frac{L \rho}{\sqrt{m}} \cdot\left\|\operatorname{loss}_{i, \ell}\right\|_{2} \tag{6.46}
\end{align*}
$$

where (1) follows by triangle inequality, (2) follows by $i, \ell=\arg \max _{i^{\prime}, \ell^{\prime}}\left\|\operatorname{loss}_{i^{\prime}, \ell^{\prime}}\right\|$, and (3) follows from event $E_{1}$. Therefore,

$$
\begin{aligned}
\sigma^{2} & \geq \frac{1}{d}\left(\left\|\operatorname{loss}_{i, \ell}\right\|_{2}-\left\|\sum_{a=\ell+1}^{L}\left(C_{a, \ell+1}\right)_{k, j} \cdot \operatorname{loss}_{i, a}\right\|_{2}\right)^{2} \\
& \geq \frac{1}{d}\left(1-\frac{L \rho}{\sqrt{m}}\right)^{2}\left\|\operatorname{loss}_{i, \ell}\right\|_{2}^{2} \geq \frac{1}{2 d}\left\|\operatorname{loss}_{i, \ell}\right\|_{2}^{2}
\end{aligned}
$$

where the first step follows by $\|a+b\|_{2} \geq\|a\|_{2}-\|b\|_{2}$, the second step follows by (6.46). In sum, we have that $v_{k}$ is a random Gaussian with $\sigma^{2} \geq \frac{1}{2 d}\left\|\operatorname{loss}_{i, \ell}\right\|_{2}^{2}$. This means,

$$
\underset{b_{N}}{\operatorname{Pr}}\left[\left|v_{k}\right| \geq \frac{1}{\mathfrak{L}_{h}}\right] \geq \underset{b_{N}}{\operatorname{Pr}}\left[\left|v_{k}\right| \geq \frac{\left\|\operatorname{loss}_{i, \ell}\right\|_{2}}{3 \sqrt{d} n L}\right] \geq 1-\frac{1}{4 n L} .
$$

according to Fact 6.10.12.

Step 2 For every $b_{1}, b_{2}, \cdots, b_{m}$ and $b_{k}^{\prime} \in \mathbb{R}^{d}$ satisfying event $E_{2}$, for every $j \neq k$, we have

$$
\begin{align*}
\left|v_{j}\left(b_{-k}, b_{k}\right)-v_{j}\left(b_{-k}, b_{k}^{\prime}\right)\right| & =\left|\left\langle b_{k}-b_{k}^{\prime}, \sum_{a=\ell+1}^{L}\left(C_{a, \ell+1}\right)_{k, k} \cdot \operatorname{loss}_{i, a}\right\rangle\right| \\
& \leq\left\|b_{k}-b_{k}^{\prime}\right\|_{2} \cdot\left\|\sum_{a=\ell+1}^{L}\left(C_{a, \ell+1}\right)_{k, k} \cdot \operatorname{loss}_{i, a}\right\|_{2} \\
& \leq \rho \cdot \frac{L \rho}{\sqrt{m}}\left\|\operatorname{loss}_{i, \ell}\right\|_{2} \leq \frac{\rho^{3}}{\sqrt{m}}\left\|\operatorname{loss}_{i, \ell}\right\|_{2} \tag{6.47}
\end{align*}
$$

where the second step follows by $|\langle a, b\rangle| \leq\|a\|_{2} \cdot\|b\|_{2}$, the third step follows by (6.46).

Step 3 Fixing $b_{-N}$ and letting function $g\left(b_{1}, b_{2}, \cdots, b_{m}\right) \stackrel{\text { def }}{=} \sum_{k \in N} h\left(v_{k}\left(b_{1}, b_{2}, \cdots, b_{m}\right)\right)$. We have

$$
\begin{aligned}
\left|g\left(b_{j}, b_{-j}\right)-g\left(b_{j}^{\prime}, b_{-j}\right)\right| & =\left|\sum_{k \in N} h\left(v_{k}\left(b_{j}, b_{-j}\right)\right)-\sum_{k \in N} h\left(v_{k}\left(b_{j}^{\prime}, b_{-j}\right)\right)\right| \\
& \leq \sum_{k \in N}\left|h\left(v_{k}\left(b_{j}, b_{-j}\right)\right)-h\left(v_{k}\left(b_{j}^{\prime}, b_{-j}\right)\right)\right| \\
& =\left|h\left(v_{j}\left(b_{j}, b_{-j}\right)\right)-h\left(v_{j}\left(b_{j}^{\prime}, b_{-j}\right)\right)\right|+\sum_{k \in N \backslash\{j\}}\left|h\left(v_{k}\left(b_{j}, b_{-j}\right)\right)-h\left(v_{k}\left(b_{j}^{\prime}, b_{-j}\right)\right)\right| \\
& \stackrel{\perp}{\leq} 1+\sum_{k \in N \backslash\{j\}}\left|h\left(v_{k}\left(b_{j}, b_{-j}\right)\right)-h\left(v_{k}\left(b_{j}^{\prime}, b_{-j}\right)\right)\right| \\
& \leq 1+\sum_{k \in N \backslash\{j\}} \mathfrak{L}_{h} \cdot\left|v_{k}\left(b_{j}, b_{-j}\right)-v_{k}\left(b_{j}^{\prime}, b_{-j}\right)\right| \\
& \stackrel{(2}{\leq} 1+|N| \cdot \frac{3 \sqrt{d} n L}{\left\|\operatorname{loss}_{i, \ell}\right\|_{2}} \cdot\left((n L d \log m)^{3} \frac{\left\|\operatorname{loss}_{i, \ell}\right\|_{2}}{\sqrt{m}}\right) \\
& \leq 1+\frac{|N| \cdot 3(n L d \log m)^{4}}{\sqrt{m}} \leq 2 .
\end{aligned}
$$

Above, inequality (1) follows by $h(t) \in[0,1]$, inequality (2) follows from (6.47). Using McDiarmid inequality (see Lemma 6.10.17), we have

$$
\begin{aligned}
& \operatorname{Pr}_{b_{N}}\left[\left|g\left(b_{1}, \cdots, b_{m}\right)-\underset{b_{N}}{\mathbb{E}}[g]\right| \geq \epsilon\right] \leq \exp \left(-\frac{\epsilon^{2}}{8|N|}\right) \\
\Longrightarrow & \operatorname{Pr}_{b_{N}}\left[\left|g\left(b_{1}, \cdots, b_{m}\right)-\underset{b_{N}}{\mathbb{E}}[g]\right| \geq \frac{|N|}{4 n L}\right]<\exp \left(-\frac{|N|}{128(n L)^{2}}\right)
\end{aligned}
$$

where we choose $\epsilon=\frac{|N|}{4 n L}$.
Finally, combining this with Step 1 -which gives $\mathbb{E}_{b_{N}}\left[g\left(b_{1}, \cdots, b_{m}\right)\right] \geq|N|\left(1-\frac{1}{4 n L}\right)$ we have with probability at least $1-e^{-\Omega\left(|N| /(n L)^{2}\right)}$ over randomness of $b_{N}$,

$$
g\left(b_{1}, \cdots, b_{m}\right) \geq|N|\left(1-\frac{1}{2 n L}\right) .
$$

This implies at least $\left(1-\frac{1}{2 n L}\right)$-fraction of $k \in N$ will have $\left|v_{k}\right| \geq \frac{1}{2 \mathfrak{R}_{h}}=\frac{\left\|\operatorname{loss}_{i, \ell}\right\|_{2}}{6 n L \sqrt{d}}$.

### 6.14 Gradient Bound at Random Initialization (Theorem 6.14.2)

The goal of this section is to understand the lower and upper bounds of gradients. Instead of analyzing the true gradient directly -where the forward and backward propagation have correlated randomness- we assume that the loss vectors $\left\{\operatorname{loss}_{i, a}\right\}_{i \in[n], a \in\{2, \ldots, L\}}$ are fixed (no randomness) in this section, as opposed to being defined as loss ${ }_{i, a}=B h_{i, a}-y_{i, a}^{*}$ which is random. We call this the "fake loss." We define a corresponding "fake gradient" with respect to this fixed loss.

Definition 6.14.1 (fake gradient, c.f. (6.13)). Given fixed vectors $\left\{\operatorname{loss}_{i, a}\right\}_{i \in[n], a \in\{2, \ldots, L\}}$, we define

$$
\begin{aligned}
& \widehat{\nabla}_{k} f(W) \stackrel{\text { def }}{=} \sum_{i=1}^{n} \sum_{a=2}^{L} \sum_{\ell=1}^{a-1}\left(\operatorname{Back}_{i, \ell+1 \rightarrow a}^{\top} \cdot \operatorname{loss}_{i, a}\right)_{k} \cdot h_{i, \ell} \cdot \mathbf{1}_{\left\langle W_{k}, h_{i, \ell}\right\rangle+\left\langle A_{k}, x_{i, \ell+1}\right\rangle \geq 0} \\
& \widehat{\nabla} f(W) \stackrel{\text { def }}{=} \sum_{i=1}^{n} \sum_{a=2}^{L} \sum_{\ell=1}^{a-1} D_{i, \ell+1}\left(\text { Back }_{i, \ell+1 \rightarrow a}^{\top} \cdot \operatorname{loss}_{i, a}\right) \cdot h_{i, \ell}^{\top} \\
& \widehat{\nabla} f_{i}(W) \stackrel{\text { def }}{=} \sum_{a=2}^{L} \sum_{\ell=1}^{a-1} D_{i, \ell+1}\left(\operatorname{Back}_{i, \ell+1 \rightarrow a}^{\top} \cdot \operatorname{loss}_{i, a}\right) \cdot h_{i, \ell}^{\top}
\end{aligned}
$$

In our analysis, we also write it as

$$
\begin{aligned}
\widehat{\nabla}_{k} f(W) & =\sum_{i=1}^{n} \sum_{\ell=1}^{L-1}\left(u_{i, \ell}\right)_{k} \cdot h_{i, \ell} \cdot \mathbf{1}_{\left(g_{i, \ell+1}\right)_{k} \geq 0} \\
\widehat{\nabla} f(W) & =\sum_{i=1}^{n} \sum_{\ell=1}^{L-1} D_{i, \ell+1} u_{i, \ell} \cdot h_{i, \ell}^{\top} \\
\widehat{\nabla} f_{i}(W) & =\sum_{\ell=1}^{L-1} D_{i, \ell+1} u_{i, \ell} \cdot h_{i, \ell}^{\top}
\end{aligned}
$$

where vectors $u_{i, \ell}, g_{i, \ell+1} \in \mathbb{R}^{m}$ are defined as

$$
u_{i, \ell} \stackrel{\text { def }}{=} \sum_{a=\ell+1}^{L} \text { Back }_{i, \ell+1 \rightarrow a}^{\top} \cdot \operatorname{loss}_{i, a} \quad g_{i, \ell+1} \stackrel{\text { def }}{=} W h_{i, \ell}+A x_{i, \ell+1}
$$

We first state a simple upper bound on the gradients.

Lemma 6.14.1. Letting $\rho=n L d \log m$, given fixed vectors $\left\{\operatorname{loss}_{i, \ell} \in \mathbb{R}^{d}\right\}_{i \in[n], \ell \in[L] \backslash\{1\}}$, with probability at least $1-e^{-\Omega\left(m / L^{2}\right)}$ over $W, A, B$, we have

$$
\begin{array}{r}
\|\widehat{\nabla} f(W)\|_{F} \leq O\left(n L^{6} \sqrt{m}\right) \cdot \max _{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|\right\} \quad \text { and } \\
\forall i \in[n]:\left\|\widehat{\nabla} f_{i}(W)\right\|_{F} \leq O\left(L^{6} \sqrt{m}\right) \cdot \max _{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|\right\}
\end{array}
$$

With probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$, we have for every $k \in[m]$ :

$$
\left\|\widehat{\nabla}_{k} f(W)\right\|_{2} \leq O\left(n \rho L^{3}\right) \cdot \max _{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|\right\}
$$

Proof of Lemma 6.14.1. For each $i, a, \ell$, we can calculate that

$$
\begin{aligned}
\left\|D_{i, \ell+1}\left(\operatorname{Back}_{i, \ell+1 \rightarrow a}^{\top} \cdot \operatorname{loss}_{i, a}\right) \cdot h_{i, \ell}^{\top}\right\|_{F} & =\left\|D_{i, \ell+1}\left(\operatorname{Back}_{i, \ell+1 \rightarrow a}^{\top} \cdot \operatorname{loss}_{i, a}\right)\right\|_{2} \cdot\left\|h_{i, \ell}^{\top}\right\|_{2} \\
& \leq\left\|\operatorname{Back}_{i, \ell+1 \rightarrow a}\right\|_{2} \cdot\left\|\operatorname{loss}_{i, a}\right\|_{2} \cdot\left\|h_{i, \ell}^{\top}\right\|_{2} \\
& \stackrel{\oplus}{\leq} O(\sqrt{m}) \cdot O\left(L^{3}\right) \cdot O(L) \cdot\left\|\operatorname{loss}_{i, a}\right\|_{2} .
\end{aligned}
$$

where inequality (1) uses Lemma 6.11.8, Lemma 6.11.1, and $\|B\|_{2} \leq O(\sqrt{m})$ with high probability. Applying triangle inequality and using the definition of fake gradient (see Definition 6.14.1), we finish the proof of the first statement.

As for the second statement, we replace the use of Lemma 6.11.8 with Corollary 6.11.16.

The rest of this section is devoted to proving a (much more involved) lower bound on this fake gradient.

Theorem 6.14.2 (gradient lower bound at random init, restated). Letting $\rho=n L d \log m$, given fixed vectors $\left\{\operatorname{loss}_{i, \ell} \in \mathbb{R}^{d}\right\}_{i \in[n], \ell \in[L] \backslash\{1\}}$, with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$ over $W, A, B$, we have

$$
\|\widehat{\nabla} f(W)\|_{F}^{2} \geq \Omega\left(\frac{\delta}{\rho^{14}}\right) \times m \times \max _{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|^{2}\right\} .
$$

In the rest of this section, we first present a elegant way to decompose the randomness in Section 6.14.1 (motivated by smooth analysis [ST04]). We give a lower bound on the expected fake gradient in Section 6.14.2. We then calculate the stability of fake gradient against rank-one perturbations in Section 6.14.3. Finally, in Section 6.14.4, we apply our extended McDiarmid's inequality to prove Theorem 6.14.2.

### 6.14.1 Randomness Decomposition

We introduce a parameter $\theta$ as follows:
Definition 6.14.2. Choose any

$$
\theta \in\left[\rho^{4} \cdot \beta_{-}, \rho^{-3} \cdot \beta_{+}\right]
$$

where the notions $\beta_{-}$and $\beta_{+}$come from Definition 6.13.1. We have $\theta \leq \frac{\delta}{\rho}$.
In the next lemma, we introduce a way to decompose the randomness of $W$ into $W=W_{2}+W^{\prime}$ for two corollated random matrices $W_{2}$ and $W^{\prime}$. We will make sure that the entries of $W_{2}$ are also i.i.d. from $N\left(0, \frac{2}{m}\right)$. This definition requires us to choose a specific pair $\left(i^{*}, \ell^{*}\right) \in[n] \times[L]$.

Lemma 6.14.3. Suppose $W$ and $A$ follow from random initialization. Given any fixed $i^{*} \in[n], \ell^{*} \in[L]$, and parameter $\theta \in\left(0, \frac{1}{2 \rho}\right]$, letting random vector $v_{i^{*}, \ell^{*}} \in \mathbb{R}^{m}$ be defined as

$$
v_{i^{*}, \ell^{*}} \stackrel{\text { def }}{=} \frac{\left(I-U_{\ell^{*}-1} U_{\ell^{*}-1}^{\top}\right) h_{i^{*}, \ell^{*}}}{\left\|\left(I-U_{\ell^{*}-1} U_{\ell^{*}-1}^{\top}\right) h_{i^{*}, \ell^{*}}\right\|_{2}},
$$

then we can write $W=W_{2}+W^{\prime}$ for two random matrices $W_{2} \in \mathbb{R}^{m \times m}$ and $W^{\prime} \in \mathbb{R}^{m \times m}$ where
(a) entries of $W_{2}$ are i.i.d. drawn from $\mathcal{N}\left(0, \frac{2}{m}\right)$ (so $W_{2}$ is in the same distribution as $W$ );
(b) $W^{\prime}$ is correlated with $W_{2}$ and can be written as $W^{\prime}=u v_{i^{*}, \ell^{*}}^{\top}$ for some $u \in \mathbb{R}^{m}$;
(c) With probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{2}$ :

$$
\begin{aligned}
\forall k: \quad \operatorname{Pr}_{W^{\prime}}\left[\left.u_{k}=\left(W^{\prime} v_{i^{*}, \ell^{*}}\right)_{k} \geq \frac{\theta}{2 \sqrt{m}} \right\rvert\, W_{2}\right] & \geq \frac{1}{4} \quad \text { and } \\
& \operatorname{Pr}_{W^{\prime}}\left[\left.u_{k}=\left(W^{\prime} v_{i^{*}, \ell^{*}}\right)_{k} \leq-\frac{\theta}{2 \sqrt{m}} \right\rvert\, W_{2}\right]
\end{aligned}
$$

(d) With probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{2}$ :

$$
\underset{W^{\prime}}{\operatorname{Pr}}\left[\left.\|u\|_{\infty}=\left\|W^{\prime} v_{i^{*}, e^{*}}\right\|_{\infty} \leq \frac{3 \theta \rho}{2 \sqrt{m}} \right\rvert\, W_{2}\right] \geq 1-e^{-\Omega\left(\rho^{2}\right)} .
$$

Proof. Since $v_{i^{*}, \ell^{*}}$ only depends on the randomness of $W U_{\ell^{*}-1}$, we have $W v_{i^{*}, \ell^{*}} \sim \mathcal{N}\left(0, \frac{2 I}{m}\right)$. As a result, letting $g_{1}, g_{2} \in \mathbb{R}^{m}$ be two independent random Gaussian vectors sampled from $\mathcal{N}\left(0, \frac{2 I}{m}\right)$, we can couple them to make sure

$$
W v_{i^{*}, \ell^{*}}=\sqrt{1-\theta^{2}} g_{1}+\theta g_{2} .
$$

We define $W_{2}, W^{\prime} \in \mathbb{R}^{n \times n}$ as follows

$$
\begin{aligned}
& W_{2} \stackrel{\text { def }}{=} W U_{\ell^{*}-1} U_{\ell^{*}-1}^{\top}+g_{1} v_{i^{*}, \ell^{*}}^{\top}+W\left(I-U_{\ell^{*}-1} U_{\ell^{*}-1}^{\top}\right)\left(I-v_{i^{*}, \ell^{*}} v_{i^{*}, \ell^{*}}^{\top}\right) \\
& W^{\prime} \stackrel{\text { def }}{=} u v_{i^{*}, \ell^{*}}^{\top} \quad \text { where } \quad u \xlongequal{\text { def }} \theta g_{2}-\left(1-\sqrt{1-\theta^{2}}\right) g_{1}
\end{aligned}
$$

- It is easy to verify that

$$
\begin{aligned}
W_{2}+W^{\prime} & =W\left(I-v_{i^{*}, \ell^{*}} v_{i^{*}, \ell^{*}}^{\top}\right)+g_{1} v_{i^{*}, \ell^{*}}^{\top}+\left(\theta g_{2}-\left(1-\sqrt{1-\theta^{2}}\right) g_{1}\right) v_{i^{*}, \ell^{*}}^{\top} \\
& =W\left(I-v_{i^{*}, \ell^{*}} v_{i^{*}, \ell^{*}}^{\top}\right)+\left(\theta g_{2}+\sqrt{1-\theta^{2}} g_{1}\right) v_{i^{*}, \ell^{*}}^{\top} \\
& =W\left(I-v_{i^{*}, \ell^{*}} v_{i^{*}, \ell^{*}}^{\top}\right)+W v_{i^{*}, \ell^{*}} v_{i^{*}, \ell^{*}}^{\top} \\
& =W .
\end{aligned}
$$

- We verify that the entries of $W_{2}$ are i.i.d. Gaussian in two steps.

On one hand, let $U=\left[\widehat{h}_{1}, \ldots, \widehat{h}_{m}\right] \in \mathbb{R}^{m \times m}$ be an arbitrary orthonormal matrix where its first $n L$ columns are identical to $U_{L} \in \mathbb{R}^{m \times n L}$. For each $k \in[n L]$, although $\widehat{h}_{k}$ may depend on the randomness of $W=W_{2}+W^{\prime}$, it can only depend on $W\left[\widehat{h}_{1}, \ldots, \widehat{h}_{k-1}\right]$.

Therefore, conditioning on the randomness of $W_{2}\left[\widehat{h}_{1}, \ldots, \widehat{h}_{k-1}\right]$, we have that $W_{2} \widehat{h_{k}}$ is still $\mathcal{N}\left(0, \frac{2 I}{m}\right)$. This shows that all entries of $W_{2} U$ still follow from $N\left(0, \frac{2}{m}\right)$.

On the other hand, let $V=W_{2} U=\left[v_{1}, \ldots, v_{m}\right] \in \mathbb{R}^{m \times m}$, then we have $v_{k}=W_{2} \widehat{h}_{k}$ is independent of $\widehat{h}_{k}$ as argued above. Since each $v_{k} \sim N\left(0, \frac{2 I}{m}\right)$, we have that all entries of $V U^{\top}=\sum_{k} v_{k} \widehat{k}_{k}^{\top}$ are i.i.d. from $N\left(0, \frac{2}{m}\right)$. Since $W_{2}=V U^{\top}$ this concludes the proof.

- By standard Gaussian tail bound, we have with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$, it satisfies $\left\|g_{1}\right\|_{\infty} \leq \frac{\rho}{\sqrt{m}}$ and therefore $\left\|\left(1-\sqrt{1-\theta^{2}}\right) g_{1}\right\|_{\infty} \leq \frac{\theta^{2} \rho}{\sqrt{m}}$. If this happens, for every $k \in[m]$, we have (see for instance Fact 6.10.12) that $\operatorname{Pr}_{g_{2}}\left[\left(\theta g_{2}\right)_{k}>\frac{\theta}{\sqrt{m}}\right] \geq \frac{1}{4}$ and $\operatorname{Pr}_{g_{2}}\left[\left(\theta g_{2}\right)_{k}<-\frac{\theta}{\sqrt{m}}\right] \geq \frac{1}{4}$. Recalling

$$
u=\left(\theta g_{2}-\left(1-\sqrt{1-\theta^{2}}\right) g_{1}\right.
$$

and using the fact that $\theta \leq \frac{1}{2 \rho}$, we conclude that

$$
\underset{W^{\prime}}{\operatorname{Pr}}\left[\left.u_{k}>\frac{\theta}{2 \sqrt{m}} \right\rvert\, W_{2}, A\right] \geq \frac{1}{4} \quad \text { and } \quad \operatorname{Pr}_{W^{\prime}}\left[\left.u_{k}<-\frac{\theta}{2 \sqrt{m}} \right\rvert\, W_{2}, A\right] \geq \frac{1}{4} .
$$

- Again by Gaussian tail bound, we have with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$, it satisfies $\left\|g_{1}\right\|_{\infty},\left\|g_{2}\right\|_{\infty} \leq \frac{\rho}{\sqrt{m}}$. Therefore, $\left\|\left(1-\sqrt{1-\theta^{2}}\right) g_{1}\right\|_{\infty} \leq \frac{\theta^{2} \rho}{\sqrt{m}}$ and using our assumption on $\theta$, we have

$$
\|u\|_{\infty}=\|\left(\theta g_{2}-\left(1-\sqrt{1-\theta^{2}}\right) g_{1} \|_{\infty} \leq \frac{3 \theta \rho}{2 \sqrt{m}} .\right.
$$

Lemma 6.14.4. Given any fixed $i^{*} \in[n], \ell^{*} \in[L]$, and parameter $\theta \in\left(0, \frac{1}{2 \rho}\right]$, and suppose suppose $W$ and $A$ are at random initialization and $W=W_{2}+W^{\prime}$ where $W^{\prime}=u v_{i^{*}, \ell^{*}}^{\top}$ (following Lemma 6.14.3).

Now, we introduce another two ways of decomposing randomness based on this definition.
(a) Given fixed $k \in[m]$, we can write $W=W_{0}+W_{k}^{\prime}$ where $W_{k}^{\prime}=u_{k} v_{i^{*}, \ell^{*}}^{\top}$ for $u_{k}=$ $\left(0, \ldots, 0, u_{k}, 0, \ldots, 0\right)$ that is non-zero only at the $k$-th coordinate. We again have the entries of $W_{0}$ are i.i.d. from $N\left(0, \frac{2}{m}\right)$.
(b) Given fixed $N \subseteq[m]$, we can write $W=W_{1}+W_{N}^{\prime}$ where $W_{N}^{\prime}=u_{N} v_{i^{*}, \ell^{*}}^{\top}$ for $u_{N}$ being the projection of $u$ onto coordinates in $N\left(s o\|u\|_{0}=|N|\right)$. We again have the entries of $W_{0}$ are i.i.d. from $N\left(0, \frac{2}{m}\right)$.
(c) Given fixed $N \subseteq[m]$, we can write $W=W_{2}+W_{N}^{\prime}+W_{-N}^{\prime}$ where $W_{N}^{\prime}=u_{N} v_{i^{*}, \ell^{*}}^{\top}$ and $W_{-N}^{\prime}=u_{-N} v_{i^{*}, \ell^{*}}^{\top}$. Here, $u_{-N}$ is the projection of $u$ onto coordinates in $[m] \backslash N$.

Proof. The proofs are analogous to Lemma 6.14.3.

### 6.14.2 Gradient Lower Bound in Expectation

The goal of this subsection is to prove Lemma 6.14.5 and then translate it into our Core Lemma A (see Lemma 6.14.6).

Lemma 6.14.5 (c.f. (6.18)). Let $\rho=n L d \log m$ and $\varrho=n L d \delta^{-1} \log (m / \epsilon)$. Fix integer $N \in\left[\rho^{6} / \beta_{-}, \varrho^{100}\right]$, $\theta$ from Definition 6.14.2, fix vectors $\left\{\operatorname{loss}_{i, \ell} \in \mathbb{R}^{d}\right\}_{i \in[n], \ell \in[L] \backslash\{1\}}$, let $\left(i^{*}, \ell^{*}\right)=$ $\operatorname{argmin}_{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|\right\}$. Suppose $W_{1}, A, B$ are at random initialization, and suppose $W_{N}^{\prime}$ is defined according to Lemma 6.14.4b so that $W=W_{1}+W_{N}^{\prime}$ is also at random initialization. With probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{1}, A, B$, we have

$$
\left|\left\{k \in N \left\lvert\, \underset{W_{N}^{\prime}}{\mathbb{E}}\left[\left\|\widehat{\nabla}_{k} f\left(W_{1}+W_{N}^{\prime}\right)\right\|_{2}\right] \geq \Omega\left(\frac{\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}}{n L \sqrt{d}}\right)\right.\right\}\right| \geq \frac{\beta_{-}|N|}{100 L}
$$

### 6.14.2.1 Proof of Lemma 6.14.5

Throughout this proof we denote by $\widetilde{W}=W_{1}$. We introduce some (old and new) notations:

$$
\begin{array}{ll}
\widetilde{\operatorname{Back}}_{i, \ell \rightarrow \ell}=B, & \text { Back }_{i, \ell \rightarrow \ell}=B \\
\widetilde{\operatorname{Back}}_{i, \ell \rightarrow a}=B \widetilde{D}_{i, a} \widetilde{W} \cdots \widetilde{D}_{i, \ell+1} \widetilde{W}, & \operatorname{Back}_{i, \ell \rightarrow a}=B D_{i, a} W \cdots D_{i, \ell+1} W,
\end{array}
$$

and we let $\operatorname{Back}_{i, \ell \rightarrow a}^{\prime}=\operatorname{Back}_{i, \ell \rightarrow a}-{\widetilde{\operatorname{Back}_{i, \ell \rightarrow a}}}^{\text {. For each } i \in[n], \ell \in[L] \text {, we define }}$

$$
\begin{array}{lll}
\widetilde{g}_{i, \ell}=A x_{i, \ell}+\widetilde{W} \widetilde{D}_{i, \ell-1} \widetilde{g}_{i, \ell-1} & g_{i, \ell}=A x_{i, \ell}+W D_{i, \ell-1} g_{i, \ell-1} & g_{i, \ell}^{\prime}=g_{i, \ell}-\widetilde{g}_{i, \ell} \\
\widetilde{h}_{i, \ell}=\widetilde{D}_{i, \ell}\left(A x_{i, \ell}+\widetilde{W} \widetilde{h}_{i, \ell-1}\right) & h_{i, \ell}=D_{i, \ell}\left(A x_{i, \ell}+W h_{i, \ell-1}\right) & h_{i, \ell}^{\prime}=h_{i, \ell}-\widetilde{h}_{i, \ell} \\
\widetilde{u}_{i, \ell}=\sum_{a=\ell}^{L}{\widetilde{\operatorname{Back}_{\ell \rightarrow a}}{ }^{\top} \operatorname{loss}_{i, a}}^{u_{i, \ell}}=\sum_{a=\ell}^{L} \operatorname{Back}_{\ell \rightarrow a}^{\top} \operatorname{loss}_{i, a} & u_{i, \ell}^{\prime}=u_{i, \ell}-\widetilde{u}_{i, \ell}
\end{array}
$$

We also let

$$
v_{i^{*}, \ell^{*}}=\frac{\left(I-U_{\ell^{*}-1} U_{\ell^{*}-1}^{\top}\right) \widetilde{h}_{i^{*}, \ell^{*}}}{\|\left(I-U_{\ell^{*}-1} U_{\ell^{*}-1}^{\top} \widetilde{h}_{i^{*}, \ell^{*}} \|_{2}\right.}
$$

and recall from Lemma 6.14.3d, we have

$$
W_{N}^{\prime}=y v_{i^{*}, \ell^{*}}^{\top} \text { for some } y \in \mathbb{R}^{m} \text { with }\|y\|_{0}=|N| \text { and }\|y\|_{\infty} \leq \frac{\tau_{0}}{\sqrt{m}} \stackrel{\text { def }}{=} \frac{3 \theta \rho}{2 \sqrt{m}} \in\left[\frac{\varrho^{-20}}{\sqrt{m}}, \frac{\varrho^{20}}{\sqrt{m}}\right]
$$

Proof of Lemma 6.14.5. We let $N_{1}=N$, apply Lemma 6.13.1 to obtain $N_{4} \subseteq N_{1}$ with $\left|N_{4}\right| \geq \frac{\beta_{-}\left|N_{1}\right|}{64 L}$, and apply Lemma 6.13 .5 to obtain $N_{5} \subseteq N_{4}$. According to the statements of these lemmas, we know that with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$, the random choice of $W_{1}, A, B$ will satisfy $\left|N_{5}\right| \geq \frac{\beta_{-}\left|N_{1}\right|}{100 L}$.

Fixing such $k \in N_{5}$, we use $u_{i, \ell}$ to denote $\left(u_{i, \ell}\right)_{k}$ and $g_{i, \ell}$ to denote $\left(g_{i, \ell}\right)_{k}$ for notational simplicity. We can rewrite $\widehat{\nabla}_{k} f(W)$ as follows and apply triangle inequality:

$$
\begin{align*}
\left\|\widehat{\nabla}_{k} f(W)\right\|= & \left\|u_{i^{*}, \ell^{*}} \cdot h_{i^{*}, \ell^{*}} \cdot \mathbf{1}_{g_{i^{*}, \ell^{*}+1} \geq 0}+\sum_{(i, \ell) \neq\left(i^{*}, \ell^{*}\right)} u_{i, \ell} \cdot h_{i, \ell} \cdot \mathbf{1}_{g_{i, \ell+1} \geq 0}\right\| \\
\geq & \underbrace{\left\|\widetilde{u}_{i^{*}, \ell^{*}} \cdot \widetilde{h}_{i^{*}, \ell^{*}} \cdot \mathbf{1}_{g_{i^{*}, \ell^{*}+1} \geq 0}+\sum_{(i, \ell) \neq\left(i^{*}, \ell^{*}\right)} \widetilde{u}_{i, \ell} \cdot \widetilde{h}_{i, \ell} \cdot \mathbf{1}_{g_{i, \ell+1} \geq 0}\right\|_{2}}_{*} \\
& -\underbrace{\sum_{i, \ell}\left\|u_{i, \ell}^{\prime} \widetilde{h}_{i, \ell}+\widetilde{u}_{i, \ell} \cdot h_{i, \ell}^{\prime}+u_{i, \ell}^{\prime} \cdot h_{i, \ell}^{\prime}\right\|_{2}}_{*} \tag{6.48}
\end{align*}
$$

where recall $g_{i, \ell+1}=\left\langle W_{k}, h_{i, \ell}\right\rangle+\left\langle A_{k}, x_{i, \ell+1}\right\rangle$.

Step 1 To bound the term on the RHS of (6.48), we consider the following bounds:

1. Using Lemma 6.12.10, we have with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$, for all $i, \ell, a$

$$
\left|\left(\widetilde{\operatorname{Back}}_{i, \ell \rightarrow a} \cdot \operatorname{loss}_{i, a}\right)_{k}-\left(\operatorname{Back}_{i, \ell \rightarrow a} \cdot \operatorname{loss}_{i, a}\right)_{k}\right| \leq O\left(\frac{\rho^{3} \tau_{0}^{1 / 3} N^{1 / 6}}{m^{1 / 6}}\right) \cdot\left\|\operatorname{loss}_{i, a}\right\|
$$

This implies, by triangle inequality and the fact that $\left\|\operatorname{loss}_{i, a}\right\| \leq\left\|\operatorname{loss}_{i^{*}, a^{*}}\right\|$,

$$
\left|u_{i, \ell}^{\prime}\right|=\left|u_{i, \ell}-\widetilde{u}_{i, \ell}\right| \leq O\left(\frac{\rho^{4} \tau_{0}^{1 / 3} N^{1 / 6}}{m^{1 / 6}}\right) \cdot\left\|\operatorname{loss}_{i^{*}, a^{*}}\right\|
$$

2. Using Corollary 6.11 .16 , we have with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$, for all $i, \ell, a$

$$
\left|\left(\widetilde{\operatorname{Back}}_{i, \ell \rightarrow a} \cdot \operatorname{loss}_{i, a}\right)_{k}\right| \leq \rho\left\|\operatorname{loss}_{i, a}\right\|
$$

This implies, by triangle inequality and the fact that $\left\|\operatorname{loss}_{i, a}\right\| \leq\left\|\operatorname{loss}_{i^{*}, a^{*}}\right\|$,

$$
\left|\widetilde{u}_{i, \ell}\right| \leq \rho^{2}\left\|\operatorname{loss}_{i^{*}, a^{*}}\right\| .
$$

3. Using Lemma 6.11.1 and Lemma 6.12.1a we have

$$
\left\|\widetilde{h}_{i, \ell}\right\|_{2} \leq O(L) \quad \text { and } \quad\left\|h_{i, \ell}^{\prime}\right\|_{2} \leq O\left(L^{6} \tau_{0} \sqrt{N}\right) / m^{1 / 2}
$$

All together, they imply

$$
\begin{equation*}
\sum_{i, \ell}\left\|u_{i, \ell}^{\prime} \cdot \widetilde{h}_{i, \ell}+\widetilde{u}_{i, \ell} \cdot h_{i, \ell}^{\prime}+u_{i, \ell}^{\prime} \cdot h_{i, \ell}^{\prime}\right\|_{2} \leq O\left(\frac{\rho^{6} \tau_{0}^{1 / 3} N^{1 / 6}}{m^{1 / 6}}\right) \cdot\left\|\operatorname{loss}_{i^{*}, a^{*}}\right\| \tag{6.49}
\end{equation*}
$$

Step 2 We next turn to the $\boldsymbol{\&}$ term on the RHS of (6.48). We divide into three cases to analyze the sign change of $g_{i, \ell+1}$ for all possible $i$ and $\ell$ in the above formula (under the randomness of $W_{N}^{\prime}$ ). Before we do so, first note that Lemma 6.12.9 implies

$$
\begin{equation*}
\left|\left(W h_{i, \ell}^{\prime}\right)_{k}\right|=\left|\left(W_{1}+W_{N}^{\prime}\right) h_{i, \ell}^{\prime}\right| \leq O\left(\frac{\rho^{8} N^{2 / 3} \tau_{0}^{5 / 6}}{m^{2 / 3}}\right) \tag{6.50}
\end{equation*}
$$

1. Consider $g_{i, \ell+1}$ for the case of $i=i^{*}, \ell=\ell^{*}$. We want to show that the sign of $g_{i, \ell+1}$ changes (at least with constant probability) with respect to the randomness of $W_{N}^{\prime}$.

One can easily check that: ${ }^{20}$

$$
\begin{equation*}
g_{i, \ell+1}=\left\langle A_{k}, x_{i, \ell+1}\right\rangle+\left(W_{1} \widetilde{h}_{i, \ell}\right)_{k}+\left(W h_{i, \ell}^{\prime}\right)_{k}+\left(W_{N}^{\prime} v_{i, \ell}\right)_{k} \cdot\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}\right\|_{2} . \tag{6.51}
\end{equation*}
$$

By Lemma 6.13.1, we have $\left|\left\langle A_{k}, x_{i, \ell}\right\rangle+\left(W_{1} \widetilde{h}_{i, \ell}\right)_{k}\right| \leq \frac{\beta_{-}}{\sqrt{m}}$ Thus, putting this and (6.50) into (6.51), we can write

$$
g_{i, \ell+1}=\Xi+\left(W_{N}^{\prime} v_{i, \ell}\right)_{k} \cdot\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) \widetilde{h}_{i, \ell}\right\|_{2}
$$

for some value $\Xi$ with $|\Xi| \leq \frac{2 \beta_{-}}{\sqrt{m}}$. By Lemma 6.11 .4 , we have $\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) \widetilde{h}_{i, \ell}\right\|_{2} \geq$ $\Omega\left(\frac{1}{L^{2} \log ^{3} m}\right)$. By Lemma 6.14.3c, we have $\left(W_{N}^{\prime} v_{i, \ell}\right)_{k}$ shall be larger than $\frac{\theta}{2 \sqrt{m}}$ and smaller than $-\frac{\theta}{2 \sqrt{m}}$ each with probability at least $1 / 4$. Since by our choice of $\theta$ (see Definition 6.14.2),

$$
\frac{\theta}{2 \sqrt{m}} \cdot \Omega\left(\frac{1}{L^{2} \log ^{3} m}\right) \gg \frac{2 \beta_{-}}{\sqrt{m}} \geq|\Xi|
$$

This concludes that, with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$ over $W_{1}$ and $A$, it satisfies

$$
\underset{W_{N}^{\prime}}{\operatorname{Pr}}\left[g_{i, \ell+1}>0 \mid W_{1}, A\right] \geq \frac{1}{5} \quad \text { and } \quad \operatorname{Pr}_{W_{N}^{\prime}}\left[g_{i, \ell+1}<0 \mid W_{1}, A\right] \geq \frac{1}{5}
$$

${ }^{20}$ We can write $\left(g_{i, \ell+1}\right)_{k}=\left\langle A_{k}, x_{i, \ell+1}\right\rangle+\left(W\left(\widetilde{h}_{i, \ell}+h_{i, \ell}^{\prime}\right)\right)_{k}=\left\langle A_{k}, x_{i, \ell+1}\right\rangle+\left(W \widetilde{h}_{i, \ell}\right)_{k}+\left(W h_{i, \ell}^{\prime}\right)_{k}$. For the second term in the above equation, we have

$$
\begin{aligned}
\left(W \widetilde{h}_{i, \ell}\right)_{k} & =\left(W U_{\ell-1} U_{\ell-1}^{\top} \widetilde{h}_{i, \ell}+\frac{W\left(I-U_{\ell-1} U_{\ell-1}^{\top} \widetilde{h}_{i, \ell}\right.}{\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) \widetilde{h}_{i, \ell}\right\|_{2}}\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) \widetilde{h}_{i, \ell}\right\|_{2}\right)_{k} \\
& \stackrel{\ominus}{=}\left(W U_{\ell-1} U_{\ell-1}^{\top} \widetilde{h}_{i, \ell}+W v_{i, \ell}\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top} \widetilde{h}_{i, \ell}\right)\right\|_{2}\right)_{k} \\
& \stackrel{(2}{=}\left(W_{1} U_{\ell-1} U_{\ell-1}^{\top} \widetilde{h}_{i, \ell}+W v_{i, \ell}\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top} \widetilde{h}_{i, \ell}\right)\right\|_{2}\right)_{k} \\
& \stackrel{\otimes}{=}\left(W_{1} U_{\ell-1} U_{\ell-1}^{\top} \widetilde{h}_{i, \ell}+W_{1} v_{i, \ell}\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top} \widetilde{h}_{i, \ell}\right)\right\|_{2}\right)_{k}+\left(W_{N}^{\prime} v_{i, \ell} \cdot\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top} \widetilde{h}_{i, \ell}\right)\right\|_{2}\right)_{k} \\
& =\left(W_{1} \widetilde{h}_{i, \ell}\right)_{k}+\left(W_{N}^{\prime} v_{i, \ell}\right)_{k} \cdot\left\|\left(I-U_{\ell-1} U_{\ell-1}^{\top} \widetilde{h}_{i, \ell}\right)\right\|_{2}
\end{aligned}
$$

where (1) follows by definition of $v_{i, \ell}$, (2) follows by definition of $W_{1}$, (3) follows by $W=W_{1}+W_{N}^{\prime}$.
2. Consider $g_{i, \ell+1}$ for the case of $i \in[n] \backslash\left\{i^{*}\right\}, \ell=\ell^{*}$ or $i \in[n], \ell>\ell^{*}$. We want to show the sign of $g_{i, \ell+1}$ is fixed, meaning with high probability independent of the choice of $W_{N}^{\prime}$. One can easily check that:

$$
\begin{equation*}
g_{i, \ell+1}=\left\langle A_{k}, x_{i, \ell+1}\right\rangle+\left(W_{1} \widetilde{h}_{i, \ell}\right)_{k}+\left(W h_{i, \ell}^{\prime}\right)_{k}+\left(W_{N}^{\prime} \widetilde{h}_{i, \ell}\right)_{k} \tag{6.52}
\end{equation*}
$$

By Lemma 6.13.1, we have $\left|\left\langle A_{k}, x_{i, \ell}\right\rangle+\left(W_{1} \widetilde{h}_{i, \ell}\right)_{k}\right| \geq \frac{\beta_{+}}{\sqrt{m}}$. Thus, putting this and (6.50) into (6.53), we can write

$$
g_{i, \ell+1}=\Xi+\left(W_{N}^{\prime} \widetilde{h}_{i, \ell}\right)_{k}=\Xi+\left(W_{N}^{\prime} v_{i^{*}, \ell^{*}}\right)_{k} \cdot\left\langle v_{i^{*}, \ell^{*}}, \widetilde{h}_{i, \ell}\right\rangle
$$

for some value $\Xi$ with $|\Xi| \geq \frac{\beta_{+}}{2 \sqrt{m}}$. By Lemma 6.11.1, we have $\left|\left\langle v_{i^{*}, \ell^{*}}, \widetilde{h}_{i, \ell}\right\rangle\right| \leq\left\|\widetilde{h}_{i, \ell}\right\|_{2} \leq$ $O(L)$. By Lemma 6.14.3d, we have $\left|\left(W_{N}^{\prime} v_{i^{*}, \ell^{*}}\right)_{k}\right| \leq \frac{2 \theta \rho}{\sqrt{m}}$ and by our choice of $\theta$ (see Definition 6.14.2),

$$
\frac{2 \theta \rho}{\sqrt{m}} \cdot O(L) \ll \frac{2 \beta_{+}}{\sqrt{m}} \leq|\Xi|
$$

This concludes that, with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$ over $W_{1}$ and $A$, it satisfies

$$
\exists s \in\{-1,1\}: \quad \operatorname{Pr}_{W_{N}^{\prime}}\left[\operatorname{sgn}\left(g_{i, \ell+1}\right)=s \mid W_{1}, A\right] \geq 1-e^{-\Omega\left(\rho^{2}\right)}
$$

3. Consider $g_{i, \ell+1}$ for the case of $i \in[n]$ and $\ell<\ell^{*}$. We want to show such $g_{i, \ell+1}$ is fixed, meaning independent of $W_{N}^{\prime}$. One can easily check that:

$$
\begin{equation*}
g_{i, \ell+1}=\left\langle A_{k}, x_{i, \ell+1}\right\rangle+\left(W_{1} \widetilde{h}_{i, \ell}\right)_{k}+\left(W h_{i, \ell}^{\prime}\right)_{k}+\left(W_{N}^{\prime} \widetilde{h}_{i, \ell}\right)_{k} . \tag{6.53}
\end{equation*}
$$

Now, since $\ell<\ell^{*}$, we know $h_{i, \ell}=\widetilde{h}_{i, \ell}$ because both of which only depend on the randomness of $W U_{\ell^{*}-1}=W_{1} U_{\ell^{*}-1}$ and $A$; this implies $h_{i, \ell}^{\prime}=0$. Also, since $W_{N}^{\prime}=$ $\left(W_{N}^{\prime} v_{i^{*}, \ell^{*}}\right) v_{i^{*}, \ell^{*}}^{\top}$ but $\left\langle v_{i^{*}, \ell^{*}}, \widetilde{h}_{i, \ell}\right\rangle=0$, we know that

$$
g_{i, \ell+1}=\left\langle A_{k}, x_{i, \ell+1}\right\rangle+\left(W_{1} \widetilde{h}_{i, \ell}\right)_{k} .
$$

This is a fixed value, independent of $W_{N}^{\prime}$.

Therefore, combining the three cases above, we conclude that with at least constant probability, the sign of $g_{i^{*}, \ell^{*}+1}$ can possibly change, but the sign of $g_{i, \ell+1}$ will not change for any other $(i, \ell) \neq\left(i^{*}, \ell^{*}\right)$. This means,

$$
\begin{aligned}
\mathbb{E}[\boldsymbol{\varphi}] & =\mathbb{E}\left[\left\|\widetilde{u}_{i^{*}, \ell^{*}} \cdot \widetilde{h}_{i^{*}, \ell^{*}} \cdot \mathbf{1}_{g_{i^{*}, \ell^{*}+1} \geq 0}+\sum_{(i, \ell) \neq\left(i^{*}, \ell^{*}\right)} \widetilde{u}_{i, \ell} \cdot \widetilde{h}_{i, \ell} \cdot \mathbf{1}_{g_{i, \ell+1} \geq 0}\right\|_{2}\right] \\
& \geq \Omega\left(\left\|\widetilde{u}_{i^{*}, \ell^{*}} \widetilde{h}_{i^{*}, \ell^{*}}\right\|_{2}\right)=\Omega\left(\left|\widetilde{u}_{i^{*}, \ell^{*}}\right| \cdot\left\|\widetilde{h}_{i^{*}, \ell^{*}}\right\|_{2}\right)
\end{aligned}
$$

Finally, for each $k \in N_{5}$, owing to Lemma 6.13.5 and Lemma 6.11.1, we have

This means,

$$
\begin{equation*}
\mathbb{E}[\boldsymbol{\varphi}] \geq \Omega\left(\frac{\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}}{\sqrt{d} n L}\right) \tag{6.54}
\end{equation*}
$$

Putting (6.49) and (6.54) back to (6.48), we conclude that

$$
\forall k \in N_{5}: \quad \mathbb{E}\left[\left\|\widehat{\nabla}_{k} f(W)\right\|\right] \geq \Omega\left(\frac{\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}}{\sqrt{d} n L}\right)
$$

### 6.14.2.2 Core Lemma A

In Lemma 6.14.5, we split $W$ into $W=W_{1}+W_{N}^{\prime}$ for a fixed $N \subseteq[m]$. In this subsection, we split $W$ into three parts $W=W_{2}+W_{N}^{\prime}+W_{-N}^{\prime}$ following Lemma 6.14.4c. Our purpose is to rewrite Lemma 6.14.5 into the following variant:

Lemma 6.14.6 (Core Lemma A). Let $\rho=n L d \log m$ and $\varrho=n L d \delta^{-1} \log (m / \epsilon)$. Fix vectors $\left\{\operatorname{loss}_{i, \ell} \in \mathbb{R}^{d}\right\}_{i \in[n], \ell \in[L] \backslash\{1\}}$, integer $|N| \in\left[\rho^{6} / \beta_{-}, \varrho^{100}\right]$, and parameter $t=m /|N|$. Let $N_{1}, \ldots, N_{t}$ be i.i.d. random subsets of $[m]$ with cardinality $\left|N_{i}\right|=|N|$. Then,

- with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{2}, A, B$ and $N_{1}, \cdots N_{t}$, the following holds:
- for every $N \in\left\{N_{1}, N_{2}, \cdots, N_{t}\right\}$, the following holds:
* with probability at $1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{-N}^{\prime}$, the following holds:

$$
\sum_{k \in N} \underset{W_{N}^{\prime}}{\mathbb{E}}\left[\left\|\widehat{\nabla}_{k} f\left(W_{2}+W_{N}^{\prime}+W_{-N}^{\prime}\right)\right\|_{2}^{2}\right] \geq \Omega\left(\frac{\beta_{-}|N|}{\rho^{2}}\right) \cdot \max _{i, \ell}\left\|\operatorname{loss}_{i, \ell}\right\|_{2}^{2}
$$

Proof of Lemma 6.14.6. We denote by $\left(i^{*}, \ell^{*}\right)=\operatorname{argmax}_{i, \ell}\left\|\operatorname{loss}_{i, \ell}\right\|_{2}$. We first note that Lemma 6.14.5 directly implies

- with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{2}, W_{-N}^{\prime}, A, B, N$, the following holds

$$
\left|\left\{k \in N \left\lvert\, \underset{W_{N}^{\prime}}{\mathbb{E}}\left[\left\|\widehat{\nabla}_{k} f\left(W_{2}+W_{-N}^{\prime}+W_{N}^{\prime}\right)\right\|_{2}\right] \geq \Omega\left(\frac{\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}}{n L \sqrt{d}}\right)\right.\right\}\right| \geq \frac{\beta_{-}}{100 L}|N|
$$

or putting it in another way, using our choice of $q$,

- with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{2}, W_{-N}^{\prime}, A, B, N$, the following holds ( $q=\frac{\beta-|N|}{c \rho^{2}}$ for some sufficiently large constant):

$$
\sum_{k \in N} \underset{W_{N}^{\prime}}{\mathbb{E}}\left[\left\|\widehat{\nabla}_{k} f\left(W_{2}+W_{-N}^{\prime}+W_{N}\right)\right\|_{2}^{2}\right] \geq q \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}^{2}
$$

Applying simple tricks to switch the ordering of randomness (see Fact 6.10.2), we have

- with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{2}, A, B$, the following holds:
- with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $N$, the following holds:
* with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over randomness of $W_{-N}^{\prime}$, the following holds:

$$
\sum_{k \in N} \underset{W_{N}^{\prime}}{\mathbb{E}}\left[\left\|\widehat{\nabla}_{k} f\left(W_{2}+W_{N}^{\prime}+W_{-N}^{\prime}\right)\right\|_{2}^{2}\right] \geq q \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}^{2}
$$

Repeating the choice of $N$ for $t$ times, and using union bound, we have

- with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{2}, A, B$, the following holds:
- with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $N_{1}, N_{2}, \cdots, N_{t}$, the following holds:
* for all $N \in\left\{N_{1}, N_{2}, \cdots, N_{t}\right\}$, the following holds:
- with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over randomness of $W_{-N}^{\prime}$, the following holds:

$$
\sum_{k \in N} \underset{W_{N}^{\prime}}{\mathbb{E}}\left[\left\|\widehat{\nabla}_{k} f\left(W_{2}+W_{N}^{\prime}+W_{-N}^{\prime}\right)\right\|_{2}^{2}\right] \geq q \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}^{2}
$$

Combining the above statement with Fact 6.10.1 (to merge randomness), we conclude the proof.

### 6.14.3 Gradient Stability

The goal of this subsection is to prove Lemma 6.14.7 and then translate it into our Core Lemma B (see Lemma 6.14.9).

Lemma 6.14.7 (c.f. (6.19)). Fix parameter $\rho=n L d \log m$, vectors $\left\{\operatorname{loss}_{i, \ell} \in \mathbb{R}^{d}\right\}_{i \in[n], \ell \in[L] \backslash\{1\}}$, $\left(i^{*}, \ell^{*}\right)=\operatorname{argmin}_{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|\right\}$, index $j \in[m]$. Let $N \subseteq[m]$ be a random subset containing $j$ of fixed cardinality $|N|$, let $W, A, B$ be at random initialization, and $W_{j}^{\prime \prime}$ be defined as Definition 6.14.4a so that $W+W_{j}^{\prime \prime}$ is also at random initialization. Then, with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$

$$
\begin{equation*}
\sum_{k \in[N]}\left|\left\|\widehat{\nabla}_{k}(W)\right\|_{2}^{2}-\left\|\widehat{\nabla}_{k}\left(W+W_{j}^{\prime \prime}\right)\right\|_{2}^{2}\right| \leq O\left(\rho^{8}+\frac{\rho^{11} \theta^{1 / 3}}{m^{1 / 6}}|N|\right) \tag{6.55}
\end{equation*}
$$

### 6.14.3.1 Proof of Lemma 6.14.7

Throughout this proof we denote by $\widetilde{W}=W$ to emphasize that $W$ is at random initialization. We first show some properties before proving Lemma 6.14.7.

Claim 6.14.8. Fix parameter $\rho=n L d \log m$, vectors $\left\{\operatorname{loss}_{i, \ell} \in \mathbb{R}^{d}\right\}_{i \in[n], \ell \in[L] \backslash\{1\}},\left(i^{*}, \ell^{*}\right)=$ $\operatorname{argmin}_{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|\right\}$, index $j \in[m]$. Let $\widetilde{W}, A, B$ be at random initialization, and $W_{j}^{\prime \prime}$ be defined as Lemma 6.14.4a so that $W=\widetilde{W}+W_{j}^{\prime \prime}$ is also at random initialization. Then, then with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over $\widetilde{W}, A, B, W_{j}^{\prime \prime}$,

- $\left|\left\{k \in[m]\left|\left|\left|\widehat{\nabla}_{k} f\left(\widetilde{W}+W_{j}^{\prime \prime}\right)\left\|_{2}^{2}-\right\| \widehat{\nabla}_{k} f(\widetilde{W}) \|_{2}^{2}\right| \leq O\left(\frac{\rho^{11} \theta^{1 / 3}}{m^{1 / 6}}\right)\right\}\right| \geq\left(1-\frac{\rho^{5} \theta^{2 / 3}}{m^{1 / 3}}\right) m\right.\right.$.
- $\left|\left\|\widehat{\nabla}_{k} f\left(\widetilde{W}+W_{j}^{\prime \prime}\right)\right\|_{2}^{2}-\left\|\widehat{\nabla}_{k} f(\widetilde{W})\right\|_{2}^{2}\right| \leq O\left(\rho^{6}\right)$ for every $k \in[m]$.

Proof of Claim 6.14.8. As before, let $\widetilde{D}_{i, \ell}$ denote the diagonal matrix consisting of the indicator function $\mathbb{1}_{\left(\widetilde{g}_{i, \ell+1}\right)_{k} \geq 0}$ for its $k$-th diagonal entry. Let $D_{i, \ell}$ be the same thing but with respect to weight matrix $\widetilde{W}+W_{j}^{\prime \prime}$. Let $D_{i, \ell}^{\prime \prime}=D_{i, \ell}-\widetilde{D}_{i, \ell}$ be the diagonal matrix of sign change.

- By Lemma 6.12 .1 (with $\tau_{0}=\frac{3 \theta \rho}{2}$ owing to Lemma 6.14.3d), we have $\left\|D_{\ell}^{\prime \prime}\right\|_{0} \leq \rho^{4} \theta^{2 / 3} m^{2 / 3}$. Therefore, letting $J \stackrel{\text { def }}{=}[m] \backslash \cup_{i, \ell} \operatorname{supp}\left(D_{i, \ell}\right)$ be the set of all indices of sign changes, we have

$$
|J| \geq m-n L \rho^{4} \theta^{2 / 3} m^{2 / 3} \geq\left(1-\frac{\rho^{5} \theta^{2 / 3}}{m^{1 / 3}}\right) \cdot m
$$

Next, for every index $k \in J$, we have

$$
\begin{aligned}
&\left\|\widehat{\nabla}_{k} f\left(\widetilde{W}+W_{j}^{\prime \prime}\right)-\widehat{\nabla}_{k} f(\widetilde{W})\right\|_{2} \\
&\left.\stackrel{(1)}{=} \| \sum_{i=1}^{n} \sum_{\ell=1}^{L}\left(\widetilde{u}_{i, \ell}\right)_{k} \cdot \widetilde{h}_{i, \ell} \cdot \mathbf{1}_{\left(\widetilde{g}_{i}, \ell+1\right.}\right)_{k} \geq 0 \\
& \stackrel{n}{=} \sum_{i=1}^{n} \sum_{\ell=1}^{L}\left(u_{i, \ell}\right)_{k} \cdot h_{i, \ell} \cdot \mathbf{1}_{\left(g_{i, \ell+1}\right)_{k} \geq 0}\left\|_{2}^{L}\left(\widetilde{u}_{i, \ell}\right)_{k} \cdot \widetilde{h}_{i, \ell} \cdot \mathbf{1}_{\left(g_{i, \ell+1}\right)_{k} \geq 0}-\sum_{i=1}^{n} \sum_{\ell=1}^{L}\left(u_{i, \ell}\right)_{k} \cdot h_{i, \ell} \cdot \mathbf{1}_{\left(g_{i, \ell+1}\right)_{k} \geq 0}\right\|_{2} \\
& \stackrel{(3)}{\leq} \sum_{i=1}^{n} \sum_{\ell=1}^{L}\left\|\left(\left(\widetilde{u}_{i, \ell}\right)_{k} \cdot \widetilde{h}_{i, \ell}-\left(u_{i, \ell}\right)_{k} \cdot h_{i, \ell}\right) \cdot \mathbf{1}_{\left(g_{i, \ell+1}\right)_{k} \geq 0}\right\|_{2} .
\end{aligned}
$$

where (1) follows by definition, (2) follows by our choice of $k \in J$, and (3) follows by triangle inequality. For each $i, \ell$, using the same proof as (6.49) (so invoking Lemma 6.12.10, Corollary 6.11.16, Lemma 6.11 .1 and Lemma 6.12.1a), we have with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}:$

$$
\begin{equation*}
\left\|\left(\widetilde{u}_{i, \ell}\right)_{k} \cdot \widetilde{h}_{i, \ell}-\left(u_{i, \ell}\right)_{k} \cdot h_{i, \ell}\right\|_{2} \leq O\left(\frac{\rho^{6} \tau_{0}^{1 / 3}}{m^{1 / 6}}\right) \cdot\left\|\operatorname{loss}_{i^{*}, a^{*}}\right\| \tag{6.56}
\end{equation*}
$$

which implies

$$
\left\|\widehat{\nabla}_{k} f\left(\widetilde{W}+W_{j}^{\prime \prime}\right)-\widehat{\nabla}_{k} f(\widetilde{W})\right\|_{2} \leq O\left(\frac{\rho^{6} \tau_{0}^{1 / 3}}{m^{1 / 6}}\right) \cdot\left\|\operatorname{loss}_{i^{*}, a^{*}}\right\|
$$

Combining this with $\left\|\widehat{\nabla}_{k} f(\widetilde{W})\right\|_{2} \leq O\left(\rho^{4}\right)$ from Lemma 6.14.1, we finish the proof of the first item.

- For every index $k \in[m]$, we have

$$
\begin{aligned}
& \left\|\widehat{\nabla}_{k} f\left(\widetilde{W}+W_{j}^{\prime \prime}\right)-\widehat{\nabla}_{k} f(\widetilde{W})\right\|_{2} \\
= & \left\|\sum_{i=1}^{n} \sum_{\ell=1}^{L}\left(\widetilde{u}_{i, \ell}\right)_{k} \cdot \widetilde{h}_{i, \ell} \cdot \mathbf{1}_{\left(\widetilde{g}_{i, \ell+1}\right)_{k} \geq 0}-\sum_{i=1}^{n} \sum_{\ell=1}^{L}\left(u_{i, \ell}\right)_{k} \cdot h_{i, \ell} \cdot \mathbf{1}_{\left(g_{i, \ell+1}\right)_{k} \geq 0}\right\|_{2} \\
\leq & \left.\sum_{i=1}^{n} \sum_{\ell=1}^{L}\left\|\left(\widetilde{u}_{i, \ell}\right)_{k} \cdot \widetilde{h}_{i, \ell}\right\|_{2}+\|\left(u_{i, \ell}\right)_{k} \cdot h_{i, \ell}\right) \|_{2} .
\end{aligned}
$$

Using (6.56) and $\left\|\left(\widetilde{u}_{i, \ell}\right)_{k} \cdot \widetilde{h}_{i, \ell}\right\|_{2} \leq O\left(\rho^{2} L\right)$ (see the proof of Lemma 6.14.1), we immediately have the desired bound.

Proof of Lemma 6.14.7. Letting $J \stackrel{\text { def }}{=}\left\{k \in[m]\left|\left|\left\|\widehat{\nabla}_{k} f\left(\widetilde{W}+W_{j}^{\prime \prime}\right)\right\|_{2}^{2}-\left\|\widehat{\nabla}_{k} f(\widetilde{W})\right\|_{2}^{2}\right| \leq O\left(\frac{\rho^{11} \theta^{1 / 3}}{m^{1 / 6}}\right)\right\}\right.$, we have $|J| \geq\left(1-\frac{\rho^{5} \theta^{2 / 3}}{m^{1 / 3}}\right) m$ according to Claim 6.14.8.

Now, for the indices in $N$ that are randomly sampled from [m], if $|N \cap J| \geq|N|-S$ for a parameter $S=\rho^{2}$, then we have

$$
\sum_{k \in N}\left|\left\|\widehat{\nabla}_{k} f(\widetilde{W})\right\|_{2}^{2}-\left\|\widehat{\nabla}_{k} f\left(\widetilde{W}+W_{j}^{\prime \prime}\right)\right\|_{2}^{2}\right| \leq|S| \cdot O\left(\rho^{6}\right)+|N| \cdot O\left(\frac{\rho^{11} \theta^{1 / 3}}{m^{1 / 6}}\right) \leq O\left(\rho^{8}+\frac{\rho^{11} \theta^{1 / 3}}{m^{1 / 6}}|N|\right)
$$

Otherwise, if $|N \cap J| \leq|N|-S$, this means at least $S$ indices that are chosen from $N$ are outside $J \subseteq[m]$. This happens with probability at most $\left(\frac{\rho^{5} \theta^{2 / 3}}{m^{1 / 3}}\right)^{S} \leq e^{-\Omega\left(\rho^{2}\right)}$.

### 6.14.3.2 Core Lemma B

In this subsection, we split $W$ into three parts $W=W_{2}+W_{N}^{\prime}+W_{-N}^{\prime}$ following def:random-decomp:N-N. Our purpose is to rewrite Lemma 6.14.7 into the following variant:

Lemma 6.14.9 (Core Lemma B). Let $\rho=n L d \log m$ and $\varrho=n L d \delta^{-1} \log (m / \epsilon)$. Fix parameter $\theta$ from Definition 6.14.2, vectors $\left\{\operatorname{loss}_{i, \ell} \in \mathbb{R}^{d}\right\}_{i \in[n], \ell \in[L] \backslash\{1\}}$, integer $|N| \in\left[1, \varrho^{100}\right]$, parameter $t=m /|N|$. Let $N_{1}, \ldots, N_{t}$ be i.i.d. random subsets of $[m]$ with cardinality $\left|N_{i}\right|=|N|$. Then, with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{2}, A, B$ and $N_{1}, \cdots N_{t}$, the following holds:

- for every $N \in\left\{N_{1}, N_{2}, \cdots, N_{t}\right\}$, the following holds:
- with probability at $1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{-N}^{\prime}$, the following holds:
* with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{N}^{\prime}$, we have :
- for all $j \in N$, and for all $W_{j}^{\prime \prime}=u_{j} v_{i^{*}, \ell^{*}}^{\top}$ satisfying

$$
\begin{equation*}
\left\|u_{j}\right\|_{0}=1, \quad\left\|u_{j}\right\|_{\infty} \leq \frac{3 \theta \rho}{\sqrt{m}}, \quad v_{i^{*}, \ell^{*}}=\frac{\left(I-U_{\ell^{*}-1} U_{\ell-1}^{\top}\right) h_{i^{*}, \ell} h_{i^{*}, \ell^{*}}}{\left\|\left(I-U_{\ell^{*}-1} U_{\ell^{*}-1}^{\top}\right) h_{i^{*}, \ell^{*}}\right\|_{2}}, \tag{6.57}
\end{equation*}
$$

we have

$$
\begin{aligned}
& \sum_{k \in N}\left|\left\|\widehat{\nabla}_{k} f\left(W_{2}+W_{N}^{\prime}+W_{-N}^{\prime}\right)\right\|_{2}^{2}-\left\|\widehat{\nabla}_{k} f\left(W_{2}+W_{N}^{\prime}+W_{-N}^{\prime}+W_{j}^{\prime \prime}\right)\right\|_{2}^{2}\right| \\
\leq & O\left(\rho^{8}\right) \cdot \max _{i, \ell}\left\|\operatorname{loss}_{i, \ell}\right\|_{2}^{2} .
\end{aligned}
$$

Proof of Lemma 6.14.9. Without loss of generality we assume $\max _{i, \ell}\left\|\operatorname{loss}_{i, \ell}\right\|_{2}^{2}=1$ in the proof. We now first rewrite Lemma 6.14.7 as follows (recalling $m$ is sufficiently large so we only need to keep the $O\left(\rho^{8}\right)$ term):

- with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over random $N \subset[m]$, random $j \in N$, random $W, A, B$, and random $W_{j}^{\prime}$, the following holds:

$$
\begin{equation*}
\sum_{k \in[N]}\left|\left\|\widehat{\nabla}_{k}(W)\right\|_{2}^{2}-\left\|\widehat{\nabla}_{k}\left(W+W_{j}^{\prime \prime}\right)\right\|_{2}^{2}\right| \leq O\left(\rho^{8}+\frac{\rho^{11} \theta^{1 / 3}}{m^{1 / 6}}|N|\right) \leq O\left(\rho^{8}\right) \tag{6.58}
\end{equation*}
$$

We can split the randomness (see Fact 6.10.1) and derive that

- With probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over $N \subseteq[m], W, A, B$, the following holds:
- With probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over $j \in N$ and $W_{j}^{\prime \prime}$, Eq. (6.58) holds.

Applying standard $\epsilon$-net argument, we derive that

- With probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over $N \subseteq[m], W, A, B$, the following holds:
- for all $j \in N$, and for all $W_{j}^{\prime \prime}=u_{j} v_{i^{*}, \ell^{*}}^{\top}$ satisfying (6.57), we have Eq. (6.58) holds.

Finally, letting $W=W_{2}+W_{N}^{\prime}+W_{-N}^{\prime}$, we have

- With probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $N \subseteq[m], W_{2}, W_{N}, W_{-N}^{\prime}, A, B$, the following holds:
- for all $j \in N$, and for all $W_{j}^{\prime \prime}=u_{j} v_{i^{*}, \ell^{*}}^{\top}$ satisfying (6.57), we have

$$
\sum_{k \in N}\left|\left\|\widehat{\nabla}_{k} f\left(W_{2}+W_{N}^{\prime}+W_{-N}^{\prime}\right)\right\|_{2}^{2}-\left\|\widehat{\nabla}_{k} f\left(W_{2}+W_{N}^{\prime}+W_{-N^{\prime}}+W_{j}^{\prime \prime}\right)\right\|_{2}^{2}\right| \leq O\left(\rho^{8}\right)
$$

Finally, splitting the randomness (using Fact 6.10.1), and applying a union bound over multiple samples $N_{1}, \ldots, N_{t}$, we finish the proof.

### 6.14.4 Main Theorem: Proof of Theorem 6.14.2

Proof of Theorem 6.14.2. Without loss of generality, we assume that $\max _{i, \ell}\left\|\operatorname{loss}_{i, \ell}\right\|_{2}^{2}=1$.
Combining Core Lemma A and B (i.e., Lemma 6.14.6 and Lemma 6.14.9), we know that if $|N|$ is appropriately chosen, with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{2}, A, B$ and $N_{1}, \cdots N_{t}$, the following holds:

- for every $N \in\left\{N_{1}, N_{2}, \cdots, N_{t}\right\}$, the following holds:
- with probability at $1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{-N}^{\prime}$, the following boxed statement holds:
- (core A)

$$
\mathbb{E}_{W_{N}^{\prime}}\left[\sum_{k \in N}\left\|\widehat{\nabla}_{k} f\left(W_{2}+W_{N}^{\prime}+W_{-N}^{\prime}\right)\right\|_{2}^{2}\right] \geq q \stackrel{\text { def }}{=} \Omega\left(\frac{\beta-|N|}{\rho^{2}}\right)
$$

- (core B) with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{N}^{\prime}$, the following holds:
- for all $j \in N$, and for all $W_{j}^{\prime \prime}=u_{j} v_{i^{*}, \ell^{*}}^{\top}$ satisfying

$$
\left\|u_{j}\right\|_{0}=1, \quad\left\|u_{j}\right\|_{\infty} \leq \frac{3 \theta \rho}{\sqrt{m}}, \quad v_{i^{*}, \ell^{*}}=\frac{\left(I-U_{\ell^{*}-1} U_{\ell-1}^{\top}\right) h_{i^{*}, \ell} h_{i^{*}, \ell^{*}}}{\left\|\left(I-U_{\ell^{*}-1} U_{\ell^{*}-1}^{\top}\right) h_{i^{*}, \ell^{*}}\right\|_{2}},
$$

we have

$$
\sum_{k \in N}\left|\left\|\widehat{\nabla}_{k} f\left(W_{2}+W_{N}^{\prime}+W_{-N}^{\prime}\right)\right\|_{2}^{2}-\left\|\widehat{\nabla}_{k} f\left(W_{2}+W_{N}^{\prime}+W_{-N}^{\prime}+W_{j}^{\prime \prime}\right)\right\|_{2}^{2}\right| \leq p \stackrel{\text { def }}{=} O\left(\rho^{8}\right)
$$

We now wish to apply our extended McDiarmid's inequality (see Lemma 6.10.18) to the boxed statement. For this goal, recalling that $W_{N}^{\prime}=u_{N} v_{i^{*}, \ell^{*}}^{\top}$ for a vector $u_{N} \in \mathbb{R}^{m}$ that is only supported on indices in $N$. That is, $\left(u_{N}\right)_{k}=0$ for all $k \notin N$. Therefore, we can define function $\digamma\left(u_{N}\right) \stackrel{\text { def }}{=} \sum_{k \in N}\left\|\widehat{\nabla}_{k} f\left(W_{2}+W_{-N}^{\prime}+W_{N}^{\prime}\right)\right\|^{2}$ where $W_{N}^{\prime}=u_{N} v_{i^{*}, \ell^{*}}^{\top}$. We emphasize
here that, inside the boxed statement, $W_{2}, W_{-N}^{\prime}$ and $v_{i^{*}, \ell^{*}}$ are already fixed, and $u_{N}$ is the only source of randomness.

Below, we condition on the high probability event (see Lemma 6.14.3d) that $\left\|u_{N}\right\|_{\infty} \leq$ $\frac{3 \theta \rho}{2 \sqrt{m}}$. The boxed statement tells us:

- $\mathbb{E}_{u_{N}}\left[\digamma\left(u_{N}\right)\right] \geq q$, and
- With probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over $u_{N}$, it satisfies

$$
\forall j \in[N], \forall u_{j}^{\prime \prime}:\left|\digamma\left(u_{-j}, u_{j}\right)-\digamma\left(u_{-j}, u_{j}^{\prime \prime}\right)\right| \leq p
$$

At the same time, we have $0 \leq \digamma\left(u_{N}\right) \leq O\left(N \rho^{8}\right)$ owing to Lemma 6.14.1. Therefore, scaling down $\digamma\left(u_{N}\right)$ by $\Theta\left(N \rho^{8}\right)$ (to make sure the function value stays in $[0,1]$ ), we can applying extended McDiarmid's inequality (see Lemma 6.10.18), we have

$$
\begin{equation*}
\operatorname{Pr}_{u_{N}}\left[\digamma\left(u_{N}\right) \geq \frac{q}{2}\right] \geq 1-N^{2} \cdot e^{-\Omega\left(\rho^{2}\right)}-\exp \left(-\Omega\left(\frac{\left(q / \rho^{8}\right)^{2}}{N\left(p / \rho^{8}\right)^{2}+N e^{-\Omega\left(\rho^{2}\right)}}\right)\right) . \tag{6.59}
\end{equation*}
$$

As long as $N \geq \frac{\rho^{22}}{\beta_{-}^{2}}$, we have that the above probability is at least $1-e^{-\Omega\left(\rho^{2}\right)}$.
Finally, we choose

$$
\begin{equation*}
N=\frac{\rho^{22}}{\beta_{-}^{2}} \tag{6.60}
\end{equation*}
$$

in order to satisfy Lemma 6.14.6. We replace the boxed statement with (6.59). This tells us

- with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{2}, A, B$ and $N_{1}, \cdots N_{t}$, the following holds:
- for every $N \in\left\{N_{1}, N_{2}, \cdots, N_{t}\right\}$, the following holds:
* with probability at $1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W_{-N}^{\prime}$ and $W_{N}^{\prime}$, the following holds:

$$
\sum_{k \in N}\left\|\widehat{\nabla}_{k} f\left(W_{2}+W_{N}^{\prime}+W_{-N}^{\prime}\right)\right\|_{2}^{2} \geq \Omega\left(\frac{\beta_{-}|N|}{\rho^{2}}\right)
$$

After rearranging randomness, and using $W=W_{2}+W_{N}^{\prime}+W_{-N}^{\prime}$, we have

- with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W, A, B$, the following holds:
- with probability $\geq 1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $N_{1}, \cdots N_{t}$, the following holds:
* for every $N \in\left\{N_{1}, N_{2}, \cdots, N_{t}\right\}$, the following holds:

$$
\sum_{k \in N}\left\|\widehat{\nabla}_{k} f(W)\right\|_{2}^{2} \geq \Omega\left(\frac{\beta_{-}|N|}{\rho^{2}}\right)
$$

Finally, since $N_{1}, \ldots, N_{t}$ are $t$ random subsets of $[m$ ], we know that with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$, for each index $k \in[m]$, it is covered by at most $\rho^{2}$ random subsets. Therefore,

$$
\sum_{k \in N}\left\|\widehat{\nabla}_{k} f(W)\right\|_{2}^{2} \geq \frac{1}{\rho^{2}} \times t \times \Omega\left(\frac{\beta_{-}|N|}{\rho^{2}}\right)=\Omega\left(\frac{\delta}{\rho^{14}}\right) \times m
$$

### 6.15 Gradient Bound After Perturbation (Theorem 6.15.2)

We use the same fake gradient notion (see Definition 6.14.1) and first derive the following result based on Lemma 6.14.1 and Theorem 6.14.2.

Lemma 6.15.1. Let $\rho=n L d \log m$ and $\varrho=n L d \delta^{-1} \log (m / \epsilon)$. Given fixed vectors $\left\{\operatorname{loss}_{i, \ell} \in\right.$ $\left.\mathbb{R}^{d}\right\}_{i \in[n], \ell \in[L] \backslash\{1\}}$, with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$ over $\widetilde{W}, A, B$, it satisfies for all $W^{\prime} \in$ $\mathbb{R}^{m \times m}$ with $\left\|W^{\prime}\right\|_{2} \leq \frac{\tau_{0}}{\sqrt{m}}$ with $\tau_{0} \leq \varrho^{100}$,

$$
\begin{aligned}
& \left\|\widehat{\nabla} f\left(\widetilde{W}+W^{\prime}\right)\right\|_{F}^{2} \geq \Omega\left(\frac{\delta}{\rho^{14}}\right) \times m \times \max _{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|^{2}\right\} \\
& \left\|\widehat{\nabla} f\left(\widetilde{W}+W^{\prime}\right)\right\|_{F}^{2} \leq O\left(\rho^{12} m\right) \times \max _{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|^{2}\right\} \\
& \left\|\widehat{\nabla} f_{i}\left(\widetilde{W}+W^{\prime}\right)\right\|_{F}^{2} \leq \frac{1}{n^{2}} O\left(\rho^{12} m\right) \times \max _{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|^{2}\right\}
\end{aligned}
$$

Proof of Lemma 6.15.1. Like before, we denote by $\widetilde{D}_{i, \ell}, \widetilde{u}_{i, \ell}, \widetilde{h}_{i, \ell}$ the corresponding matrices or vectors when the weight matrix is $\widetilde{W}$, and by $D_{i, \ell}, u_{i, \ell}, h_{i, \ell}$ the corresponding terms when when the weight matrix is $W=\widetilde{W}+W^{\prime}$.

Lower bound on $\left\|\widehat{\nabla} f\left(\widetilde{W}+W^{\prime}\right)\right\|_{F} \quad$ Recall

$$
\widehat{\nabla} f\left(\widetilde{W}+W^{\prime}\right)=\sum_{i=1}^{n} \sum_{\ell=1}^{L} D_{i, \ell+1}\left(\sum_{a=\ell}^{L} \operatorname{Back}_{i, \ell \rightarrow a}^{\top} \operatorname{loss}_{i, a}\right) h_{i, \ell}^{\top}=\sum_{i=1}^{n} \sum_{\ell=1}^{L} D_{i, \ell+1} u_{i, \ell} h_{i, \ell}^{\top}
$$

In Theorem 6.14.2 of the previous section, we already have a lower bound on $\|\widehat{\nabla} f(\widetilde{W})\|_{F}$.

We just need to upper bound $\left\|\widehat{\nabla} f\left(\widetilde{W}+W^{\prime}\right)-\widehat{\nabla} f(\widetilde{W})\right\|_{F}$.

$$
\begin{aligned}
\left\|\widehat{\nabla} f\left(\widetilde{W}+W^{\prime}\right)-\widehat{\nabla} f(\widetilde{W})\right\|_{F} & \stackrel{(1)}{=}\left\|\sum_{i=1}^{n} \sum_{\ell=1}^{L} D_{i, \ell+1} u_{i, \ell} h_{i, \ell}^{\top}-\widetilde{D}_{i, \ell+1} \widetilde{u}_{i, \ell} \widetilde{h}_{i, \ell}^{\top}\right\|_{F} \\
& \stackrel{(2)}{\leq} \sum_{i=1}^{n} \sum_{\ell=1}^{L}\left\|D_{i, \ell+1} u_{i, \ell} h_{i, \ell}^{\top}-\widetilde{D}_{i, \ell+1} \widetilde{u}_{i, \ell} \widetilde{h}_{i, \ell}^{\top}\right\|_{F} \\
& \stackrel{(3)}{\leq} \sum_{i=1}^{n} \sum_{\ell=1}^{L}\left\|D_{i, \ell+1}^{\prime} \widehat{u}_{i, \ell} \hat{h}_{i, \ell}^{\top}\right\|_{F}+\left\|\widetilde{D}_{i, \ell+1} u_{i, \ell}^{\prime} \widehat{h}_{i, \ell}^{\top}\right\|_{F}+\left\|\widetilde{D}_{i, \ell+1} \widetilde{u}_{i, \ell} h_{i, \ell}^{\prime}\right\|_{F}+o\left(m^{1 / 3}\right) .
\end{aligned}
$$

where (1) follows by definition and (2) and (3) follow by the triangle inequality. Note that in inequality (3) we have hidden four more higher order terms in $o\left(m^{1 / 3}\right)$. We ignore the details for how to bound them, for the ease of presentation. We bound the three terms separately:

- Using Corollary 6.11 .15 (with $s^{2}=O\left(L^{10 / 3} \tau_{0}^{2 / 3}\right)$ from Lemma 6.12.1b) and Lemma 6.11.1 we have

$$
\begin{aligned}
\left\|D_{i, \ell+1}^{\prime} \widehat{u}_{i, \ell} \widehat{h}_{i, \ell}^{\top}\right\|_{F} & \leq\left\|D_{i, \ell+1}^{\prime} \widehat{u}_{i, \ell}\right\|_{2} \cdot\left\|\widehat{h}_{i, \ell}\right\|_{2} \\
& \leq O\left(L \cdot L^{4} \tau_{0}^{1 / 3} m^{1 / 3} \log m \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}\right) \cdot L
\end{aligned}
$$

- Using $\left\|\widetilde{D}_{i, \ell+1}\right\|_{2} \leq 1$, Lemma 6.12.8 and Lemma 6.11 .1 we have

$$
\begin{aligned}
\left\|\widetilde{D}_{i, \ell+1} u_{i, \ell}^{\prime} \widehat{h}_{i, \ell}^{\top}\right\|_{F} & \leq\left\|\widetilde{D}_{i, \ell+1}\right\| \cdot\left\|u_{i, \ell}^{\prime}\right\|_{2} \cdot\left\|\widehat{h}_{i, \ell}\right\|_{2} \\
& \leq O\left(L \cdot L^{6} \tau_{0}^{1 / 3} m^{1 / 3} \log m \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}\right) \cdot L
\end{aligned}
$$

- Using $\left\|\widetilde{D}_{i, \ell+1}\right\|_{2} \leq 1,\left\|\widetilde{u}_{i, \ell}\right\|_{2} \leq O\left(L^{4}\right) \sqrt{m}\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}$ (implied by Lemma 6.11.8) and Lemma 6.12.1a, we have

$$
\begin{aligned}
\left\|\widetilde{D}_{i, \ell+1} \widetilde{u}_{i, \ell} h_{i, \ell}^{\prime}{ }^{\top}\right\|_{F} & \leq\left\|\widetilde{D}_{i, \ell+1}\right\| \cdot\left\|\widetilde{u}_{i, \ell}\right\|_{2} \cdot\left\|h_{i, \ell}^{\prime}\right\|_{2} \\
& \leq O\left(L^{4} \sqrt{m} \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2}\right) \cdot\left(L^{6} \tau_{0}^{1 / 2} \frac{1}{\sqrt{m}}\right)
\end{aligned}
$$

Putting it all together, we have

$$
\begin{equation*}
\left\|\widehat{\nabla} f\left(\widetilde{W}+W^{\prime}\right)-\widehat{\nabla} f(\widetilde{W})\right\|_{F} \leq O\left(\rho^{8} \tau_{0}^{1 / 3} m^{1 / 3}\right) \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{F} \tag{6.61}
\end{equation*}
$$

Finally, using $(a-b)^{2} \geq \frac{1}{2} a^{2}-b^{2}$, we have

$$
\begin{aligned}
\left\|\widehat{\nabla} f\left(\widetilde{W}+W^{\prime}\right)\right\|_{F}^{2} & \geq\left(\|\widehat{\nabla} f(\widetilde{W})\|_{F}-\left\|\widehat{\nabla} f\left(\widetilde{W}+W^{\prime}\right)-\widehat{\nabla} f(\widetilde{W})\right\|_{F}\right)^{2} \\
& \geq \frac{1}{2}\|\widehat{\nabla} f(\widetilde{W})\|_{F}^{2}-\left\|\widehat{\nabla} f\left(\widetilde{W}+W^{\prime}\right)-\widehat{\nabla} f(\widetilde{W})\right\|_{F}^{2} \\
& \geq \Omega\left(\frac{\delta}{\rho^{14}}\right) \times m \times\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|^{2} .
\end{aligned}
$$

where the last step follows by (6.61) (with our sufficiently large choice of $m$ ) and Theorem 6.14.2.

Upper bound on $\left\|\widehat{\nabla} f\left(\widetilde{W}+W^{\prime}\right)\right\|_{F} \quad$ Using $(a+b)^{2} \leq 2 a^{2}+2 b^{2}$, we have

$$
\begin{aligned}
\left\|\widehat{\nabla} f\left(\widetilde{W}+W^{\prime}\right)\right\|_{F}^{2} & \leq 2\|\widehat{\nabla} f(\widetilde{W})\|_{F}^{2}+2\left\|\widehat{\nabla} f\left(\widetilde{W}+W^{\prime}\right)-\widehat{\nabla} f(\widetilde{W})\right\|_{F}^{2} \\
& \stackrel{\otimes}{\leq} O\left(\rho^{12} m\right) \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{F}^{2}+2\left\|\widehat{\nabla} f\left(\widetilde{W}+W^{\prime}\right)-\widehat{\nabla} f(\widetilde{W})\right\|_{F}^{2} \stackrel{(2}{\leq} O\left(\rho^{12} m\right) \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{F}^{2}
\end{aligned}
$$

where (1) follows by Lemma 6.14 .1 and (2) follows by (6.61).

Upper bound on $\left\|\widehat{\nabla} f_{i}\left(\widetilde{W}+W^{\prime}\right)\right\|_{F} \quad$ This is completely analogous so we do not replicate the proofs here.

By applying an $\epsilon$-net argument over all possible $\operatorname{loss}_{i, \ell} \in \mathbb{R}^{d}$, we can modify Lemma 6.15.1 from "for fixed loss" to "for all loss." This allows us to plug in the true loss $\operatorname{loss}_{i, \ell}=B h_{i, \ell}-y_{i, \ell}^{*}$ and derive that:

Theorem 6.15.2 (restated). Let $\rho=n L d \log m$ and $\varrho=n L d \delta^{-1} \log (m / \epsilon)$. With probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$ over $W, A, B$, it satisfies for all $W^{\prime} \in \mathbb{R}^{m \times m}$ with $\left\|W^{\prime}\right\|_{2} \leq \frac{\varrho^{100}}{\sqrt{m}}$,

$$
\begin{aligned}
& \left\|\nabla f\left(\widetilde{W}+W^{\prime}\right)\right\|_{F}^{2} \geq \Omega\left(\frac{\delta}{\rho^{14}}\right) \times m \times \max _{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|^{2}\right\} \\
& \left\|\nabla f\left(\widetilde{W}+W^{\prime}\right)\right\|_{F}^{2} \leq O\left(\rho^{12} m\right) \times \max _{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|^{2}\right\} \\
& \left\|\nabla f_{i}\left(\widetilde{W}+W^{\prime}\right)\right\|_{F}^{2} \leq \frac{1}{n^{2}} O\left(\rho^{12} m\right) \times \max _{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|^{2}\right\}
\end{aligned}
$$

### 6.16 Objective Semi-Smoothness (Theorem 6.16.1)

The purpose of this section is to prove
Theorem 6.16.1 (objective semi-smoothness, restated). Let $\rho=n L d \log m, \varrho=n L d \delta^{-1} \log (m / \epsilon)$, $\tau_{0} \in\left[\varrho^{-100}, \varrho^{100}\right]$, and $\widetilde{W}, A, B$ be at random initialization. With probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $\widetilde{W}, A, B$, we have for every $\breve{W} \in \mathbb{R}^{m \times m}$ with $\|\breve{W}-\widetilde{W}\| \leq \frac{\tau_{0}}{\sqrt{m}}$, and for every $W^{\prime} \in \mathbb{R}^{m \times m}$ with $\left\|W^{\prime}\right\| \leq \frac{\tau_{0}}{\sqrt{m}}$, and letting $\left(i^{*}, \ell^{*}\right)=\operatorname{argmax}_{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}\right\|_{2}\right\}$,
$f\left(\breve{W}+W^{\prime}\right) \leq f(\breve{W})+\left\langle\nabla f(\breve{W}), W^{\prime}\right\rangle+O\left(\rho^{11} \tau_{0}^{1 / 3} m^{1 / 3}\right) \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}\right\|_{2} \cdot\left\|W^{\prime}\right\|_{2}+O\left(L^{18} n m\right)\left\|W^{\prime}\right\|_{2}^{2}$

We introduce the following notations before we go to proofs.
Definition 6.16.1. For $i \in[n]$ and $\ell \in[L]$ :

$$
\begin{array}{lll}
\widetilde{g}_{i, 0}=\widetilde{h}_{i, 0}=0 & \breve{g}_{i, 0} & =\hbar_{i, 0}=0 \\
\widetilde{g}_{i, \ell} & =\widetilde{W} \widetilde{h}_{i, \ell-1}+A x_{i, \ell} & \breve{g}_{i, \ell}
\end{array}=\breve{W} \hbar_{i, \ell-1}+A x_{i, \ell} \quad g_{i, \ell}=\left(\breve{W}+W^{\prime}\right) h_{i, \ell-1}+A x_{i, \ell} .
$$

Define diagonal matrices $\widetilde{D}_{i, \ell}$ and $\breve{D}_{i, \ell}$ respectively by letting

$$
\left(\widetilde{D}_{i, \ell}\right)_{k, k}=\mathbb{1}_{\left(\widetilde{g}_{i, \ell}\right)_{k} \geq 0} \text { and }\left(\breve{D}_{i, \ell}\right)_{k, k}=\mathbb{1}_{\left(\breve{g}_{i, \ell}\right)_{k} \geq 0} .
$$

The following claim gives rise to a new recursive formula to calculate $h_{i, \ell}-\hbar_{i, \ell}$.
Claim 6.16.2 (c.f. (6.21)). There exist diagonal matrices $D_{i, \ell}^{\prime \prime} \in \mathbb{R}^{m \times m}$ with entries in $[-1,1]$ such that,

$$
\begin{equation*}
\forall i \in[n], \forall \ell \in[L]: \quad h_{i, \ell}-\hbar_{i, \ell}=\sum_{a=1}^{\ell-1}\left(\breve{D}_{i, \ell}+D_{i, \ell}^{\prime \prime}\right) \breve{W} \cdots \breve{W}\left(\breve{D}_{i, a+1}+D_{i, a+1}^{\prime \prime}\right) W^{\prime} h_{i, a} \tag{6.62}
\end{equation*}
$$

Furthermore, we have $\left\|h_{i, \ell}-\hbar_{i, \ell}\right\| \leq O\left(L^{9}\right)\left\|W^{\prime}\right\|_{2}$ and $\left\|D_{i, \ell}^{\prime \prime}\right\|_{0} \leq O\left(L^{10 / 3} \tau_{0}^{2 / 3} m^{2 / 3}\right)$.

Proof of Claim 6.16.2. We ignore the subscript in $i$ for cleanness, and calculate that

$$
\begin{aligned}
h_{\ell}-\hbar_{\ell} & \stackrel{(1)}{=} \phi\left(\left(\breve{W}+W^{\prime}\right) h_{\ell-1}+A x_{\ell}\right)-\phi\left(\breve{W} \hbar_{\ell-1}+A x_{\ell}\right) \\
& \stackrel{(2)}{=}\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right)\left(\left(\breve{W}+W^{\prime}\right) h_{\ell-1}-\breve{W} \hbar_{\ell-1}\right) \\
& =\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right) \breve{W}\left(h_{\ell-1}-\hbar_{\ell-1}\right)+\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right) W^{\prime} h_{\ell-1} \\
& \stackrel{(3)}{=} \sum_{a=1}^{\ell-1}\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right) \breve{W} \cdots \breve{W}\left(\breve{D}_{a+1}+D_{a+1}^{\prime \prime}\right) W^{\prime} h_{a}
\end{aligned}
$$

Above, (1) is by the recursive definition of $h_{\ell}$ and $\hbar_{\ell}$; (2) is by Proposition 6.16 .3 and $D_{\ell}^{\prime \prime}$ is defined according to Proposition 6.16.3; and inequality (3) is by recursively computing $h_{\ell-1}-\hbar_{\ell-1}$. As for the two properties:

- We have $\left\|h_{\ell}-\hbar_{\ell}\right\| \leq O\left(L^{9}\right)\left\|W^{\prime}\right\|_{2}$. This is because we have
$-\left\|\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right) \breve{W} \cdots \breve{W}\left(\breve{D}_{a+1}+D_{a+1}^{\prime \prime}\right)\right\|_{2} \leq O\left(L^{7}\right)$ by Lemma 6.12.6;
- $\left\|h_{a}\right\| \leq O(L)$ (by $\left\|\widetilde{h}_{a}\right\| \leq O(L)$ from Lemma 6.11.1 and $\left\|\widetilde{h}_{a}-h_{a}\right\| \leq o(1)$ from Lemma 6.12.1a); and

$$
-\left\|W^{\prime} h_{a}\right\| \leq\left\|W^{\prime}\right\|_{2}\left\|h_{a}\right\|
$$

- We have $\left\|D_{\ell}^{\prime \prime}\right\|_{0} \leq O\left(L^{10 / 3} \tau_{0}^{2 / 3} m^{2 / 3}\right)$.

This is because, $\left(D_{\ell}^{\prime \prime}\right)_{k, k}$ is non-zero only at the coordinates $k \in[m]$ where the signs of $\breve{g}_{\ell}$ and $g_{\ell}$ are opposite (by Proposition 6.16.3). Such a coordinate $k$ must satisfy either $\left(\widetilde{D}_{\ell}\right)_{k, k} \neq\left(\breve{D}_{\ell}\right)_{k, k}$ or $\left(\widetilde{D}_{\ell}\right)_{k, k} \neq\left(D_{\ell}\right)_{k, k}$, and therefore by Lemma $6.12 .1-$ with probability $\geq 1-e^{-\Omega\left(L^{6} \tau_{0}^{4 / 3} m^{1 / 3}\right)}$ - there are at most $O\left(L^{10 / 3} \tau_{0}^{2 / 3} m^{2 / 3}\right)$ such coordinates $k$.

Proof of Theorem 6.16.1. First of all, since
$\frac{1}{2}\left\|B h_{i, \ell}-y_{i, \ell}^{*}\right\|^{2}=\frac{1}{2}\left\|\operatorname{loss}_{i, \ell}+B\left(h_{i, \ell}-\hbar_{i, \ell}\right)\right\|^{2}=\frac{1}{2} \| \breve{\left.\operatorname{loss}_{i, \ell}\left\|^{2}+\operatorname{loss}_{i, \ell}^{\top} B\left(h_{i, \ell}-\hbar_{i, \ell}\right)+\frac{1}{2}\right\| B\left(h_{i, \ell}-\hbar_{i, \ell}\right) \|^{2}{ }^{2}{ }^{2}\right)}$
we can write

$$
\begin{align*}
& f\left(\breve{W}+W^{\prime}\right)-f(W)-\left\langle\nabla f(W), W^{\prime}\right\rangle \\
& \stackrel{(1)}{=}-\left\langle\nabla f(\breve{W}), W^{\prime}\right\rangle+\frac{1}{2} \sum_{i=1}^{n} \sum_{\ell=2}^{L}\left\|B h_{i, \ell}-y_{i, \ell}^{*}\right\|^{2}-\left\|B \hbar_{i, \ell}-y_{i, \ell}^{*}\right\|^{2} \\
& \stackrel{(2)}{=}-\left\langle\nabla f(\breve{W}), W^{\prime}\right\rangle+\sum_{i=1}^{n} \sum_{\ell=2}^{L} \operatorname{loss}_{i, \ell}^{\top} B\left(h_{i, \ell}-\hbar_{i, \ell}\right)+\frac{1}{2}\left\|B\left(h_{i, \ell}-\hbar_{i, \ell}\right)\right\|^{2} \\
& \stackrel{(3)}{=} \sum_{i=1}^{n} \sum_{\ell=2}^{L} \breve{\operatorname{loss}}_{i, \ell}^{\top} B\left(\left(h_{i, \ell}-\hbar_{i, \ell}\right)-\sum_{a=1}^{\ell-1} \breve{D}_{i, \ell} \breve{W} \cdots \breve{W} \breve{D}_{i, a+1} W^{\prime} \hbar_{i, a}\right)+\frac{1}{2}\left\|B\left(h_{i, \ell}-\hbar_{i, \ell}\right)\right\|^{2} \\
& \stackrel{\oplus}{=} \sum_{i=1}^{n} \sum_{\ell=2}^{L} \operatorname{loss}_{i, \ell}^{\top} B\left(\sum_{a=1}^{\ell-1}\left(\breve{D}_{i, \ell}+D_{i, \ell}^{\prime \prime}\right) \breve{W} \cdots \breve{W}\left(\breve{D}_{i, a+1}+D_{i, a+1}^{\prime \prime}\right) W^{\prime} h_{i, a}-\breve{D}_{i, \ell} \breve{W} \cdots \breve{W}_{D_{i, a+1}} W^{\prime} \hbar_{i, a}\right) \\
& +\frac{1}{2}\left\|B\left(h_{i, \ell}-\hbar_{i, \ell}\right)\right\|^{2} \tag{6.64}
\end{align*}
$$

Above, (1) is by the definition of $f(\cdot)$; (2) is by (6.63); (3) is by the definition of $\nabla f(\cdot)$ (see Fact 6.2.3 for an explicit form of the gradient); (4) is by Claim 6.16.2.

We next bound the RHS of (6.64). We drop subscripts in $i$ for notational simplicity.

We first use (6.62) in Claim 6.16.2 to calculate

$$
\begin{align*}
\left\|B\left(h_{\ell}-\hbar_{\ell}\right)\right\| & \leq \sum_{a=1}^{\ell-1}\left\|B\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right) \breve{W} \cdots \breve{W}\left(\breve{D}_{a+1}+D_{a+1}^{\prime \prime}\right) W^{\prime} h_{a}\right\| \\
& \leq \sum_{a=1}^{\ell-1}\|B\|_{2}\left\|\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right) \breve{W} \cdots \breve{W}\left(\breve{D}_{a+1}+D_{a+1}^{\prime \prime}\right)\right\|_{2}\left\|W^{\prime}\right\|_{2}\left\|h_{a}\right\| \\
& \leq \sum_{a=1}^{\ell-1} O(\sqrt{m}) \cdot O\left(L^{7}\right) \cdot O(L) \cdot\left\|W^{\prime}\right\|_{2} \leq O\left(L^{9} \sqrt{m}\right) \cdot\left\|W^{\prime}\right\|_{2} . \tag{6.65}
\end{align*}
$$

In the last inequality (1) above, we have used $\left\|\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right) \breve{W} \cdots \breve{W}\left(\breve{D}_{a+1}+D_{a+1}^{\prime \prime}\right)\right\|_{2} \leq O\left(L^{7}\right)$ from Lemma 6.12 .6 ; we have used $\|B\|_{2} \leq O(\sqrt{m})$ with high probability; and we have used $\left\|h_{a}\right\| \leq O(L)$ (by $\left\|\widetilde{h}_{a}\right\| \leq O(L)$ from Lemma 6.11.1 and $\left\|\widetilde{h}_{a}-h_{a}\right\| \leq o(1)$ from Lemma 6.12.1a).

Since for each $D_{\ell}^{\prime \prime}$, we can write it as $D_{\ell}^{\prime \prime}=D_{\ell}^{0 / 1} D_{\ell}^{\prime \prime} D_{\ell}^{0 / 1}$, where each $D_{\ell}^{0 / 1}$ is a diagonal matrix satisfying

$$
\left(D_{\ell}^{0 / 1}\right)_{k, k}=\left\{\begin{array}{cc}
1, & \left(D_{\ell}^{\prime \prime}\right)_{k, k} \neq 0 ; \\
0, & \left(D_{\ell}^{\prime \prime}\right)_{k, k}=0 .
\end{array} \quad \text { and } \quad\left\|D_{\ell}^{0 / 1}\right\|_{0} \leq O\left(L^{10 / 3} \tau_{0}^{2 / 3} m^{2 / 3}\right)\right.
$$

Therefore,

$$
\begin{align*}
& \quad\left|\operatorname{loss}_{\ell}^{\top} B\left(\breve{D}_{\ell}+D_{\ell}^{\prime \prime}\right) \breve{W} \cdots \breve{W}\left(\breve{D}_{a+1}+D_{a+1}^{\prime \prime}\right) W^{\prime} h_{a}-\operatorname{loss}_{\ell}^{\top} B \breve{D}_{\ell} \breve{W} \cdots \breve{D}_{a+1} W^{\prime} h_{a}\right| \\
& \stackrel{(1)}{\leq}\left\|\operatorname{loss}_{\ell}\right\|_{2} \cdot \sum_{b=1}^{\ell-a}\binom{\ell-a}{b}\left\|B \breve{D} \breve{W} \cdots \breve{D} \breve{W} D^{0 / 1}\right\|_{2} \cdot\left\|D^{0 / 1} \breve{W} \breve{D} \cdots \breve{D} \breve{W} D^{0 / 1}\right\|_{2}^{b-1} \cdot\left\|D^{0 / 1} \breve{W} \cdots \breve{W} \breve{D} W^{\prime} h_{a}\right\|_{2} \\
& \stackrel{(2)}{\leq}\left\|\operatorname{loss}_{\ell}\right\|_{2} \cdot \sum_{b=1}^{\ell-a}\binom{\ell-a}{b} \cdot O\left(\rho^{2} \tau_{0}^{1 / 3} m^{1 / 3}\right) \cdot\left(O\left(\frac{\rho^{2}}{m^{1 / 6}}\right)\right)^{b-1} \cdot O\left(L^{7}\right) \cdot\left\|W^{\prime}\right\|_{2} \cdot O(L) \\
& \stackrel{(2}{\leq}\left\|\operatorname{loss}_{\ell}\right\|_{2} \cdot O\left(\rho^{11} \tau_{0}^{1 / 3} m^{1 / 3}\right) \cdot\left\|W^{\prime}\right\|_{2} \tag{6.66}
\end{align*}
$$

Above, (1) is an abbreviation and we have dropped the subscripts for the easy of presentation; (2) uses Corollary 6.11.15, Corollary 6.11.12, Lemma 6.12.6 and Lemma 6.11.1.

Finally, we also have

$$
\begin{align*}
\left|\breve{\operatorname{loss}}{ }_{\ell}^{\top} B \breve{D}_{\ell} \breve{W} \cdots \breve{D}_{a+1} W^{\prime}\left(h_{a}-\hbar_{a}\right)\right| & \stackrel{\oplus}{\leq}\left\|\operatorname{loss}_{\ell}\right\|_{2} \cdot O(\sqrt{m}) \cdot O\left(L^{7}\right) \cdot\left\|W^{\prime}\right\|_{2} \cdot\left\|h_{a}-\hbar_{a}\right\|_{2} \\
& \stackrel{(2}{\leq} O\left(\rho^{16} \sqrt{m}\right) \cdot\left\|\operatorname{loss}_{\ell}\right\|_{2} \cdot\left\|W^{\prime}\right\|_{2}^{2} \tag{6.67}
\end{align*}
$$

where (1) uses Lemma 6.12 .6 and (2) uses Claim 6.16 .2 to bound $\left\|h_{a}-\hbar_{a}\right\|_{2}$.
Putting (6.65), (6.66) and (6.67) back to (6.64), and using triangle inequality, we have the desired result.

### 6.16.1 Tool

Proposition 6.16.3. Given vectors $a, b \in \mathbb{R}^{m}$ and $D \in \mathbb{R}^{m \times m}$ the diagonal matrix where $D_{k, k}=\mathbb{1}_{a_{k} \geq 0}$. Then, then there exists a diagonal matrix $D^{\prime \prime} \in \mathbb{R}^{m \times m}$ with

- $\left|D_{k, k}^{\prime \prime}\right| \leq 1$ for every $k \in[m]$,
- $D_{k, k}^{\prime \prime} \neq 0$ only when $\mathbb{1}_{a_{k} \geq 0} \neq \mathbb{1}_{b_{k} \geq 0}$, and
- $\phi(a)-\phi(b)=D(a-b)+D^{\prime \prime}(a-b)$

Proof. We verify coordinate by coordinate for each $k \in[m]$.

- If $a_{k} \geq 0$ and $b_{k} \geq 0$, then $(\phi(a)-\phi(b))_{k}=a_{k}-b_{k}=(D(a-b))_{k}$.
- If $a_{k}<0$ and $b_{k}<0$, then $(\phi(a)-\phi(b))_{k}=0-0=(D(a-b))_{k}$.
- If $a_{k} \geq 0$ and $b_{k}<0$, then $(\phi(a)-\phi(b))_{k}=a_{k}=\left(a_{k}-b_{k}\right)+\frac{b_{k}}{a_{k}-b_{k}}\left(a_{k}-b_{k}\right)=$ $\left(D(a-b)+D^{\prime \prime}(a-b)\right)_{k}$, if we define $\left(D^{\prime \prime}\right)_{k, k}=\frac{b_{k}}{a_{k}-b_{k}} \in[-1,0]$.
- If $a_{k}<0$ and $b_{k} \geq 0$, then $(\phi(a)-\phi(b))_{k}=-b_{k}=0 \cdot\left(a_{k}-b_{k}\right)-\frac{b_{k}}{b_{k}-a_{k}}\left(a_{k}-b_{k}\right)=$ $\left(D(a-b)+D^{\prime \prime}(a-b)\right)_{k}$, if we define $\left(D^{\prime \prime}\right)_{k, k}=\frac{b_{k}}{b_{k}-b_{a}} \in[0,1]$.


### 6.17 Convergence Rate of Gradient Descent (Theorem 6.17.1)

Theorem 6.17.1 (gradient descent, restated). There exists some absolute constant $C>1$ such that the following holds. Let $\rho=n L d \log m, \epsilon \in(0,1], \varrho=n L d \delta^{-1} \log (m / \epsilon)$, and $\eta \stackrel{\text { def }}{=} \frac{\delta}{\rho^{44} m}$. Suppose $m \geq C \varrho^{C}$, and $W^{(0)}, A, B$ be at random initialization. Then, with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W^{(0)}, A, B$, suppose we start at $W^{(0)}$ and for each $t=0,1, \ldots, T-1$,

$$
W^{(t+1)}=W^{(t)}-\eta \nabla f\left(W^{(t)}\right)
$$

Then, it satisfies

$$
f\left(W^{(T)}\right) \leq \epsilon \quad \text { for all } \quad T \in\left[\frac{\rho^{59}}{\delta^{2}} \log \frac{1}{\epsilon}, \frac{\rho^{59} \varrho^{28}}{\delta^{2}} \log \frac{1}{\epsilon}\right] .
$$

In other words, the training loss of the recurrent neural network drops to $\epsilon$ in a linear convergence speed.

To present the simplest possible result, we have not tried to tighten the polynomial dependency with respect to $n, d$ and $L$. We only tightened the dependency with respect to $\delta$ and $\epsilon$. In fact, a more involved analysis can also get ride of the $\log (1 / \epsilon)$ dependency in $m$ [AZLS19].

Proof of Theorem 6.17.1. Using Lemma 6.11.1 and the randomness of $B$, it is easy to show that $\left\|B h_{i, \ell}-y_{i, \ell}^{*}\right\|^{2} \leq O\left(\rho^{2} L^{2}\right)$ with at least $1-e^{-\Omega\left(\rho^{2}\right)}$ (where $h_{i, \ell}$ is defined with respect to $\left.W^{(0)}\right)$, and therefore

$$
f\left(W^{(0)}\right) \leq O\left(n \rho^{2} L^{3}\right)
$$

In the rest of the proof, we first assume that for every $t=0,1, \ldots, T-1$, the following holds

$$
\begin{equation*}
\left\|W^{(t)}-W^{(0)}\right\|_{F} \leq \frac{\tau_{0}}{\sqrt{m}} \stackrel{\text { def }}{=} \frac{\varrho^{50}}{\sqrt{m}} \tag{6.68}
\end{equation*}
$$

We shall prove the convergence of gradient descent assuming (6.68), so that previous statements such as Theorem 6.16.1 and Theorem 6.15.2 can be applied. At the end of the proof, we shall verify that (6.68) is satisfied throughout the gradient descent process.

For each $t=0,1, \ldots, T-1$, we let $\left\{\operatorname{loss}_{i, \ell}^{(t)}\right\}_{i, \ell}$ denote the loss vectors with respect to the current point $W^{(t)}$. We denote by $\left(i^{*}, \ell^{*}\right)=\operatorname{argmax}_{i, \ell}\left\{\left\|\operatorname{loss}_{i, \ell}^{(t)}\right\|_{2}\right\}$ and $\nabla_{t}=\nabla f\left(W^{(t)}\right)$.

We calculate that

$$
\begin{aligned}
f\left(W^{(t+1)}\right) & \stackrel{(1)}{\leq} f\left(W^{(t)}\right)-\eta\left\|\nabla_{t}\right\|_{F}^{2}+O\left(\rho^{11} \tau_{0}^{1 / 3} m^{1 / 3}\right) \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}^{(t)}\right\|_{2} \cdot \eta\left\|\nabla_{t}\right\|_{2}+O\left(L^{18} n m \eta^{2}\right)\left\|\nabla_{t}\right\|_{2}^{2} \\
& \stackrel{(2)}{\leq} f\left(W^{(t)}\right)-\eta\left\|\nabla_{t}\right\|_{F}^{2}+O\left(\rho^{30} \eta^{2} m^{2}\right) \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}^{(t)}\right\|_{2}^{2} \\
& \stackrel{\circledR}{\leq} f\left(W^{(t)}\right)-\left(\Omega\left(\frac{\eta \delta}{\rho^{14}} m\right)-O\left(\rho^{30} \eta^{2} m^{2}\right)\right) \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}^{(t)}\right\|_{2}^{2} \\
& \stackrel{\oplus}{\leq} f\left(W^{(t)}\right)-\Omega\left(\frac{\eta \delta}{\rho^{14}} m\right) \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}^{(t)}\right\|_{2}^{2} \\
& \stackrel{\text { © }}{\leq}\left(1-\Omega\left(\frac{\eta \delta}{\rho^{15}} m\right)\right) f\left(W^{(t)}\right)
\end{aligned}
$$

Above, (1) uses Theorem 6.16.1; (2) uses Theorem 6.15.2 (which gives $\left\|\nabla_{t}\right\|_{2} \leq\left\|\nabla_{t}\right\|_{F} \leq$ $\left.O\left(\rho^{6} \sqrt{m}\right) \times\left\|\operatorname{loss}_{i^{*},,^{*}}^{(t)}\right\|\right)$, and our choices of $\tau_{0}$ and $m$; (3) uses Theorem 6.15.2; (4) uses our choice of $\eta$; and (5) uses $f\left(W^{(t)}\right) \leq n L\left\|\operatorname{loss}_{i^{*}, \ell^{*}}^{(t)}\right\|^{2}$. In other words, after $T=\Omega\left(\frac{\rho^{15}}{\eta \delta m}\right) \log \frac{n L^{2}}{\epsilon}$ iterations we have $f\left(W^{(T)}\right) \leq \epsilon$.

We need to verify for each $t,\left\|W^{(t)}-W^{(0)}\right\|_{F}$ is small so that (6.68) holds. By Theorem 6.15.2,

$$
\begin{aligned}
\left\|W^{(t)}-W^{(0)}\right\|_{F} \leq \sum_{i=0}^{t-1}\left\|\eta \nabla f\left(W^{(i)}\right)\right\|_{F} \leq O\left(\eta \rho^{6} \sqrt{m}\right) \cdot \sum_{i=0}^{t-1} \sqrt{f\left(W^{(i)}\right)} & \leq O\left(\eta \rho^{6} \sqrt{m}\right) \cdot O\left(T \cdot \sqrt{n \rho^{2} L^{3}}\right) \\
& \leq \eta T \cdot O\left(\rho^{8.5} \sqrt{m}\right) \leq \frac{\varrho^{50}}{\sqrt{m}}
\end{aligned}
$$

where the last step follows by our choice of $T$.

### 6.18 Convergence Rate of Stochastic Gradient Descent (Theorem 6.18.1)

Theorem 6.18.1 (stochastic gradient descent, stated). There exists some absolute constant $C>1$ such that the following holds. Let $\rho=n L d \log m, \epsilon \in(0,1], \varrho=n L d \delta^{-1} \log (m / \epsilon)$, and $\eta \stackrel{\text { def }}{=} \frac{\delta}{\rho^{42} m}$. Suppose $m \geq C \varrho^{C}$, and $W^{(0)}, A, B$ be at random initialization. Then, with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $W^{(0)}, A, B$, suppose we start at $W^{(0)}$ and for each $t=0,1, \ldots, T-1$,

$$
W^{(t+1)}=W^{(t)}-\eta \cdot \frac{n}{\left|S_{t}\right|} \sum_{i \in S_{t}} \nabla f\left(W^{(t)}\right) \quad \text { (for a random subset } S_{t} \subseteq[n] \text { of fixed cardinality.) }
$$

Then, it satisfies with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$ over the randomness of $S_{1}, \ldots, S_{T}$ :

$$
f\left(W^{(T)}\right) \leq \epsilon \quad \text { for all } \quad T \in\left[\frac{\rho^{57}}{\delta^{2}} \log \frac{1}{\epsilon}, \frac{\rho^{57} \varrho^{28}}{\delta^{2}} \log \frac{1}{\epsilon}\right] .
$$

(To present the simplest possible result, we have not tried to tighten the polynomial dependency with respect to $n, d$ and $L$. We only tightened the dependency with respect to $\delta$ and $\epsilon$. In fact, a more involved analysis can also get ride of the $\log (1 / \epsilon)$ dependency in $m$.)

Proof of Theorem 6.18.1. The proof is almost identical to that of Theorem 6.17.1. We again have with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$

$$
f\left(W^{(0)}\right) \leq O\left(n \rho^{2} L^{3}\right)
$$

Again, we first assume for every $t=0,1, \ldots, T-1$, the following holds

$$
\begin{equation*}
\left\|W^{(t)}-W^{(0)}\right\|_{F} \leq \frac{\tau_{0}}{\sqrt{m}} \stackrel{\text { def }}{=} \frac{\varrho^{50}}{\sqrt{m}} \tag{6.69}
\end{equation*}
$$

We shall prove the convergence of SGD assuming (6.69), so that previous statements such as Theorem 6.16.1 and Theorem 6.15.2 can be applied. At the end of the proof, we shall verify that (6.69) is satisfied throughout the SGD with high probability.

For each $t=0,1, \ldots, T-1$, using the same notation as Theorem 6.17.1, except that we choose $\nabla_{t}=\frac{n}{\left|S_{t}\right|} \sum_{i \in S_{t}} \nabla f_{i}\left(W^{(t)}\right)$. We have $\mathbb{E}_{S_{t}}\left[\nabla_{t}\right]=\nabla f\left(W^{(t)}\right)$ and therefore

$$
\begin{align*}
& \quad \underset{S_{t}}{\mathbb{E}}\left[f\left(W^{(t+1)}\right)\right] \\
& \stackrel{(1)}{\leq} f\left(W^{(t)}\right)-\eta\left\|\nabla f\left(W^{(t)}\right)\right\|_{F}^{2}+O\left(\rho^{11} \tau_{0}^{1 / 3} m^{1 / 3}\right) \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}^{(t)}\right\|_{2} \cdot \eta \underset{S_{t}}{\mathbb{E}}\left[\left\|\nabla_{t}\right\|_{2}\right]+O\left(L^{18} n m \eta^{2}\right) \underset{S_{t}}{\underset{E}{E}}\left[\left\|\nabla_{t}\right\|_{2}^{2}\right] \\
& \stackrel{\text { (2) }}{\leq} f\left(W^{(t)}\right)-\eta\left\|\nabla_{t}\right\|_{F}^{2}+O\left(\rho^{30} \eta^{2} m^{2}\right) \cdot\left\|\operatorname{loss}_{i^{*},,^{*}}^{(t)}\right\|_{2}^{2} \\
& \stackrel{\text { ® }}{\leq}\left(1-\Omega\left(\frac{\eta \delta}{\rho^{15}} m\right)\right) f\left(W^{(t)}\right) . \tag{6.70}
\end{align*}
$$

Above, (1) uses Theorem 6.16.1 and $\mathbb{E}_{S_{t}}\left[\nabla_{t}\right]=\nabla f\left(W^{(t)}\right)$; (2) uses Theorem 6.15.2 (which gives $\left.\left\|\nabla_{t}\right\|_{2}^{2} \leq \frac{n^{2}}{\left|S_{t}\right|^{2}} \sum_{i \in S_{t}}\left\|\nabla f_{i}\left(W^{(t)}\right)\right\|_{F}^{2} \leq O\left(\rho^{12} m\right) \times\left\|\operatorname{loss}_{i^{*}, \ell^{*}}^{(t)}\right\|^{2}\right)$; (3) is identical to the proof of Theorem 6.17.1.

At the same time, we also have the following absolute value bound:

$$
\begin{align*}
& \quad f\left(W^{(t+1)}\right) \\
& \stackrel{(1)}{\leq} f\left(W^{(t)}\right)+\eta\left\|\nabla f\left(W^{(t)}\right)\right\|_{F} \cdot\left\|\nabla_{t}\right\|_{F}+O\left(\rho^{11} \tau_{0}^{1 / 3} m^{1 / 3}\right) \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}^{(t)}\right\|_{2} \cdot \eta\left\|\nabla_{t}\right\|_{2}+O\left(L^{18} n m \eta^{2}\right)\left\|\nabla_{t}\right\|_{2}^{2} \\
& \stackrel{(2)}{\leq} f\left(W^{(t)}\right)+O\left(\rho^{12} \eta m+\rho^{30} \eta^{2} m^{2}\right) \cdot\left\|\operatorname{loss}_{i^{*}, \ell^{*}}^{(t)}\right\|_{2}^{2} \\
& \stackrel{\circledR}{\leq}\left(1+O\left(\rho^{12} \eta m\right)\right) f\left(W^{(t)}\right) . \tag{6.71}
\end{align*}
$$

Above, (1) uses Theorem 6.16.1 and Cauchy-Shwartz $\langle A, B\rangle \leq\|A\|_{F}\|B\|_{F}$, and (2) uses Theorem 6.15.2 and the derivation from (6.70).

Next, taking logarithm on both sides of (6.70) and (6.71), and using Jensen's inequality $\mathbb{E}[\log X] \leq \log \mathbb{E}[X]$, we have
$\mathbb{E}\left[\log f\left(W^{(t+1)}\right)\right] \leq \log f\left(W^{(t)}\right)-\Omega\left(\frac{\eta \delta}{\rho^{15}} m\right) \quad$ and $\quad \log f\left(W^{(t+1)}\right) \leq \log f\left(W^{(t)}\right)+O\left(\rho^{12} \eta m\right)$

By one-sided Azuma's inequality (a.k.a. martingale concentration), we have with probability at least $1-e^{-\Omega\left(\rho^{2}\right)}$, for every $t=1,2, \ldots, T$ :

$$
\log f\left(W^{(t)}\right)-\mathbb{E}\left[\log f\left(W^{(t)}\right)\right] \leq \sqrt{t} \cdot O\left(\rho^{12} \eta m\right) \cdot \rho
$$

This implies two things.

- On one hand, after $T=\Omega\left(\frac{\rho^{15} \log \left(n L^{2} / \epsilon\right)}{\eta \delta m}\right)$ iterations we have

$$
\begin{aligned}
\log f\left(W^{(T)}\right) & \leq \sqrt{T} \cdot O\left(\rho^{12} \eta m\right) \cdot \rho+\log f\left(W^{(0)}\right)-\Omega\left(\frac{\eta \delta}{\rho^{15}} m\right) T \\
& \leq \log f\left(W^{(0)}\right)-\Omega\left(\frac{\eta \delta}{\rho^{15}} m\right) T \leq \log O\left(n \rho^{2} L^{3}\right)-\Omega\left(\log \frac{\rho^{5}}{\epsilon}\right) \leq \log \epsilon
\end{aligned}
$$

Therefore, we have $f\left(W^{(T)}\right) \leq \epsilon$.

- On the other hand, for every $t=1,2, \ldots, T$, we have

$$
\begin{aligned}
\log f\left(W^{(t)}\right) & \leq \sqrt{t} \cdot O\left(\rho^{12} \eta m\right) \cdot \rho+\log f\left(W^{(0)}\right)-\Omega\left(\frac{\eta \delta}{\rho^{15}} m\right) t \\
& \stackrel{(1)}{=} \log f\left(W^{(0)}\right)-\left(\sqrt{\frac{\eta \delta m}{\rho^{15}}} \cdot \Omega(\sqrt{t})-\sqrt{\frac{\rho^{15}}{\eta \delta m}} \cdot O\left(\rho^{13} \eta m\right)\right)^{2}+O\left(\frac{\rho^{41} \eta m}{\delta}\right) \\
& \stackrel{(2)}{\leq} \log f\left(W^{(0)}\right)+1
\end{aligned}
$$

where in (1) we have used $2 a \sqrt{t}-b^{2} t=-(b \sqrt{t}-a / b)^{2}+a^{2} / b^{2}$, and in (2) we have used $\eta \leq O\left(\frac{\delta}{\rho^{42} m}\right)$. This implies $f\left(W^{(t)}\right) \leq O\left(n \rho^{2} L^{3}\right)$. We can now verify for each $t$, $\left\|W^{(t)}-W^{(0)}\right\|_{F}$ is small so that (6.69) holds. By Theorem 6.15.2,

$$
\begin{aligned}
\left\|W^{(t)}-W^{(0)}\right\|_{F} \leq \sum_{i=0}^{t-1}\left\|\eta \nabla_{t}\right\|_{F} \leq O\left(\eta \rho^{6} \sqrt{m}\right) \cdot \sum_{i=0}^{t-1} \sqrt{f\left(W^{(i)}\right)} & \leq O\left(\eta \rho^{6} \sqrt{m}\right) \cdot O\left(T \sqrt{n \rho^{2} L^{3}}\right) \\
& \leq \eta T \cdot O\left(\rho^{8.5} \sqrt{m}\right) \leq \frac{\varrho^{50}}{\sqrt{m}}
\end{aligned}
$$

where the last step follows by our choice of $T$.

## Part II

## Matrix Concentrations

## Chapter 7

## Matrix Chernoff Bound on Expander

We prove a Chernoff-type bound for sums of matrix-valued random variables sampled via a random walk on an expander, confirming a conjecture due to Wigderson and Xiao [WX06]. Our proof is based on a new multi-matrix extension of the Golden-Thompson inequality which improves in some ways the inequality in [SBT17] and may be of independent interest, as well as an adaptation of an argument for the scalar case due to Healy [Hea08]. Secondarily, we also provide a generic reduction showing that any concentration inequality for vector-valued martingales implies a concentration inequality for the corresponding expander walk, with a weakening of parameters proportional to the squared mixing time.

### 7.1 Introduction

The Chernoff Bound [Che52] is one of the most widely used probabilistic results in computer science. It states that a sum of independent bounded random variables exhibits subgaussian concentration around its mean. In particular, when the random variables are i.i.d. samples from a fixed distribution, it implies that the empirical mean of $k$ samples is $\epsilon$-close to the true mean with exponentially small deviation probability proportional to $e^{-\Omega\left(k \epsilon^{2}\right)}$.

An important generalization of this bound was achieved by Gillman [Gi198] (with refinements later by [Lez98, Kah97, LP04, WX05, Hea08, Wag08, CLLM12, RR17]), who significantly relaxed the independence assumption to Markov dependence. In particular, suppose $G$ is a regular graph with vertex set $V=[n], X: V \rightarrow \mathbb{C}$ is a bounded function, and $v_{1}, \ldots, v_{k}$ is a stationary random walk ${ }^{1}$ of length $k$ on $G$. Then, even though the random variables $X\left(v_{i}\right)$ are in not independent (except when $G$ is the complete graph with self loops), it is shown that:

$$
\begin{equation*}
\operatorname{Pr}\left[\left|\frac{1}{k} \sum_{i=1}^{k} X\left(v_{i}\right)-\mathbb{E}[X]\right|>\epsilon\right] \leq 2 \cdot \exp \left(-\Omega\left((1-\lambda) k \epsilon^{2}\right)\right) \tag{7.1}
\end{equation*}
$$

where $1-\lambda$ is the spectral gap of the transition matrix of the random walk. The gain here is that sampling a stationary random walk of length $k$ on a constant degree graph with constant spectral gap requires $\log (n)+O(k)$ random bits, which is much less than the $k \log (n)$ bits required to produce $k$ independent samples. Since such graphs can be explicitly constructed, this leads to a generic "derandomization" of the Chernoff bound, which has

[^47]had several important applications (see [WX05] for a detailed discussion). In particular, it leads to the following randomness efficient sampler for scalar-valued functions ([Gil98]) using known strongly explicit constructions of expander graphs [RVW00, LPS88]:

Theorem 7.1.1 ([Gil98]). For any $\epsilon>0$ and $k \geq 1$, there is a poly $(r)$-time computable sampler $\sigma:\{0,1\}^{r} \rightarrow[n]^{k}$, where $r=\log (n)+O(k)$ s.t. for all functions $f:[n] \rightarrow[-1,1]$ satisfying $\mathbb{E} f=0$, we have that

$$
\operatorname{Pr}_{w \in R\{0,1\}^{r}}\left[\left|\frac{1}{k} \sum_{i=1}^{k} f\left(\sigma(w)_{i}\right)\right| \geq \epsilon\right] \leq 2 \exp \left(-\Omega\left(\epsilon^{2} k\right)\right) .
$$

In many applications of interest $k$ is about $\log (n)$, and going from $O\left(\log ^{2}(n)\right)$ to $O(\log (n))$ random bits leads to a complete derandomization by cycling over all seeds $w \in$ $\{0,1\}^{r}$.

A different generalization of the Chernoff bound appeared in the works of Rudelson [Rud99], Ahlswede-Winter [AW02], and Tropp [Tro12], who showed that a similar concentration phenomenon is true for matrix-valued random variables. In particular, if $X_{1}, \ldots, X_{k}$ are independent $d \times d$ complex Hermitian random matrices with $\left\|X_{i}\right\| \leq 1$, then the following is true:

$$
\begin{equation*}
\operatorname{Pr}\left[\left\|\frac{1}{k} \sum_{i=1}^{k} X_{i}-\mathbb{E}[X]\right\|>\epsilon\right] \leq 2 d \cdot \exp \left(-\Omega\left(k \epsilon^{2}\right)\right) \tag{7.2}
\end{equation*}
$$

The only difference between this and the usual Chernoff bound is the factor of $d$ in front of the deviation probability; to see that it is necessary, notice that the diagonal case simply corresponds to a direct sum of $d$ arbitrarily correlated instances of the scalar Chernoff bound, so by the union bound the probability should be $d$ times as large in the worst case. This so called "Matrix Chernoff Bound" has seen several applications as well, notably in quantum
information theory, numerical linear algebra, and spectral graph theory; the reader may consult e.g. the book [Tro15] for many examples.

We present two different extensions of the above results in this chapter.

### 7.1.1 A Matrix Expander Chernoff Bound

It is natural to wonder whether there is a common generalization of (7.1) and (7.2), i.e., a "Matrix Expander Chernoff Bound". Such a result was conjectured by Wigderson and Xiao in [WX06] — in fact, [WX05] contained a proof of it, but the authors later discovered a gap in the proof. In this chapter, we prove the Wigderson and Xiao conjecture, namely:

Theorem 7.1.2. Let $G=(V, E)$ be a regular undirected graph whose transition matrix has second eigenvalue $\lambda$, and let $f: V \rightarrow \mathbb{C}^{d \times d}$ be a function such that:

1. For each $v \in V, f(v)$ is Hermitian and $\|f(v)\| \leq 1$.
2. $\sum_{v \in V} f(v)=0$.

Then, for a stationary random walk $v_{1}, \ldots, v_{k}$ with $\epsilon \in(0,1)$ we have:

$$
\begin{aligned}
& \operatorname{Pr}\left[\lambda_{\max }\left(\frac{1}{k} \sum_{j=1}^{k} f\left(v_{j}\right)\right) \geq \epsilon\right] \leq d \cdot \exp \left(-\Omega\left(\epsilon^{2}(1-\lambda) k\right)\right), \\
& \operatorname{Pr}\left[\lambda_{\min }\left(\frac{1}{k} \sum_{j=1}^{k} f\left(v_{j}\right)\right) \leq-\epsilon\right] \leq d \cdot \exp \left(-\Omega\left(\epsilon^{2}(1-\lambda) k\right)\right) .
\end{aligned}
$$

This theorem adds to the amazingly long list of pseudorandom properties of expander graphs. By applying the theorem with a strongly explicit bounded degree expander, one obtains the following randomness-efficient sampler for matrix-calued functions conjectured in [WX06].

Theorem 7.1.3. For any $\epsilon>0, k \geq 1$ and $d \geq 1$, there is a poly $(r)$-time computable sampler $\sigma:\{0,1\}^{r} \rightarrow[n]^{k}$, where $r=\log (n)+O(k)$ s.t. for all functions $f:[n] \rightarrow \mathbb{C}^{d \times d}$ satisfying $\mathbb{E} f=0$ and for each $v \in[n], f(v)$ is Hermitian and $\|f(v)\| \leq 1$, we have that

$$
\operatorname{Pr}_{w \in_{R}\{0,1\}^{r}}\left[\left\|\frac{1}{k} \sum_{i=1}^{k} f\left(\sigma(w)_{i}\right)\right\| \geq \epsilon\right] \leq 2 d \exp \left(-\Omega\left(-\epsilon^{2} k\right)\right) .
$$

We remark that while the derandomization applications studied in [WX05] were later recovered in [WX08] using the method of pessimistic estimators, that method requires additional assumptions to be efficiently implementable (specifically, computability of the matrix moment generating function, which is problem-dependent) and therefore does not constitute a truly black box derandomization of the matrix Chernoff bound, whereas Theorem 7.1.3 does. Given the increasing ubiquity of applications of this bound, we therefore suspect that it will find further applications in the study of derandomization and expander graphs, beyond the ones mentioned in [WX05].

## Techniques

To describe the ideas that go into the proof of Theorem 7.1.2, let us begin by recalling how the usual scalar Chernoff bound is proved, in the case when the random variables have mean zero. The key observation is that if $X_{1}, \ldots, X_{k}$ are independent random variables, then the moment generating function of the sum is equal to the product of the moment generating functions:

$$
\mathbb{E}\left[\exp \left(t \sum_{i=1}^{k} X_{i}\right)\right]=\prod_{i=1}^{k} \mathbb{E}\left[\exp \left(t X_{i}\right)\right] .
$$

This is no longer true in case where the $X_{i}$ come from a random walk, but we still have the algebraic fact that

$$
\begin{equation*}
\exp \left(t \sum_{i=1}^{k} X_{i}\right)=\prod_{i=1}^{k} \exp \left(t X_{i}\right) \tag{7.3}
\end{equation*}
$$

which allows one to decompose the sum as a product. The latter allows one to consider the steps of the random walk separately and analyze the change in the expectation inductively.

The analogue of the moment generating function in the matrix setting is

$$
\mathbb{E}\left[\operatorname{tr}\left[\exp \left(t \sum_{i=1}^{k} X_{i}\right)\right]\right]
$$

and the main difficulty is that (7.3) no longer holds if the matrices $X_{i}$ do not commute. A substitute for this fact is given by the Golden-Thompson inequality [Gol65, Tho65], which states that for any Hermitian $A, B$ :

$$
\begin{equation*}
\operatorname{tr}[\exp (A+B)] \leq \operatorname{tr}[\exp (A) \exp (B)] \tag{7.4}
\end{equation*}
$$

The latter expression may further be bounded by $\|\exp (A)\| \operatorname{tr}[\exp (B)]$, and this is sufficient to prove (7.2) in the independent case as is done in [AW02], where an inductive application of it yields

$$
\mathbb{E}\left[\operatorname{tr}\left[\exp \left(t \sum_{i=1}^{k} X_{i}\right)\right]\right] \leq \operatorname{tr}[I] \cdot \prod_{i=1}^{k}\left\|\mathbb{E}\left[\exp \left(t X_{i}\right)\right]\right\|
$$

However, this approach is too crude to handle the Markov case, roughly because in the absence of inependence, passing to the norm makes it difficult to utilize the fact that the expectation of each $X_{i}$ is zero.

The original proof of Wigderson-Xiao was based on the following plausible multimatrix generalization of (7.4):

$$
\operatorname{tr}\left[\exp \left(\sum_{i=1}^{k} A_{i}\right)\right] \leq \operatorname{tr}\left[\prod_{i=1}^{k} \exp \left(A_{i}\right)\right],
$$

which turns out to be false for $k>2$. To see why, observe that the left hand side is always nonnegative, whereas the right hand side can be the trace of a product of any three positive semidefinite matrices, which can be negative (and this is not the case for two matrices). This led to a fatal gap in their proof.

The main ingredient in our proof is a new multi-matrix generalization of (7.4), which is inspired by the following statement that was recently proven in [SBT17] (see also [HKT16]).

Theorem 7.1.4 (Corollary 3.3 in [SBT17]). Let $H_{1}, \ldots, H_{k} \in \mathbb{C}^{d \times d}$ be Hermitian matrices. Then
$\log \left[\operatorname{tr}\left(\exp \left(\sum_{j=1}^{k} H_{j}\right)\right)\right] \leq \int_{-\infty}^{\infty} \log \left[\operatorname{tr}\left(\prod_{j=1}^{k} \exp \left(\frac{H_{j}(1+\mathbf{i} b)}{2}\right) \prod_{j=k}^{1} \exp \left(\frac{H_{j}(1-\mathbf{i} b)}{2}\right)\right)\right] \mathrm{d} \mu(b)$
where $\mu$ is some probability distribution on $(-\infty, \infty)$.

The above inequality successfully relates the matrix exponential of a sum to a product of matrix exponential, but is not adequate for proving an optimal Chernoff bound. The reason is that all known arguments require a Taylor expansion, and Theorem 7.1.4 involves integration over an unbounded region (this region can be truncated, but this introduces a loss which leads to a suboptimal bound). To remedy this, we prove a new multi-matrix Golden-Thompson inequality, which only involves integration over a bounded region instead of a line.

Theorem 7.1.5 (Bounded Multi-matrix Golden-Thompson inequality). Let $H_{1}, \ldots, H_{k} \in$ $\mathbb{C}^{d \times d}$ be Hermitian matrices. Then

$$
\log \left(\operatorname{tr}\left[\exp \left(\sum_{j=1}^{k} H_{j}\right)\right]\right) \leq \frac{4}{\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \log \left(\operatorname{tr}\left[\prod_{j=1}^{k} \exp \left(\frac{e^{\mathbf{i} \phi}}{2} H_{j}\right) \prod_{j=k}^{1} \exp \left(\frac{e^{-\mathrm{i} \phi}}{2} H_{j}\right)\right]\right) \mathrm{d} \mu(\phi)
$$

where $\mu$ is some probability distribution on $\left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$.

We present the proof in Section 7.3. Theorem 7.1.5 is likely to be of independent interest and could have further applications, e.g. in quantum information theory. We draw attention to the following notable features of the above two theorems:
(a) Since $\exp (H(1+\mathbf{i} b))=\exp (H(1-\mathbf{i} b))^{*}$ and $\exp \left(\frac{e^{\mathbf{i} \phi}}{2} H_{j}\right)=\exp \left(\frac{e^{-\mathbf{i} \phi}}{2} H_{j}\right)^{*}$ for Hermitian $H$, the right hand always considers the trace of a matrix times its adjoint, which is always positive semidefinite, ruling out the bad example described above.
(b) They are average case inequalities, where the averaging is done over specific distributions. We remark that the first inequality is known to be false in the worst case (i.e., with $b=0$ and with other small values of $b$; see [SBT17] for a discussion).

The main point is that Theorem 7.1.5 allows one to relate the exponential of a sum of matrices to a (two-sided) product of bounded $d \times d$ matrices and their adjoints. In order to prove Theorem 7.1.2, we rewrite this as a one-sided matrix product of $d^{2} \times d^{2}$ matrices acting on $\mathbb{C}^{d \times d}$, by encoding left and right multiplication on this space via a tensor product. However, these matrices are no longer Hermitian (or even normal), so it is difficult to analyze the moment generating function of their product over the random walk using the perturbation-theoretic approach of [WX05]. We surmount this difficulty by employing a
variant of the more robust linear algebraic proof technique of Healy [Hea08]. The proof of Theorem 7.1.2 is presented in Section 7.4.

### 7.1.2 Martingale Approximation of Expander Walks

While Theorem 7.1.2 provides a satisfactory generalization of the expander Chernoff bound to the case when one is interested in the spectral norm of a matrix-valued function on $V$, one could ask what happens for other matrix norms (such as Schatten norms), or even more generally, for functions taking values in an arbitrary Banach space. Our second contribution is a generic reduction from this problem, of proving concentration for random variables sampled using a Markov chain, to the much more well-studied problem (see e.g. [CL06]) of concentration for sums of martingale random variables. We remark that the technique used in the proof of Theorem 7.1.6 is not new, and was introduced in a slightly different context in the paper [NPSS06].

Theorem 7.1.6. Suppose $G=(V, E)$ is a regular graph whose transition matrix has second eigenvalue $\lambda$ and $f: V \rightarrow \mathbb{R}^{N}$ is a vector-valued function satisfying $\sum_{v \in V} f(v)=0$ with $F:=\sqrt{\sum_{v \in V}\|f(v)\|_{2}^{2}}$, where $\|\cdot\|_{2}$ denotes the Frobenius norm. If $v_{1}, \ldots, v_{k}$ is a stationary random walk on $G$, then for every $\epsilon>0$, there is a martingale difference sequence $Z_{1}, \ldots, Z_{k}$ with respect to the filtration generated by initial segments of $v_{1}, \ldots, v_{k}$ such that

$$
\frac{1}{k} \sum_{i=1}^{k} f\left(v_{i}\right)=W+\frac{1}{k} \sum_{i=1}^{k} Z_{i}
$$

where

1. $W$ is a random vector satisfying $\|W\|_{2} \leq \epsilon$.
2. Each term $Z_{i}$ satisfies

$$
\left\|Z_{i}\right\|_{*} \leq \frac{2 \log (F / \epsilon)}{1-\lambda} \cdot \max _{v \in V}\|f(v)\|_{*}
$$

for every norm $\|\cdot\|_{*}$.

Thus, the empirical sums of any bounded (in any norm) function on a graph are well-approximated by a martingale whose increments are also bounded, with a loss in the bound depending on the $\ell_{2}$ norm $F$ of the function and the spectral gap of the graph. Since $F$ will typically scale with the number of vertices, the ratio above is typically comparable to the mixing time.

To see the theorem in action, consider the case when $f(v)$ is matrix-valued in $d \times d$ Hermitian matrices and $\|\cdot\|_{*}$ is the operator norm. If $\|f(v)\| \leq 1$ then we have the bound

$$
F^{2}=\sum_{v \in V}\|f(v)\|_{F}^{2} \leq d n
$$

Suppose we are interested in obtaining an estimate on the probability:

$$
\begin{equation*}
\operatorname{Pr}\left[\left\|\frac{1}{k} \sum_{v \in V} f(v)\right\|>\epsilon\right] . \tag{7.5}
\end{equation*}
$$

Applying Theorem 7.1.6 with parameter $\epsilon / 2$ and noting that $\|W\| \leq\|W\|_{F}$, we have that (7.5) is at most

$$
\operatorname{Pr}\left[\left\|\frac{1}{k} \sum_{i=1}^{k} Z_{i}\right\|>\epsilon / 2\right],
$$

where $Z_{i}$ is a martingale with bound

$$
\left\|Z_{i}\right\| \leq \frac{\log (n d / \epsilon)}{1-\lambda}
$$

We now appeal to the existing martingale generalization of (7.2) (see e.g. [Tro12]) and find that this probability is at most

$$
2 d \cdot \exp \left(-\Omega\left(\frac{k \epsilon^{2}(1-\lambda)^{2}}{\log ^{2}(n d)}\right)\right)
$$

While this theorem is much weaker than the previous one in terms of parameters (depending on the square of the mixing time rather than on the spectral gap), it shows qualitatively that concentration for Markov chains is a generic phenomenon rather than something specific to matrices. It also allows one to instantly import the wealth of results regarding concentration for martingales in various Banach spaces (see e.g., [LT13]) to the random walk setting, albeit with suboptimal parameters. The simple proof of Theorem 7.1.6 is presented in Section 7.5.

### 7.2 Preliminaries

For an $n \in \mathbb{N}_{+}$, let $[n]$ denote the set $\{1,2, \cdots, n\}$. Let $\mathbf{i}$ denote $\sqrt{-1}$. For $z=a+\mathbf{i} b$, where $a, b \in \mathbb{R}$, we define the complex conjugate of $z$ to be $\bar{z}=a-\mathbf{i} b$ and $|z|=\sqrt{a^{2}+b^{2}}$. Then $|z|^{2}=z \bar{z}$. We define real part $\operatorname{Re}(z)=a$ and imaginary part $\operatorname{Im}(\mathrm{z})=b$. Then $\operatorname{Re}(z)=\frac{z+\bar{z}}{2}$ and $\operatorname{Im}(z)=\frac{z-\bar{z}}{2 \mathbf{i}}$.

We will be working with $D$-regular undirected graphs $G=(V, E)$. The number of vertices of the graph, $V$ will be denoted by $n$. $A$ will denote the adjacency matrix of the graph and $P=A / D$ will denote its normalized adjacency matrix. A regular graph $G$ will be called a $\lambda$-expander $(0<\lambda<1)$ if $\|P x\| \leq \lambda \cdot\|x\|$ for all vectors $x \in \mathbb{C}^{n}$ s.t. $\sum_{i=1}^{n} x_{i}=0$.

We will use $e_{i} \in \mathbb{C}^{d}$ to denote the standard basis vector with 1 in $i^{\text {th }}$ position and 0 everywhere else.

### 7.2.1 Linear Algebra

Matrices and Norms. For matrix $A$, we use $A^{\top}$ to denote the transpose of $A$, we use $\bar{A}$ to denote the entry-wise complex conjugate of $A$. For square matrix $A \in \mathbb{C}^{n \times n}$, we use $A^{*}$ to denote the conjugate transpose of matrix $A$. It is obvious that $A^{*}=\bar{A}^{\top}=\overline{A^{\top}}$. We say a complex square matrix $A$ is Hermitian, if $A=A^{*}$, unitary if $A A^{*}=A^{*} A=I$, and positive-semidefinite (psd) if $A=A^{*}$ and $x^{*} A x \geq 0$ for all $x \in \mathbb{C}^{n}$. We use $\succeq, \preceq$ to denote the semidefinite ordering, e.g. $A \succeq 0$ means that $A$ is psd.

For $p \in[1, \infty)$, define the Schatten $p$-norm of $A$ as

$$
\|A\|_{p}=\left(\sum_{n \geq 1} s_{n}^{p}(A)\right)^{1 / p}
$$

for $s_{1}(A) \geq s_{2}(A) \geq \cdots \geq s_{n}(A) \geq \cdots \geq 0$ the singular values of $A$, i.e., the eigenvalues of the Hermitian matrix $|A|=\sqrt{\left(A^{*} A\right)}$. Then $\|A\|_{p}^{p}=\operatorname{tr}\left[|T|^{p}\right]$.

For matrix $A \in \mathbb{C}^{n \times n}$, we define $\|A\|$ to be the spectral norm of $A$, i.e.,

$$
\|A\|=\max _{\|x\|_{2}=1, x \in \mathbb{C}^{n}} x^{*} A x
$$

Tensor Products. Given two vectors $v \in \mathbb{C}^{d_{1}}$ and $w \in \mathbb{C}^{d_{2}}$, their tensor product $v \otimes w \in$ $\mathbb{C}^{d_{1} d_{2}}$ is the vector whose $(i, j)^{\text {th }}$ entry is $v(i) w(j)$ (for concreteness, assume the entries are in lexicographic order). Given two matrices $A_{1} \in \mathbb{C}^{d_{1} \times d_{1}}$ and $A_{2} \in \mathbb{C}^{d_{2} \times d_{2}}$, their tensor product $A_{1} \otimes A_{2} \in \mathbb{C}^{d_{1} d_{2} \times d_{1} d_{2}}$ is the matrix whose $((i, k),(j, l))^{\text {th }}$ entry is $A_{1}(i, j) A_{2}(k, l)$. It is easy to see that

$$
(A \otimes B)(v \otimes w)=A v \otimes B w
$$

and

$$
(A \otimes B)(C \otimes D)=A B \otimes C D
$$

For a matrix $X \in \mathbb{C}^{d \times d}, \operatorname{vec}(X) \in \mathbb{C}^{d^{2}}$ will denote the vectorized version of the matrix $X$. That is

$$
\operatorname{vec}(X)=\sum_{i, j=1}^{d} X(i, j) e_{i} \otimes e_{j}
$$

We have the following relationship between matrix multiplication and the tensor product:

$$
\operatorname{vec}(A X B)=\left(A \otimes B^{\top}\right) \operatorname{vec}(X)
$$

Exponential and Logarithm. All logarithms will be taken with the base $e$, and $\exp (x)$ will denote $e^{x}$. The matrix exponential of a complex matrix $A \in \mathbb{C}^{d \times d}$ is defined by the

Taylor expansion:

$$
\exp (A)=\sum_{j=0}^{\infty} \frac{A^{j}}{j!}
$$

which converges for all matrices $A$. We will use the fact that

$$
\exp (A) \otimes \exp (B)=\exp (A \otimes I+I \otimes B)
$$

which may be checked by expanding both sides and comparing terms.
The matrix logarithm of a positive definite matrix $A=U D U^{*}$ with $D$ diagonal and positive is defined by

$$
\log (A):=U \log (D) U^{*}
$$

where the logarithm of $D$ is taken entrywise. For such matrices we have

$$
\log (\exp (A))=\exp (\log (A))=A
$$

For positive definite $A$ and complex $z$, we define $A^{z}:=\exp (z \log (A))$.

Polar Decomposition. The polar decomposition of a square complex matrix $A$ is a matrix decomposition of the form

$$
A=U V
$$

where $U$ is a unitary matrix and $V$ is a psd matrix. The polar decomposition separates matrix $A$ into a component that stretches the space along a set of orthogonal axes, represented by $V$, and a rotation (with possible reflection) represented by $U$. The decomposition of the complex conjugate of matrix $A$ can written as

$$
\bar{A}=\overline{U V}
$$

The decomposition of the conjugate transpose of matrix $A$ can be written as

$$
A^{*}=V^{*} U^{*}
$$

The following simple proposition will be useful in our proofs.

Proposition 7.2.1. Let $A$ and $B$ be Hermitian psd matrices. Then

$$
\operatorname{tr}[A B] \leq\|A\| \cdot \operatorname{tr}[B]
$$

Proof. Let $B=\sum_{j=1}^{n} \sigma_{j} v_{j} v_{j}^{\dagger}$ be the eigenvalue decomposition of $B$. Then

$$
\begin{aligned}
\operatorname{tr}[A B] & =\sum_{j=1}^{n} \sigma_{j} \operatorname{tr}\left[A v_{j} v_{j}^{\dagger}\right] \\
& =\sum_{j=1}^{n} \sigma_{j} v_{j}^{\dagger} A v_{j} \\
& \leq \sum_{j=1}^{n} \sigma_{j} \cdot\|A\| \\
& =\|A\| \cdot \operatorname{tr}[B] .
\end{aligned}
$$

### 7.2.2 Complex analysis

A function $f: U \rightarrow \mathbb{C}$ on a domain $U \subseteq \mathbb{C}$ is holomorphic if it has a complex derivative in a neighborhood of every point $z \in U$. The existence of a complex derivative in a neighborhood is a very strong condition, for it implies that any holomorphic function is actually infinitely differentiable and equal to its own Taylor series (i.e., analytic) at every point in $U$. A biholomorphic function is a bijective holomorphic function whose inverse is
also holomorphic. It follows from the definition that sums, products, and compositions of holomorphic functions are holomorphic. We will also talk about matrix-valued holomorphic functions $f: U \rightarrow \mathbb{C}^{d \times d}$, which just means that every entry is holomorphic.

The main property that we will use is that the value of a holomorphic function at a point $z \in U$ can be related to values that it takes on the boundary of $U$, in the following way. A function $f: U \rightarrow \mathbb{R} \cup\{-\infty\}$ is called subharmonic if it is upper semicontinuous and

$$
f(z) \leq \frac{1}{2 \pi} \int_{0}^{2 \pi} f\left(z+r e^{i \theta}\right) d \theta
$$

for all $z \in U$ and $r>0$ such that the closed disk $D(z, r)$ is contained in $U$, and all of the above integrals converge.

We will make frequent use of the following standard fact.

Proposition 7.2.2. If $f$ is analytic on a domain $U \subset \mathbb{C}$ then $\log |f(z)|$ is subharmonic on $U$.

Our main tool will be the Poisson Integral Formula for subharmonic functions.

Lemma 7.2.3 (Poisson integral formula on unit disk [Gao08, Eq 1.3.35]). For any subharmonic function $U$ defined on the unit disk $\{z \in \mathbb{C}:|z| \leq 1\}$, we have that

$$
U(z) \leq \frac{1}{2 \pi} \int_{-\pi}^{\pi} U\left(e^{\mathbf{i} \varphi}\right) \frac{1-|z|^{2}}{\left|e^{\mathbf{i} \varphi}-z\right|^{2}} \mathrm{~d} \varphi, \quad \forall|z|<1
$$

### 7.3 New Golden-Thompson inequality

We begin by giving an outline of the proof of Theorem 7.1.5. The proof of Theorem 7.1.4 in [SBT17] relies on the multivariate Lie-Trotter product formula (e.g. see [Bha97]), which states that:

$$
\exp \left(\sum_{j=1}^{k} L_{j}\right)=\lim _{\theta \rightarrow 0^{+}}\left(\prod_{j=1}^{k} \exp \left(\theta L_{j}\right)\right)^{\frac{1}{\theta}}
$$

For Hermitian $L_{j}$, a judicious application of the above allows one to rewrite the trace of the exponential as a limit of Schatten norms:

$$
\log \left(\operatorname{tr}\left[\exp \left(\sum_{j=1}^{k} L_{j}\right)\right]\right)=\lim _{\theta \rightarrow 0^{+}} \frac{2}{\theta} \log \left\|\prod_{j=1}^{k} \exp \left(\frac{\theta}{2} L_{j}\right)\right\|_{2 / \theta}
$$

Thus, understanding the matrix exponential of a sum is the same as understanding the behavior of a certain norm of the product as $\theta \rightarrow 0$. The idea of [SBT17] is to use complex interpolation, along the lines of the Stein-Hirschman theorem in complex analysis: for every fixed real $\theta$ near zero, find a complex function $F_{\theta}(z)$ that agrees with the right hand side at $z=\theta$ and is holomorphic on the strip $\{0 \leq \Re(z) \leq 1\}$. Since the value of a holomorphic function at any point can be related to an integral of its values on the boundary, this allows one to relate $F_{\theta}(\theta)$ to its integrals on $\{\Re(z)=0\}$ and $\{\Re(z)=1\}$, which are easy to understand. Taking the limit in $\theta$ yields Theorem 7.1.4.

To avoid integration on the whole vertical line $\{1+\mathbf{i} b: b \in \mathbb{R}\}$, we observe that the above strategy only relies on the fact that $\theta$ is enclosed by the two vertical lines $\{\mathbf{i} b: b \in \mathbb{R}\}$ and $\{1+\mathbf{i} b: b \in \mathbb{R}\}$, and we could have used any other region enclosing a neighborhood of real positive $\theta$ near zero, provided we can define the required holomorphic functions $F_{\theta}$. We choose the half-circle (which is easy to work with because the Riemann map to the unit


Figure 7.1: The function $h(z)=-\frac{1+z}{1-z}+\sqrt{\left(\frac{1+z}{1-z}\right)^{2}+1}$ maps the unit disk $\{z \in \mathbb{C}:|z| \leq 1\}$ to the half disk $\{z \in \mathbb{C}:|z| \leq 1$ and $\operatorname{Re}(z) \geq 0\}$.
disk is explicit) and use it to derive a variant of the Riesz-Thorin theorem (Theorem 7.3.3), from which our new multi-matrix Golden-Thompson inequality follows by mimicking the remainder of the proof of 7.1.4 given in [SBT17].

### 7.3.1 Complex estimate on the half disk

In general, we can upper bound the value of any subharmonic function on a simply connected domain by mapping the domain to the unit disk via Riemann mapping theorem and applying the Poisson integral formula. In this section, we will give such estimate on a unit half disk.

The following lemma follows from the biholomorphic map from unit disk onto the half disk defined in Figure 7.1.

Lemma 7.3.1 (Poisson Integral Formula on the Half-Disk). For any analytic function $F$ on
the half disk $\{z \in \mathbb{C}:|z| \leq 1$ and $\operatorname{Re}(z) \geq 0\}$, we have that

$$
\begin{equation*}
\log |F(x)| \leq \frac{1}{2 \pi} \int_{-\pi}^{\pi} \log \left|F \circ h\left(e^{\mathrm{i} \varphi}\right)\right| \frac{1-\rho^{2}}{1-2 \rho \cos (\varphi)+\rho^{2}} \mathrm{~d} \varphi \tag{7.6}
\end{equation*}
$$

for any $0 \leq x \leq 1$ where

$$
h(z)=-\frac{1+z}{1-z}+\sqrt{\left(\frac{1+z}{1-z}\right)^{2}+1} \quad \text { and } \quad \rho=\frac{x^{2}+2 x-1}{x^{2}-2 x-1}
$$

Proof. Note that the function $h(z)$ is a biholomorphic map from the unit disk $\{z \in \mathbb{C}:|z| \leq$ $1\}$ to the half disk $\{z \in \mathbb{C}:|z| \leq 1$ and $\operatorname{Re}(z) \geq 0\}$. (We provide a proof for completeness, see Lemma 7.6.1 ${ }^{2}$ )

Since $F$ and $h$ are holomorphic, $\log |F \circ h(z)|$ is subharmonic. Therefore, we can apply the Poisson integral formula for subharmonic functions (Lemma 7.2.3) and get

$$
\begin{aligned}
& \log |F \circ h(z)| \\
\leq & \frac{1}{2 \pi} \int_{-\pi}^{\pi} \log \left|F \circ h\left(e^{\mathbf{i} \varphi}\right)\right| \frac{1-|z|^{2}}{\left|e^{\mathbf{i} \varphi}-z\right|^{2}} \mathrm{~d} \varphi \\
= & \frac{1}{2 \pi} \int_{-\pi}^{\pi} \log \left|F \circ h\left(e^{\mathbf{i} \varphi}\right)\right| \frac{1-\rho^{2}}{1-2 \rho \cos (\theta-\varphi)+\rho^{2}} \mathrm{~d} \varphi
\end{aligned}
$$

for $z=\rho e^{\mathrm{i} \theta}$.
Setting $x=h(z)$, we obtain that

$$
z=\frac{x^{2}+2 x-1}{x^{2}-2 x-1} .
$$

[^48]Therefore, we have that

$$
\theta=0 \text { and } \rho=\frac{x^{2}+2 x-1}{x^{2}-2 x-1} .
$$

This gives the desired result.

The following lemma allows us to conveniently study the behavior of certain analytic functions near zero. The idea is that when $|F(z)|$ is at most 1 on the imaginary axis, $\log |F(\theta)|$ should be close to 0 for small $\theta$, and the value of $\log |F(\theta)| / \theta$ can be upper bounded by a suitable average of the values on the boundary of the half disk.

Lemma 7.3.2. Given any analytic function $F$ on the half disk $\{z \in \mathbb{C}:|z| \leq 1$ and $\operatorname{Re}(z) \geq$ $0\}$. Suppose that $|F(\mathbf{i} y)| \leq 1$ for all $y \in[-1,1]$. Then, for any $0 \leq \theta \leq 1 / 4$, we have

$$
\log |F(\theta)| \leq\left(\frac{4 \theta}{\pi}+O\left(\theta^{2}\right)\right) \int_{-\pi / 2}^{\pi / 2} \log \left|F\left(e^{\mathbf{i} \phi}\right)\right| \mathrm{d} \mu_{\theta}(\phi)
$$

where $\mu_{\theta}$ is some probability distribution on $\left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$ depending only on $\theta$, and $\mu_{\theta} \rightarrow$ some probability distribution as $\theta \rightarrow 0^{+}$.

Proof. Since $\log |F(\mathbf{i} y)| \leq 0$ for all $y \in[-1,1]$ and since $h\left(e^{\mathbf{i} \varphi}\right)$ is imaginary with modulus at most 1 whenever $|\varphi| \leq \pi / 2$, we can ignore these $\varphi$ in the integral (7.6), namely,

$$
\begin{align*}
& \log |F(x)| \\
\leq & \frac{1}{2 \pi} \int_{-\pi}^{\pi} \log \left|F \circ h\left(e^{\mathbf{i} \varphi}\right)\right| \frac{1-\rho^{2}}{1-2 \rho \cos (\varphi)+\rho^{2}} \mathrm{~d} \varphi \\
\leq & \frac{1}{2 \pi} \int_{\pi / 2 \leq|\varphi| \leq \pi} \log \left|F \circ h\left(e^{\mathbf{i} \varphi}\right)\right| \frac{1-\rho^{2}}{1-2 \rho \cos (\varphi)+\rho^{2}} \mathrm{~d} \varphi . \tag{7.7}
\end{align*}
$$

To bound the right hand side, for $\pi / 2 \leq|\varphi| \leq \pi$ and $0 \leq \rho \leq 1$, we can prove the following statement using elementary calculations (see Lemma 7.6.4),

$$
\frac{1-\rho^{2}}{1-2 \rho \cos (\varphi)+\rho^{2}}=\frac{1-\rho}{1-\cos (\varphi)} \pm O(1-\rho)^{2} .
$$

Note that $\rho=\frac{\theta^{2}+2 \theta-1}{\theta^{2}-2 \theta-1} \geq 1-4 \theta$ for all $0 \leq \theta \leq 1 / 4$. Therefore, we have that $0 \leq \rho \leq 1$ and

$$
\frac{1-\rho^{2}}{1-2 \rho \cos (\varphi)+\rho^{2}} \leq \frac{4 \theta}{1-\cos (\varphi)}+O\left(\theta^{2}\right)
$$

Putting this inequality into (7.7), we have that

$$
\begin{aligned}
& \log |F(x)| \\
\leq & \frac{1}{2 \pi} \int_{\pi / 2 \leq|\varphi| \leq \pi}\left(\frac{4 \theta}{1-\cos (\varphi)}+O\left(\theta^{2}\right)\right) \log \left|F \circ h\left(e^{\mathrm{i} \varphi}\right)\right| \mathrm{d} \varphi .
\end{aligned}
$$

We now observe that

$$
\begin{aligned}
& \int_{\pi / 2 \leq|\varphi| \leq \pi} \frac{1}{1-\cos (\varphi)} \mathrm{d} \varphi \\
= & \int_{-\pi}^{-\pi / 2} \frac{1}{1-\cos (\varphi)} \mathrm{d} \varphi+\int_{\pi / 2}^{\pi} \frac{1}{1-\cos (\varphi)} \mathrm{d} \varphi \\
= & 2 .
\end{aligned}
$$

Note that $h$ maps $e^{\mathrm{i} \varphi}$ for $\pi / 2 \leq|\varphi| \leq \pi$ to the boundary of the half disk $([-\pi / 2, \pi / 2])$, and let $\mu_{\theta}$ be some probability distribution on $\left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$ depending only on $\theta$. Then we can have

$$
\log |F(\theta)| \leq\left(\frac{4 \theta}{\pi}+O\left(\theta^{2}\right)\right) \int_{-\pi / 2}^{\pi / 2} \log \left|F\left(e^{\mathbf{i} \phi}\right)\right| \mathrm{d} \mu_{\theta}(\phi)
$$

Note that $\mu_{\theta} \rightarrow$ some probability distribution as $\theta \rightarrow 0^{+}$.

### 7.3.2 Bounded Multimatrix Golden-Thompson type inequality

Plugging our new complex estimate into the proof of Theorem 3.1 in [SBT17], we obtain the following Riesz-Thorin-type inequality. We give a complete proof for completeness.

Theorem 7.3.3 (Riesz-Thorin-type inequality). Let $S=\{z \in \mathbb{C}:|z| \leq 1$ and $\operatorname{Re}(z) \geq 0\}$ and let $G$ be a holomorphic map from $S$ to square matrices. Let $p_{0} \geq p_{1} \in[1, \infty]$, for $\theta \in(0,1)$, define $p_{\theta}$ by

$$
\frac{1}{p_{\theta}}=\frac{1-\theta}{p_{0}}+\frac{\theta}{p_{1}}
$$

If $z \rightarrow\|G(z)\|_{p_{\operatorname{Re}(z)}}$ is uniformly bounded on $S$ and $\|G(\mathbf{i} t)\|_{p_{0}} \leq 1$, then for any $0 \leq \theta \leq 1 / 4$,

$$
\log \|G(\theta)\|_{p_{\theta}} \leq\left(\frac{4 \theta}{\pi}+O\left(\theta^{2}\right)\right) \int_{-\pi / 2}^{\pi / 2} \log \left\|G\left(e^{\mathbf{i} \phi}\right)\right\|_{p_{1}} \mathrm{~d} \mu_{\theta}(\phi)
$$

where $\mu_{\theta}$ is some probability distribution on $\left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$ depending only on $\theta$.

Proof. To apply Lemma 7.3.2, we want to define a holomorphic function $F(z)$ such that

$$
|F(\mathbf{i} t)| \leq 1,|F(z)| \leq\|G(z)\|_{p_{1}} \text { and } F(\theta)=\|G(\theta)\|_{p_{\theta}}
$$

Now, we describe how to define such $F(z)$. For $x \in[0,1]$, define $q_{x}$ as the Hölder conjugate of $p_{x}$ such that $p_{x}^{-1}+q_{x}^{-1}=1$. Hence, using the definition of $p_{x}$ in the statement, we have

$$
\frac{1}{q_{x}}=\frac{1-x}{q_{0}}+\frac{x}{q_{1}} .
$$

Now for our fixed $\theta \in(0,1)$, let $G(\theta)=U V$ be the polar decomposition of $G(\theta)$, where $V$ is positive definite since $G(\theta)$ is always invertible, and $U$ is unitary. Finally, we define $X(z)$
and $F(z)$ by

$$
\begin{aligned}
X(z)^{*} & =(V / c)^{p_{\theta}\left(\frac{1-z}{q_{0}}+\frac{z}{q_{1}}\right)} U^{*}, \text { where } c=\|V\|_{p_{\theta}}=\|G(\theta)\|_{p_{\theta}} \\
F(z) & =\operatorname{tr}\left[X(z)^{*} G(z)\right] .
\end{aligned}
$$

Note that $F(z)$ is holomorphic. Due to the renormalization $c$, we can show $\| X(x+$ $\mathbf{i} y)) \|_{q_{x}}^{q_{x}}=1$ for all $x \in[0,1]:$

$$
\begin{aligned}
\|X(x+\mathbf{i} y)\|_{q_{x}}^{q_{x}} & =\operatorname{tr}\left[\sqrt{X^{*}(x+\mathbf{i} y) X(x+\mathbf{i} y)^{q_{x}}}\right] \\
& =\operatorname{tr}\left[(V / c)^{q_{x} p_{\theta}\left(\frac{1-x}{q_{0}}+\frac{x}{q_{1}}\right)}\right] \\
& =\operatorname{tr}\left[(V / c)^{p_{\theta}}\right] \\
& =1
\end{aligned}
$$

where the first step follows by definition of $\|\cdot\|_{p}$, the second step follows by $U^{*} U=I$, the third step follows by $\frac{1}{q_{x}}=\frac{1-x}{q_{0}}+\frac{x}{q_{1}}$, and the last step follows by $c=\|V\|_{p_{\theta}}$.

Therefore, $F(z)$ is bounded on $S$ as follows

$$
|F(x+\mathbf{i} y)| \leq\|X(x+\mathbf{i} y)\|_{q_{x}} \cdot\|G(x+\mathbf{i} y)\|_{p_{x}} \leq\|G(x+\mathbf{i} y)\|_{p_{x}}
$$

Using this, it is obvious that

$$
|F(\mathbf{i} t)| \leq\|G(\mathbf{i} t)\|_{p_{0}} \leq 1, \text { and }|F(z)| \leq\|G(z)\|_{p_{\operatorname{Re}(z)}} \leq\|G(z)\|_{p_{1}}
$$

for all $z \in S$ where we used that $p_{0} \geq p_{\operatorname{Re}(z)} \geq p_{1}$.

Finally, we verify that $F(\theta)=\|G(\theta)\|_{p_{\theta}}$ :

$$
\begin{aligned}
F(\theta) & =\operatorname{tr}\left[X(\theta)^{*} G(\theta)\right] \\
& =\operatorname{tr}\left[(V / c)^{p_{\theta}\left(\frac{1-\theta}{q_{0}}+\frac{\theta}{q_{1}}\right)} U^{*} \cdot U V\right] \\
& =\operatorname{tr}\left[(V / c)^{p_{\theta}\left(\frac{1-\theta}{q_{0}}+\frac{\theta}{q_{1}}\right)} V\right] \\
& =\operatorname{tr}\left[c^{-p_{\theta} / q_{\theta}} V^{1+p_{\theta} / q_{\theta}}\right] \\
& =\operatorname{tr}\left[c^{1-p_{\theta}} V^{p_{\theta}}\right] \\
& =c^{1-p_{\theta}} c^{p_{\theta}} \\
& =\|G(\theta)\|_{p_{\theta}},
\end{aligned}
$$

where the first step follows by definition of $X$ and $G$, the second step follows by $U^{*} U=I$, the third step follows by $\frac{1}{q_{\theta}}=\frac{1-\theta}{q_{0}}+\frac{\theta}{q_{1}}$, the fourth step follows by $p_{\theta} / q_{\theta}=p_{\theta}\left(1-1 / p_{\theta}\right)=p_{\theta}-1$, the fifth step follows by $\left(\operatorname{tr}\left[V^{p_{\theta}}\right]\right)^{1 / p_{\theta}}=c$, and the last step follows by $c=\|G(\theta)\|_{p_{\theta}}$.

Hence, the statement follows from Lemma 7.3.2.

Now, we are ready to prove our variant of multimatrix Golden-Thompson inequality. It follows from plugging in Theorem 7.3.3 into the proof of Theorem 3.5 in [SBT17]. We give a complete proof for completeness.

Theorem 7.3.4 (Multimatrix Golden-Thompson inequality). For any $k$ Hermitian matrices $H_{1}, \cdots, H_{k}$, we have:

$$
\log \left(\operatorname{tr}\left[\exp \left(\sum_{j=1}^{k} H_{j}\right)\right]\right) \leq \frac{4}{\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \log \left(\operatorname{tr}\left[\prod_{j=1}^{k} \exp \left(\frac{e^{\mathbf{i} \phi}}{2} H_{j}\right) \prod_{j=k}^{1} \exp \left(\frac{e^{-\mathbf{i} \phi}}{2} H_{j}\right)\right]\right) \mathrm{d} \mu(\phi)
$$

where $\mu$ is some probability distribution on $\left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$.

Proof. Define

$$
G(z)=\prod_{j=1}^{k} \exp \left(\frac{z}{2} H_{j}\right)
$$

Note that $\|G(\mathbf{i} y)\|_{\infty}=1$ for all $y \in \mathbb{R}$. We now have:

$$
\begin{aligned}
& \log \left(\operatorname{tr}\left[\exp \left(\sum_{j=1}^{k} H_{j}\right)\right]\right) \\
= & \log \left(\operatorname{tr}\left[\exp \left(\sum_{j=1}^{k} H_{j} / 2+\sum_{j=k}^{1} H_{j} / 2\right)\right]\right) \\
= & \log \left(\operatorname{tr}\left[\lim _{\theta \rightarrow 0^{+}}\left(G(\theta) G(\theta)^{*}\right)^{1 / \theta}\right]\right) \\
= & \lim _{\theta \rightarrow 0^{+}} \frac{2}{\theta} \log \left(\operatorname{tr}\left[\left(G(\theta) G(\theta)^{*}\right)^{1 / \theta}\right]^{\theta / 2}\right) \\
= & \lim _{\theta \rightarrow 0^{+}} \frac{2}{\theta} \log \|G(\theta)\|_{2 / \theta} \\
\leq & \lim _{\theta \rightarrow 0^{+}} 2\left(\frac{4}{\pi}+O(\theta)\right) \int_{-\pi / 2}^{\pi / 2} \log \left\|G\left(e^{\mathbf{i} \phi}\right)\right\|_{2} \mathrm{~d} \mu_{\theta}(\phi) \\
= & \lim _{\theta \rightarrow 0^{+}}\left(\frac{4}{\pi}+O(\theta)\right) \int_{-\pi / 2}^{\pi / 2} \log \left(\operatorname{tr}\left[G\left(e^{\mathbf{i} \phi}\right) G\left(e^{\mathbf{i} \phi}\right)^{*}\right]\right) \mathrm{d} \mu_{\theta}(\phi)
\end{aligned}
$$

where the first step follows from by the Lie-Trotter formula, the second step follows from $H_{j}$ are Hermitian, the third step follows from continuity of $\log$ away from 0 , last inequality follows from Theorem 7.3.3 with $p_{0}=\infty$ and $p_{1}=2$.

When $\theta \rightarrow 0^{+}, \mu_{\theta}(\phi)$ converges to some probability distribution $\mu$. This completes the proof.

Remark 7.3.1. We suspect the constant $4 / \pi$ is tight and that any constant larger than one is unavoidable if we consider the maximum over a bounded domain inside the strip $\{z \in \mathbb{C}: 0 \leq \operatorname{Re}(z) \leq 1\}$.

### 7.4 Proof of Theorem 7.1.2

In this section we present the proof of Theorem 7.1.2. We restate it here (with explicit constants) for convenience.

Theorem 7.4.1. Let $G$ be a regular $\lambda$-expander on $V$ and let $f$ be a function $f: V \rightarrow \mathbb{C}^{d \times d}$ s.t.

1. For each $v \in V, f(v)$ is Hermitian and $\|f(v)\| \leq 1$.
2. $\sum_{v \in V} f(v)=0$.

If $v_{1}, \ldots, v_{k}$ is a stationary random walk on $G$, and $\epsilon \in(0,1)$,

$$
\begin{aligned}
& \operatorname{Pr}\left[\lambda_{\max }\left(\sum_{j=1}^{k} f\left(v_{j}\right)\right) \geq+k \epsilon\right] \leq d^{2-\pi / 4} \cdot \exp \left(-\epsilon^{2}(1-\lambda) k / 80\right), \\
& \operatorname{Pr}\left[\lambda_{\min }\left(\sum_{i=1}^{k} f\left(v_{j}\right)\right) \leq-k \epsilon\right] \leq d^{2-\pi / 4} \cdot \exp \left(-\epsilon^{2}(1-\lambda) k / 80\right) .
\end{aligned}
$$

Remark 7.4.1. Depsite the exponent of $d$ is different from Theorem 7.1.2, we note that since the left hand side (the probability) is at most 1, one can prove the same statement with any positive exponent by changing the constant 80 .

Proof. Due to symmetry, it suffices to prove just one of the statements. Let $t>0$ be a
parameter to be chosen later. Then

$$
\begin{align*}
& \operatorname{Pr}\left[\lambda_{\max }\left(\sum_{j=1}^{k} f\left(v_{j}\right)\right) \geq k \epsilon\right] \\
\leq & \operatorname{Pr}\left[\operatorname{tr}\left[\exp \left(t \sum_{j=1}^{k} f\left(v_{j}\right)\right)\right] \geq \exp (t k \epsilon)\right] \\
\leq & \frac{\mathbb{E}\left[\operatorname{tr}\left[\exp \left(t \sum_{j=1}^{k} f\left(v_{j}\right)\right)\right]\right]}{\exp (t k \epsilon)} \tag{7.8}
\end{align*}
$$

The second inequality follows from Markov's inequality.
Now the question is how to bound
$\mathbb{E}_{v_{1}, \cdots, v_{k}}\left[\operatorname{tr}\left[\exp \left(t \sum_{j=1}^{k} f\left(v_{j}\right)\right)\right]\right]$. Using Theorem 7.3.4 and note that $\mu(\phi)$ is a probability distribution on $\left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$, we have

$$
\begin{aligned}
& \log \left(\operatorname{tr}\left[\exp \left(t \sum_{j=1}^{k} f\left(v_{j}\right)\right)\right]\right) \\
\leq & \frac{4}{\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \log \operatorname{tr}\left[\prod_{j=1}^{k} \exp \left(\frac{e^{\mathbf{i} \phi}}{2} t f\left(v_{j}\right)\right) \prod_{j=k}^{1} \exp \left(\frac{e^{-\mathbf{i} \phi}}{2} t f\left(v_{j}\right)\right)\right] \mathrm{d} \mu(\phi) \\
\leq & \frac{4}{\pi} \log \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \operatorname{tr}\left[\prod_{j=1}^{k} \exp \left(\frac{e^{\mathbf{i} \phi}}{2} t f\left(v_{j}\right)\right) \prod_{j=k}^{1} \exp \left(\frac{e^{-\mathbf{i} \phi}}{2} t f\left(v_{j}\right)\right)\right] \mathrm{d} \mu(\phi),
\end{aligned}
$$

where the the second step follows by concavity of log function. This implies that

$$
\begin{align*}
& \operatorname{tr}\left[\exp \left(t \sum_{j=1}^{k} f\left(v_{j}\right)\right)\right] \\
\leq & \left(\int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \operatorname{tr}\left[\prod_{j=1}^{k} \exp \left(\frac{e^{\mathbf{i} \phi}}{2} t f\left(v_{j}\right)\right) \prod_{j=k}^{1} \exp \left(\frac{e^{-\mathbf{i} \phi}}{2} t f\left(v_{j}\right)\right)\right] \mathrm{d} \mu(\phi)\right)^{\frac{4}{\pi}} . \tag{7.9}
\end{align*}
$$

Note that $\|x\|_{p} \leq d^{1 / p-1}\|x\|_{1}$ for $p \in(0,1)$, choosing $p=\pi / 4$ we have

$$
\begin{equation*}
\left(\operatorname{tr}\left[\exp \left(\frac{\pi}{4} t \sum_{j=1}^{k} f\left(v_{j}\right)\right)\right]\right)^{\frac{4}{\pi}} \leq d^{4 / \pi-1} \operatorname{tr}\left[\exp \left(t \sum_{j=1}^{k} f\left(v_{j}\right)\right)\right] \tag{7.10}
\end{equation*}
$$

Combining Eq. (7.9) and Eq. (7.10), we have
$\operatorname{tr}\left[\exp \left(\frac{\pi}{4} t \sum_{j=1}^{k} f\left(v_{j}\right)\right)\right] \leq d^{1-\pi / 4} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \operatorname{tr}\left[\prod_{j=1}^{k} \exp \left(\frac{e^{\mathbf{i} \phi}}{2} t f\left(v_{j}\right)\right) \prod_{j=k}^{1} \exp \left(\frac{e^{-\mathbf{i} \phi}}{2} t f\left(v_{j}\right)\right)\right] \mathrm{d} \mu(\phi)$.

The core of the proof is the following bound on the moment generating function-like expression that appears above, thinking of $\mathbf{i} \phi$ as $\gamma+\mathbf{i} b$ with $\gamma^{2}+b^{2}=1$ :

Lemma 7.4.2. Let $G$ be a regular $\lambda$-expander on $V$, let $f$ be a function $f: V \rightarrow \mathbb{C}^{d \times d}$ and $\sum_{v \in V} f(v)=0$, let $v_{1}, \cdots, v_{k}$ be a stationary random walk on $G$, for any $t>0, \gamma \geq 0$, $b>0, t^{2}\left(\gamma^{2}+b^{2}\right) \leq 1$, and $t \gamma \leq \frac{1-\lambda}{4 \lambda}$ we have

$$
\begin{aligned}
& \mathbb{E}\left[\operatorname{tr}\left[\prod_{j=1}^{k} \exp \left(\frac{t f\left(v_{j}\right)(\gamma+\mathbf{i} b)}{2}\right) \prod_{j=k}^{1} \exp \left(\frac{t f\left(v_{j}\right)(\gamma-\mathbf{i} b)}{2}\right)\right]\right] \\
\leq & d \cdot \exp \left(k t^{2}\left(\gamma^{2}+b^{2}\right)\left(1+\frac{8}{1-\lambda}\right)\right)
\end{aligned}
$$

Assuming this lemma, we can easily complete the proof of the theorem as:

$$
\begin{align*}
& \underset{v_{1}, \cdots, v_{k}}{\mathbb{E}}\left[\operatorname{tr}\left[\exp \left(\frac{\pi}{4} t \sum_{j=1}^{k} f\left(v_{j}\right)\right)\right]\right] \\
\leq & d^{1-\pi / 4} \underset{v_{1}, \cdots, v_{k}}{\mathbb{E}}\left[\int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \operatorname{tr}\left[\prod_{j=1}^{k} \exp \left(\frac{e^{\mathbf{i} \phi}}{2} t f\left(v_{j}\right)\right) \prod_{j=k}^{1} \exp \left(\frac{e^{-\mathbf{i} \phi}}{2} t f\left(v_{j}\right)\right)\right] \mathrm{d} \mu(\phi)\right] \\
= & d^{1-\pi / 4} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \underset{v_{1}, \cdots, v_{k}}{\mathbb{E}}\left[\operatorname{tr}\left[\prod_{j=1}^{k} \exp \left(\frac{e^{\mathbf{i} \phi}}{2} t f\left(v_{j}\right)\right) \prod_{j=k}^{1} \exp \left(\frac{e^{-\mathbf{i} \phi}}{2} t f\left(v_{j}\right)\right)\right]\right] \mathrm{d} \mu(\phi) \\
\leq & d^{1-\pi / 4} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} d \exp \left(k t^{2}\left|e^{\mathbf{i} \phi}\right|^{2}\left(1+\frac{8}{1-\lambda}\right)\right) \mathrm{d} \mu(\phi) \\
= & d^{2-\pi / 4} \exp \left(k t^{2}\left(1+\frac{8}{1-\lambda}\right)\right) \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \mathrm{~d} \mu(\phi) \\
= & d^{2-\pi / 4} \exp \left(k t^{2}\left(1+\frac{8}{1-\lambda}\right)\right), \tag{7.12}
\end{align*}
$$

where the first step follows by the Equation (7.11), the second step follows by swapping $\mathbb{E}$ and $\int$, the third step follows by Lemma 7.4.2, the fourth step follows by $\left|e^{\mathbf{i} \phi}\right|=1$, and the last step follows by $\int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \mathrm{~d} \mu(\phi)=1$.

Finally, putting it all together,

$$
\begin{aligned}
& \operatorname{Pr}_{v_{1}, \cdots, v_{k}}\left[\lambda_{\max }\left(\sum_{j=1}^{k} f\left(v_{j}\right)\right) \geq k \epsilon\right] \\
\leq & d^{2-\pi / 4} \cdot \exp \left((4 / \pi)^{2} k t^{2} \frac{9}{1-\lambda}-k t \epsilon\right) \\
= & d^{2-\pi / 4} \cdot \exp \left((4 / \pi)^{2} k \epsilon^{2}(1-\lambda)^{2} \frac{1}{36^{2}} \frac{9}{1-\lambda}-k \frac{(1-\lambda) \epsilon}{36} \epsilon\right) \\
\leq & d^{2-\pi / 4} \cdot \exp \left(-k \epsilon^{2}(1-\lambda) / 72\right) .
\end{aligned}
$$

where the first step follows by Eq. (7.12) the second step follows by choosing $t=(1-$ $\lambda) \epsilon / 36$.

We now give the proof of Lemma 7.4.2

Proof of Lemma 7.4.2. We start by writing the expected trace expression in terms of the transition matrix of the random walk. This is an analogue of a step which is common to most of the expander chernoff bound proofs in the scalar case. Let $P$ be the normalized adjacency matrix of $G$ and let $\widetilde{P}=P \otimes I_{d^{2}}$. Let $E$ denote the $n d^{2} \times n d^{2}$ block diagonal matrix where the $v^{\text {th }}$ diagonal block is the matrix

$$
\begin{equation*}
M_{v}=\exp \left(\frac{t f(v)(\gamma+\mathbf{i} b)}{2}\right) \otimes \exp \left(\frac{t f(v)(\gamma-\mathbf{i} b)}{2}\right) \tag{7.13}
\end{equation*}
$$

Then $(E \widetilde{P})^{k}$ is an $n d^{2} \times n d^{2}$ block matrix whose $(u, v)^{\text {th }}\left(d^{2} \times d^{2}\right)$ block is given by the matrix

$$
\begin{equation*}
\sum_{v_{1}, \ldots, v_{k-1}} P_{u, v_{1}} \cdot\left(\prod_{j=1}^{k-2} P_{v_{j}, v_{j+1}}\right) \cdot P_{v_{k-1}, v} \cdot M_{u} \cdot\left(\prod_{j=1}^{k-1} M_{v_{j}}\right) . \tag{7.14}
\end{equation*}
$$

Let $z_{0} \in \mathbb{C}^{n d^{2}}$ be the vector $\frac{1}{\sqrt{n}} \otimes \operatorname{vec}\left(I_{d}\right)$. Here $\mathbf{1}$ is the all 1 's vector and $\operatorname{vec}\left(I_{d}\right)$ is the vector form of the identity matrix. Then, by applying

$$
\left\langle\operatorname{vec}\left(I_{d}\right), A_{1} \otimes A_{2} \operatorname{vec}\left(I_{d}\right)\right\rangle=\operatorname{tr}\left[A_{1} A_{2}^{\top}\right]
$$

it follows that for a stationary random walk $v_{1}, \ldots, v_{k}$ :

$$
\begin{aligned}
& \mathbb{E}\left[\operatorname{tr}\left[\prod_{j=1}^{k} \exp \left(\frac{t f\left(v_{j}\right)(\gamma+\mathbf{i} b)}{2}\right) \prod_{j=k}^{1} \exp \left(\frac{t f\left(v_{j}\right)(\gamma-\mathbf{i} b)}{2}\right)\right]\right] \\
= & \mathbb{E}\left[\left\langle\operatorname{vec}\left(I_{d}\right), \prod_{i=1}^{k} M_{v_{i}} \operatorname{vec}\left(I_{d}\right)\right\rangle\right] \\
= & \left\langle z_{0},(E \widetilde{P})^{k} z_{0}\right\rangle .
\end{aligned}
$$

Hence we can focus our attention on $\left\langle z_{0},(E \widetilde{P})^{k} z_{0}\right\rangle$.
Let 1 be the all 1 's vector. For a vector $z \in \mathbb{C}^{n d^{2}}$, let $z^{\| l}$ denote the component of $z$ which lies in the subspace spanned by the $d^{2}$ vectors $\mathbf{1} \otimes e_{i}, 1 \leq i \leq d^{2}$. Let $z^{\perp}$ denote the component in the orthogonal space. Letting $z_{j}=(E \widetilde{P})^{j} z_{0}$, we are interested in bounding

$$
\left\langle z_{0}, z_{k}\right\rangle=\left\langle z_{0}, z_{k}^{\|}\right\rangle \leq\left\|z_{0}\right\| \cdot\left\|z_{k}^{\|}\right\|=\sqrt{d} \cdot\left\|z_{k}^{\|}\right\| .
$$

The following lemma is the analogue of the main lemma in Healy's proof for the scalar valued expander Chernoff bound [Hea08]. Roughly speaking, it tracks how much a vector can move in and out of the subspace we are interested in as the operator $E \widetilde{P}$ is applied.

Lemma 7.4.3. Given four parameters $\lambda \in[0,1], \gamma \geq 0, \ell \geq 0$, and $t>0$. Let $G$ be a regular $\lambda$-expander on $V$. Suppose each vertex $v$ is assigned a matrix $H_{v} \in \mathbb{C}^{d^{2} \times d^{2}}$ s.t. $\left\|H_{v}\right\| \leq \ell$ and $\sum_{v} H_{v}=0$. Let $P$ be the normalized adjacency matrix of $G$ and let $\widetilde{P}=P \otimes I_{d^{2}}$. Let $E$ denote the $n d^{2} \times n d^{2}$ block diagonal matrix where the $v$-th diagonal block is the matrix $\exp \left(t H_{v}\right)$. Also suppose that $\|E\|=\max _{v \in V}\left\|\exp \left(t H_{v}\right)\right\| \leq \exp (\gamma t)$. Then for any $z \in \mathbb{C}^{n d^{2}}$, we have:

1. $\left\|\left(E \widetilde{P} z^{\|}\right)^{\|}\right\| \leq \alpha_{1}\left\|z^{\|}\right\| \quad$ where $\alpha_{1}=\exp (t \ell)-t \ell$,
2. $\left\|\left(E \widetilde{P} z^{\|}\right)^{\perp}\right\| \leq \alpha_{2}\left\|z^{\|}\right\| \quad$ where $\alpha_{2}=\exp (t \ell)-1$,
3. $\left\|\left(E \widetilde{P} z^{\perp}\right)^{\|}\right\| \leq \alpha_{3}\left\|z^{\perp}\right\| \quad$ where $\alpha_{3}=\lambda \cdot(\exp (t \ell)-1)$,
4. $\left\|\left(E \widetilde{P} z^{\perp}\right)^{\perp}\right\| \leq \alpha_{4}\left\|z^{\perp}\right\| \quad$ where $\alpha_{4}=\lambda \cdot \exp (t \gamma)$.

Proof of Lemma 7.4.3. Part 1.

Note that

$$
\left(E \widetilde{P} z^{\|}\right)^{\|}=\left(E z^{\|}\right)^{\|} .
$$

Let $\mathbf{1} \in \mathbb{R}^{n}$ denote all ones vector, suppose $z^{\|}=\mathbf{1} \otimes \mathbf{w}$ for some $\mathbf{w} \in \mathbb{C}^{d^{2}}$. Then $\left\|z^{\|}\right\|=$ $\sqrt{n} \cdot\|\mathbf{w}\|$ and

$$
\left(E z^{\|}\right)^{\|}=\mathbf{1} \otimes\left(\frac{1}{n} \sum_{v \in V} \exp \left(H_{v}\right) \mathbf{w}\right)
$$

We can upper bound $\left\|\frac{1}{n} \sum_{v \in V} \exp \left(t H_{v}\right)\right\|$ in the following way,

$$
\begin{aligned}
\left\|\frac{1}{n} \sum_{v \in V} \exp \left(t H_{v}\right)\right\| & =\left\|\frac{1}{n} \sum_{v \in V} \sum_{j=0}^{\infty} \frac{t^{j} H_{v}^{j}}{j!}\right\| \\
& =\left\|I+\frac{1}{n} \sum_{v \in V} \sum_{j=2}^{\infty} \frac{t^{j} H_{v}^{j}}{j!}\right\| \\
& \leq 1+\frac{1}{n} \sum_{v \in V} \sum_{j \geq 2} \frac{t^{j}}{j!}\left\|H_{v}\right\|^{j} \\
& =1+\sum_{j \geq 2} \frac{(t \ell)^{j}}{j!} \\
& =\exp (t \ell)-t \ell,
\end{aligned}
$$

where the first step follows by Taylor expansion, the second step follows by $\sum_{v \in V} H_{v}=0$, the third step follows by triangle inequality, the fourth step follows by $|V|=n$ and $\left\|H_{v}\right\| \leq \ell$, and last step follows by Taylor expansion.

Thus,

$$
\begin{aligned}
\left\|\left(E z^{\|}\right)^{\|}\right\| & =\sqrt{n}\left\|\frac{1}{n} \sum_{v \in V} \exp \left(t H_{v}\right) w\right\| \\
& \leq \sqrt{n}\|w\|(\exp (t \ell)-t \ell) \\
& =\left\|z^{\|}\right\|(\exp (t \ell)-t \ell)
\end{aligned}
$$

Part 2. Note that $\left(E \widetilde{P} z^{\|}\right)^{\perp}=\left(E z^{\|}\right)^{\perp}=\left((E-I) z^{\|}\right)^{\perp}$. and $\left(z^{\|}\right)^{\perp}=0$. We can upper bound $\left\|\left((E-I) z^{\|}\right)^{\perp}\right\|$ in the following way,

$$
\begin{aligned}
\left\|\left((E-I) z^{\|}\right)^{\perp}\right\| & \leq\left\|(E-I) z^{\|}\right\| \\
& \leq\|E-I\| \cdot\left\|z^{\|}\right\| \\
& =\max _{v \in V}\left\|\exp \left(t H_{v}\right)-I\right\| \cdot\left\|z^{\|}\right\| \\
& =\max _{v \in V}\left\|\sum_{j=1}^{\infty} \frac{t^{j}}{j!} H_{v}^{j}\right\| \cdot\left\|z^{\|}\right\| \\
& \leq\left(\sum_{j=1}^{\infty} \frac{t^{j} \ell^{j}}{j!}\right) \cdot\left\|z^{\|}\right\| \\
& =(\exp (t \ell)-1) \cdot\left\|z^{\|}\right\|
\end{aligned}
$$

where the second step follows by $\|A x\| \leq\|A\| \cdot\|x\|$, the third step follows by definition of $E$, the fourth step follows by Taylor expansion, the fifth step follows by triangle inequality and $\left\|H_{v}\right\| \leq \ell$, and the last step follows by Taylor expansion.

Part 3. Note that $\left(E \widetilde{P} z^{\perp}\right)^{\|}=\left((E-I) \widetilde{P} z^{\perp}\right)^{\|}$. This is because $\left(\widetilde{P} z^{\perp}\right)^{\|}=0$ since $\widetilde{P}$ preserves the property of being orthogonal to the space spanned by the vectors $\mathbf{1} \otimes e_{i}$ (these are the top eigenvectors of $\widetilde{P})$. Hence we can bound

$$
\begin{aligned}
\left\|\left((E-I) \widetilde{P} z^{\perp}\right)^{\|}\right\| & \leq\left\|(E-I) \widetilde{P} z^{\perp}\right\| \\
& \leq\|E-I\| \cdot\left\|\widetilde{P} z^{\perp}\right\| \\
& \leq(\exp (t \ell)-1) \cdot \lambda \cdot\left\|z^{\perp}\right\| .
\end{aligned}
$$

Third inequality follows from the fact that $\|E-I\| \leq \exp (t \ell)-1$ and that $G$ is a $\lambda$-expander.

Part 4. We can bound

$$
\left\|\left(E \widetilde{P} z^{\perp}\right)^{\perp}\right\| \leq\left\|E \widetilde{P} z^{\perp}\right\| \leq \exp (t \gamma) \cdot \lambda\left\|z^{\perp}\right\|
$$

where the second step follows by $\|E\| \leq \exp (\gamma t)$ and $G$ is a $\lambda$-expander.

We now use the above Lemma 7.4.3 to analyze the evolution of $z_{j}^{\|}$and $z_{j}^{\perp}$. Recall the definition of $H_{v}$,

$$
H_{v}=\frac{f(v)(\gamma+\mathbf{i} b)}{2} \otimes I_{d}+I_{d} \otimes \frac{f(v)(\gamma-\mathbf{i} b)}{2}
$$

which means

$$
\begin{aligned}
\exp \left(t H_{v}\right) & =\exp \left(\frac{t f(v)(\gamma+\mathbf{i} b)}{2} \otimes I_{d}+I_{d} \otimes \frac{t f(v)(\gamma-\mathbf{i} b)}{2}\right) \\
& =\exp \left(\frac{t f(v)(\gamma+\mathbf{i} b)}{2}\right) \otimes \exp \left(\frac{t f(v)(\gamma-\mathbf{i} b)}{2}\right) \\
& =M_{v}
\end{aligned}
$$

where the the first step follows by definition of $H_{v}$, the second step follows by $\exp \left(A \otimes I_{d}+\right.$ $\left.I_{d} \otimes B\right)=\exp (A) \otimes \exp (B)$, and the last step follows by Eq. (7.13).

We can upper bound $\left\|H_{v}\right\| \leq \sqrt{\gamma^{2}+b^{2}}$ and then set $\ell=\sqrt{\gamma^{2}+b^{2}}$. We can also upper bound $\left\|\exp \left(t H_{v}\right)\right\|$,

$$
\begin{aligned}
\left\|\exp \left(t H_{v}\right)\right\| & =\left\|\exp \left(t \cdot \operatorname{Re}\left(H_{v}\right)\right)\right\| \\
& =\left\|\exp \left(\gamma t\left(\frac{f(v)}{2} \otimes I_{d}+I_{d} \otimes \frac{f(v)}{2}\right)\right)\right\| \\
& \leq \exp (\gamma t)
\end{aligned}
$$

Note that $\sum_{v \in V} H_{v}=0$ since $\sum_{v \in V} f(v)=0$.

Claim 7.4.4. $\left\|z_{i}^{\perp}\right\| \leq \frac{\alpha_{2}}{1-\alpha_{4}} \max _{j<i}\left\|z_{j}^{\|}\right\|$.

Proof.

$$
\begin{aligned}
\left\|z_{i}^{\perp}\right\| & =\left\|\left(E \widetilde{P} z_{i-1}\right)^{\perp}\right\| \\
& \leq\left\|\left(E \widetilde{P} z_{i-1}^{\|}\right)^{\perp}\right\|+\left\|\left(E \widetilde{P} z_{i-1}^{\perp}\right)^{\perp}\right\| \\
& \leq \alpha_{2}\left\|z_{i-1}^{\|}\right\|+\alpha_{4}\left\|z_{i-1}^{\perp}\right\| \\
& \leq\left(\alpha_{2}+\alpha_{2} \alpha_{4}+\alpha_{2} \alpha_{4}+\cdots\right) \cdot \max _{j<i}\left\|z_{j}^{\|}\right\| \\
& \leq \frac{\alpha_{2}}{1-\alpha_{4}} \max _{j<i}\left\|z_{j}^{\|}\right\|,
\end{aligned}
$$

where the first step follows by definition of $z_{i}$, the second step follows by triangle inequality, the third step follows by part 2 and 4 of Lemma 7.4.3.

Claim 7.4.5. $\left\|z_{i}^{\|}\right\| \leq\left(\alpha_{1}+\frac{\alpha_{2} \alpha_{3}}{1-\alpha_{4}}\right) \max _{j<i}\left\|z_{j}^{\|}\right\|$.

Proof.

$$
\begin{aligned}
\left\|z_{i}^{\|}\right\| & =\left\|\left(E \widetilde{P} z_{i-1}\right)^{\|}\right\| \\
& \leq\left\|\left(E \widetilde{P} z_{i-1}^{\|}\right)^{\|}\right\|+\left\|\left(E \widetilde{P} z_{i-1}^{\perp}\right)^{\|}\right\| \\
& \leq \alpha_{1}\left\|z_{i-1}^{\|}\right\|+\alpha_{3}\left\|z_{i-1}^{\perp}\right\| \\
& \leq \alpha_{1}\left\|z_{i-1}^{\|}\right\|+\alpha_{3} \frac{\alpha_{2}}{1-\alpha_{4}} \max _{j<i-1}\left\|z_{j}^{\|}\right\| \\
& \leq\left(\alpha_{1}+\frac{\alpha_{2} \alpha_{3}}{1-\alpha_{4}}\right),
\end{aligned}
$$

where the first step follows by definition of $z_{i}$, the second step follows by triangle inequality, the third step follows by part 1 and 3 of Lemma 7.4.3, the fourth step follows by Claim 7.4.4.

Combining Claim 7.4.4 and Claim 7.4.5 gives

$$
\left\|z_{k}^{\|}\right\| \leq\left(\alpha_{1}+\frac{\alpha_{2} \alpha_{3}}{1-\alpha_{4}}\right)^{k}\left\|z_{0}^{\|}\right\|=\sqrt{d} \cdot\left(\alpha_{1}+\frac{\alpha_{2} \alpha_{3}}{1-\alpha_{4}}\right)^{k}
$$

which implies that

$$
\left\langle z_{0},(E \widetilde{P})^{k} z_{0}\right\rangle \leq d \cdot\left(\alpha_{1}+\frac{\alpha_{2} \alpha_{3}}{1-\alpha_{4}}\right)^{k}
$$

Now the question is how to bound $\left(\alpha_{1}+\frac{\alpha_{2} \alpha_{3}}{1-\alpha_{4}}\right)^{k}$.
We can upper bound $\alpha_{1}, \alpha_{2} \alpha_{3}$ and $\alpha_{4}$ in the following sense,

$$
\alpha_{1}=\exp (t \ell)-t \ell \leq 1+t^{2} \ell^{2}=1+t^{2}\left(\gamma^{2}+b^{2}\right),
$$

and

$$
\alpha_{2} \alpha_{3}=\lambda(\exp (t \ell)-1)^{2} \leq \lambda(2 t \ell)^{2}=4 \lambda t^{2}\left(\gamma^{2}+b^{2}\right)
$$

where the second step follows by $t \ell<1$ (because $\exp (x) \leq 1+2 x, \forall x \in[0,1])$,

$$
\alpha_{4}=\lambda \cdot \exp (t \gamma) \leq \lambda(1+2 t \gamma) \leq \frac{1}{2}+\frac{1}{2} \lambda,
$$

where the second step follows by $t \gamma<1$, and the third step follows by $t \gamma \leq(1-\lambda) / 4 \lambda$.
Thus,

$$
\begin{aligned}
\left(\alpha_{1}+\frac{\alpha_{2} \cdot \alpha_{3}}{1-\alpha_{4}}\right)^{k} & \leq\left(1+t^{2}\left(\gamma^{2}+b^{2}\right)+\frac{4 \lambda t^{2}\left(\gamma^{2}+b^{2}\right)}{\frac{1}{2}-\frac{1}{2} \lambda}\right)^{k} \\
& \leq \exp \left(k t^{2}\left(\gamma^{2}+b^{2}\right)\left(1+\frac{8}{1-\lambda}\right)\right)
\end{aligned}
$$

Remark 7.4.2. As is the case with Healy's proof [Hea08], our proof also works for the case when there are different mean zero functions $f_{1}, \ldots, f_{k}$ for the different steps of the walk and also when there are $k \lambda$-expanders $G_{1}, \ldots, G_{k}$ and the $j^{\text {th }}$ step of the walk is taken according to $G_{j}$.

Remark 7.4.3. We suspect that with appropriate modifications, our proof should generalize to random walks on irregular undirected graphs (or reversible Markov chains) as was done for Healy's proof in [CLLM12].

Remark 7.4.4. Although we have stated the theorem for Hermitian matrices, the same result can be obtained for general matrices by a standard dilation trick, namely replacing every $d \times d$ matrix $M$ that appears with the $2 d \times 2 d$ Hermitian matrix

$$
\left[\begin{array}{cc}
0 & M \\
M^{*} & 0
\end{array}\right]
$$

whose norm is always within a factor of two of $M$.

### 7.5 Proof of Theorem 7.1.6

Proof. Observe that for every $i=2, \ldots, k$ we have

$$
\mathbb{E}\left[f\left(v_{i}\right) \mid v_{i-1}\right]=\sum_{u \sim v_{i-1}} P\left(v_{i-1}, u\right) f(u)=\operatorname{Pf}\left(v_{i-1}\right),
$$

whence the random vectors

$$
Y_{i}^{(1)}:=f\left(v_{i}\right)-P f\left(v_{i-1}\right)
$$

satisfy

$$
\mathbb{E}\left[Y_{i}^{(1)} \mid v_{1}, \ldots, v_{i-1}\right]=0
$$

and thus form a martingale difference sequence with respect to the filtration generated by initial segments of $v_{1}, \ldots, v_{k}$. Denoting $Y_{1}^{(1)}:=f\left(v_{1}\right)$, we can write the sum of interest as a martingale part plus a remainder, which is a sum of $k-1$ (i.e., one fewer) random variables:

$$
S=\sum_{i=1}^{k} Y_{i}^{(1)}+\sum_{i=1}^{k-1} \operatorname{Pf}\left(v_{i}\right) .
$$

Notice that $P f$ is also a mean zero function on $G$, and by Jensen's inequality we have

$$
\|(P f)(v)\|_{*} \leq M:=\max _{v \in V}\|f(v)\|_{*} \quad \text { for all } v
$$

The key point is that the remainder terms $P f\left(v_{i}\right)$ are smaller on average than the original terms $f\left(v_{i}\right)$ in squared Euclidean norm, because $P$ is a contraction orthogonal to the constant vector; in particular, by considering the action of $P$ on each coordinate of $f$ separately, we have:

$$
\begin{equation*}
\sum_{v}\|P f(v)\|_{2}^{2} \leq \lambda \cdot \sum_{v}\|f(v)\|_{2}^{2} \tag{7.15}
\end{equation*}
$$

Iterating this construction on the remainder a total of $T \leq k$ times, we obtain a sequence of martingales $1 \leq t \leq T$ :

$$
\begin{aligned}
& Y_{1}^{(t)}:=P^{t-1} f\left(v_{1}\right) \\
& Y_{i}^{(t)}:=P^{t-1} f\left(v_{i}\right)-P^{t} f\left(v_{i-1}\right), i=1, \ldots, k-(t-1)
\end{aligned}
$$

which are related to the original sum as:

$$
S=\sum_{t=1}^{T} \sum_{i=1}^{k-t+1} Y_{i}^{(t)}+\sum_{i=1}^{k-T} P^{T} f\left(v_{i}\right)
$$

Interchanging the order of summation, we find that the random matrices

$$
Z_{i}:=\sum_{t=1}^{\min \{(k+1-i), T\}} Y_{i}^{(t)}
$$

themselves form a martingale difference sequence, with each

$$
\left\|Z_{i}\right\|_{*} \leq \sum_{t}\left\|Y_{i}^{(t)}\right\|_{*} \leq T M
$$

We bound the error $W:=\frac{1}{k} \sum_{i=1}^{k-T}\left(P^{T} f\right)\left(v_{i}\right)$ crudely as:

$$
\begin{aligned}
\|W\|_{2} & \leq \frac{1}{k} \sum_{i=1}^{k-T}\left\|P^{T} f\left(v_{i}\right)\right\|_{2} \\
& \leq \frac{k-T}{k} \sum_{v \in V}\left\|\left(P^{T} f\right)(v)\right\|_{2} \\
& \leq \lambda^{T / 2} F \\
& \leq \exp (-(1-\lambda) T / 2) F
\end{aligned}
$$

where $F:=\left(\sum_{v \in V}\|f(v)\|_{2}^{2}\right)^{1 / 2}$, by applying (7.15). Rearranging and setting $T=2 \log (F / \epsilon) /(1-$ $\lambda)$ yields the advertised bound on $\left\|Z_{i}\right\|_{*}$.

### 7.6 Elementary calculations

Lemma 7.6.1. We define function $h(z): \mathbb{C} \rightarrow \mathbb{C}$ as follows

$$
h(z)=-\frac{1+z}{1-z}+\sqrt{\left(\frac{1+z}{1-z}\right)^{2}+1} .
$$

Then function $h(z)$ maps the unit disk $\{z \in \mathbb{C}:|z| \leq 1\}$ to the half disk $\{z \in \mathbb{C}:|z| \leq$ 1 and $\operatorname{Re}(z) \geq 0\}$.

Proof. We first compute the inverse of function $h(z)$, let $f=h^{-1}$. By definition of $h(z)$, we can do the following elementary calculations,

$$
\begin{aligned}
& h(z)+\frac{1+z}{1-z}=\sqrt{\left(\frac{1+z}{1-z}\right)^{2}+1} \\
\Longrightarrow & \left(h(z)+\frac{1+z}{1-z}\right)^{2}=\left(\frac{1+z}{1-z}\right)^{2}+1 \\
\Longrightarrow & (h(z))^{2}+2 h(z) \frac{1+z}{1-z}=1 \\
\Longrightarrow & (1-z)(h(z))^{2}+2 h(z)(1+z)=(1-z) \\
\Longrightarrow & (h(z))^{2}+2 h(z)-1=z\left((h(z))^{2}-2 h(z)-1\right),
\end{aligned}
$$

Thus, we obtain that

$$
f(z)=\frac{z^{2}+2 z-1}{z^{2}-2 z-1}
$$

To finish the proof, we need Claim 7.6.2 and Claim 7.6.3.
Claim 7.6.2. For all $z \in\{z \in \mathbb{C}:|z| \leq 1$ and $\operatorname{Re}(z) \geq 0\}$,

$$
|f(z)| \leq 1
$$

Proof. Let $z=r e^{\mathbf{i} \phi}$, where $r \in[0,1]$ and
$\phi \in[-\pi / 2, \pi / 2]$. It is easy to observe that $\cos (\phi) \in[0,1]$ We have

$$
\begin{aligned}
|f(z)| & =\left|\frac{z^{2}+2 z-1}{z^{2}-2 z-1}\right| \\
& =\left|\frac{r^{2} e^{\mathbf{i} 2 \phi}+2 r e^{\mathbf{i} \phi}}{r^{2} e^{\mathbf{i} 2 \phi}-2 r e^{\mathbf{i} \phi}-1}\right| \\
& =\frac{\left|r^{2} e^{\mathbf{i} 2 \phi}+2 r e^{\mathbf{i} \phi}-1\right|}{\left|r^{2} e^{\mathbf{i} 2 \phi}-2 r e^{\mathbf{i} \phi}-1\right|} .
\end{aligned}
$$

We can compute the numerator,

$$
\begin{aligned}
& \left|r^{2} e^{\mathbf{i} 2 \phi}+2 r e^{\mathbf{i} \phi}-1\right|^{2} \\
= & \left|r^{2} \cos 2 \phi+\mathbf{i} r^{2} \sin 2 \phi+2 r \cos \phi+2 r \mathbf{i} \sin \phi-1\right|^{2} \\
= & \left(r^{2} \cos 2 \phi+2 r \cos \phi-1\right)^{2}+\left(r^{2} \sin 2 \phi+2 r \sin \phi\right)^{2} \\
= & r^{4}+4 r^{2}+1-2 r^{2} \cos 2 \phi+4 r^{3} \cos 2 \phi \cos \phi \\
& -4 r \cos \phi+4 r^{3} \sin 2 \phi \sin \phi \\
= & r^{4}+4 r^{2}+1-2 r^{2} \cos 2 \phi+4 r^{3} \cos \phi-4 r \cos \phi .
\end{aligned}
$$

We can compute the denominator,

$$
\begin{aligned}
& \left|e^{\mathbf{i} 2 \phi}-2 e^{\mathbf{i} \phi}-1\right|^{2} \\
= & \left|r^{2} \cos 2 \phi+\mathbf{i} r^{2} \sin 2 \phi-2 r \cos \phi-2 r \mathbf{i} \sin \phi-1\right|^{2} \\
= & \left(r^{2} \cos 2 \phi-2 r \cos \phi-1\right)^{2}+\left(r^{2} \sin 2 \phi-2 r \sin \phi\right)^{2} \\
= & r^{4}+4 r^{2}+1-2 r^{2} \cos 2 \phi-4 r^{3} \cos 2 \phi \cos \phi \\
& +4 r \cos \phi-4 r^{3} \sin 2 \phi \sin \phi \\
= & r^{4}+4 r^{2}+1-2 r^{2} \cos 2 \phi-4 r^{3} \cos \phi+4 r \cos \phi .
\end{aligned}
$$

Note that, in order to show $|f(z)| \leq 1$, it is sufficient to prove

$$
4 r^{3} \cos \phi-4 r \cos \phi \leq-4 r^{3} \cos \phi+4 r \cos \phi
$$

which is equivalent to

$$
8 r\left(1-r^{2}\right) \cos \phi \geq 0
$$

It follows by definition of $r$ and $\phi$. Thus, we complete the proof.

Next, we can show that
Claim 7.6.3. For all $z$ is on the boundary of half disk,

$$
|f(z)|=1
$$

Proof. First, we want to show that $\forall z \in[-\mathbf{i}, \mathbf{i}],|f(z)|=1$. Let $b \in[0,1]$, let $z=\mathbf{i} b$, then we have

$$
|f(z)|=|f(\mathbf{i} b)|=\left|\frac{-b^{2}+2 b \mathbf{i}-1}{-b^{2}-2 b \mathbf{i}-1}\right|=1 .
$$

Second, we want to show that for all $z$ on half circle, $|f(z)|=1$. We replace $z$ by $e^{\mathbf{i} \phi}$, where $\phi \in[-\pi / 2, \pi / 2]$. Then we have

$$
\begin{aligned}
|f(z)| & =\left|\frac{z^{2}+2 z-1}{z^{2}-2 z-1}\right| \\
& =\left|\frac{e^{\mathbf{i} 2 \phi}+2 e^{\mathbf{i} \phi}-1}{e^{\mathbf{i} 2 \phi}-2 e^{\mathbf{i} \phi}-1}\right| \\
& =\frac{\left|e^{\mathbf{i} 2 \phi}+2 e^{\mathbf{i} \phi}-1\right|}{\left|e^{\mathbf{i} 2 \phi}-2 e^{\mathbf{i} \phi}-1\right|} .
\end{aligned}
$$

We can compute the numerator,

$$
\begin{aligned}
& \left|e^{\mathbf{i} 2 \phi}+2 e^{\mathbf{i} \phi}-1\right| \\
= & |\cos 2 \phi+\mathbf{i} \sin 2 \phi+2 \cos \phi+2 \mathbf{i} \sin \phi-1| \\
= & \left((\cos 2 \phi+2 \cos \phi-1)^{2}+(\sin 2 \phi+2 \sin \phi)^{2}\right)^{1 / 2} \\
= & (6-2 \cos 2 \phi+4 \cos 2 \phi \cos \phi-4 \cos \phi+4 \sin 2 \phi \sin \phi)^{1 / 2} \\
= & (6-2 \cos 2 \phi+4 \cos \phi-4 \cos \phi)^{1 / 2} \\
= & (6-2 \cos 2 \phi)^{1 / 2} .
\end{aligned}
$$

We can compute the denominator,

$$
\begin{aligned}
& \left|e^{\mathbf{i} 2 \phi}-2 e^{\mathbf{i} \phi}-1\right| \\
= & |\cos 2 \phi+\mathbf{i} \sin 2 \phi-2 \cos \phi-2 \mathbf{i} \sin \phi-1| \\
= & \left((\cos 2 \phi-2 \cos \phi-1)^{2}+(\sin 2 \phi-2 \sin \phi)^{2}\right)^{1 / 2} \\
= & (6-2 \cos 2 \phi-4 \cos 2 \phi \cos \phi+4 \cos \phi-4 \sin 2 \phi \sin \phi)^{1 / 2} \\
= & (6-2 \cos 2 \phi+4 \cos \phi-4 \cos \phi)^{1 / 2} \\
= & (6-2 \cos 2 \phi)^{1 / 2} .
\end{aligned}
$$

Thus, we have $|f(z)|=1$.

Note the biholomorphic is basically follows from the formula of $f$ and $h$, because they are composition of holomorphic function.

Lemma 7.6.4. For any $\rho \in[0,1]$ and $\cos \varphi \in[-1,0]$, we have

$$
\frac{1-\rho}{1-\cos \varphi}-(1-\rho)^{2} \leq \frac{1-\rho^{2}}{1-2 \rho \cos \varphi+\rho^{2}} \leq \frac{1-\rho}{1-\cos \varphi}+2(1-\rho)^{2}
$$

Proof. This directly follows by combining Claim 7.6.5 and Claim 7.6.6

Claim 7.6.5. There exists some sufficiently large constant $c \geq 1$ such that for any $\rho \in[0,1]$ and $\cos \varphi \in[-1,0]$, we have

$$
\frac{1-\rho^{2}}{1-2 \rho \cos \varphi+\rho^{2}} \leq \frac{1-\rho}{1-\cos \varphi}+c(1-\rho)^{2} .
$$

Proof. It is equivalent to

$$
\begin{aligned}
\frac{1+\rho}{1-2 \rho \cos \varphi+\rho^{2}} & \leq \frac{1}{1-\cos \varphi}+c(1+\rho) \\
(1+\rho)(1-\cos \varphi) & \leq 1-2 \rho \cos \varphi+\rho^{2}+c(1+\rho)(1-\cos \varphi)\left(1-2 \rho \cos \varphi+\rho^{2}\right) \\
\rho-\cos \varphi & \leq-\rho \cos \varphi+\rho^{2}+c(1+\rho)(1-\cos \varphi)\left(1-2 \rho \cos \varphi+\rho^{2}\right)
\end{aligned}
$$

which is equivalent to,

$$
-\rho \cos \varphi+\rho^{2}-\rho+\cos \varphi+c(1+\rho)(1-\cos \varphi)\left(1-2 \rho \cos \phi+\rho^{2}\right) \geq 0
$$

Since $1-2 \rho \cos \varphi+\rho^{2} \geq 1$, thus it suffices to show

$$
(\rho-1)(\rho-\cos \varphi)+c(1+\rho)(1-\cos \varphi) \geq 0 .
$$

Note that $(\rho-1)(\rho-\cos \varphi) \geq-2$, by choosing $c \geq 2$, we complete the proof.

Claim 7.6.6. There exists some sufficiently large constant $c \geq 1$ such that for any $\rho \in[0,1]$ and $\cos \varphi \in[-1,0]$, we have

$$
\frac{1-\rho^{2}}{1-2 \rho \cos \varphi+\rho^{2}} \geq \frac{1-\rho}{1-\cos \varphi}-c(1-\rho)^{2} .
$$

Proof. It is equivalent to

$$
\begin{aligned}
\frac{1+\rho}{1-2 \rho \cos \varphi+\rho^{2}} & \geq \frac{1}{1-\cos \varphi}+c(1+\rho) \\
(1+\rho)(1-\cos \varphi) & \geq 1-2 \rho \cos \varphi+\rho^{2}-c(1+\rho)(1-\cos \varphi)\left(1-2 \rho \cos \varphi+\rho^{2}\right) \\
\rho-\cos \varphi & \geq-\rho \cos \varphi+\rho^{2}-c(1+\rho)(1-\cos \varphi)\left(1-2 \rho \cos \varphi+\rho^{2}\right)
\end{aligned}
$$

which is equivalent to,

$$
-\rho \cos \varphi+\rho^{2}-\rho+\cos \varphi-c(1+\rho)(1-\cos \varphi)\left(1-2 \rho \cos \phi+\rho^{2}\right) \leq 0
$$

which is equivalent to

$$
(\rho-1)(\rho-\cos \varphi) \leq c(1+\rho)(1-\cos \varphi)\left(1-2 \rho \cos \varphi+\rho^{2}\right)
$$

It suffices to choose $c=1$.

## Chapter 8

## Matrix Chernoff Bound for Random Spanning Trees

Strongly Rayleigh distributions are a class of negatively dependent distributions of binary-valued random variables [Borcea, Brändén, Liggett JAMS 09]. Recently, these distributions have played a crucial role in the analysis of algorithms for fundamental graph problems, e.g. Traveling Salesman Problem [Gharan, Saberi, Singh FOCS 11]. We prove a new matrix Chernoff bound for Strongly Rayleigh distributions.

As an immediate application, we show that adding together the Laplacians of $\epsilon^{-2} \log ^{2} n$ random spanning trees gives an ( $1 \pm \epsilon$ ) spectral sparsifiers of graph Laplacians with high probability. Thus, we positively answer an open question posed in [Baston, Spielman, Srivastava, Teng JACM 13]. Our number of spanning trees for spectral sparsifier matches the number of spanning trees required to obtain a cut sparsifier in [Fung, Hariharan, Harvey, Panigraphi STOC 11]. The previous best result was by naively applying a classical matrix Chernoff bound which requires $\epsilon^{-2} n \log n$ spanning trees. For the tree averaging procedure to agree with the original graph Laplacian in expectation, each edge of the tree should be reweighted by the inverse of the edge leverage score in the original graph. We also show that when using this reweighting of the edges, the Laplacian of single random tree is bounded above in the PSD order by the original graph Laplacian times a factor $\log n$ with high probability, i.e. $L_{T} \preceq O(\log n) L_{G}$.

We show a lower bound that almost matches our last result, namely that in some graphs, with high probability, the random spanning tree is not bounded above in the spectral order by $\frac{\log n}{\log \log n}$ times the original graph Laplacian.

We also show a lower bound that in $\epsilon^{-2} \log n$ spanning trees are necessary to get a $(1 \pm \epsilon)$ spectral sparsifier.

### 8.1 Introduction

The study of concentration of sums of random variables dates back to Central Limit Theorems, and hence de Moivre and Laplace [Tij], while modern concentration bounds for sums of random variables were perhaps first established by Bernstein [Ber24], and a popular variant now known as Chernoff bounds was introduced by Rubin and published by Chernoff [Che52].

Concentration of measure for matrix-valued random variables is the phenomenon that many matrix valued distributions are to close their mean with high probability, closeness usually being measured by spectral norm. Modern quantitative bounds of the form often used in theoretical computer science were derived by Rudelson [Rud99], while Ahlswede and Winter [AW02] established a useful matrix-version of the Laplace transform that plays a central role in scalar concentration results such as those of Bernstein. [AW02] combined this with the Golden-Thompson trace inequality to prove matrix concentration results. Tropp refined this approach, and by replacing the use of Golden-Thompson with deep a theorem on concavity of certain trace functions due to Lieb, Tropp was able to recover strong versions of a wide range of scalar concentration results, including matrix Chernoff bounds, Azuma and Freedman's inequalities for matrix martingales [Tro12].

Matrix concentration results have had an enormous range of applications in computer science, and are ubiquitous throughout spectral graph theory [ST04, SS11, CKP ${ }^{+}$17], sketching [Coh16a], approximation algorithms [HSSS16], and deep learning [ZSJ ${ }^{+}$17, ZSD17, SY19]. Most applications are based on results for independent random matrices, but more flexible bounds, such as Tropp's Matrix Freedman Inquality [Tro11a], have been used to greatly simplify algorithms, e.g. for solving Laplacian linear equations [KS16] and for semi-streaming
graph sparsification [AG09, KPPS17]. Matrix concentration results are also closely related to other popular tools sampling tools, such as Karger's techniques for generating sparse graphs that approximately preserve the cuts of denser graphs [BK96].

Negative dependence of random variables is an appealing property that intuition suggests should help with concentration of measure. Notions of negative dependence can be formalized in many ways. Roughly speaking, these notions characterize distributions where where some event occurring ensures that other events of interest become less likely. A simple example is the distribution of a sequence of coin flips, conditioned on the total number of heads in the outcome. In this distribution, conditioning on some coin coming out heads makes all other coins less likely to come out heads. Unfortunately, negative dependence phenomena are not as robust as positive association which can be established from local conditions using the powerful FKG theorem [FKG71].

Strongly Rayleigh distributions were introduced recently by Borcea, Brändén, and Liggett [BBL09] as a class of negatively dependent distributions of binary-valued random variables with many useful properties. Strongly Rayleigh distributions satisfy useful negative dependence properties, and retain these properties under natural conditioning operations. Strongly Rayleigh distributions also satisfy a powerful stability property under conditioning known as Stochastic Covering [PP14], which is useful for analyzing them through martingale techniques. A measure on $\{0,1\}^{n}$ is said to be Strongly Rayleigh if its generating polynomial is real stable [BBL09]. There are many interesting examples of Strongly Rayleigh distributions [PP14]: The example mentioned earlier of heads of independent coin flips conditional on the total number of heads in the outcome; symmetric exclusion processes; determinental point processes and determinental measures on a boolean lattice. An example of particu-
lar interest to us is the edges of uniform or weighted random spanning trees, which form a Strongly Rayleigh distribution.

We prove a Matrix Chernoff bound for the case of $k$-homogeneous Strongly Rayleigh distributions. Our bound is slightly weaker than the bound for independent variables. We give lower bounds that show our bounds are close to tight in some regimes, but importantly, our lower bounds do not establish separation from the behaviour of indepedent random matrices, leaving open the question of whether the true bound should match the independent case in all regimes - which seems plausible. We use our bound to show new concentration results related to random spanning trees of graphs. An open question is to find other interesting applications of our concentration result, e.g. by analyzing concentration for matrices generated by exclusion processes.

Random spanning trees are one among the most well-studied probabilistic objects in graph theory, going back to the work of Kirchoff [Kir47] in 1847, who gave formula relating the number of spanning trees in a graph to the determinant of the Laplacian of the same graph.

Algorithms for sampling of random spanning trees have been studied extensively, [Gue83, Bro89, Ald90, Kul90, Wil96, CMN96, KM09, MST15, HX16, DKP ${ }^{+}$17, DPPR17, Sch18], and a random spanning tree can now be sampled in almost linear time [Sch18].

In theoretical computer science, random spanning trees have found a number of applications, most notably in breakthrough results on approximating the traveling salesperson problem with symmetric [GSS11] and asymmetric costs [AGM ${ }^{+} 10$ ]. Goyal et al. [GRV09] demonstrated that adding just two random spanning trees sampled from a bounded degree
graph gives a $O(\log n)$ cut sparsifier with probability $1-o(1)$. Later, it was shown by Fung, Hariharan, Harvey, Panigraphi [FHHP11], that if we sample $O\left(\epsilon^{-2} \log ^{2} n\right)$ random spanning trees from a graph, reweight the tree edges by the inverse of their leverage scores in the original graph, and average them together, then whp. we get a graph where every the weight of edges crossing every cut is approximately the same in as in the original graph, up to a factor $(1 \pm \epsilon)$. We refer to this as an $\epsilon$-cut sparsifier. The techniques of Fung et al. unfortunately do not extend to proving spectral sparsifiers.

Spectral graph sparsifiers were introduced by Spielman and Teng [ST04], who for any graph $G$ showed how to construct a another graph $H$ with $\epsilon^{-2} n$ poly $\log n$ edges s.t. $(1-\epsilon) L_{G} \preceq L_{H} \preceq(1+\epsilon) L_{G}$, which we refer to as an $\epsilon$-spectral sparsifier. The construction was refined by Spielman and Srivastava [SS11], who suggested sampling edges independently ${ }^{1}$ with probability proportional to their leverage scores, and brought the number of required samples down to $\epsilon^{-2} n \log n$. This analysis is tight in the sense that if fewer than $o\left(\epsilon^{-2} n \log n\right)$ samples are used, there will be at least a $1 / \operatorname{poly}(n)$ probability of failure. Meanwhile, $\epsilon^{-2} n \frac{\log n}{\log \log n}$ independent samples in a union of cliques can be shown whp. to fail to give a cut sparsifier. This can be observed directly from the degree distribution of a single vertex in the complete graph. For a variant of [SS11] sampling based on flipping a single coin for each edge to decide whether to keep it or not, it can also be shown that when the expected number of edges is $\epsilon^{-2} n \frac{\log n}{\log \log n}$, whp. the procedure fails to give a cut sparsifier. For arbitrary sparsification schemes, bounds in [BSS12] show that $\Theta\left(\epsilon^{-2} n\right)$ edges are necessary and sufficient to give an $\epsilon$-spectral sparsifier.

[^49]The marginal probability of an edge being present in a random spanning tree is exactly the leverage score of the edge. This seems to suggest that combining $\epsilon^{-2}$ poly $\log n$ spanning trees might give a spectral sparsifier, but the lack of independence between the sampled edges means the process cannot be analyzed using existing techniques. Observing this, Baston, Spielman, Srivastava, Teng [BSST13] in their excellent 2013 survey on sparsification noted that "it remains to be seen if the union of a small number of random spanning trees can produce a spectral sparsifier." We answer this question in the affirmative. In particular, we show that adding together $O\left(\epsilon^{-2} \log ^{2} n\right)$ spanning trees with edges scaled proportional to inverse leverage scores in the original graph leads to an $\epsilon$-spectral sparsifier. This matches the bound obtained for cut sparsifiers in [FHHP11]. Our result also implies their earlier bound since a spectral sparsifier is always a cut sparsifier with the same approximation quality. Before our result, only a trivial bound on the number of spanning trees required to build a spectral sparsifier was known. In particular standard matrix concentration arguments like those in [SS11] prove that $O\left(\epsilon^{-2} n \log n\right)$ spanning trees suffice. Lower bounds in [FHHP11] show that whp. $\Omega(\log n)$ random spanning trees are required to give a constant factor spectral sparsifier.

We show that whp. $\epsilon^{-2} \frac{\log n}{\log \log n}$ random spanning trees do not give an $\epsilon$-spectral sparsifier. We also show that the Laplacian of a single random tree with edges weighted as above satisfies $L_{T} \preceq O(\log n) L_{G}$ whp., and we give an almost matching lower bound, showing that in some graphs whp. $L_{T} \nprec \frac{1}{8} \frac{\log n}{\log \log n} L_{G}$. Before our work, the main result known about approximating graphs using $O(1)$ random spanning trees is due to Goyal, Rademacher, Vempala [GRV09], who showed that surprisingly, when the original graph has bounded degree, adding two random spanning trees gives a graph whose cuts approximate the cuts in the
original graph up to a factor $O(\log n)$ with good probability. As our result for a single tree establishes only a one-sided bound an interesting open question remains: Does sampling $O(1)$ random spanning trees give a $\log n$-factor spectral sparsifier with, say, constant probability?

### 8.1.1 Previous work

Chernoff-type bound for matrices. Chernoff-like bounds for matrices appear in Rudelson [Rud99] and Ahlswede and Winter [AW02]. The latter introduced a useful matrixvariant of the Laplace transform that is central in concentration bounds for scalar-valued matrices. Their bounds restricted to iid random matrices, an artifact of their use of the Golden-Thompson inequality for bounding traces. In contrast, Tropp obtained more flexible concentration bounds for random matrices by using a result of Lieb to bound the expected trace of various operators [Tro12], including bounds for matrix martingales [Tro11c].

In a recent work by Garg, Lee, Song and Srivastava [GLSS18], they show a Chernoff bound for sums of matrix-valued random variables sampled via a random walk on an expander graph. This work confirms a conjecture due Wigderson and Xiao. The proof of Garg et al. is also concerned with matrices that are not fully independent. In this case the matrices are generated from random walks on an expander graph. The main idea to deal with dependence issue is using a new multi-matrix extension of the Golden-Thompson inequality and an adaptation of Healy's proof of the expander Chernoff bound in the scalar's case [Hea08] to matrix case. Their techniques deal with fairly generic types of dependence, and cannot leverage the very strong stability properties that arise from the negative dependence and stochastic covering properties of Strongly Rayleigh distributions. Harvey and Olver [HO14] proved a matrix concentration result for randomized pipage rounding, which can be used to
show concentration results for random spanning trees obtained from pipage rounding, but not for (weighted) uniformly random spaning trees. The central technical element of their proof is a new variant of a theorem of Lieb on concavity of certain matrix trace functions.

Matrix martingales have played a central role in a number of algorithmic results in theoretical computer science [KS16, CMP16, KPPS17], but beyond a reliance on Tropp's Matrix Freedman Inequality, these works have little in common with our approach. However, our bound does share a technical similarity with [KS16], namely that a sequence of increasingly restricted random choices in a martingale process lead to a $\log n$ factor in a variance bound.

Strongly Rayleigh Distributions in Theoretical Computer Science. Perhaps the most prominent result on Strongly Rayleigh distributions in theoretical computer science is the generalization of [MSS15b] to Strongly Rayleigh distributions.

The central technical result of [MSS15b] essentially shows that given a collection of independent random vectors $v_{1}, \ldots, v_{m}$ with finite support in $\mathbb{C}^{n}$ s.t. $\sum_{i=1}^{m} \mathbb{E}\left[v_{i} v_{i}^{*}\right]=I$ and for all $i,\left\|v_{i}\right\|^{2} \leq \epsilon$, then $\operatorname{Pr}\left[\left\|\sum_{i=1}^{m} v_{i} v_{i}^{*}\right\| \leq(1+\sqrt{\epsilon})^{2}\right]>0$. [AG15] establishes a related result for $k$-homogeneous Strongly Rayleigh distributions, though they require an additional constraint on the marginal probability that any given random variable is non-zero being bounded above by $\delta$, and then establish $\operatorname{Pr}\left[\left\|\sum_{i=1}^{m} v_{i} v_{i}^{*}\right\| \leq 4(\epsilon+\delta)+2(\epsilon+\delta)^{2}\right]>0$. Based on this, [AG15] shows ${ }^{2}$ that given an unweighted $k$-edge connected graph $G$ where every edge has leverage score at most $\epsilon$, there exists an unweighted spanning tree s.t. $L_{T} \preceq O\left(\frac{1}{k}+\epsilon\right) \cdot L_{G}$. This is referred to as a spectrally thin tree with parameter $O\left(\frac{1}{k}+\epsilon\right)$.

[^50][AGR16] showed how to algorithmically sample from $k$-homogeneous Determinental Point Process in time poly $(k) n \log (n / \epsilon)$, where $n$ is the dimension of matrix giving rise to the determinental point process and $\epsilon$ is the allowed total variation distance. Their techniques are based on generalization proofs of expansion in the base graph associated with a balanced matroid, a result first established by [FM92].

Random spanning trees. Algorithms for sampling random spanning trees have a long history, but only recently have they explicitly used matrix concentration [DKP ${ }^{+} 17$, DPPR17, Sch18]. The matrix concentration arguments in these papers, however, deal mostly with how modifying a graph results in changes to the distribution of random spanning trees in the graph. We instead study how closely random spanning trees resemble the graph they were initially sampled from. Whether our result in turn has applications for improving sampling algorithms for random spanning trees is unclear.

The fact that spanning tree edges exhibit negative dependence has been used strikingly in concentration arguments by Goyal et al. [GRV09] to show that two random spanning trees gives $O(\log n)$-factor approximate cut sparsifier in bounded degree graphs, with good probability. This is clearly false when sampling the same number of edges independently, because this graph has large probability of having isolated vertices. Goyal et al. improve over independent sampling by leveraging the fact that for a fixed tree, in some sense, very few cuts of a given size exist. This is a variant of Karger's famous cut-counting techniques [Kar93, KS96] specialized to unweighted trees.

Uses of negatively dependent Chernoff bounds applied to tree edges also appeared in works on approximation algorithms for TSP problems [GSS11, AGM ${ }^{+}$10], where additionally
the connectivity properties of the tree play an important role
In contrast, the techniques of Fung et al. [FHHP11] show that $O\left(\epsilon^{-2} \log ^{2} n\right)$ spanning trees suffice to give a $(1 \pm \epsilon)$-cut sparsifier, but they do not show that tree-based sparsifiers improve over independent sampling, The focus of their paper is to establish that wide range of different techniques for choosing sampling probabilities all give cut sparsifiers, by establishing a more flexible framework than the original cut-sparsifier results of Benczur-Karger [BK96], using related cut-counting techniques (see [Kar93, KS96]). To extend their results to spanning trees, they simply observe that the (scalar-valued) Chernoff bounds they use directly apply to negatively dependent variables, and hence edges in spanning trees.

Fung et at. [FHHP11] also establish a lower bound, showing that for any constant $c$, there exists a graph for which obtaining a factor $c$-cut sparsifier by averaging trees requires using at least $\Omega(\log n)$ trees to succeed with constant probability.

### 8.1.2 Our results and techniques

Theorem 8.1.1 (First main result, a Matrix Chernoff Bound $k$-homogeneous Strongly Rayleigh Distributions). Suppose $\left(\xi_{1}, \ldots, \xi_{m}\right) \in\{0,1\}^{m}$ is a random vector of $\{0,1\}$ variables whose distribution is $k$-homogeneous and Strongly Rayleigh.

Given a collection of PSD matrices $A_{1}, \ldots A_{m} \in \mathbb{R}^{n \times n}$ s.t. for all $e \in[m]$ we have $\left\|A_{e}\right\| \leq R$ and $\left\|\mathbb{E}\left[\sum_{e} \xi_{e} A_{e}\right]\right\| \leq \mu$.

Then for any $\epsilon>0$,

$$
\operatorname{Pr}\left[\left\|\sum_{e} \xi_{e} A_{e}-\mathbb{E}\left[\sum_{e} \xi_{e} A_{e}\right]\right\| \geq \epsilon \mu\right] \leq n \exp \left(-\frac{\epsilon^{2} \mu}{R(\log k+\epsilon)} \Theta(1)\right)
$$

This Matrix Chernoff bound matches the bounds due to Tropp [Tro12], up to the $\log k$ factor in the exponent. Our lower bounds rule out that a much stronger bound is true in the $\epsilon \approx \log k$ regime, since the level of concentration we prove at this $\epsilon$ is stronger than what is ruled out for $\epsilon \approx \frac{\log k}{\log \log k}$ by the lower bound in Theorem 8.1.3. However, this does not rule out that the $\log k$ factor in our bound is unnecessary. I.e. we cannot rule out that a stronger concentration statement might match the bound for the independent case given by Tropp [Tro12].

Remark 8.1.1. When the Strongly Rayleigh distribution is in fact a product distribution on a collection of $l$ Strongly Rayleigh distributions that are each $t$-homogeneous, then the joint distribution is $l t$-homogeneous Strongly Rayleigh, but in Theorem 8.1.1, the factor $\log (l t)$ can be replaced by a factor $\log t$. This applies to the case of independent random spanning trees, but only gives a constant factor improvement in the number of trees required.

Our work is related to the concentration inequality of Peres and Pemantle [PP14], who showed a concentration result for scalar-valued Lipschitz functions of Strongly Rayleigh distributions. They used Doob martingales (martingales constructed from sequences of conditional expectations) to prove their result. We use a similar approach for matrices, constructing Doob matrix martingales from our Strongly Rayleigh distributions. In addition, we use the stochastic covering property of Strongly Rayleigh distributions observed by Peres and Pemantle, but implicitly derived in [BBL09]. This property leads to bounded differences in Doob martingale sequences for scalars. As in the scalar setting, it is possible to show concentration results for matrix-valued martingales. We use the Matrix Freedman inequality ${ }^{3}$

[^51]of Tropp. This inequality allows makes it possible to establish strong concentration bounds based on control of sample norms and control of the predictable quadratic variation process of the martingale, a matrix-valued object that is used to measure variance (see [Tro11c]). We show that as in the scalar setting, the stochastic covering property of Strongly Rayleigh distributions leads to bounded differences for Doob matrix martingales. But, we also combine the stochastic covering property with deceptively simple matrix martingale properties and a negative dependence condition to derive additional bounds on the predictable quadratic variation process of the martingale. The key negative dependence property we use is a simple observation that generalizes, to $k$-homogeneous Strongly Rayleigh distributions, the fact that in a random spanning tree, conditioning on the presence of a set of edges lowers the marginal probability of every other graph edge being in the tree (see Lemma 8.1.7). While we frame it differently, it is essentially an immediate consequence of statements in [BBL09]. The surprise here is how useful this simple observation is for removing issues with characterizing conditional $k$-homogeneous Strongly Rayleigh distributions. As a corollary we get our second main result.

Theorem 8.1.2 (Second main result, concentration bound of a batch of independent random spanning trees). Given as input a weighted graph $G$ with $n$ vertices and a parameter $\epsilon>0$, let $T_{1}, T_{2}, \cdots, T_{t}$ denote $t$ independent inverse leverage score weighted random spanning trees, if we choose $t=O\left(\epsilon^{-2} \log ^{2} n\right)$ then with probability $1-1 / \operatorname{poly}(n)$,

$$
(1-\epsilon) L_{G} \preceq \frac{1}{t} \sum_{i=1}^{t} L_{T_{i}} \preceq(1+\epsilon) L_{G} .
$$

process, which means the bound we use is more analogous to a matrix version Bernstein's inequality, adapted to martingales. We resort to the more complicated Freedman's inequality only because it gives a directly applicable statement that is known in the literature.

Prior to our work, only a trivial bound on the number of spanning trees required to build a spectral sparsifier was known, namely that standard matrix concentration arguments like those in [SS11] prove that $O\left(\epsilon^{-2} n \log n\right)$ spanning trees suffice. Note that, the number of spanning trees required to build spectral sparsifier in our Theorem 8.1.2 matches the number of spanning trees required to construct cut sparsifier in previous best result [FHHP11]. The total edge count we require is $\Theta\left(\epsilon^{-2} n \log ^{2} n\right)$, worse by a factor $\log n$ than the bound for independent edge sampling obtained in [SS11]. It is not clear whether this factor in necessary. Remark 8.1.2. Suppose we apply our Theorem 8.1.3 to show that any single random spanning tree satisfies $L_{T} \preceq O(\log n) \cdot L_{G}$ whp. This is tight up a $\log \log n$ factor. Then, one can from use this to derive Theorem 8.1.2 based on a standard (and tight) Matrix Chernoff bound, and a (laborious) combination of Doob martingales and stopping time arguments similar to those found in [Tro12, KS16]. This line of reasoning will lead to the same bounds as Theorem 8.1.3. Thus, unless one proves more than just a norm bound for each individual tree, it is not possible improve over our result, except for $\log \log n$ factors.

Like the work of Fung et al. our results for spanning trees do not improve over the independent case. Fung et al. achieved their result by combining cut counting techniques with Chernoff bounds for scalar-valued negatively dependent variables. In our random matrix setting, there are no clear candidates for a Chernoff bound for negatively dependent random matrices that we can adopt, and this type of bound is exactly what we develop in the Strongly Rayleigh case.

We establish a one-sided concentration result for a single tree, namely that whp. $L_{T} \preceq O(\log n) \cdot L_{G}$. Again, this is a direct application of Theorem 8.1.1.

Theorem 8.1.3 (Third main result, upper bound for the concentration of one random spanning tree). Given a graph $G$, let $T$ be a random spanning tree, then with probability at least $1-1 / \operatorname{poly}(n)$

$$
L_{T} \preceq O(\log n) \cdot L_{G} .
$$

This upper bound is tight up to a factor $\log \log n$ as shown by our almost matching lower bound stated below.

Theorem 8.1.4 (Lower bound for the concentration of one random spanning tree). For any $n \geq 2^{2^{6}}$, there is an unweighted graph $G$ with $n$ nodes, s.t. if we sample an inverse leverage score weighted random spanning tree $T$, then with probability at least $1-e^{-n^{4}}$,

$$
L_{T} \nprec \frac{\log n}{8 \log \log n} \cdot L_{G} .
$$

Trivially, the presence of degree one nodes in $L_{T}$ means that in a complete graph, $L_{G} \nprec L_{T}$. So choosing any other scaling of the tree will make at least one of the inequalities $L_{T} \nprec \frac{1}{8} \log n / \log \log n \cdot L_{G}$ and $L_{G} \nprec L_{T}$ true with a larger gap. Note that in the complete unweighted graph the trees we consider have weight $\Theta(n)$ on each edge. A random spanning tree in the complete graph has diameter about $\sqrt{n}[\operatorname{RS} 67]$. This can be shown to imply that for an unweighted random tree $\widehat{T}, L_{G} \nprec \sqrt{n} L_{\widehat{T}}$. But once we scale up every edge of the tree by a factor $\Theta(n)$, the diameter bound no longer directly implies a spectral gap of the form $L_{G} \nprec \alpha L_{T}$ for some $\alpha$. In a ring graph, we get $L_{G} \nprec(n-2) L_{T}$ and $L_{T} \nprec L_{G}$.

We can also show in general that $L_{T} \nprec 10 \log n \cdot L_{G}$, but with a much smaller probability.

Theorem 8.1.5 (Lower bound for the concentration of one random spanning tree). For any $n \geq 4$, there is an unweighted graph $G$ with $n$ nodes, s.t. if we sample an inverse leverage score weighted random spanning tree $T$, then with probability at least $2^{-150 \log n \log \log n}$,

$$
L_{T} \nprec 10 \log n \cdot L_{G} .
$$

And we show a lower bound for $\epsilon$-spectral sparsifiers for random spanning trees.

Theorem 8.1.6 (Lower bound for the concentration of multiple random spanning trees). For any $n \geq 2^{100}$, there is an unweighted graph $G$ with $n$ nodes, s.t. for any accuracy parameter $\epsilon \in\left(5 n^{-0.1}, 1 / 2\right)$, if we sample $t=0.05 \epsilon^{-2} \log n$ independent random spanning trees with edges weighted by inverse leverage score, then with probability at least $1-e^{-n^{39}}$,

$$
(1-\epsilon) L_{G} \nprec \frac{1}{t} \sum_{i=1}^{t} L_{T_{i}} \text { and } \frac{1}{t} \sum_{i=1}^{t} L_{T_{i}} \nprec(1+\epsilon) L_{G} \text {. }
$$

Our lower bound is incomparable with that of Fung et al.[FHHP11], who showed that for any constant $c$, there exists a graph for obtaining a factor $c$-cut sparsifier by averaging trees requires using at least $\Omega(\log n)$ trees to succeed with constant probability. Where Fung et. al [FHHP11] used triangles in their lower bound construction, our bad examples are based on collections of small cliques, which lets us ensure cut differences in even a single tree, by giving longer-tailed degree distributions. All of our lower bounds are based on simple constructions from collections of edge disjoint cliques, and use the fact that the exact distribution of degrees of a fixed vertex in a random spanning tree of the complete graph is known. Note that a lower bound for cut approximation implies a lower-bound for spectral approximation, because the contrapositive statement is true: spectral approximation implies cut approximation.

Remark 8.1.3. In fact, all our lower bounds also directly apply for cut approximation, which is a strictly stronger result. For example, there is an unweighted graph $G$, s.t. if we sample an inverse leverage score weighted random spanning tree $T$, then with probability at least $1-e^{-n^{4}}, T$ has a cut which is larger than the corresponding cut in $G$ by a factor $\frac{\log n}{8 \log \log n}$.

Connection to Spectrally Thin Trees Using their MSS-type existence proof for "small norm outcomes" of homogeneous Strongly Rayleigh distributions, [AG15] showed that in an unweighted $k$-edge connected graph $G$ where every edge has leverage score at most $\epsilon$, there exists an unweighted spanning tree $\widehat{T}$ s.t. $L_{\widehat{T}} \preceq O\left(\frac{1}{k}+\epsilon\right) \cdot L_{G}$. This is referred to a spectrally thin tree with parameter $O\left(\frac{1}{k}+\epsilon\right)$.

In contrast, applying our $k$-homogeneous Strongly Rayleigh Matrix Chernoff bound to an unweighted graph $G$ where every edge has leverage score at most $\epsilon$, we can show that an unweighted random spanning tree satisfies $L_{T} \preceq O(\epsilon \log n) L_{G}$ with high probability. This follows immediately from our Theorem 8.1.3, because if we let $T$ denote the unweighted spanning tree and $\widehat{T}$ corresponding spanning tree with edges weighted by inverse leverage scores, then

$$
\begin{aligned}
L_{\widehat{T}} & =\sum_{e \in T} w(e) b_{e} b_{e}^{\top} \\
& \preceq \epsilon \sum_{e \in T} \frac{1}{l(e)} w(e) b_{e} b_{e}^{\top} \\
& =\epsilon L_{\widehat{T}} \\
& \preceq O(\epsilon \log n) L_{G}
\end{aligned}
$$

whp.

The proof in [AG15] is based on an adaptation of the [MSS15b] proof, and does not have clear parallels with our approach. Whereas the key properties of Strongly Rayleigh distributions that we use are stochastic covering (a property that limits change in a distribution under conditioning) and conditional negative dependence, the central element of their approach is a proof that certain mixed characteristic polynomials associated with $k$ homogeneous distributions are real stable when the original distribution is Strongly Rayleigh.

The following Lemma captures a simple but crucial property of Strongly Rayleigh distributions.

Lemma 8.1.7 (Shrinking Marginals). Suppose $\left(\xi_{1}, \ldots, \xi_{m}\right) \in\{0,1\}^{m}$ is a random vector of $\{0,1\}$ variables whose distribution is $k$-homogeneous and Strongly Rayleigh, then any set $S \subseteq[m]$ with $|S| \leq k$ for all $j \in[m] \backslash S$

$$
\operatorname{Pr}\left[\xi_{j}=1 \mid \xi_{S}=\mathbf{1}_{S}\right] \leq \operatorname{Pr}\left[\xi_{j}=1\right] .
$$

We provide a proof in Section 8.7.

### 8.2 Notation

We use $[n]$ to denote set $\{1,2, \cdots, n\}$. Given a vector $x$, we use $\|x\|_{0}$ to denote the number of non-zero entries in the vector.

Matrix and norms. For a matrix $A$, we use $A^{\top}$ to denote the transpose of $A$. We say matrix $A$ is positive semi-definite (PSD) if $A=A^{\top}$ and $x^{\top} A x \geq 0$ for all $x \in \mathbb{R}^{n}$. We use $\succeq, \preceq$ to denote the semidefinite ordering, e.g. $A \succeq 0$ denotes that $A$ is PSD, and $A \succeq B$ means $A-B \succeq 0$. We say matrix $A$ is positive definite ( PD ) if $A=A^{\top}$ and $x^{\top} A x>0$ for all $x \in \mathbb{R}^{n}-\{0\} . A \succ B$ means $A-B$ is PD .

For matrix $A \in \mathbb{R}^{n \times n}$, we define $\|A\|$ to be the spectral norm of $A$, i.e.,

$$
\|A\|=\max _{\|x\|_{2}=1, x \in \mathbb{R}^{n}} x^{\top} A x
$$

Let $\operatorname{tr}(A)$ denote the trace of a square matrix $A$. We use $\lambda_{\max }(A)$ to denote the largest eigenvalue of matrix $A$. For symmetric matrix $A \in \mathbb{R}^{n \times n}, \lambda_{\max }(A)=\|A\|$ and $\operatorname{tr}(A) \leq n\|A\|$.

The Laplacian matrix-related definitions. Let $G=(V, E, w)$ be a connected weighted undirected graph with $n$ vertices and $m$ edges and edge weights $w_{e}>0$. If we orient the edges of $G$ arbitrarily, we can write its Laplacian as $L=B^{\top} W B$, where $B \in \mathbb{R}^{m \times n}$ is the signed edge-vertex incidence matrix and defined as follows

$$
B(e, v)= \begin{cases}1, & \text { if } v \text { is } e^{\prime} \text { s head } \\ -1, & \text { if } v \text { is } e^{\prime} \text { s tail } \\ 0, & \text { otherwise }\end{cases}
$$

and $W \in \mathbb{R}^{m \times m}$ is the diagonal matrix with $W(e, e)=w_{e}$.

### 8.3 Preliminaries

### 8.3.1 Useful facts and tools

This section, we provide some useful tools.

Fact 8.3.1. For any two square matrices $A$ and $B$, we have

$$
\frac{1}{2}(A+B)^{2}+\frac{1}{2}(A-B)^{2}=A^{2}+B^{2}
$$

Proof.

$$
\begin{aligned}
\frac{1}{2}(A+B)^{2}+\frac{1}{2}(A-B)^{2} & =\frac{1}{2}(A+B)(A+B)+\frac{1}{2}(A-B)(A-B) \\
& =\frac{1}{2}\left(A^{2}+B A+A B+B^{2}\right)+\frac{1}{2}\left(A^{2}-B A-A B+B^{2}\right) \\
& =A^{2}+B^{2}
\end{aligned}
$$

which completes the proof.

Fact 8.3.2. For any two symmetric matrices

$$
(A-B)^{2} \preceq 2 A^{2}+2 B^{2} .
$$

Proof. Using Fact 8.3.1, we have

$$
(A+B)^{2}+(A-B)^{2}=2 A^{2}+2 B^{2}
$$

Because $A$ and $B$ are symmetric matrices, then $(A+B)^{2} \succeq 0$. It implies that

$$
(A-B)^{2} \preceq 2 A^{2}+2 B^{2},
$$

which completes the proof.

### 8.3.2 Strongly Rayleigh distributions

This section provides definitions related to Strongly Rayleigh distributions. For more details, we refer the readers to [BBL09, PP14].

Let $\mu: 2^{[n]} \rightarrow \mathbb{R}_{\geq 0}$ denote a probability distribution over $2^{[n]}$, and $\sum_{S \subseteq[n]} \mu(S)=1$.
Let $x_{1}, x_{2}, \cdots, x_{n}$ denote $n$ variables, we use $x$ to denote $\left(x_{1}, x_{2}, \cdots, x_{n}\right)$. For each set $S \subseteq[n]$, we define $x_{S}=\prod_{i \in S} x_{i}$. We define the generating polynomial for $\mu$ as follows

$$
f_{\mu}(x)=\sum_{S \subseteq[n]} \mu(S) \cdot x_{S}
$$

We say distribution $\mu$ is $k$-homogeneous if the polynomial $f_{\mu}$ is a homogeneous polynomial of degree $k$. In other words, for each $S \in \operatorname{supp}(\mu),|S|=k$.

We say a polynomial $p\left(x_{1}, x_{2}, \cdots, x_{n}\right)$ is stable, if $\operatorname{Im}\left(x_{i}\right)>0, \forall i \in[n]$, then $p\left(x_{1}, \cdots, x_{n}\right) \neq$ 0 . We say polynomial $p$ is real stable, it is stable and all of its coefficients are real. We say $\mu$ is a Strongly Rayleigh distribution if $f_{\mu}$ is a real stable polynomial.

Fact 8.3.3 (Conditioning on subset of coordinates). Consider a random vector $\left(\xi_{1}, \ldots, \xi_{m}\right) \in$ $\{0,1\}^{m}$ whose distribution is $k$-homogeneous Strongly Rayleigh. Suppose we get a binary vector $\mathbf{b}=\left(b_{1}, \ldots, b_{t}\right) \in\{0,1\}^{t}$ with $\|\mathbf{b}\|_{0}=l \leq k$, and we get a set $S \subset[m]$ with $|S|=$ $t$. Then conditional on $\xi_{S}=\mathbf{b}$, the distribution of $\xi_{[m] \backslash S}$ is $(k-l)$-homogeneous Strongly Rayleigh.

This fact tells us that if we condition on the value of some entries in the vector, the remaining coordinates still have a Strongly Rayleigh distribution.

Fact 8.3.4 (Stochastic Covering Property). Consider a random vector $\left(\xi_{1}, \ldots, \xi_{m}\right) \in\{0,1\}^{m}$ whose distribution is $k$-homogeneous Strongly Rayleigh. Suppose we are given an index $i \in$ $[m]$. Let $\xi^{\prime}=\xi_{[m] \backslash\{i\}}$ be the distribution on entries of $\xi$ except $i$. Let $\xi^{\prime \prime}$ be the distribution of $\xi_{[m] \backslash\{i\}}$ conditional on $\xi_{i}=1$. Then, there exists a coupling between $\xi^{\prime}$ and $\xi^{\prime \prime}$ (i.e. a joint distribution the two vectors), s.t. in every outcome of the coupling the value of $\xi^{\prime}$ can be obtained from the value of $\xi^{\prime \prime}$ by either changing a single from 0 to 1 or by leaving all entries unchanged.

This fact is known as the Stochastic Covering Property (see [PP14]). It gives us a convenient tool for relating the conditional distribution of a subset of the coordinates of the vector to the unconditional distribution.

Note that by Fact 8.3.3, the distribution of $\xi^{\prime \prime}$ used in Fact 8.3.4 is $k-1$ homogeneous. In contrast, the outcomes of $\xi^{\prime}$ may have $k$ or $k-1$ ones. Fact 8.3.4 tells us that we can pair up all the outcomes of the conditional distribution $\xi^{\prime \prime}$ with outcome of the unconditional distribution $\xi^{\prime}$ s.t. only a small change is required to make them equal. This tells us that the distribution is in some sense not changing too quickly under conditioning.

### 8.3.3 Random spanning trees

We provide the formal definition of random spanning tree in this section.
We use the same definitions about spanning trees as $\left[\mathrm{DKP}^{+} 17\right]$. Let $\mathcal{T}_{G}$ denote the set of all spanning subtrees of $G$. We now define a probability distribution on these trees.

Definition 8.3.1 ( $w$-uniform distribution on trees). Let $\mathcal{D}_{G}$ be a probability distribution
on $\mathcal{T}_{G}$ such that

$$
\operatorname{Pr}_{X \sim \mathcal{D}_{G}}[X=T] \propto \prod_{e \in T} w_{e}
$$

We refer to $\mathcal{D}_{G}$ as the $w$-uniform distribution on $\mathcal{T}_{G}$. When the graph $G$ is unweighted, this corresponds to the uniform distribution on $\mathcal{T}_{G}$. Crucially, random spanning tree distributions are Strongly Rayleigh, as shown in [BBL09].

Fact 8.3.5 (Spanning Trees are Strongly Rayleigh). In a connected weighted graph $G$, the $w$-uniform distribution on spanning trees is $(n-1)$-homogeneous Strongly Rayleigh.

Definition 8.3.2 (Effective Resistance). The effective resistance of a pair of vertices $u, v \in$ $V_{G}$ is defined as

$$
R_{\mathrm{eff}}(u, v)=b_{u, v}^{\top} L^{\dagger} b_{u, v},
$$

where $b_{u, v}$ is an all zero vector corresponding to $V_{G}$, except for entries of 1 at $u$ and -1 at $v$.

The a reference for following standard fact about random spanning trees can be found in $\left[\mathrm{DKP}^{+} 17\right]$.

Definition 8.3.3 (Leverage Score). The statistical leverage score, which we will abbreviate to leverage score, of an edge $e=(u, v) \in E_{G}$ is defined as

$$
l_{e}=w_{e} R_{\mathrm{eff}}(u, v)
$$

Fact 8.3.6 (Spanning Tree Marginals). The probability $\operatorname{Pr}[e]$ that an edge $e \in E_{G}$ appears in a tree sampled $w$-uniformly randomly from $\mathfrak{T}_{G}$ is given by

$$
\operatorname{Pr}[e]=l_{e},
$$

where $l_{e}$ is the leverage score of the edge $e$.

### 8.4 A Matrix Chernoff Bound for Strongly Rayleigh Distributions

We first define a mapping which maps an element into a psd matrix.

Definition 8.4.1 ( $Y$-operator). We use $\Gamma$ to denote $\left[m\right.$ ], we define a mapping $Y: \Gamma \rightarrow \mathbb{R}^{n \times n}$ such that $Y_{e}$ is a psd matrix and $\left\|Y_{e}\right\| \leq R$.

Throughout this section, we will use $\xi \in\{0,1\}^{m}$ to denote a random length $m$ boolean vector whose distribution is $k$-homogeneous Strongly Rayleigh. For any set $S \subseteq[m$ ], we use $\xi_{S}$ to denote the length $|S|$ vector that only chooses the entry from indices in $S$.

We will frequently need to work with a different representation of the random variable $\xi$. We use $\gamma$ to denote this second representation. The random variable $\gamma$ is composed of a sequence of $k$ random indices $\gamma_{1}, \gamma_{2}, \cdots, \gamma_{k}$, each of which takes a value $e_{1}, e_{2}, \cdots, e_{k} \in[m]$. The indices give the locations of the ones in $\xi$, i.e. in an outcome of the two variables $(\xi, \gamma)$, we always have $\xi_{\left\{\gamma_{1}, \gamma_{2}, \cdots, \gamma_{k}\right\}}=\mathbf{1}$ and $\xi_{[m] \backslash\left\{\gamma_{1}, \gamma_{2}, \cdots, \gamma_{k}\right\}}=\mathbf{0}$. Additionally, we want to ensure that the distribution of $\gamma$ is invariant under permutation: This can clearly be achieved by starting with any distribution for $\gamma$ that satisfies the coupling with $\xi$ and the applying a uniformly random permutation to reorder the $k$ indices of $\gamma$ (see [PP14] for a further discussion).

For convenience, for each $i \in[k]$, we define $\gamma_{\leq i}$ and $\gamma_{\geq i}$ as abbreviated notation for

$$
\gamma_{1}, \gamma_{2}, \cdots, \gamma_{i} \text { and } \gamma_{i}, \gamma_{i+1}, \cdots, \gamma_{k}
$$

respectively. Let $S=\left\{e_{1}, e_{2}, \cdots, e_{i}\right\} \subset[m]$ be one possible assignment for indices of a subset of the ones in $\xi$, (we require $i \leq k$ ). Then the distribution of $\xi_{[m] \backslash S}$ conditional on $\xi_{S}=\mathbf{1}$ is the same as the the distribution of $\xi_{[m] \backslash S}$ conditional on $\left(\gamma_{1}, \gamma_{2}, \cdots, \gamma_{i}\right)=\left(e_{1}, e_{2}, \cdots, e_{i}\right)$. In
other words, in terms of the resulting distribution of $\xi_{[m] \backslash S}$, it is equivalent to condition on either $\gamma_{\leq i}$ or $\xi_{\gamma_{\leq i}}=\mathbf{1}$. We define matrix $Z \in \mathbb{R}^{n \times n}$ as follows.

Definition 8.4.2 (Z). Let $Z$ denote $\sum_{e \in \Gamma} \xi_{e} \cdot Y_{e}$ where $\|\xi\|_{0}=k$. Due to the relationship between $\xi$ and $\gamma$, we can also write $Z$ as

$$
Z=\sum_{i=1}^{k} Y_{\gamma_{i}}
$$

For simplicity, for each $i \in[k]$, we define $Z_{\leq i}$ and $Z_{\geq i}$ as follows,

$$
Z_{\leq i}=\sum_{j=1}^{i} Y_{\gamma_{j}} \text { and } Z_{\geq i}=\sum_{j=i}^{k} Y_{\gamma_{j}}
$$

We define a series of matrices $M_{i} \in \mathbb{R}^{n \times n}$ as follows
Definition 8.4.3 ( $M_{i}$, martingale). We define $M_{0}=\mathbb{E}[Z]$. For each $i \in\{1,2, \cdots, k-1\}$, we define $M_{i}$ as follows

$$
M_{i}=\underset{\gamma \geq i+1}{\mathbb{E}}\left[Z \mid \gamma_{1}, \cdots, \gamma_{i}\right]
$$

It is easy to see that

$$
\underset{\gamma_{i+1}}{\mathbb{E}}\left[M_{i+1}\right]=M_{i},
$$

which implies

$$
\begin{equation*}
\underset{\gamma_{i+1}}{\mathbb{E}}\left[M_{i+1}-M_{i} \mid \gamma_{1}, \cdots, \gamma_{i}\right]=0 \tag{8.1}
\end{equation*}
$$

Note that we can split $Z$ up as

$$
\begin{align*}
Z & =\sum_{j=i+2}^{k} Y_{\gamma_{j}}+Y_{\gamma_{i+1}}+\sum_{j=1}^{i} Y_{\gamma_{j}} \\
& =Z_{\geq i+2}+Y_{\gamma_{i+1}}+Z_{\leq i} \tag{8.2}
\end{align*}
$$

And similarly $Z=Z_{\geq i+1}+Z_{\leq i}$.
In order to relate $M_{i}$ and $M_{i+1}$, we will consider a fresh copy of $\gamma_{\geq i+1}$ which we denote by $\widehat{\gamma}_{\geq i+1}$. We denote the corresponding fresh copy of $Z_{\geq i+1}$, by $\widehat{Z}_{\geq i+1}$. We can now give an equivalent definition of $M_{i}$ in terms of the expectation over $\widehat{\gamma}_{\geq i+1}$, while $M_{i+1}$ is still defined in terms of the expectation over $\gamma_{\geq i+2}$, so that

$$
\begin{equation*}
M_{i}=\underset{\widehat{\gamma} \geq i+1}{\mathbb{E}}\left[\widehat{Z}_{\geq i+1}+Z_{\leq i} \mid \gamma_{1}, \cdots, \gamma_{i}\right] \text { and } M_{i+1}=\underset{\gamma \geq i+2}{\mathbb{E}}\left[Z_{\geq i+2}+Y_{\gamma_{i+1}}+Z_{\leq i} \mid \gamma_{1}, \cdots, \gamma_{i}, \gamma_{i+1}\right] \tag{8.3}
\end{equation*}
$$

Note that both still depend on the same $\gamma_{\leq i}$ vector, and $M_{i+1}$ depends on $\gamma_{i+1}$, but $M_{i}$ does not. So far, we have simply introduced a sligtly different notation for $M_{i}$, since the expectation operation ensures that the value of $M_{i}$ is unchanged.

We let $\xi \in\{0,1\}^{m}$ denote the binary vector indicating the positions of indices $\gamma_{1}, \cdots, \gamma_{i}, \gamma_{i+1}, \cdots, \gamma_{k}$. while letting $\widehat{\xi} \in\{0,1\}^{m}$ indicate the positions of indices of $\gamma_{1}, \cdots, \gamma_{i}$, $\widehat{\gamma}_{i+1}, \cdots, \widehat{\gamma}_{k}$.

Note that $k$-homogeneous Strongly Rayleigh implies the stochastic covering property. By Fact 8.3.4, the stochastic covering property implies that a coupling exists s.t. adding either one or no extra "ones" to the vector $\xi_{[m] \backslash \gamma_{i+1}}$ results in the vector $\widehat{\xi}_{[m] \backslash \gamma_{i+1}}$. But, since $\xi_{[m] \backslash \gamma_{i+1}}$ is $k-1$ homogenous and $\widehat{\xi}_{[m]}$ is $k$-homogenous, we can conclude that $\widehat{\xi}_{[m] \backslash \gamma_{i+1}}$ is obtained from $\xi_{[m] \backslash \gamma_{i+1}}$ by adding no ones, if and only $\widehat{\xi}_{[m]}$ has a one at index $\gamma_{i+1}$. From this we conclude a more helpful form of stochastic covering: we can construct an index $\widetilde{\gamma}_{i+1}$ and a coupling s.t. conditional on $\gamma_{i+1}$, the indices $\widetilde{\gamma}_{i+1}, \gamma_{i+2}, \cdots, \gamma_{k}$ have the same distribution as $\widehat{\gamma}_{i+1}, \widehat{\gamma}_{i+2}, \cdots, \widehat{\gamma}_{k}$. Thus

$$
\begin{equation*}
Z_{\geq i+2}+Y_{\widetilde{\gamma}_{i+1}}=\widehat{Z}_{\geq i+1} \tag{8.4}
\end{equation*}
$$

We define $X_{i+1}=M_{i+1}-M_{i}$ and then by Eq. (8.1)

$$
\begin{equation*}
\underset{\gamma_{i+1}}{\mathbb{E}}\left[X_{i+1} \mid \gamma_{1}, \cdots, \gamma_{i}\right]=0 \tag{8.5}
\end{equation*}
$$

Then we can rewrite $X_{i+1}$ in the following way,
Claim 8.4.1. Let $\mathcal{D}_{\gamma_{\leq i+1}}$ denote the coupling distribution between $Z_{\geq i+2}$ and $Y_{\widetilde{\gamma}_{i+1}}$ such that Eq. (8.4) holds. Then

$$
X_{i+1}=Y_{\gamma_{i+1}}-\underset{\left(Z_{\geq i+2}, Y_{\gamma_{i+1}}\right) \sim \mathcal{D}_{\gamma_{\leq i+1}}}{\mathbb{E}}\left[Y_{\widetilde{\gamma}_{i+1}} \mid \gamma_{\leq i+1}\right] .
$$

Proof. Note that in the following proof, we should think of $\gamma_{\leq i+1}$ as fixed.

$$
\begin{align*}
X_{i+1} & =M_{i+1}-M_{i} \\
& =\underset{\gamma \geq i+2}{\mathbb{E}}\left[Z \mid \gamma_{\leq i+1}\right]-\underset{\gamma \geq i+1}{\mathbb{E}}\left[Z \mid \gamma_{\leq i}\right] \\
& =\underset{\gamma \geq i+2}{\mathbb{E}}\left[Z_{\geq i+2}+Y_{\gamma_{i+1}}+Z_{\leq i} \mid \gamma_{\leq i+1}\right]-\underset{\widehat{\gamma} \geq i+1}{\mathbb{E}}\left[\widehat{Z}_{\geq i+1}+Z_{\leq i} \mid \gamma_{\leq i}\right] \\
& =\underset{\gamma \geq i+2}{\mathbb{E}}\left[Z_{\geq i+2}+Y_{\gamma_{i+1}}+Z_{\leq i} \mid \gamma_{\leq i+1}\right]-\underset{\left(Z_{\geq i+2}, Y_{\gamma_{i+1}}\right) \sim \mathcal{D}_{\gamma_{\leq i+1}}}{\mathbb{E}}\left[Z_{\geq i+2}+Y_{\widetilde{\gamma}_{i+1}}+Z_{\leq i} \mid \gamma_{\leq i+1}\right] \\
& =\underset{\left(Z_{\geq i+2}, Y_{\widetilde{\gamma}_{i+1}}\right) \sim \mathcal{D}_{\gamma_{\leq i+1}}}{\mathbb{E}}\left[Y_{\gamma_{i+1}}-Y_{\widetilde{\gamma}_{i+1}} \mid \gamma_{\leq i+1}\right] \tag{8.6}
\end{align*}
$$

where the first equality follows by definition of $X_{i+1}$, the second equality follows by Definition 8.4.3, the third equality follows by Eq. (8.3), the fourth equality follows by Eq. (8.4), and the fifth equality is by linearity of expectation and cancellation of terms that agree.

Once we condition on $\gamma_{\leq i}$ and $\gamma_{i+1}$ being fixed, then $Y_{\gamma_{i+1}}$ is also fixed. Thus in Eq. (8.6), we can move $Y_{\gamma_{i+1}}$ out of Expectation, so that the right hand side of Eq. (8.6) becomes

$$
\begin{equation*}
\underset{\left(Z_{\geq i+2}, Y_{\gamma_{i+1}}\right) \sim \mathcal{D}_{\gamma_{\leq i+1}}}{\mathbb{E}}\left[Y_{\gamma_{i+1}}-Y_{\widetilde{\gamma}_{i+1}} \mid \gamma_{\leq i+1}\right]=Y_{\gamma_{i+1}}-\underset{\left(Z_{\geq i+2}, Y_{\gamma_{i+1}}\right) \sim \mathcal{D}_{\gamma_{\leq i+1}}}{\mathbb{E}}\left[Y_{\widetilde{\gamma}_{i+1}} \mid \gamma_{\leq i+1}\right] . \tag{8.7}
\end{equation*}
$$

Fact 8.4.2. We condition on $\gamma_{\leq i+1}$. Let $\mathcal{D}_{\gamma_{\leq i+1}}$ denote the coupling distribution such that $Z_{\geq i+2}+Y_{\widetilde{\gamma}_{i+1}}=\widehat{Z}_{\geq i+1}$ holds. We define $U_{\gamma_{i+1}}$ as follows

$$
U_{\gamma_{i+1}}=\underset{\left(Z_{\geq i+2}, Y_{\gamma_{i+1}}\right) \sim \mathcal{D}_{\gamma_{\leq i+1}}}{\mathbb{E}}\left[Y_{\widetilde{\gamma}_{i+1}} \mid \gamma_{\leq i+1}\right] .
$$

Then, we have the following four properties,

$$
\begin{aligned}
& \text { (I) } \underset{\gamma_{i+1}}{\mathbb{E}}\left[U_{\gamma_{i+1}} \mid \gamma_{\chi_{i}}\right]=\underset{\gamma_{i+1}}{\mathbb{E}}\left[Y_{\gamma_{i+1}} \mid \gamma_{\leq i}\right], \\
& \text { (II) }\left\|Y_{\gamma_{i+1}}\right\| \leq R,\left\|U_{\gamma_{i+1}}\right\| \leq R, \\
& \text { (III) }\left\|Y_{\gamma_{i+1}}-U_{\gamma_{i+1}}\right\| \leq R, \\
& \text { (IV) } Y_{\gamma_{i+1}}^{2} \preceq R \cdot Y_{\gamma_{i+1}}, U_{\gamma_{i+1}}^{2} \preceq R \cdot U_{\gamma_{i+1}} .
\end{aligned}
$$

Proof. Proof of (I). We have

$$
\begin{aligned}
\underset{\gamma_{i+1}}{\mathbb{E}}\left[U_{\gamma_{i+1}} \mid \gamma_{\leq i}\right] & =\underset{\gamma_{i+1}}{\mathbb{E}}\left[\underset{\left(Z_{\geq i+2}, Y_{\gamma_{i+1}}\right) \sim \mathcal{D}_{\gamma_{\leq i+1}}}{\mathbb{E}}\left[Y_{\widetilde{\gamma}_{i+1}} \mid \gamma_{\leq i+1}\right] \mid \gamma_{\leq i}\right] \\
& =\underset{\gamma_{i+1}}{\mathbb{E}}\left[\underset{\left(Z_{\geq i+2}, Y_{\gamma_{\gamma_{i+1}}}\right) \sim \mathcal{D}_{\gamma_{\leq i+1}}}{\mathbb{E}}\left[Y_{\widetilde{\gamma}_{\gamma_{i+1}}} \mid \gamma_{\leq i+1}\right]-Y_{\gamma_{i+1}}+Y_{\gamma_{i+1}} \mid \gamma_{\leq i}\right] \\
& =\underset{\gamma_{i+1}}{\mathbb{E}}\left[-X_{i+1}+Y_{\gamma_{i+1}} \mid \gamma_{\leq i}\right] \\
& =\underset{\gamma_{i+1}}{\mathbb{E}}\left[-X_{i+1} \mid \gamma_{\leq i}\right]+\underset{\gamma_{i+1}}{\mathbb{E}}\left[Y_{\gamma_{i+1}} \mid \gamma_{\leq i}\right] \\
& =\underset{\gamma_{i+1}}{\mathbb{E}}\left[Y_{\gamma_{i+1}} \mid \gamma_{\leq i}\right]
\end{aligned}
$$

where the third step follows by Eq. (8.7) and Eq. (8.6), the fourth step follows by linearity of expectation, and the last step follows by $\mathbb{E}_{\gamma_{i+1}}\left[-X_{i+1} \mid \gamma_{\leq i}\right]=0$.

Proof of (II).

By definition of $Y$, we have $\left\|Y_{\gamma_{i+1}}\right\| \leq R$.

$$
\begin{aligned}
\left\|U_{\gamma_{i+1}}\right\| & =\left\|\underset{\left(Z_{\geq i+2}, Y_{\widetilde{\gamma}_{i+1}}\right) \sim \mathcal{D}_{\gamma_{\leq i+1}}}{\mathbb{E}}\left[Y_{\widetilde{\gamma}_{i+1}} \mid \gamma_{\leq i+1}\right]\right\| \\
& \leq \underset{\left(Z_{\geq i+2}, Y_{\gamma_{i+1}}\right) \sim \mathcal{D}_{\gamma_{\leq i+1}}}{\mathbb{E}}\left[\left\|Y_{\widetilde{\gamma}_{i+1}}\right\| \mid \gamma_{\leq i+1}\right] \\
& \leq R .
\end{aligned}
$$

Proof of (III). For any two PSD matrices $A$ and $B$, we have $\|A-B\| \leq \max (\|A\|,\|B\|)$. Because both $Y_{\gamma_{i+1}}$ and $U_{\gamma_{i+1}}$ are PSD matrices and $\max \left(\left\|Y_{\gamma_{i+1}}\right\|,\left\|U_{\gamma_{i+1}}\right\|\right) \leq R$, we get the desired property.

Proof of (IV).
It follows by (II) and that $Y_{\gamma_{i+1}}$ and $U_{\gamma_{i+1}}$ are both PSD matrices.

We can show

## Claim 8.4.3.

$$
\underset{\gamma_{i+1}}{\mathbb{E}}\left[\left(Y_{\gamma_{i+1}}-U_{\gamma_{i+1}}\right)^{2} \mid \gamma_{\leq i}\right] \preceq 4 R \cdot \underset{\gamma_{i+1}}{\mathbb{E}}\left[Y_{\gamma_{i+1}} \mid \gamma_{\leq i}\right]
$$

Proof.

$$
\begin{aligned}
& \underset{\gamma_{i+1}}{\mathbb{E}}\left[\left(Y_{\gamma_{i+1}}-U_{\gamma_{i+1}}\right)^{2} \mid \gamma_{\leq i}\right] \\
\preceq & \underset{\gamma_{i+1}}{\mathbb{E}}\left[2 Y_{\gamma_{i+1}}^{2}+2 U_{\gamma_{i+1}}^{2} \mid \gamma_{\leq i}\right] \\
\preceq & \underset{\gamma_{i+1}}{\mathbb{E}}\left[2 R \cdot Y_{\gamma_{i+1}}+2 R \cdot U_{\gamma_{i+1}} \mid \gamma_{\leq i}\right] \\
\preceq & \underset{\gamma_{i+1}}{\mathbb{E}}\left[4 R \cdot Y_{\gamma_{i+1}} \mid \gamma_{\leq i}\right],
\end{aligned}
$$

where the first step follows by Fact 8.3.2, and the second step follows by $U_{\gamma_{i+1}}^{2} \preceq R \cdot U_{\gamma_{i+1}}$ and $Y_{\gamma_{i+1}}^{2} \preceq R \cdot Y_{\gamma_{i+1}}$.

Lemma 8.4.4. Let $\mathbb{E}\left[\sum_{e \in \Gamma} \xi_{e} Y_{e}\right] \preceq \mu I$. For each $i \in\{1,2, \cdots, k\}$, we have

$$
\underset{\gamma_{i}}{\mathbb{E}}\left[Y_{\gamma_{i}} \mid \gamma_{\leq i-1}\right] \preceq \frac{1}{k+1-i} \mu I .
$$

Proof. We use 1 to denote a length $i-1$ vector where each entry is one. We can think of $\gamma_{\leq i-1}$ as having its values already set to some edges in $\Gamma$, for example $\gamma_{1}=e_{1}, \cdots, \gamma_{i-1}=e_{i-1}$. Note that all of the $e_{1}, \cdots, e_{i-1}$ must be distinct. Then we use $\Gamma \backslash \gamma_{\leq i-1}$ to denote $\Gamma \backslash\left\{e_{1}, \cdots, e_{i-1}\right\}$.

$$
\begin{aligned}
\mathbb{E}\left[Y_{\gamma_{i}} \mid \gamma_{\leq i-1}\right] & =\sum_{e \in \Gamma \backslash \gamma_{\leq i-1}} \operatorname{Pr}\left[\gamma_{i}=e \mid \gamma_{\leq i-1}\right] \cdot Y_{e} \\
& =\sum_{e \in \Gamma \backslash \gamma_{\leq i-1}} \frac{\operatorname{Pr}\left[\xi_{e}=1 \mid \gamma_{\leq i-1}\right]}{k-(i-1)} \cdot Y_{e} \\
& =\sum_{e \in \Gamma \backslash \gamma_{\leq i-1}} \frac{\operatorname{Pr}\left[\xi_{e}=1 \mid \xi_{\gamma_{\leq i-1}}=1\right]}{k-(i-1)} \cdot Y_{e} \\
& \preceq \sum_{e \in \Gamma \backslash \gamma_{\leq i-1}} \frac{\operatorname{Pr}\left[\xi_{e}=1\right]}{k-(i-1)} \cdot Y_{e} \\
& \preceq \sum_{e \in \Gamma} \frac{\operatorname{Pr}\left[\xi_{e}=1\right]}{k-(i-1)} \cdot Y_{e} \\
& =\frac{1}{k-(i-1)} \mathbb{E}\left[\sum_{e} \xi_{e} Y_{e}\right] \\
& \preceq \frac{1}{k+1-i} \mu I
\end{aligned}
$$

where the first step follows by definition of expectation, the second step follows by $\operatorname{Pr}\left[\gamma_{i}=\right.$ $\left.e \mid \gamma_{\leq i-1}\right]=\operatorname{Pr}\left[\xi_{e}=1 \mid \gamma_{\leq i-1}\right] /(k-(i-1))$, the third step follows because $\left[\cdot \mid \gamma_{\leq i-1}\right]$ is equivalent to $\left[\cdot \mid \xi_{\gamma \leq i-1}=\mathbf{1}\right]$, the fourth step follows by $\left(\operatorname{Pr}\left[\xi_{e}=1 \mid \xi_{\gamma \geq i-1}=\mathbf{1}\right] \leq \operatorname{Pr}\left[\xi_{e}=1\right]\right)$
from the Shrinking Marginals Lemma 8.1.7, the fifth step follows by relaxing $\Gamma \backslash \gamma_{\leq i-1}$, the sixth step follows by $\operatorname{Pr}\left[\xi_{e}=1\right]=\mathbb{E}\left[\xi_{e}\right]$ and linearity of expectation, and the last step follows by $\mathbb{E}\left[\sum_{e \in \Gamma} \xi_{e} Y_{e}\right] \preceq \mu I$.

Lemma 8.4.5. For each $i \in\{1,2, \cdots, k\}$

$$
\underset{\gamma_{i}}{\mathbb{E}}\left[X_{i}^{2} \mid \gamma_{\leq i-1}\right] \preceq 4 \mu R \frac{1}{k+1-i} I .
$$

Proof. It follows by combining Claim 8.4.3 and Lemma 8.4.4 directly.

The above lemma implies this corollary directly

## Corollary 8.4.6.

$$
\sum_{i=1}^{k} \underset{\gamma_{i}}{\mathbb{E}}\left[X_{i}^{2} \mid \gamma_{\leq i-1}\right] \preceq 10 \mu R \log k \cdot I .
$$

### 8.4.1 Main result

To prove our main theorem 8.1.1, we need a useful tool for random matrices: Freedman's inequality for matrices (Lemma A.2.1).

Proof. of Theorem 8.1.1.
We use $Y$ to denote $A$ and $\Gamma$ to denote $[m]$.
In order to use Lemma A.2.1, we first we define $W_{i}$ as follows

$$
W_{i}=\sum_{j=1}^{i} \underset{\gamma_{i}}{\mathbb{E}}\left[X_{i}^{2} \mid \gamma_{\leq i-1}\right]
$$

According to definition of $M_{i},\left\{M_{0}, M_{1}, M_{2} \cdots\right\}$ is a matrix martingale and $M_{k}-M_{0}=$ $\sum_{e} \xi_{e} A_{e}-\mathbb{E}\left[\sum_{e} \xi_{e} A_{e}\right]$.

We have proved the following facts,
The first one is, $\mathbb{E}_{\gamma_{i}}\left[X_{i} \mid \gamma_{\leq i-1}\right]=0$. It follows by Eq. (8.5)
The second one is

$$
\lambda_{\max }\left(X_{i}\right) \leq R
$$

It follows by combining Property (III) of Fact 8.4.2 and Claim 8.4.1.
The third one is

$$
\left\|W_{i}\right\| \leq \sigma^{2}, \forall i \in[k]
$$

where $\sigma^{2}=10 \mu R \log k$. It follows by Corollary 8.4.6.
Thus,

$$
\operatorname{Pr}\left[\lambda_{\max }\left(M_{k}-M_{0}\right) \geq \epsilon \mu\right] \leq n \exp \left(-\frac{(\epsilon \mu)^{2} / 2}{\sigma^{2}+R(\epsilon \mu) / 3}\right) .
$$

We have

$$
\begin{array}{rlr}
\frac{t^{2} / 2}{\sigma^{2}+R t / 3} & =\frac{\epsilon^{2} \mu^{2} / 2}{10 \mu R \log k+R \epsilon \mu / 3} \quad \quad \text { by choosing } t=\epsilon \mu \\
& =\frac{3 \epsilon^{2} \mu}{(60 \log k+2 \epsilon) R}
\end{array}
$$

Thus we prove one side of the bound. Since $\mathbb{E}_{\gamma_{i}}\left[-X_{i} \mid \gamma_{\leq i-1}\right]=0$ and $\mathbb{E}_{\gamma_{i}}\left[\left(-X_{i}\right)^{2} \mid \gamma_{\leq i-1}\right]=$ $\mathbb{E}_{\gamma_{i}}\left[X_{i}^{2} \mid \gamma_{\leq i-1}\right]$, then following the similar procedure as proving $\lambda_{\max }$, we have bound for $\lambda_{\min }$

$$
\operatorname{Pr}\left[\lambda_{\min }\left(M_{k}-M_{0}\right) \leq-\epsilon \mu\right] \leq n \exp \left(-\frac{3 \epsilon^{2} \mu}{(60 \log k+2 \epsilon) R}\right)
$$

Putting two sides of the bound together, we complete the proof.

### 8.5 Applications to Random Spanning Trees

In this section, we show how to use Theorem 8.1.1 to prove the bound for one random spanning and also summation of random spanning trees.

Proof. of Theorem 8.1.3.
Let $G=(V, E, w)$ be a undirected weighted graph, $w: E \rightarrow R$, which is connected. The Laplacian of $G$ is $L_{G}=\sum_{e \in E} w(e) b_{e} b_{e}^{\top}$.

Let $T \subseteq E$ be a random spanning tree of $G$ in the sense of Definition 8.3.1. Let the weights of the edges in $T$ be given by $w^{\prime}: T \rightarrow R$ where $w^{\prime}(e)=w(e) / l_{e}$, where $l_{e}$ is the leverage score of $e$ in $G$. Thus the Laplacian of the tree is $L_{T}=\sum_{e \in T} w^{\prime}(e) b_{e} b_{e}^{\top}=$ $\sum_{e \in T} \frac{w(e)}{l_{e}} b_{e} b_{e}^{\top}$. Then by Fact 8.3.6, $\operatorname{Pr}[e \in T]=l_{e}$, and hence $\mathbb{E}\left[L_{T}\right]=L_{G}$.

Note also that for all $e \in E,\left\|\left(L_{G}^{\dagger}\right)^{1 / 2} w(e) b_{e} b_{e}^{\top}\left(L_{G}^{\dagger}\right)^{1 / 2}\right\|=l_{e}$. Consider the random matrix $\left(L_{G}^{\dagger}\right)^{1 / 2} L_{T}\left(L_{G}^{\dagger}\right)^{1 / 2}$. The distribution of edge in the spanning tree can be seen as an $n-1$ homogeneous vector in $\{0,1\}^{m}$ where $m=|E|$. To apply Theorem 8.1.1, let $\xi_{e}$ be the $e$ th entry of this random vector, and

$$
A_{e}=\left(L_{G}^{\dagger}\right)^{1 / 2} w^{\prime}(e) b_{e} b_{e}^{\top}\left(L_{G}^{\dagger}\right)^{1 / 2}
$$

Note $A_{e} \succeq 0$. Now $\left\|A_{e}\right\|=1$ and $\mathbb{E}\left[\sum_{e} \xi_{e} A_{e}\right]=\mathbb{E}\left[\left(L_{G}^{\dagger}\right)^{1 / 2} L_{T}\left(L_{G}^{\dagger}\right)^{1 / 2}\right]=\left(L_{G}^{\dagger}\right)^{1 / 2} L_{G}\left(L_{G}^{\dagger}\right)^{1 / 2}=$ $\Pi=I-\frac{1}{n} \mathbf{1 1} \mathbf{1}^{\top}$, where we used in the last equality that the null space of the Laplacian of a connected graph is the span of the all ones vector. Thus, as each we get $\left\|\mathbb{E}\left[\sum_{e} \xi_{e} A_{e}\right]\right\|=1$ This means we can apply Theorem 8.1 .1 with $R=1, \mu=1$ and $\epsilon=100 \log n$ to whp. $\left\|\left(L_{G}^{\dagger}\right)^{1 / 2} L_{T}\left(L_{G}^{\dagger}\right)^{1 / 2}-\Pi\right\| \leq 100 \log n$.

As $L_{T}$ is a Laplacian, it has $\mathbf{1}$ in the null space, so can conclude that $\left(L_{G}^{\dagger}\right)^{1 / 2} L_{T}\left(L_{G}^{\dagger}\right)^{1 / 2} \preceq$ $100 \log n \Pi$. Hence $L_{T} \preceq \log n L_{G}$.

Proof. of Theorem 8.1.2.
The proof is similar to the proof of Theorem 8.1.3. Now we view the edges of $t=O\left(\epsilon^{-2} \log ^{2} n\right)$ independent random spanning trees as a $t(n-1)$-homogeneous Strongly Rayleigh Distribution a vector in $\{0,1\}^{t|E|}$. Note that the product of independent Strongly Rayleigh distributions is Strongly Rayleigh [BBL09]. Again we get $\left\|\mathbb{E}\left[\sum_{e} \xi_{e} A_{e}\right]\right\|=1$, but now we can take $R=\frac{1}{t}$, and hence we obtain the desired result by plugging into Theorem 8.1.1.

### 8.6 Lower bounds

### 8.6.1 Single spanning tree, low probability

The goal of this section is to prove Theorem 8.1.5. First, we recall a helpful fact estbliashed by Prüfer [Prü18].

Fact 8.6.1. If $T$ is a uniformly random spanning tree of the complete graph $G$ on $n$ vertices, the degree distribution of a fixed node $v$ in $T$ is $1+\operatorname{Binomial}(n-2,1 / n)$.

We now prove two claims that will serve as helpful tools, Claims 8.6.2 and 8.6.3.

Claim 8.6.2. Let $G$ be complete graph $K_{n}$ with $n \geq 4$, let $T$ denote a random spanning tree, the probability that at least one node of the $T$ has degree at least $b \log n$ is at least $2^{-b \log n \log (b \log n)-3}$.

Proof. By Fact 8.6.1, the degree distribution of a fixed node in $T$ is, $1+\operatorname{Binomial}(n-2,1 / n)$.
For a random variable $x$ sampled from $\operatorname{Binomial}(n-2,1 / n)$, we use $q_{i}$ to denote the probability that $x=i$.

Let $p=1 / n$. We consider $q_{b \log n}$, which is

$$
\begin{array}{rlr}
q_{b \log n} & =\binom{n-2}{b \log n} \cdot p^{b \log n} \cdot(1-p)^{n-2-b \log n} \\
& =\binom{n-2}{b \log n} \cdot(1 / n)^{b \log n} \cdot(1-1 / n)^{n-2-b \log n} \quad \text { by } p=1 / n \\
& \geq((n-2) /(b \log n))^{b \log n} \cdot(1 / n)^{b \log n} \cdot(1-1 / n)^{n-2-b \log n} \\
& =(b \log n)^{-b \log n} \cdot((n-2) / n)^{b \log n} \cdot(1-1 / n)^{n-2-b \log n} \\
& \geq(b \log n)^{-b \log n} \cdot(1-2 / n)^{b \log n} \cdot(1-2 / n)^{n-2-b \log n} \\
& =(b \log n)^{-b \log n} \cdot(1-2 / n)^{n-2} \\
& \geq(b \log n)^{-b \log n} \cdot \frac{1}{e^{2}} \\
& \geq 2^{-b \log n \cdot \log (b \log n)-3} .
\end{array}
$$

where the seventh step follows by $(1-2 / n)^{n-2} \geq 1 / e^{2}$ when $n \geq 4$. Then the desired probability is

$$
\sum_{i=b \log n}^{n} q_{i} \geq 2^{-b \log n \cdot \log (b \log n)-3}
$$

Claim 8.6.3. Let $G$ be a complete graph $K_{n}$, let $T$ denote a random spanning tree, if $T$ has a node with degree at least d, then the inverse leverage score weighted Laplacian of the tree satisfies

$$
L_{T} \npreceq(d / 2) \cdot L_{G} .
$$

Proof. There are $n(n-1) / 2$ edges in the graph $G$. Let $l_{e}$ denote the leverage of each $e \in G$. The the sum of the leverage scores is $\sum_{e \in G} l_{e}=n-1$, e.g. see [SS11]. Since all the edges in the graph $G$ are symmetric, we have $l_{e}=2 / n$ for all edge $e$ in $G$.

Let $v$ denote a fixed node in graph $G$ and let $d$ be the degree of $v$. Let $L_{v}$ denote the Laplacian matrix of the subgraph of $T$ consisting of edges incident on $v$, i.e. the star of $d+1$ nodes with $v$ at the center, and with edge weights as in $T$ (which differ from those in $G)$. We should think of $L_{v}$ as a $n \times n$ matrix with only $d+1$ nonzeros on the diagonal.

Observe that $\lambda_{\max }\left(L_{v}\right) \leq \frac{n}{2} \lambda_{\max }\left(L_{K_{d+1}}\right) \leq \frac{n}{2}(d+1)$. We can also exhibit a unit vector $x=\frac{1}{\sqrt{d^{2}+d}}(d,-1, \cdots,-1)$, for which $x^{\top} L_{v} x=\frac{n}{2}(d+1)$ which implies that $\lambda_{\max }\left(L_{v}\right) \geq$ $\frac{n}{2}(d+1)$. Therefore $\lambda_{\max }\left(L_{v}\right)=\frac{n}{2}(d+1)$.

We can split the $L_{T}$ into two parts,

$$
L_{T}=L_{v}+L_{T \backslash v}
$$

and both parts are PSD matrices. We also know that $\lambda_{\max }\left(L_{G}\right)=n$. Thus,

$$
L_{T} \npreceq(d / 2) \cdot L_{G} .
$$

Proof. of Theorem 8.1.5.
The proof is a direct combination of Claim 8.6.2 and Claim 8.6.3.
The approximation ratio is

$$
\frac{d}{2}=\frac{b \log n}{2}=10 \log n
$$

where the last step follows by choosing $b=20$.

Then the probability is

$$
\begin{aligned}
2^{-b \log n \log (b \log n)-3} & =2^{-b \log n \log \log n-b \log n \log b-3} \\
& \geq 2^{-20 \log n \log \log n-20 \log n \log 20-3} \\
& \geq 2^{-150 \log n \log \log n}
\end{aligned}
$$

where the last step follows by $\log \log n \geq 1$ when $n \geq 4$.

### 8.6.2 Single spanning tree, high probability

Proof. of Theorem 8.1.4.
Let $C=4$, and let $\delta=1 /(C \log \log n)$. Note we have assumed $n \geq 2^{2^{6}}$, which ensures that $\delta<0.05$.

We constuct a graph of size $n$ as a union of $n^{1-\delta}$ cliques of size $n^{1-\delta}$ that are disjoint except they all share one central vertex. Applying Claim 8.6.2 with $n$ replaced by $n^{\delta}$, and $b=1$, and assuming $n^{\delta} \geq 8$, we get that for each clique, the probability that at least one node has degree at least $\log n^{\delta}$ in $T$ is at least

$$
2^{-\delta c_{0} \log n \cdot \log (\delta \log n)}
$$

where $c_{0}=2$. We can lower bound this probability:

$$
\begin{aligned}
2^{-\delta c_{0} \log n \cdot \log (\delta \log n)} & =2^{-\frac{c_{0} \log n}{C \log \log n} \log \left(\frac{1}{C \log \log n} \log n\right)} \\
& \geq 2^{-\frac{c_{0} \log n}{C \log \log n} \log \log n} \\
& =2^{-\left(c_{0} \log n\right) / C} \\
& =n^{-c_{0} / C}
\end{aligned}
$$

where the first step follows by $\delta=1 /(C \log \log n)$.
The probability that at least one node in $G$ has degree at least $\log n^{\delta}$ in $T$ is at least

$$
\begin{aligned}
1-\left(1-2^{-\delta c_{0} \log n \log (\delta \log n)}\right)^{n^{1-\delta}} & =1-\left(1-\frac{1}{n^{c_{0} / C}}\right)^{n^{1-\delta}} \\
& =1-\left(1-\frac{1}{n^{c_{0} / C}}\right)^{n^{c_{0} / C} \cdot \frac{n^{1-\delta}}{n_{0} / C}} \\
& \geq 1-(1 / e)^{n^{1-\delta-c_{0} / C}} \\
& \geq 1-(1 / e)^{n^{0.4}}
\end{aligned}
$$

where the last step follows by $\delta \leq 0.05, c_{0}=2$ and $C=4$. Thus, we have the desired probability.

Using Claim 8.6.3, we have the approximation ratio

$$
\frac{d}{2}=\frac{\log \left(n^{\delta}\right)}{2}=\frac{\delta \log n}{2}=\frac{\log n}{2 C \log \log n}=\frac{\log n}{8 \log \log n}
$$

Note that we still need to make sure $n^{\delta} \geq 8$, which is implied by $2^{\log n / 4 \log \log n} \geq 2^{3}$ which is equivalent to $\log n \geq 12 \log \log n$. This holds for all $n \geq 2^{2^{6}}$ : At $n=2^{2^{6}}, \log n=$ $2^{6}>60=12 \cdot 5=12 \log \log n$, and as $n$ grows, the left hand side grows faster than the right hand side.

### 8.6.3 Sum of a batch of spanning trees

Proof. of Theorem 8.1.6.
We constuct a graph of size $n$ as a union of $n^{1-\delta}$ cliques of size $n^{1-\delta}$ that are disjoint except they all share one central vertex. The parameter $\delta \in(0,1)$ will be decided later.

We use $H$ to denote the graph formed by a collection of trees $T_{1}, T_{2}, \cdots, T_{t}$. Let $L_{H}$ denote the Laplacian matrix of new graph $H$.

We use $\operatorname{deg}(v)$ to denote the degree of a vertex $v$. We use wdeg $(v)$ to denote weighted degree (after re-weighting). In the original graph $G$ and the new graph $H$, we have for each vertex $v$, that

$$
\operatorname{wdeg}_{H}(v)=\frac{d}{2 t} \operatorname{deg}_{H}(v), \text { and } \operatorname{wdeg}_{G}(v)=\operatorname{deg}_{G}(v)
$$

By our construction of the graph, it is easy to see that $\operatorname{deg}_{G}(v)=n^{\delta}$ for all vertices $v$ except the special central vertex that appears in all the cliques. Let $d=n^{\delta}$. Let $\xi_{1}$ denote the event that there exists a vertex $x \in V$ such that

$$
\operatorname{wdeg}_{H}(x)>(1+\epsilon) \operatorname{wdeg}_{G}(x)
$$

and let $\xi_{2}$ denote that there exists a vertex $y \in V$, such that

$$
\operatorname{wdeg}_{H}(y)<(1-\epsilon) \operatorname{wdeg}_{G}(y)
$$

We want to show that events $\xi_{1}$ and $\xi_{2}$ both occur simultaneously with probability at least $1-e^{-n^{0.39}}$, which implies absence of spectral $(1 \pm \epsilon)$ approximation, as desired. We first bound the probability of $\xi_{1}$. Note that

$$
\begin{aligned}
\operatorname{Pr}\left[\operatorname{wdeg}_{H}(v) \geq(1+\epsilon) \operatorname{wdeg}_{G}(v)\right] & =\operatorname{Pr}\left[\frac{d}{2 t} \operatorname{deg}_{H}(v) \geq(1+\epsilon) \operatorname{deg}_{G}(v)\right] \\
& \geq \operatorname{Pr}\left[\frac{d}{2 t} \operatorname{deg}_{H}(v) \geq(1+\epsilon) d\right] \\
& =\operatorname{Pr}\left[\operatorname{deg}_{H}(v) \geq 2 t(1+\epsilon)\right] \\
& =\operatorname{Pr}\left[\operatorname{deg}_{H}(v)-t \geq t(1+2 \epsilon)\right] .
\end{aligned}
$$

By Fact 8.6.1, the degree of a fixed node in $T_{i}$ is distributed as $1+\operatorname{Binomial}(d-2,1 / d)$. Then as $H$ is a union of independent spanning trees, the degree in $H$ of a fixed node is distributed as $t+\operatorname{Binomial}(t(d-2), 1 / d)$.

For a random variable $x$ sampled from $\operatorname{Binomial}(t(d-2), 1 / d)$, we know that $\mathbb{E}[x]=$ $t(1-2 / d)$. For $\epsilon>5 / d$

$$
\begin{array}{rlrl}
t(1+2 \epsilon) & =t(1-2 / d)+(2 t / d-2 \epsilon t+8 \epsilon t / d)+4 \epsilon t(1-2 / d) & \\
& \leq t(1-2 / d)+(10 t / d-2 \epsilon t)+4 \epsilon t(1-2 / d) & & \text { by } \epsilon \leq 1 \\
& \leq t(1-2 / d)+4 \epsilon t(1-2 / d) & \text { by } \epsilon>5 / d .
\end{array}
$$

So it suffices to calculate the probability that

$$
\begin{equation*}
x \geq t(1-2 / d)+4 \epsilon t(1-2 / d) \tag{8.8}
\end{equation*}
$$

For any $k \geq 10, p \in(0,1 / 2), \epsilon \in(0,1 / 2)$ with $\epsilon^{2} p k \geq 3$, using Lemma A.1.8, we can prove the probability that Eq. (8.8) holds is at least $2^{-c \epsilon^{2} k p}$, where $c=9$. We choose $k=t(d-2)$ and $p=1 / d$, and get that this probability is at least $2^{-c \epsilon^{2} t(d-2) / d}$. Now, the probability that event $\xi_{1}$ holds is at least

$$
1-\left(1-2^{-c \epsilon^{2} t(1-2 / d)}\right)^{n^{1-\delta}}
$$

We have

$$
\begin{aligned}
2^{-c \epsilon^{2} t(1-2 / d)} & \geq 2^{-c \epsilon^{2} t} \\
& \geq 1 / n^{0.5}
\end{aligned}
$$

where the last step follows by $t \leq 0.5 \log n /\left(c \epsilon^{2}\right)$.

Thus, we have

$$
\begin{aligned}
1-\left(1-2^{-c \epsilon^{2} t(1-2 / d)}\right)^{n^{1-\delta}} & \geq 1-\left(1-1 / n^{0.5}\right)^{n^{1-\delta}} \\
& \geq 1-e^{-n^{1-\delta-0.5}}
\end{aligned}
$$

$$
=1-e^{-n^{0.4}} \quad \text { by } \delta=0.1
$$

We summarize the conditions for $\epsilon$ :

$$
\begin{array}{rlr}
\epsilon & \geq \max (5 / d, 3 / \sqrt{p n}) & \\
& =\max \left(5 / n^{0.1}, 3 / \sqrt{p n}\right) & \text { by } d=n^{0.1} \\
& \geq \max \left(5 / n^{0.1}, 5 / \sqrt{t}\right) &
\end{array}
$$

Since we choose $t=0.05 \epsilon^{-2} \log n$, then as long as $\log n \geq 100$ we have $\epsilon \geq 5 / \sqrt{t}$.
Similarly, we can control the probability of event $\xi_{2}$ similarly, completing the proof.

### 8.7 Shrinking Marginals Lemma

Lemma 8.7.1 (Restatement of Lemma 8.1.7, Shrinking Marginals). Suppose $\left(\xi_{1}, \ldots, \xi_{m}\right) \in$ $\{0,1\}^{m}$ is a random vector of $\{0,1\}$ variables whose distribution is $k$-homogeneous and Strongly Rayleigh, then any set $S \subseteq[m]$ with $|S| \leq k$ for all $j \in[m] \backslash S$

$$
\operatorname{Pr}\left[\xi_{j}=1 \mid \xi_{S}=\mathbf{1}_{S}\right] \leq \operatorname{Pr}\left[\xi_{j}=1\right]
$$

Proof. Note that by an immediate consequence of negative association, for any pair $i, j \in[m]$, with $i \neq j$,

$$
\operatorname{Pr}\left[\xi_{j}=1 \mid \xi_{i}=1\right] \leq \operatorname{Pr}\left[\xi_{j}=1 \mid \xi_{i}=0\right] .
$$

Hence

$$
\begin{aligned}
\operatorname{Pr}\left[\xi_{j}=1 \mid \xi_{i}=1\right] & \leq \operatorname{Pr}\left[\xi_{j}=1 \mid \xi_{i}=1\right] \cdot \operatorname{Pr}\left[\xi_{i}=1\right]+\operatorname{Pr}\left[\xi_{j}=1 \mid \xi_{i}=0\right] \cdot\left(1-\operatorname{Pr}\left[\xi_{i}=1\right]\right) \\
& =\operatorname{Pr}\left[\xi_{j}=1\right]
\end{aligned}
$$

By [BBL09], the distribution of $\xi_{[m] \backslash\{i\}} \in\{0,1\}^{m-1}$ conditional on $\xi_{i}=1$ is Strongly Rayleigh.
With loss of generality, let us order the indices s.t. $S=\{1, \ldots, s\}$, where $s \leq k$. We use $[i]$ to denote $\{1,2, \cdots, i\}$. Using the above observations, we can now prove the lemma by induction. The induction hypothesis at the $i$-th step (where $i \leq k$ ), is that the following two statements are true.

1. $\forall j \in\{i+1, \ldots, m\} \cdot \operatorname{Pr}\left[\xi_{j}=1 \mid \xi_{[i]}=1\right] \leq \operatorname{Pr}\left[\xi_{j}=1\right]$
2. The distribution of the vector of random variables $\xi_{m \backslash[i]} \in\{0,1\}^{m-i}$ conditional on $\xi_{[i]}=1$ is Strongly Rayleigh.

## Chapter 9

## Four Derivation Suffice for Rank 1 Matrix

We prove a matrix discrepancy bound that strengthens the famous Kadison-Singer result of Marcus, Spielman, and Srivastava. Consider any independent scalar random variables $\xi_{1}, \ldots, \xi_{n}$ with finite support, e.g. $\{ \pm 1\}$ or $\{0,1\}$-valued random variables, or some combination thereof. Let $u_{1}, \ldots, u_{n} \in \mathbb{C}^{m}$ and

$$
\sigma^{2}=\left\|\sum_{i=1}^{n} \mathbb{V}\left[\xi_{i}\right]\left(u_{i} u_{i}^{*}\right)^{2}\right\|
$$

Then there exists a choice of outcomes $\epsilon_{1}, \ldots, \epsilon_{n}$ in the support of $\xi_{1}, \ldots, \xi_{n}$ s.t.

$$
\left\|\sum_{i=1}^{n} \mathbb{E}\left[\xi_{i}\right] u_{i} u_{i}^{*}-\sum_{i=1}^{n} \epsilon_{i} u_{i} u_{i}^{*}\right\| \leq 4 \sigma
$$

A simple consequence of our result is an improvement of a Lyapunov-type theorem of Akemann and Weaver.

### 9.1 Introduction

Discrepancy theory is an area of combinatorics that studies how well continuous objects can be approximated by discrete ones. It lies at the heart of numerous problems in mathematics and computer science [Cha00]. Although closely tied to probability theory, direct randomized approaches rarely yield the best bounds. In a classical formulation in discrepancy theory, we have $n$ sets on $n$ elements, and would like to two-color the elements so that each set has roughly the same number of elements of each color. Using a simple random coloring, it is an easy consequence of Chernoff's bound that there exists a coloring such that the discrepancy in all $n$ sets is $O(\sqrt{n \log n})$ [AS16]. However, in a celebrated result, Spencer showed that in fact there is a coloring with discrepancy at most $6 \sqrt{n}$ [Spe85].

Recently, there has been significant success in generalizing Chernoff [Che52], Hoeffding, Bernstein, and Bennett-type concentration bounds for scalar random variables to matrix-valued random variables [Rud99, AW02, Tro12]. Consider the following matrix concentration bound, which is a direct consequence of a matrix Hoeffding bound.

Theorem 9.1.1 ([Tro12]). Let $\xi_{i} \in\{ \pm 1\}$ be independent, symmetric random signs and $A_{1}, \ldots, A_{n} \in \mathbb{C}^{m \times m}$ be positive semi-definite matrices. Suppose $\max _{i \in[n]}\left\|A_{i}\right\| \leq \epsilon$ and $\left\|\sum_{i=1}^{n} A_{i}\right\| \leq 1$. Then,

$$
\operatorname{Pr}\left[\left\|\sum_{i=1}^{n} \xi_{i} A_{i}\right\| \geq t \sqrt{\epsilon}\right] \leq 2 m \exp \left(-t^{2} / 2\right)
$$

A consequence of this theorem is that with high probability

$$
\begin{equation*}
\left\|\sum_{i=1}^{n} \xi_{i} A_{i}\right\|=O(\sqrt{\log m}) \sqrt{\epsilon} \tag{9.1}
\end{equation*}
$$

The Kadison-Singer theorem of [MSS15b] is essentially equivalent to the following statement (which can readily be derived from the bipartition statement in [MSS15b]).

Theorem 9.1.2 ([MSS15b]). Let $u_{1}, \ldots, u_{n} \in \mathbb{C}^{m}$ and suppose $\max _{i \in[n]}\left\|u_{i} u_{i}^{*}\right\| \leq \epsilon$ and $\sum_{i=1}^{n} u_{i} u_{i}^{*}=I$. Then, there exists signs $\xi_{i} \in\{ \pm 1\}$ s.t.

$$
\left\|\sum_{i=1}^{n} \epsilon_{i} u_{i} u_{i}^{*}\right\| \leq O(\sqrt{\epsilon})
$$

Thus, for rank 1 matrices the theorem improves on the norm bound in Equation (9.1), by a factor $\sqrt{\log m}$, in a manner analogous to the improvement of Spencer's theorem over the bound based on the scalar Chernoff bound.

For random signings of matrices one can establish bounds in some cases that are much stronger than Theorem 9.1.1.

Theorem 9.1.3 ([Tro12]). Let $\xi_{i} \in\{ \pm 1\}$ be independent random signs, and let $A_{1}, \ldots, A_{n} \in$ $\mathbb{C}^{m \times m}$ be Hermitian matrices. Let $\sigma^{2}=\left\|\sum_{i=1}^{n} \mathbb{V}\left[\xi_{i}\right] A_{i}^{2}\right\|$. Then,

$$
\operatorname{Pr}\left[\left\|\sum_{i=1}^{n} \mathbb{E}\left[\xi_{i}\right] A_{i}-\sum_{i=1}^{n} \xi_{i} A_{i}\right\| \geq t \cdot \sigma\right] \leq 2 m \exp \left(-t^{2} / 2\right)
$$

From this theorem we deduce that with high probability

$$
\begin{equation*}
\left\|\sum_{i=1}^{n} \mathbb{E}\left[\xi_{i}\right] A_{i}-\sum_{i=1}^{n} \xi_{i} A_{i}\right\|=O(\sqrt{\log m}) \sigma \tag{9.2}
\end{equation*}
$$

Of course, this implies that there exists a choice of signs $\epsilon_{1}, \ldots, \epsilon_{n} \in\{ \pm 1\}$ such that

$$
\left\|\sum_{i=1}^{n} \mathbb{E}\left[\xi_{i}\right] A_{i}-\sum_{i=1}^{n} \epsilon_{i} A_{i}\right\|=O(\sqrt{\log m}) \sigma
$$

Our main result demonstrates that for rank-1 matrices, there exists a choice of signs with a stronger guarantee.

Theorem 9.1.4 (Main Theorem). Consider any independent scalar random variables $\xi_{1}, \ldots, \xi_{n}$ with finite support. Let $u_{1}, \ldots, u_{n} \in \mathbb{C}^{m}$ and

$$
\sigma^{2}=\left\|\sum_{i=1}^{n} \mathbb{V}\left[\xi_{i}\right]\left(u_{i} u_{i}^{*}\right)^{2}\right\|
$$

Then there exists a choice of outcomes $\epsilon_{1}, \ldots, \epsilon_{n}$ in the support of $\xi_{1}, \ldots, \xi_{n}$

$$
\left\|\sum_{i=1}^{n} \mathbb{E}\left[\xi_{i}\right] u_{i} u_{i}^{*}-\sum_{i=1}^{n} \epsilon_{i} u_{i} u_{i}^{*}\right\| \leq 4 \sigma .
$$

For rank 1 matrices our theorem improves on the norm bound in Equation (9.2), by a factor $\sqrt{\log m}$. Note for example, if $\xi_{i}$ is $\{ \pm 1\}$-valued, then $\mathbb{V}\left[\xi_{i}\right]=1-\mathbb{E}\left[\xi_{i}\right]^{2}$.

Specializing to centered random variables, we obtain the following corollary, which in a simple way generalizes the Kadison-Singer theorem, although with a slightly worse constant.

Corollary 9.1.5. Let $u_{1}, \ldots, u_{n} \in \mathbb{C}^{m}$ and

$$
\sigma^{2}=\left\|\sum_{i=1}^{n}\left(u_{i} u_{i}^{*}\right)^{2}\right\| .
$$

There exists a choice of signs $\epsilon_{i} \in\{ \pm 1\}$ such that

$$
\left\|\sum_{i=1}^{n} \epsilon_{i} u_{i} u_{i}^{*}\right\| \leq 4 \sigma
$$

Notice that if $\max _{i}\left\|u_{i} u_{i}^{*}\right\| \leq \epsilon$ and $\sum_{i=1}^{n} u_{i} u_{i}^{*}=I$, then $\sigma^{2} \leq \epsilon$, and we obtain Theorem 9.1.2 as consequence of the corollary.

Our Theorem 9.1.4 has multiple advantages over Theorem 9.1.2. It allows us to show existence of solutions close to the mean of arbitrarily biased $\pm 1$ random variables, instead
of only zero mean distributions. If some variable $\xi_{i}$ is extremely biased, its variance is correspondingly low, as $\mathbb{V}\left[\xi_{i}\right]=1-\mathbb{E}\left[\xi_{i}\right]^{2}$.

If only a small subset of the rank 1 matrices obtain the $\epsilon$ norm bound, and the rest are significantly smaller in norm, then we can have $\sigma^{2} \approx \epsilon^{2}$, so we prove a bound of $O(\epsilon)$ instead of the $O(\sqrt{\epsilon})$ bound of the Kadison-Singer theorem. On the other hand, when the problem is appropriately scaled, we always have $\sigma \geq \epsilon^{2}$, so the gap between the two results can never be more than a square root. Additionally, note that although the Kadison-Singer theorem requires $\sum_{i=1}^{n} u_{i} u_{i}^{*}=I$, a multi-paving $\operatorname{argument}^{1}$ can be used to instead relax this to $\left\|\sum_{i=1}^{n} u_{i} u_{i}^{*}\right\| \leq 1$.

Approximate Lyapunov Theorems. Marcus, Spielman and Srivastava resolved the Kadison-Singer problem by proving Weaver's conjecture [Wea04], which was shown to imply the Kadison-Singer conjecture. In [AW14], Akemann and Weaver prove a generalization of Weaver's conjecture [Wea04].

Theorem 9.1.6 ([AW14]). Let $u_{1}, \ldots, u_{n} \in \mathbb{C}^{m}$ such that $\left\|\sum_{i=1}^{n} u_{i} u_{i}^{*}\right\| \leq 1$ and $\max _{i}\left\|u_{i} u_{i}^{*}\right\| \leq$ $\epsilon$. For any $t_{i} \in[0,1]$ and $1 \leq i \leq n$, there exists a set of indices $S \subset\{1,2, \ldots, n\}$ such that

$$
\left\|\sum_{i \in S} u_{i} u_{i}^{*}-\sum_{i=1}^{n} t_{i} u_{i} u_{i}^{*}\right\|=O\left(\epsilon^{1 / 8}\right)
$$

Due to the classical Lyapunov theorem [Lya40] and its equivalent versions [Lin66], in their study of operator algebras, Akemann and Anderson [AA91] refer to a result as a Lyapunov theorem if the result states that for a convex set $C$, the image of $C$ under an affine

[^52]map is equal to the image of the extreme points of $C$. Theorem 9.1.6 is an approximate Lyapunov theorem as it can be interpreted as saying that the image of $[0,1]^{n}$ under the map
$$
f:\left(t_{1}, \ldots, t_{n}\right) \rightarrow \sum_{i=1}^{n} t_{i} u_{i} u_{i}^{*}
$$
can be approximated by the image of one of the vertices of the hypercube $[0,1]^{n}$. A corollary of our main result, Theorem 9.1.4, is the following strengthening of Theorem 9.1.6. This result greatly improves the $\epsilon$ dependence of the original Lyapunov-type theorem and provides a small explicit constant.

Corollary 9.1.7. Let $u_{1}, \ldots, u_{n} \in \mathbb{C}^{m}$ such that $\left\|\sum_{i=1}^{n} u_{i} u_{i}^{*}\right\| \leq 1$ and $\max _{i}\left\|u_{i} u_{i}^{*}\right\| \leq \epsilon$. For any $t_{i} \in[0,1]$ and $1 \leq i \leq n$, there exists a set of indices $S \subset\{1,2, \ldots, n\}$ such that

$$
\left\|\sum_{i \in S} u_{i} u_{i}^{*}-\sum_{i=1}^{n} t_{i} u_{i} u_{i}^{*}\right\|=2 \epsilon^{1 / 2}
$$

The corollary follows immediately from Theorem 9.1 .4 by choosing as the $\xi_{i}$ a set of independent $\{0,1\}$-valued random variables with means $t_{i} \in[0,1]$. Note then that $\mathbb{V}\left[\xi_{i}\right]=$ $t_{i}\left(1-t_{i}\right) \leq 1 / 4$, and so by the assumptions $\max _{i}\left\|u_{i} u_{i}^{*}\right\| \leq \epsilon$ and $\sum_{i=1}^{n} u_{i} u_{i}^{*} \preceq I$, we have $\sigma^{2} \leq \epsilon / 4$, and we obtain Theorem 9.1.2 as consequence of the corollary.

### 9.1.1 Related work and open questions

Our work is an extension of the the Kadison-Singer theorem of [MSS15b]. This result has a rich history of connections with theoretical computer science, and, we hope, may represent a step toward getting polynomial time algorithms for finding Kadison-Singer partitions.

Expanders are sparse graphs that exhibit good connectivity. One view on this is that an expander graph approximates a complete graph in some sense, e.g. the spectrum of the adjacency matrix or Laplacian matrix associated with the graph resembles that of the complete graph. Ramanujan graphs [Mar73, LPS88] are sparse $d$-regular graphs whose adjacency matrix eigenvalues essentially optimally approximates those of a complete graph among all $d$-regular graphs for a given degree $d$. These initial constructions of Ramanujan graphs relied on number-theoretic techniques and could not show existence of Ramanujan for all degrees and sizes.

Using the method of interlacing families of polynomials, [MSS15a, MSS18], finally showed the existence of bipartite ${ }^{2}$ Ramanujan graphs of all degrees and of all sizes. The existential result of [MSS15a] was turned into a polynomial time algorithm in [Coh16b]. The restriction to bipartite Ramanujan graphs in these papers arose from the fact that interlacing families most naturally control only the largest eigenvalue of a matrix, while Ramanujan graphs require both the largest and the smallest eigenvalue to be controlled. [MSS15a] overcame this obstacle by studying only bipartite graphs, whose adjacency matrix eigenvalues are symmetric around zero and hence controlling the largest eigenvalue ${ }^{3}$ implies also controlling the smallest. Our techniques for simultaneously controlling the smallest and largest eigenvalues of a matrix are very different from earlier interlacing family-based methods, and an intriguing open question is whether similar ideas can be used to construct non-bipartite Ramanujan graphs of all degrees and sizes.

[^53][ST04] gave nearly-linear time algorithms for producing sparse graphs (with nonuniform weights) that closely approximate a given input graph, in the sense that the spectral behavior of the Laplacian matrices associated with the graphs is very similar. A greatly simplified and stronger result was given in [SS11], still with a nearly linear time algorithm. When used to construct a sparse graph that spectrally approximates a complete graph of size $n$, the result of [SS11] has a worse average degree than a corresponding Ramanujan graph by a factor $\Theta(\log n)$. Later, this was improved to within a factor $O(1)$, still with a polynomial time algorithm, in [BSS12]. Finally, a series of papers turned this latter result into a nearly-linear time algorithm [LS18, LS17]. The resolution of the Kadison-Singer theorem was motivated by $[\mathrm{BSS} 12]^{4}$. A major open question in the area is whether a polynomial time algorithm for finding explicit solutions to the Kadison-Singer problem exists, and we hope that our strengthened form of the result may make the problem more tractable.

The Kadison-Singer theorem and our result can be interpreted as analogous to Spencer's theorem in a matrix setting, in recent years that has been tremendous progress in obtaining polynomial time algorithms for Spencer's theorem and related problems in other settings [Bou10, LM15, Rot17, BDG16, BG17, BDGL18, NDTTJ18]. Techniques developed in these papers have also found algorithmic uses in approximation algorithms beyond typical discrepancy statements [Rot13]. While interlacing families have notably been used to to bound integrality gaps for ATSP [AG14] algorithmic results based on these have been limited, with [Coh16b] providing the most notable example of a polynomial time algorithm based on interlacing families. It seems likely, however, that a deeper understanding of the limits of matrix rounding could have many applications in approximation algorithms.

[^54]
### 9.1.2 Our techniques

Our proof is based on the method of interlacing polynomials introduced in [MSS15a, MSS15b]. One difficulty in applying the method of interlacing polynomials is the inability to control both the largest and smallest eigenvalues of a matrix simultaneously. Various techniques and restrictions are used to overcome this problem in [MSS15a, MSS15b] (studying bipartite graphs, assuming isotropic position). We develop a seemingly more natural approach for simultaneously controlling both the largest and smallest roots of the matrices we consider. We study polynomials that can be viewed as expected characteristic polynomials, but are more easily understood as the expectation of a product of multiple determinants. The using a product of two determinants helps us bound the upper and lower eigenvalues both at the same time.

We introduce an analytic expression for the expected polynomials in terms of linear operators that use second order derivatives. This has the advantage of allowing us to gain stronger control over the movement of roots of polynomials under the linear operators we apply than those used by [MSS15b]. This is because movement of the roots now depends on the curvature of our polynomials in favorable way. Interestingly, linear operators containing second order derivatives also appear in the work of [AG14] but for different reasons. This lets us reuse one of their lemmas for bounding position of the roots of a real stable polynomial.

### 9.2 Preliminaries

We gather several basic linear algebraic and analytic facts in the following sections.

### 9.2.1 Linear Algebra

Lemma 9.2.1. Let $x \in \mathbb{R}^{n}$. Then

$$
\operatorname{det}\left(I-t x x^{*}\right)=1-t x^{*} x
$$

Fact 9.2.2 (Jacobi's Formula). For $A(t) \in \mathbb{R}^{n \times n}$ a function of $t$,

$$
\frac{d}{d t} \operatorname{det}(A(t))=\operatorname{det}(A(t)) \operatorname{tr}\left[A^{-1}(t) \frac{d}{d t} A(t)\right]
$$

### 9.2.2 Real Stability

Definition 9.2.1. A multivariate polynomial $p\left(z_{1}, \ldots, z_{n}\right) \in \mathbb{C}\left[z_{1}, \ldots, z_{n}\right]$ is stable if it has no zeros in the region $\left\{\left(z_{1}, \ldots, z_{n}\right): \Im\left(z_{i}\right)>0\right.$ for all $\left.1 \leq i \leq n\right\}$. $p$ is real stable if $p$ is stable and the coefficients of $p$ are real.

Lemma 9.2.3 (Corollary 2.8, [AOGSS18]). If $p \in \mathbb{R}\left[z_{1}, \ldots, z_{n}\right]$ is real stable, then for any $c>0$, so is

$$
\left(1-c \partial_{z_{i}}^{2}\right) p\left(z_{1}, \ldots, z_{n}\right)
$$

for all $1 \leq i \leq n$.
Lemma 9.2.4 (Proposition 2.4, $[\mathrm{BB} 08]$ ). If $A_{1}, \ldots, A_{n}$ are positive semidefinite symmetric matrices, then the polynomial

$$
\operatorname{det}\left(\sum_{i=1}^{n} z_{i} A_{i}\right)
$$

is real stable.

We also need that real stability is preserved under fixing variables to real values (see [Wag11, Lemma 2.4(d)]).

Proposition 9.2.5. If $p \in \mathbb{R}\left[z_{1}, \ldots, z_{m}\right]$ is real stable and $a \in \mathbb{R}$, then $\left.p\right|_{z_{1}=a}=p\left(a, z_{2}, \ldots, z_{m}\right) \in$ $\mathbb{R}\left[z_{2}, \ldots, z_{m}\right]$ is real stable.

### 9.2.3 Interlacing Families

We recall the definition and consequences of interlacing families from [MSS15a].
Definition 9.2.2. We say a real rooted polynomial $g(x)=C \prod_{i=1}^{n-1}\left(x-\alpha_{i}\right)$ interlaces the real rooted polynomial $f(x)=C^{\prime} \prod_{i=1}^{n}\left(x-\beta_{i}\right)$ if

$$
\beta_{1} \leq \alpha_{1} \leq \cdots \leq \alpha_{n-1} \leq \beta_{n} .
$$

Polynomials $f_{1}, \ldots, f_{k}$ have a common interlacing if there is a polynomial $g$ that interlaces each of the $f_{i}$.

The following lemma relates the roots of a sum of polynomials to those of a common interlacer.

Lemma 9.2.6 (Lemma 4.2, [MSS15a]). Let $f_{1}, \ldots, f_{k}$ be degree $d$ real rooted polynomials with positive leading coefficients. Define

$$
f_{\emptyset}:=\sum_{i=1}^{k} f_{i} .
$$

If $f_{1}, \ldots, f_{k}$ have a common interlacing then there exists an $i$ for which the largest root of $f_{i}$ is upper bounded by the largest root of $f_{\emptyset}$.

Definition 9.2.3. Let $S_{1}, \ldots, S_{n}$ be finite sets. For every $S \in \mathcal{F}$, we let $f_{S}(x)$ be a real rooted polynomial of degree $d$ with positive leading coefficient. For a choice of assignment $s_{1}, \ldots, s_{n} \in S_{1} \times \cdots \times S_{n}$, let $f_{s_{1}, \ldots, s_{n}}(x)$ be a real rooted degree $d$ polynomial with positive leading coefficient. For a partial assignment $s_{1}, \ldots, s_{k} \in S_{1} \times \cdots \times S_{k}$ for $k<n$, we define

$$
\begin{equation*}
f_{s_{1}, \ldots, s_{k}}:=\sum_{s_{k+1} \in S_{k+1}, \ldots, s_{n} \in S_{n}} f_{s_{1}, \ldots, s_{k}, s_{k+1}, \ldots, s_{n}} \tag{9.3}
\end{equation*}
$$

Note that this is compatible with our definition of $f_{\emptyset}$ from Definition 9.2.2. We say that the polynomials $\left\{f_{s_{1}, \ldots, s_{n}}\right\}$ form an interlacing family if for all $k=0, \ldots, n-1$ and all $s_{1}, \ldots, s_{k} \in S_{1} \times \cdots \times S_{k}$, the polynomials have a common interlacing.

The following lemma relates the roots of the interlacing family to those of $f_{\emptyset}$.

Lemma 9.2.7 (Theorem 4.4, [MSS15a]). Let $S_{1}, \ldots, S_{n}$ be finite sets and let $\left\{f_{s_{1}, \ldots, s_{n}}\right\}$ be an interlacing family. Then there exists some $s_{1}, \ldots, s_{n} \in S_{1} \times \cdots \times S_{n}$ so that the largest root of $f_{s_{1}, \ldots, s_{n}}$ is upper bounded by the largest root of $f_{\emptyset}$.

Finally, we recall a relationship between real-rootedness and common interlacings which has been discovered independently several times [DG94, Fel80, CS07].

Lemma 9.2.8. Let $f_{1}, \ldots, f_{k}$ be univariate polynomials of the same degree with positive leading coefficient. Then $f_{1}, \ldots, f_{k}$ have a common interlacing if and only if $\sum_{i=1}^{k} \alpha_{i} f_{i}$ is real rooted for all nonnegative $\alpha_{i}$ such that $\sum_{i} \alpha_{i}=1$.

### 9.3 Expected Characteristic Polynomial

Instead of working with the random polynomial $\operatorname{det}\left(x I-\sum_{i=1}^{n} \xi_{i} u_{i} u_{i}^{*}\right)$, we consider

$$
\operatorname{det}\left(x^{2} I-\left(\sum_{i=1}^{n} \xi_{i} u_{i} u_{i}^{*}\right)^{2}\right)=\operatorname{det}\left(x I-\sum_{i=1}^{n} \xi_{i} u_{i} u_{i}^{*}\right) \operatorname{det}\left(x I+\sum_{i=1}^{n} \xi_{i} u_{i} u_{i}^{*}\right) .
$$

Observe that the largest root $\lambda_{\max }$ of this polynomial is

$$
\begin{equation*}
\lambda_{\max }\left(\operatorname{det}\left(x^{2} I-\left(\sum_{i=1}^{n} \xi_{i} u_{i} u_{i}^{*}\right)^{2}\right)\right)=\left\|\sum_{i=1}^{n} \xi_{i} u_{i} u_{i}^{*}\right\| . \tag{9.4}
\end{equation*}
$$

We gather some results that will allow us to extract an analytic expression for the expected characteristic polynomial.

Lemma 9.3.1. For positive semidefinite ( $P S D$ ) matrices $M, N \in \mathbb{R}^{m \times m}, v \in \mathbb{R}^{m}$ and $\xi$ a random variable with zero mean and variance $\tau^{2}$,

$$
\begin{equation*}
\underset{\xi}{\mathbb{E}}\left[\operatorname{det}\left(M-\xi v v^{*}\right) \operatorname{det}\left(N+\xi v v^{*}\right)\right]=\left.\left(1-\frac{1}{2} \frac{d^{2}}{d t^{2}}\right)\right|_{t=0} \operatorname{det}\left(M+t \tau v v^{*}\right) \operatorname{det}\left(N+t \tau v v^{*}\right) \tag{9.5}
\end{equation*}
$$

Proof. We can assume that both $M$ and $N$ are positive definitive and hence invertible. The argument for the the positive semi-definite case follows by a continuity argument (using Hurwitz's theorem from complex analysis, see also [MSS15b]).

We show that the two sides of (9.5) are equivalent to the same expression. Beginning on the left hand side,

$$
\begin{aligned}
\mathbb{E}\left[\operatorname{det}\left(M-\xi v v^{*}\right)\right. & \left.\operatorname{det}\left(N+\xi v v^{*}\right)\right] \\
& =\operatorname{det}(M) \operatorname{det}(N) \mathbb{E}\left[\operatorname{det}\left(I-\xi M^{-1 / 2} v v^{*} M^{-1 / 2}\right) \operatorname{det}\left(I+\xi N^{-1 / 2} v v^{*} N^{-1 / 2}\right)\right] \\
& =\operatorname{det}(M) \operatorname{det}(N) \mathbb{E}\left[1+\xi b^{*} b-\xi a^{*} a-\xi^{2} a^{*} a b^{*} b\right] \\
& =\operatorname{det}(M) \operatorname{det}(N)\left(1-\tau^{2} a^{*} a b^{*} b\right)
\end{aligned}
$$

where $a:=M^{-1 / 2} v$ and $b:=N^{-1 / 2} v$. For the right hand side of (9.5),

$$
\begin{aligned}
\left.\left(1-\frac{1}{2} \frac{d^{2}}{d t^{2}}\right)\right|_{t=0} & \operatorname{det}\left(M+t \tau v v^{*}\right) \operatorname{det}\left(N+t \tau v v^{*}\right) \\
& =\left.\operatorname{det}(M) \operatorname{det}(N)\left(1-\frac{1}{2} \frac{d^{2}}{d t^{2}}\right)\right|_{t=0} \operatorname{det}\left(I+t \tau a a^{*}\right) \operatorname{det}\left(I+t \tau b b^{*}\right) \\
& =\operatorname{det}(M) \operatorname{det}(N)\left(1-\tau^{2} a^{*} a b^{*} b\right)
\end{aligned}
$$

where the last line follows from Lemma 9.2.1.

Centering our random variables and applying the previous lemma leads to the following corollary for non-centered random variables.

Corollary 9.3.2. Let $M, N \in \mathbb{R}^{m \times m}$ be arbitrary PSD matrices, $v \in \mathbb{R}^{m}$ and $\xi$ a random variable with expectation $\mu$ and variance $\tau^{2}$.
$\underset{\xi}{\mathbb{E}}\left[\operatorname{det}\left(M-(\xi-\mu) v v^{*}\right) \operatorname{det}\left(N+(\xi-\mu) v v^{*}\right)\right]=\left.\left(1-\frac{1}{2} \frac{d^{2}}{d t^{2}}\right)\right|_{t=0} \operatorname{det}\left(M+t \tau v v^{*}\right) \operatorname{det}\left(N+t \tau v v^{*}\right)$.
We can now derive an expression for the expected characteristic polynomial.

Proposition 9.3.3. Let $u_{1}, \ldots, u_{n} \in \mathbb{R}^{m}$. Consider independent random variables $\xi_{i}$ with means $\mu_{i}$ and variances $\tau_{i}^{2}$. Let $Q \in \mathbb{R}^{m \times m}$ be a symmetric matrix.

$$
\begin{array}{r}
\underset{\xi}{\mathbb{E}}\left[\operatorname{det}\left(x^{2} I-\left(Q+\sum_{i=1}^{n}\left(\xi_{i}-\mu_{i}\right) u_{i} u_{i}^{*}\right)^{2}\right)\right]=\left.\prod_{i=1}^{n}\left(1-\frac{\partial_{z_{i}}^{2}}{2}\right)\right|_{z_{i}=0} \\
\operatorname{det}\left(x I-Q+\sum_{i=1}^{n} z_{i} \tau_{i} u_{i} u_{i}^{*}\right) \\
\times \operatorname{det}\left(x I+Q+\sum_{i=1}^{n} z_{i} \tau_{i} u_{i} u_{i}^{*}\right),
\end{array}
$$

and this is a real rooted polynomial in $x$.

Proof. For each $i$, let $\beta_{i}$ denote the maximum value of $\left|\xi_{i}\right|$ among outcomes in the (finite) support of $\xi_{i}$. We begin by restricting the domain of our polynomials to $x>\|Q\|+$ $2 \sum_{i=1}^{n} \beta_{i}\left\|u_{i} u_{i}^{*}\right\|$. We then proceed by induction. Our induction hypothesis will be that for $0 \leq k \leq n$

$$
\begin{align*}
& \mathbb{E}\left[\operatorname{det}\left(x^{2} I-\left(Q+\sum_{i=1}^{n}\left(\xi_{i}-\mu_{i}\right) u_{i} u_{i}^{*}\right)^{2}\right)\right] \\
= & \left.\underset{\xi_{k+1}, \ldots, \xi_{n}}{\mathbb{E}} \prod_{i=1}^{k}\left(1-\frac{\partial_{z_{i}}^{2}}{2}\right)\right|_{z_{i}=0} \operatorname{det}\left(x I-Q-\sum_{i=k+1}^{n}\left(\xi_{i}-\mu_{i}\right) u_{i} u_{i}^{*}+\sum_{j=1}^{k} z_{j} \tau_{j} u_{j} u_{j}^{*}\right) \\
& \times \operatorname{det}\left(x I+Q+\sum_{i=k+1}^{n}\left(\xi_{i}-\mu_{i}\right) u_{i} u_{i}^{*}+\sum_{j=1}^{k} z_{j} \tau_{j} u_{j} u_{j}^{*}\right) \tag{9.6}
\end{align*}
$$

The base case, $k=0$ is trivially true as we get the same formula on both sides after recalling that for any matrix $Y$

$$
\operatorname{det}\left(x^{2} I-Y^{2}\right)=\operatorname{det}(x I-Y) \operatorname{det}(x I+Y)
$$

By our assumption that $x>\|Q\|+2 \sum_{i=1}^{n} \beta_{i}\left\|u_{i} u_{i}^{*}\right\|$,

$$
x I-\sum_{i=k+2}^{n}\left(\xi_{i}-\mu_{i}\right) u_{i} u_{i}^{*}+\sum_{j=1}^{k} z_{j} \tau_{j} u_{j} u_{j}^{*}
$$

is PSD for any realization of $\xi_{k+2}, \ldots, \xi_{n}$ and in a neighborhood of zero for each $z_{j}$. Applying
Corollary 9.3.2 to the right hand side of (9.6), yields

$$
\begin{aligned}
& \mathbb{E}\left[\operatorname{det}\left(x^{2} I-\left(Q+\sum_{i=1}^{n}\left(\xi_{i}-\mu_{i}\right) u_{i} u_{i}^{*}\right)^{2}\right)\right] \\
&=\left.\underset{\xi_{k+2}, \ldots, \xi_{n}}{\mathbb{E}} \prod_{i=1}^{k+1}\left(1-\frac{\partial_{z_{i}}^{2}}{2}\right)\right|_{z_{i}=0} \operatorname{det}\left(x I-Q-\sum_{i=k+2}^{n}\left(\xi_{i}-\mu_{i}\right) u_{i} u_{i}^{*}+\sum_{j=1}^{k+1} z_{j} \tau_{j} u_{j} u_{j}^{*}\right) \\
& \times \operatorname{det}\left(x I+Q+\sum_{i=k+2}^{n}\left(\xi_{i}-\mu_{i}\right) u_{i} u_{i}^{*}+\sum_{j=1}^{k+1} z_{j} \tau_{j} u_{j} u_{j}^{*}\right)
\end{aligned}
$$

which completes the induction. To extend the proof to all $x$, we remark that we have shown the equivalence of two polynomials in the interval $x>\|Q\|+2 \sum_{i=1}^{n} \beta_{i}\left\|u_{i} u_{i}^{*}\right\|$, and two polynomials that agree on an interval are identical.

Real-rootedness of the right hand side follows by Lemma 9.2.4, Lemma 9.2.3, and that by Proposition 9.2.5 restriction to $z=\mathbf{0}$ preserves real-stability, and a univariate real stable polynomial is real rooted.

### 9.4 Defining the Interlacing Family

The next proposition establishes that it suffices to bound the largest root of $q_{\emptyset}$.
Proposition 9.4.1. There exists a choice of outcomes $\epsilon_{1}, \ldots, \epsilon_{n}$ in the finite support of $\xi_{1}, \ldots, \xi_{n}$, s.t.

$$
\left\|\sum_{i=1}^{n} \epsilon_{i} u_{i} u_{i}^{*}-\sum_{i=1}^{n} \mu_{i} u_{i} u_{i}^{*}\right\|
$$

is less than the largest root of

$$
\underset{\xi_{1}, \cdots, \xi_{n}}{\mathbb{E}}\left[\operatorname{det}\left(x^{2} I-\left(\sum_{i=1}^{n}\left(\xi_{i}-\mu_{i}\right) u_{i} u_{i}^{*}\right)^{2}\right)\right] .
$$

where $\mathbb{E}\left[\xi_{i}\right]=\mu_{i}, \forall i \in[n]$.

Proof. For a vector of independent random variables $\left(\xi_{1}, \ldots, \xi_{n}\right)$ with finite support, let $p_{i, x}$ be the probability that $\xi_{i}=x$. For $s=\left(\epsilon_{1}, \ldots, \epsilon_{n}\right)$ in the support of $\xi_{1}, \ldots, \xi_{n}$, we define

$$
q_{s}(x):=\prod_{i=1}^{n} p_{i, \epsilon_{i}} \operatorname{det}\left(x^{2} I-\left(\sum_{i=1}^{n}\left(\epsilon_{i}-\mu_{i}\right) u_{i} u_{i}^{*}\right)^{2}\right)
$$

Let $t$ be a vector of $k$ outcomes in the support of $\xi_{1}, \ldots, \xi_{k}$, i.e. a partial assignment of assignment of outcomes. Then we consider the conditional expected polynomial

$$
q_{t}(x):=\left(\prod_{i=1}^{k} p_{i, t_{i}}\right) \underset{\xi_{k+1}, \ldots, \xi_{n}}{\mathbb{E}}\left[\operatorname{det}\left(x^{2} I-\left(\sum_{i=1}^{k}\left(t_{i}-\mu_{i}\right) u_{i} u_{i}^{*}+\sum_{j=k+1}^{n}\left(\xi_{j}-\mu_{j}\right) u_{j} u_{j}^{*}\right)^{2}\right)\right]
$$

which coincides with (9.3). We show that $q_{s}$ is an interlacing family. Let $r_{k+1}^{(1)} \ldots r_{k+1}^{(l)}$ be the outcomes in the support of $\xi_{k+1}$. For a given $t$ and $r \in\left\{r_{k+1}^{(1)}, \ldots, r_{k+1}^{(l)}\right\}$, let $(t, r)$ denote the vector $\left(t_{1}, \ldots, t_{k}, r\right)$. By Lemma 9.2.8, it suffices to show that for any choice of non-negative numbers $\alpha_{1}, \ldots, \alpha_{l}$ such that $\sum_{j} \alpha_{j}=1$, the polynomial

$$
\sum_{j} \alpha_{j} q_{t, r_{k+1}^{(j)}}(x)
$$

is real rooted. We can consider these $\alpha$ 's as a probability distribution and define $p_{k+1, r_{k+1}^{(j)}}=$ $\alpha_{j}$. Therefore,

$$
\begin{aligned}
q_{t}(x)= & \left(\prod_{i=1}^{k} p_{i, t_{i}}\right) \underset{\xi_{k+1}, \ldots, \xi_{n}}{\mathbb{E}}\left[\operatorname{det}\left(x^{2} I-\left(\sum_{i=1}^{k}\left(t_{i}-\mu_{i}\right) u_{i} u_{i}^{*}+\sum_{j=k+1}^{n}\left(\xi_{j}-\mu_{j}\right) u_{j} u_{j}^{*}\right)^{2}\right)\right] \\
= & \sum_{j}\left(\prod_{i=1}^{k} p_{i, t_{i}}\right) p_{k+1, r_{k+1}^{(j)}} \\
& \quad \times \underset{\xi_{k+2}, \ldots, \xi_{n}}{\mathbb{E}}\left[\operatorname{det}\left(x^{2} I-\left(\sum_{i=1}^{k}\left(t_{i}-\mu_{i}\right) u_{i} u_{i}^{*}-\left(r_{k+1}^{(j)}-\mu_{k+1}\right) u_{k+1} u_{k+1}^{*}-\sum_{j=k+2}^{n}\left(\xi_{j}-\mu_{j}\right) u_{j} u_{j}^{*}\right)^{2}\right)\right] \\
= & \sum_{j} \alpha_{j} q_{t, r_{k+1}^{(j)}}(x) .
\end{aligned}
$$

By Proposition 9.3.3, $q_{t}$ is real rooted, which completes the proof that $q_{s}$ is an interlacing family. Finally, by Lemma 9.2.7, there exists a choice of $t$ so that the largest root of $q_{t}$ is upperbounded by $q_{\emptyset}$.

### 9.5 Largest Root of the Expected Characteristic Polynomial

We use the barrier method approach to control the largest root [MSS15b].

Definition 9.5.1. For a multivariate polynomial $p\left(z_{1}, \ldots, z_{n}\right)$, we say $z \in \mathbb{R}^{n}$ is above all the roots of $p$ if for all $t \in \mathbb{R}_{+}^{n}$,

$$
p(z+t)>0 .
$$

We use $\mathbf{A} \mathbf{b}_{p}$ to denote the set of points that are above all the roots of $p$.

We use the same barrier function as in [BSS12, MSS15b].

Definition 9.5.2. For a real stable polynomial $p$ and $z \in \mathbf{A b}_{p}$, the barrier function of $p$ in direction $i$ at $z$ is

$$
\Phi_{p}^{i}(z):=\frac{\partial_{z_{i}} p(z)}{p(z)} .
$$

We will also make use of the following lemma that controls the deviation of the roots after applying a second order differential operator. The lemma is a slight variation of Lemma 4.8 in [AG14].

Lemma 9.5.1. Suppose that $p$ is real stable and $z \in \mathbf{A} \mathbf{b}_{p}$.

$$
\begin{aligned}
& \text { If } \Phi_{p}^{j}(z)<\sqrt{2} \text {, then } z \in \mathbf{A b}_{\left(1-\frac{1}{2} \partial_{z_{j}}^{2}\right) p} . \text { If additionally for } \delta>0 \\
& \qquad \frac{1}{\delta} \Phi_{p}^{j}(z)+\frac{1}{2} \Phi_{p}^{j}(z)^{2} \leq 1,
\end{aligned}
$$

then and for all $i$,

$$
\Phi_{\left(1-\frac{1}{2} \partial \bar{z}_{j}\right) p}^{i}\left(z+\delta \cdot \mathbf{1}_{j}\right) \leq \Phi_{p}^{i}(z)
$$

We provide a proof in the Appendix 9.6 for completeness.
We can now bound the largest root of the expected characteristic polynomial for subisotropic vectors.

Proposition 9.5.2. Let $\xi_{1}, \ldots, \xi_{n}$ be independent scalar random variables with finite support, with $\mathbb{E}\left[\xi_{i}\right]=\mu_{i}$, and let $\tau_{i}^{2}=\mathbb{E}\left[\left(\xi_{i}-\mu_{i}\right)^{2}\right]$.

Let $v_{1}, \ldots, v_{n} \in \mathbb{R}^{n}$ such that $\sum_{i=1}^{n} \tau_{i}^{2}\left(v_{i} v_{i}^{*}\right)^{2} \preceq I$. Then the largest root of

$$
p(x):=\underset{\xi_{1}, \cdots, \xi_{n}}{\mathbb{E}}\left[\operatorname{det}\left(x^{2} I-\left(\sum_{i=1}^{n}\left(\xi_{i}-\mu_{i}\right) v_{i} v_{i}^{*}\right)^{2}\right)\right]
$$

is at most 4.

Proof. Note that we must have

$$
\begin{equation*}
\max _{i \in[n]} \tau_{i} v_{i}^{*} v_{i} \leq 1 \tag{9.7}
\end{equation*}
$$

since otherwise $\sum_{i=1}^{n} \tau_{i}^{2}\left(v_{i} v_{i}^{*}\right)^{2} \preceq I$ is false.
Let

$$
Q(x, z)=\left(\operatorname{det}\left(x I+\sum_{i=1}^{n} z_{i} \tau_{i} v_{i} v_{i}^{*}\right)\right)^{2} .
$$

For $t>0$, define $\delta_{i}=t \tau_{i} v_{i}^{*} v_{i}$. For $t<\alpha(t)$ a parameter to be chosen later, we evaluate
our polynomial to find that

$$
\begin{aligned}
Q\left(\alpha,-\delta_{1}, \ldots,-\delta_{n}\right) & =\left(\operatorname{det}\left(\alpha I-\sum_{i=1}^{n} \delta_{i} \tau_{i} v_{i} v_{i}^{*}\right)\right)^{2} \\
& =\left(\operatorname{det}\left(\alpha I-t \sum_{i=1}^{n} \tau_{i}^{2}\left(v_{i}^{*} v_{i}\right) v_{i} v_{i}^{*}\right)\right)^{2} \\
& =\left(\operatorname{det}\left(\alpha I-t \sum_{i=1}^{n} \tau_{i}^{2}\left(v_{i} v_{i}^{*}\right) v_{i} v_{i}^{*}\right)\right)^{2} \\
& \geq(\operatorname{det}((\alpha-t) I))^{2} \\
& >0
\end{aligned}
$$

where the second last step follows by $\sum_{i=1}^{n} \tau_{i}^{2}\left(v_{i} v_{i}^{*}\right) v_{i} v_{i}^{*} \preceq I$.
This implies that $(\alpha,-\delta) \in \mathbb{R}^{n+1}$ is above the roots of $Q(x, z)$. We can upper bound $\Phi_{Q}^{i}(\alpha,-\delta)$ via Fact 9.2.2 as follows

$$
\begin{aligned}
\Phi_{Q}^{i}(\alpha,-\delta) & =\left.\frac{\partial_{z_{i}} Q}{Q}\right|_{x=\alpha, z=-\delta} \\
& =\left.\frac{2 \operatorname{det}\left(x I+\sum_{i=1}^{n} z_{i} \tau_{i} v_{i} v_{i}^{*}\right) \partial_{z_{i}} \operatorname{det}\left(x I+\sum_{i=1}^{n} z_{i} \tau_{i} v_{i} v_{i}^{*}\right)}{\left(\operatorname{det}\left(x I+\sum_{i=1}^{n} z_{i} \tau_{i} v_{i} v_{i}^{*}\right)\right)^{2}}\right|_{x=\alpha, z=-\delta} \\
& =\left.2 \operatorname{tr}\left[\left(x I+\sum_{i=1}^{n} z_{i} \tau_{i} v_{i} v_{i}^{*}\right)^{-1} \tau_{i} v_{i} v_{i}^{*}\right]\right|_{x=\alpha, z=-\delta} \\
& =2 \operatorname{tr}\left[\left(\alpha I-t \sum_{i=1}^{n} \tau_{i}^{2}\left(v_{i} v_{i}^{*}\right)^{2}\right)^{-1} \tau_{i} v_{i} v_{i}^{*}\right] \\
& \leq 2 \operatorname{tr}\left[(\alpha-t)^{-1} \tau_{i} v_{i} v_{i}^{*}\right] \\
& =\frac{2 \tau_{i} v_{i}^{*} v_{i}}{\alpha-t}
\end{aligned}
$$

where the fifth step follows by $\sum_{i=1}^{n} \tau_{i}^{2}\left(v_{i} v_{i}^{*}\right) v_{i} v_{i}^{*} \preceq I$.
Choosing $\alpha=2 t$ and $t=2$, and recalling Condition (9.7) we get

$$
\Phi_{Q}^{i}(\alpha,-\delta) \leq \frac{2 \tau_{i} v_{i}^{*} v_{i}}{\alpha-t} \leq 1<\sqrt{2}
$$

By Lemma 9.2.3, $\left(1-\frac{1}{2} \partial_{z_{i}}^{2}\right) Q$ is real stable. By Lemma 9.5.1, and $(4, z) \in \mathbf{A b}_{\left(1-\frac{1}{2} \partial_{z_{i}}^{2}\right) Q}$ and hence $\left(4, z+\delta_{i} \mathbf{1}_{i}\right) \in \mathbf{A b}_{\left(1-\frac{1}{2} \partial_{z_{i}}^{2}\right) Q}$. Also

$$
\begin{aligned}
\frac{1}{\delta_{i}} \Phi_{Q}^{i}(\alpha,-\delta)+\frac{1}{2} \Phi_{Q}^{i}(\alpha,-\delta)^{2} & \leq \frac{1}{t \tau_{i} v_{i}^{*} v_{i}} \frac{2 \tau_{i} v_{i}^{*} v_{i}}{t}+\frac{1}{2}\left(\frac{2 \tau_{i} v_{i}^{*} v_{i}}{t}\right)^{2} \\
& \leq \frac{4}{t^{2}} \\
& =1
\end{aligned}
$$

Therefore, by Lemma 9.5.1, for all $j$

$$
\Phi_{\left(1-\frac{1}{2} \partial_{z_{i}}^{2}\right) Q}^{j}\left(4, z+\delta_{i} \mathbf{1}_{i}\right) \leq \Phi_{Q}^{j}(4, z)
$$

Repeating this argument for each $i \in[n]$ demonstrates that $(4,0, \ldots, 0)$ lies above the roots of

$$
\prod_{i=1}^{n}\left(1-\frac{\partial_{z_{i}}^{2}}{2}\right)\left(\operatorname{det}\left(x I+\sum_{i=1}^{n} z_{i} \tau_{i} u_{i} u_{i}^{*}\right)\right)\left(\operatorname{det}\left(x I+\sum_{i=1}^{n} z_{i} \tau_{i} u_{i} u_{i}^{*}\right)\right)
$$

After restricting to $z_{i}=0$ for all $i$, we then conclude by Proposition 9.3.3 is equivalent to a bound on the largest root of the expected characteristic polynomial.

Having developed the necessary machinery, we now prove our main theorem.

Proof of Theorem 9.1.4. Define $v_{i}=\frac{u_{i}}{\sqrt{\sigma}}$. Then, $\left\|\sum_{i=1}^{n} \tau_{i}^{2}\left(v_{i} v_{i}^{*}\right)^{2}\right\|=1$. Applying Proposition 9.5.2 and Proposition 9.4.1, we conclude that there exists a choice of outcomes $\epsilon_{i}$ such
that

$$
\left\|\sum_{i=1}^{n}\left(\epsilon_{i}-\mu_{i}\right) v_{i} v_{i}^{*}\right\| \leq 4 .
$$

From this, we conclude that

$$
\left\|\sum_{i=1}^{n}\left(\epsilon_{i}-\mu_{i}\right) u_{i} u_{i}^{*}\right\| \leq 4 \sigma .
$$

### 9.6 Omitted Proofs

### 9.6.1 Proof of Lemma 9.5.1

Choosing $c=1 / 2$ in Lemma 9.6.1 gives a proof of Lemma 9.5.1.

Lemma 9.6.1 (Generalization of Lemma 4.8 in [AG14]). Suppose that $p\left(z_{1}, \cdots, z_{m}\right)$ is real stable and $z \in A b_{p}$. For any $c \in[0,1]$.

$$
\begin{aligned}
& \text { If } \Phi_{p}^{j}(z)<\sqrt{1 / c}, \text { then } z \in \mathbf{A b}_{\left(1-c \partial_{z_{j}}\right) p} . \text { If additionally for } \delta>0, \\
& c \cdot\left(\frac{2}{\delta} \Phi_{p}^{j}(z)+\left(\Phi_{p}^{j}(z)\right)^{2}\right) \leq 1,
\end{aligned}
$$

then, for all $i \in[m]$,

$$
\Phi_{\left(1-c \partial_{j}^{2}\right) p}^{i}\left(z+\delta \mathbf{1}_{j}\right) \leq \Phi_{p}^{i}(z) .
$$

Proof. We write $\partial_{i}$ instead of $\partial_{z_{i}}$ for ease of notation. By Lemma 9.2.3, $\left(1-c \partial_{j}^{2}\right) p$ is real stable. Recall the definitions of $\Phi_{p}^{j}(z)$ and $\Psi_{p}^{j}(z)$

$$
\Phi_{p}^{j}(z)=\frac{\partial_{z_{j}} \operatorname{det}(M)}{\operatorname{det}(M)}=\frac{\partial_{j} p}{p} \text { and } \Psi_{p}^{j}(z)=\frac{\partial_{j}^{2} \operatorname{det}(M)}{\operatorname{det}(M)}=\frac{\partial_{j}^{2} p}{p}
$$

Consider a non-negative vector $t$. By Lemmas 4.6 and 4.5 in [AG14], we have

$$
\begin{equation*}
\Phi_{p}^{j}(z+t) \leq \Phi_{p}^{j}(z), \quad \Psi_{p}^{j}(z+t) \leq \Phi_{p}^{j}(z+t)^{2} \leq \Phi_{p}^{j}(z)^{2} . \tag{9.8}
\end{equation*}
$$

Thus, $\Psi_{p}^{j}(z+t) \leq \Phi_{p}^{j}(z)^{2}<1 / c$ so $c \cdot \partial_{j}^{2} p(z+t)<p(z+t)$, i.e. $\left(1-c \partial_{j}^{2}\right) p(z+t)>0$. Thus $z \in \mathbf{A b}_{\left(1-c \partial_{z_{j}}^{2}\right) p}$.

Next, we write $\Phi_{p-c \cdot \partial_{j}^{2} p}^{i}$ in terms of $\Phi_{p}^{i}$ and $\Psi_{p}^{j}$ and $\partial_{i} \Psi_{p}^{j}$.

$$
\begin{aligned}
\Phi_{p-c \cdot \partial_{j}^{2} p}^{i} & =\frac{\partial_{i}\left(p-c \cdot \partial_{j}^{2} p\right)}{p-c \cdot \partial_{j}^{2} p} \\
& =\frac{\partial_{i}\left(\left(1-c \cdot \Psi_{p}^{j}\right) p\right)}{\left(1-c \cdot \Psi_{p}^{j}\right) p} \\
& =\frac{\left(1-c \cdot \Psi_{p}^{j}\right)\left(\partial_{i} p\right)}{\left(1-c \cdot \Psi_{p}^{j}\right) p}+\frac{\left(\partial_{i}\left(1-c \cdot \Psi_{p}^{j}\right)\right) p}{\left(1-c \cdot \Psi_{p}^{j}\right) p} \\
& =\Phi_{p}^{i}-\frac{c \cdot \partial_{i} \Psi_{p}^{j}}{1-c \cdot \Psi_{p}^{j}} .
\end{aligned}
$$

We would like to show that $\Phi_{p-c \cdot \partial_{j}^{2} p}^{i}\left(z+\delta \mathbf{1}_{j}\right) \leq \Phi_{p}^{i}(z)$. Equivalently, it is enough to show that

$$
-\frac{c \cdot \partial_{i} \Psi_{p}^{j}\left(z+\delta \mathbf{1}_{j}\right)}{1-c \cdot \Psi_{p}^{j}\left(z+\delta \mathbf{1}_{j}\right)} \leq \Phi_{p}^{i}(z)-\Phi_{p}^{i}\left(z+\delta \mathbf{1}_{j}\right)
$$

By convexity, it is enough to show that

$$
-\frac{c \cdot \partial_{i} \Psi_{p}^{j}\left(z+\delta \mathbf{1}_{j}\right)}{1-c \cdot \Psi_{p}^{j}\left(z+\delta \mathbf{1}_{j}\right)} \leq \delta \cdot\left(-\partial_{j} \Phi_{p}^{i}\left(z+\delta \mathbf{1}_{j}\right)\right)
$$

By monotonicity, $\delta \cdot\left(-\partial_{j} \Phi_{p}^{i}\left(z+\delta \mathbf{1}_{j}\right)\right)>0$ so we may divide both sides of the above inequality by this term and obtain that the above is equivalent to

$$
\frac{-c \cdot \partial_{i} \Psi_{p}^{j}\left(z+\delta \mathbf{1}_{j}\right)}{-\delta \cdot \partial_{i} \Phi_{p}^{j}\left(z+\delta \mathbf{1}_{j}\right)} \cdot \frac{1}{1-c \Psi_{p}^{j}\left(z+\delta \mathbf{1}_{j}\right)} \leq 1
$$

where we also used $\partial_{j} \Phi_{p}^{i}=\partial_{i} \Phi_{p}^{j}$. By Lemma 4.10 in [AG14], $\frac{\partial_{i} \Psi_{p}^{j}}{\partial_{i} \Phi_{p}^{j}} \leq 2 \Phi_{p}^{j}$. So, we can write,

$$
\frac{2 c}{\delta} \Phi_{p}^{j}\left(z+\delta \mathbf{1}_{j}\right) \cdot \frac{1}{1-c \cdot \Psi_{p}^{j}\left(z+\delta \mathbf{1}_{j}\right)} \leq 1
$$

By Equation (9.8), with $t=\delta \mathbf{1}_{j}$,

$$
\Phi_{p}^{j}\left(z+\delta \mathbf{1}_{j}\right) \leq \Phi_{p}^{j}(z), \quad \Psi_{p}^{j}\left(z+\delta \mathbf{1}_{j}\right) \leq \Phi_{p}^{j}\left(z+\delta \mathbf{1}_{j}\right)^{2} \leq \Phi_{p}^{j}(z)^{2}
$$

So, it is enough to show that

$$
\frac{2 c}{\delta} \Phi_{p}^{j}(z) \cdot \frac{1}{1-c \cdot \Phi_{p}^{j}(z)^{2}} \leq 1
$$

Using $\Phi_{p}^{j}(z)<1$ and $c \in[0,1]$ we know that $1-c \cdot \Phi_{p}^{j}(z)^{2}<1$. We can multiply both sides with $1-c \cdot \Phi_{p}^{j}(z)^{2}$ and we obtain that it suffices to have

$$
c \cdot\left(\frac{2}{\delta} \Phi_{p}^{j}(z)+\Phi_{p}^{j}(z)^{2}\right) \leq 1
$$

which is true by assumption.

## Part III

## Compressive Sensing and Sparse Fourier Transform

## Chapter 10

## Compressive Sensing

We consider the extensively studied problem of $\ell_{2} / \ell_{2}$ compressed sensing. The main contribution of our work is an improvement over [Gilbert, Li, Porat and Strauss, STOC 2010] with faster decoding time and significantly smaller column sparsity, answering two open questions of the aforementioned work.

Previous work on sublinear-time compressed sensing employed an iterative procedure, recovering the heavy coordinates in phases. We completely depart from that framework, and give the first sublinear-time $\ell_{2} / \ell_{2}$ scheme which achieves the optimal number of measurements without iterating; this new approach is the key step to our progress. Towards that, we satisfy the $\ell_{2} / \ell_{2}$ guarantee by exploiting the heaviness of coordinates in a way that was not exploited in previous work. Via our techniques we obtain improved results for various sparse recovery tasks, and indicate possible further applications to problems in the field, to which the aforementioned iterative procedure creates significant obstructions.

### 10.1 Introduction

Compressed Sensing, or sparse recovery, is a powerful mathematical framework the goal of which is to reconstruct an approximately $k$-sparse vector $x \in \mathbb{R}^{n}$ from linear measurements $y=\Phi x$, where $\Phi \in \mathbb{R}^{m \times n}$. The most important goal is to reduce the number of measurements $m$ needed to approximate the vector $x$, avoiding the linear dependence on $n$. In discrete signal processing, where this framework was initiated [CRT06b, Don06], the core principle that the sparsity of a signal can be exploited to recover it using much fewer samples than the Shannon-Nyquist Theorem. We refer to the matrix $\Phi$ as the sketching or sensing matrix, and $y=\Phi x$ as the sketch of vector $x$.

Sparse recovery is the primary task of interest in a number of applications, such as image processing $\left[\mathrm{TLW}^{+} 06, \mathrm{LDP} 07 \mathrm{a}, \mathrm{DDT}^{+} 08\right]$, design pooling schemes for biological tests [ $\mathrm{ECG}^{+} 09, \mathrm{DWG}^{+} 13$ ], pattern matching [CEPR07], combinatorial group testing [SAZ09, ESAZ09, $\mathrm{KBG}^{+} 10$ ], localizing sources in sensor networks [ZBSG05, ZPB06], as well as neuroscience [GS12]. Furthermore, not surprisingly, tracking heavy hitters in data streams, also known as frequent items, can be captured by the sparse recovery framework [Mut05a, CH09, KSZC03, Ind07]. In practice, streaming algorithms for detecting heavy hitters have been used to find popular destination addresses and heavy bandwidth users by AT\&T [CJK $\left.{ }^{+} 04\right]$ or answer "iceberg queries" in databases [FSGM ${ }^{+} 99$ ].

Sparse recovery attracts researchers from different communities, from both theoretical and practical perspective. During the last ten years, hundreds of papers have been published by theoretical computer scientists, applied mathematicians and electrical engineers that specialize in compressed sensing. While numerous algorithms using space linear in the universe size $n$ are known, [Don06, CRT06b, IR08a, NT09a, BIR08a, BD09a, SV16]
to name a few, our goal is to obtain algorithms that are sublinear, something that is crucial in many applications.

The desirable quantities we want to optimize may vary depending on the application. For example, in network management, $x_{i}$ could denote the total number of packets with destination $i$ passing through a network router. In such an application, storing the sketching matrix explicitly is typically not a tenable solution, since this would lead to an enormous space consumption; the number of possible IP addresses is $2^{32}$. Moreover, both the query and the update time should be very fast, in order to avoid congestion on the network. Incremental updates to $x$ come rapidly, and the changes to the sketch should also be implemented very fast; we note that in this case, even poly-logarithmic factors might be prohibitive. Interested readers can refer to [KSZC03, EV03] for more information about streaming algorithms for network management applications.
"The goal of that research is to obtain encoding and recovery schemes with good compression rate (i.e., short sketch lengths) as well as good algorithmic properties (i.e., low encoding, update and recovery times)." - Anna Gilbert and Piotr Indyk [GI10]

Sparse recovery schemes that are optimal across all axis are a challenge and an important theoretical and practical problem. For most sparse recovery tasks, we have algorithms that achieve different trade-offs for the various parameters of interest. One exception is the $\ell_{\infty} / \ell_{2}$ guarantee, for which the breakthrough work of Larsen, Nelson, Nguyên and Thorup [LNNT16] shows that this trade-off is unnecessary.

### 10.1.1 Previous work

Since compressed sensing has been extensively studied in the literature for more than a decade, different guarantees of interest have been suggested ( $x_{-k}$ is the vector that occurs after zeroing out every $i$ that does not belong among the largest $k$ coordinates). In what follows $x \in \mathbb{R}^{n}$ is the vector we want to sketch, $x^{\prime}$ is the approximation to $x, k$ is the sparsity and $\epsilon$ is the fineness of the approximation.

There are two different guarantees researchers consider in compressed sensing, one is the for-all guarantee and the other is the for-each guarantee. In the for-all guarantee, one wants to design a sketch that gives the desired result for all vectors $x \in \mathbb{R}^{n}$. In the for-each guarantee, one wants to design a distribution over sketches that gives the desired result for a fixed vector $x \in \mathbb{R}^{n}$. We note that $\ell_{\infty} / \ell_{2}, \ell_{2} / \ell_{2}$ are impossible in the for-all model, unless $\Omega(n)$ measurements are used [CDD09]. The standard approach for the for-all guarantee is via RIP matrices, satisfying the so-called Restricted Isometry Property. In what follows, we will refer to the for-each model, unless stated otherwise.


The first set of schemes that initiated the research on compressed sensing are given in [CRT06b, Don06]. There the authors show, for any $x \in \mathbb{R}^{n}$, given $y=\Phi x$, it is possible to satisfy the $\ell_{2} / \ell_{1}$ guarantee for all vectors, if $\Phi$ is a Gaussian matrix with $O(k \log (n / k))$ rows. The schemes in [CM06, CCF02] achieve the $\ell_{\infty} / \ell_{2}$ guarantee with $O(k \log n)$ measurementz, matching known lower bounds [JST11], $O(n \log n)$ decoding time and $O(\log n)$
update time. The state of the art for $\ell_{\infty} / \ell_{2}$ is [LNNT16], which gives optimal number of measurements, sublinear decoding time, $O(\log n)$ update time and $1 / \operatorname{poly}(n)$ failure probability. Price and Woodruff [PW11] show that in order to get $\ell_{2} / \ell_{2}$ with constant failure probability $<1 / 2$ with the output being exactly $k$-sparse output requires $\Omega\left(\epsilon^{-2} k\right)$ measurements. They also showed non- $k$-sparse output requires $\Omega\left(\epsilon^{-1} k \log (n / k)\right)$ measurements in the regime $\epsilon>\sqrt{k \log n / n}$, and gave an upper bound of $O\left(\epsilon^{-1} k \log n\right)$ measurements, showing thus a separation in the measurement complexity between $k$-sparse and $O(k)$-sparse output. Later, in the breakthrough work of Gilbert, Li, Porat and Strauss [GLPS10] an algorithm that runs in sublinear time, and has $O\left(\log (n / k) \log ^{2} k\right)$ column sparsity, was devised. On generic norms, nearly optimal bounds have been given by Backurs, Indyk, Razenshteyn and Woodruff [BIRW16]. We note, however, that their schemes are not computationally efficient: they have exponential running time, except in the case of Earth-Mover-Distance, which has time polynomial in $n$ and $\log ^{k} n$.

Measurements. The number of measurements corresponds to physical resources: memory in monitoring devices of data streams, number of screens in biological applications, or number of filters in dynamic spectrum access (DSA) of radio signal [HMT $\left.{ }^{+} 13\right]$.

In applications such as medical imaging, it is crucial to reduce the number of measurements, since the radiation used in CT scans could potentially increase cancer risks for patients. For instance, $\left[\mathrm{PSL}^{+} 12\right]$ showed that a positive association between radiation exposure from CT scans in childhood and subsequent risk of leukemia and brain tumors.

For more applications, we refer the readers to $\left[\mathrm{QBI}^{+} 13\right]$.

Encoding Time. Designing algorithms with fast update/encoding time is a well-motivated task for streaming algorithms, since the packets arrive at an extremely fast rate [TZ12]; even logarithmic factors are crucial in these applications. Also in digital signal processing applications, in the design of cameras or satellites which demand rapid imaging, when we observe a sequence of images that are close to each other, we may not need to encode the new signal from the beginning, rather than encode only that part which differs from the current signal; the delay is then defined by the update time of our scheme. Moreover, in Magnetic Resonance Imaging (MRI) update time or encoding time defines the time the patient waits for the scan to happen. Improvement of the runtime has benefits both for patients and for healthcare economics [LDSP08].

A natural question is the following: what are the time limitations of our data structures, regarding update time? Regarding the streaming setting, the first lower bounds are given in [LNN15] for non-adaptive algorithms. An algorithm is called non-adaptive if, during updates, the memory cells are written and read depend only on the index being updated and the random coins tossed before the stream is started to being processed. The lower bounds given concern both randomized and deterministic algorithms; the relevant bounds to sparse recovery are for $\ell_{p} / \ell_{q}$ estimation. However, for constant failure probability their results do not give anything useful, since their lower bounds start to kick in when the failure probability becomes very small, namely $o\left(2^{-\sqrt{m \cdot \log n}}\right)$.

For the column sparsity (which could be smaller than update time, and hence the lower bounds in [LNN15] might not apply ${ }^{1}$ ), the only known lower bounds are known for

[^55]RIP matrices, which are used in the for-all setting. To the best of our knowledge, the first non-trivial lower bounds were given by Nachin [Nac10], and then extended by Indyk and Razenshteyn in [IR13] for RIP-1 model-based compressed sensing matrices. Lower bounds for the column sparsity of RIP-2 matrices were given in Nelson and Nguyên [NN13b], and then to RIP-p matrices in Allen-Zhu, Gelashvili and Razenshteyn [AGR16]. Roughly speaking, the lower bounds for $\ell_{2}$ indicate that if one aims for optimal measurements, $m=k \log (n / k)$, in the regime $k<n / \log ^{3} n$, one cannot obtain column sparsity better than $\Omega(m)$. This indicates that the for-all case should be significantly worse, in terms of column sparsity, than the for-each case.

Decoding Time. Another very important quantity we want to minimize is the time needed to reconstruct the approximation of $x$ from its compressed version. This quantity is of enormous significance in cases where the universe size is huge and we cannot afford to iterate over it. This is often the case in networking applications, where the universe size is the number of distinct IP addresses. In MRI applications the decoding time corresponds to the time needed to reconstruct the image after the scan has been performed. Decoding time is highly important also in satellite systems, modern radars and airspace surveillance, where compressed sensing have found extensive application [End10].

The sparse Fourier transform can be regarded as another variant of the sparse recovery problem, the results developed along that line being incomparable with traditional sparse recovery results, where we have freedom over the design of the sketching matrix. Sparse Fourier transforms can be divided into two categories, one being sparse discrete Fourier transform, and the other being sparse continuous Fourier transform [BCG+12, Moi15, PS15, CKPS16].

The sparse discrete Fourier transform also can be split into two lines, the first category of results [GGI+ 02, GMS05, HIKP12a, HIKP12b, Iwe13, IKP14, IK14, Kap16, CKSZ17, Kap17, KVZ19] carefully choose measurements that allow for sublinear recovery time, and the second category of results [CRT06a, RV08, Bou14, HR16] focus proving Restricted Isometry Property. The techniques used in sparse Fourier transforms are related to standard compressed sensing, since the main approaches try to implement arbitrary linear measurements via sampling Fourier coefficients.

### 10.1.2 Our result

Our main result is a novel scheme for $\ell_{2} / \ell_{2}$ sparse recovery. Our contribution lies in obtaining better decoding time, and $O(\log (n / k))$ column sparsity via new techniques. The problem of improving the column sparsity to $O(\log (n / k))$ was explicitly stated in [GLPS10] as an open problem. We completely resolve the open problem. Moreover, as an important technical contribution, we introduce a different approach for sublinear-time optimalmeasurement sparse recovery tasks. Since this iterative loop is a crucial component of almost all algorithms in sublinear-time compressed sensing [IPW11, PS12, HIKP12a, GNP ${ }^{+}$13, IKP14, Kap16, GLPS17, CKSZ17, Kap17, LNW18, NSWZ18], we believe our new approach and ideas will appear useful in the relevant literature, as well as be a starting point for re-examining sparse recovery tasks under a different lens, and obtaining improved bounds.

### 10.1.3 Notation

For $x \in \mathbb{R}^{n}$ we let $H(x, k)$ to be the set of the largest $k$ in magnitude coordinates of $x$. We also write $x_{S}$ for the vector obtained after zeroing out every $x_{i}, \notin S$, and $x_{-k}=x_{[n \backslash \backslash H(x, k)}$.

| Reference | Measurements | Decoding Time | Encoding Time |
| :--- | :--- | :--- | :--- |
| [Don06, CRT06b] | $k \log (n / k)$ | LP | $k \log (n / k)$ |
| [CCF02, CM06] | $\epsilon^{-2} k \log n$ | $\epsilon^{-1} n \log n$ | $\log n$ |
| [NT09a] | $k \log (n / k)$ | $n k \log (n / k)$ | $\log (n / k)$ |
| [CM04] | $\epsilon^{-2} k \log ^{2} n$ | $\epsilon^{-1} k \log ^{c} n$ | $\log ^{2} n$ |
| [CCF02, CM06] | $\epsilon^{-2} k \log ^{c} n$ | $\epsilon^{-1} k \log ^{2} n$ | $\log ^{c} n$ |
| [GLPS10] | $\epsilon^{-1} k \log (n / k)$ | $\epsilon^{-1} k \log ^{c} n$ | $\log (n / k) \cdot \log ^{2} k$ |
| Our result | $\epsilon^{-1} k \log (n / k)$ | $\epsilon^{-1} k \log ^{2}(n / k)$ | $\log (n / k)$ |

Table 10.1: (A list of $\ell_{2} / \ell_{2}$-sparse recovery results). We ignore the " $O$ " for simplicity. LP denotes the time of solving Linear Programs [CLS19], and the state-of-the-art algorithm takes $n^{\omega}$ time where $\omega$ is the exponent of matrix multiplication. The results in [Don06, CRT06b, NT09a] do not explicitly state the $\ell_{2} / \ell_{2}$ guarantee, but their approach obtains it by an application of the Johnson-Lindenstrauss Lemma; they also cannot facilitate $\epsilon<1$, obtaining thus only a 2 -approximation. The $c$ in previous work is a sufficiently large constant, not explicitly stated, which is defined by probabilistically picking an error-correcting code of short length and iterating over all codewords. We estimate $c \geq 4$. We note that our runtime is (almost) achieved

We use $\|\cdot\|_{p}$ to denote the $\ell_{p}$ norm of a vector, i.e. $\|x\|_{p}=\left(\sum_{i=1}^{n}\left|x_{i}\right|^{p}\right)^{1 / p}$.

### 10.1.4 Technical statements

We proceed with the definition of the $\ell_{2} / \ell_{2}$ sparse recovery problem.

Problem 10.1.1 ( $\ell_{2} / \ell_{2}$ sparse recovery). Given parameters $\epsilon, k$, $n$, and a vector $x \in \mathbb{R}^{n}$. The goal is to design some matrix $\Phi \in \mathbb{R}^{m \times n}$ and a recovery algorithm $\mathcal{A}$ such that we can output a vector $x^{\prime}$ based on measurements $y=\Phi x$,

$$
\left.\left\|x^{\prime}-x\right\|_{2} \leq(1+\epsilon) \min _{k \text {-sparse }} z \in \mathbb{R}^{n}\right] z-x \|_{2} .
$$

We primarily want to minimize $m$ (which is the number of measurements), the running time of $\mathcal{A}$ (which is the decoding time) and column sparsity of $\Phi$.

In table 10.1, we provide a list of the previous results and compare with ours. Here, we formally present our main result.

Theorem 10.1.2 (stronger $\ell_{2} / \ell_{2}$ sparse recovery). There exists a randomized construction of a linear sketch $\Phi \in \mathbb{R}^{m \times n}$ with $m=O\left(\epsilon^{-1} k \log (n / k)\right)$ and column sparsity $O(\log (n / k))$, such that given $y=\Phi x$, we can find an $O(k)$-sparse vector $x^{\prime} \in \mathbb{R}^{n}$ in $O(m \cdot \log (n / k))$ time such that

$$
\left\|x^{\prime}-x\right\|_{2} \leq(1+\epsilon) \min _{k \text {-sparse }}^{z \in \mathbb{R}^{n}} \mid\|z-x\|_{2} .
$$

holds with $9 / 10$ probability.

Remark 10.1.1. In the regime where $k$ is very close to $n$, for example $k=n / \operatorname{poly}(\log n)$, we get an exponential improvement on the column sparsity over [GLPS10]. In many applications of compressed sensing, this is the desired regime of interest, check for example Figure 8 from [BI08]: $n=71,542$ while $m \geq 10,000$, which corresponds to $k$ being very close to $n$.

Remark 10.1.2. As can be inferred from the proof, our algorithms runs in time $O\left((k / \epsilon) \log ^{2}(\epsilon n / k)+\right.$ $(k / \epsilon) \log (1 / \epsilon))$, which is slightly better than the one stated in Theorem 10.1. The algorithm in [HIKP12a] achieves also the slightly worse running time of $O((k / \epsilon) \log n \log (n / k))$. That algorithm was the first algorithm that achieved running time $O(n \log n)$ for all values of $k, \epsilon$ for which the measurement complexity remained sublinear, smaller than $\gamma n$, for some absolute constant $\gamma$. A careful inspection shows that our algorithm achieves running time that is always sublinear, as long as the measurement complexity remains smaller than $\gamma n$.

### 10.1.5 Overview of techniques and difference with previous work

This subsection is devoted to highlighting the difference between our approach and the approach of [GLPS10]. We first give a brief high-level description of the state of the art algorithm before our work, then discuss our techniques, and try to highlight why the previous approach could not obtain the stronger result we present in this paper. Lastly, we show how our ideas can be possibly applied to other contexts.

### 10.1.5.1 Summary of [GLPS10].

The algorithm of [GLPS10] consists of $O(\log k)$ rounds: in the $r$-th round the algorithm finds a constant fraction of the remaining heavy hitters. Beyond this iterative loop lies the following idea about achieving the $\ell_{2} / \ell_{2}$ guarantee: in order to achieve it, you can find all but $\frac{k}{3^{r}}$ heavy hitters $i$ such that $\left|x_{i}\right|^{2}=\Omega\left(\frac{2^{r} \epsilon}{k}\left\|x_{-k}\right\|_{2}^{2}\right)$. This means that the algorithm is allowed to "miss" a small fraction of the heavy hitters, depending on their magnitude. For example, if all heavy hitters are as small as $\Theta\left(\sqrt{\epsilon / k}\left\|x_{-k}\right\|_{2}\right)$, a correct algorithm may even not find any of them. This crucial observation leads naturally to the main iterative loop of combinatorial compressed sensing, which, as said before, loop proceeds in $O(\log k)$ rounds. Every round consists of an identification and an estimation step: in the identification step most heavy hitters are recognized, while in the estimation step most of them are estimated correctly. Although in the estimation step some coordinates might have completely incorrect estimates, this is guaranteed (with some probability) be fixed in a later round. The reason why this will be fixed is the following. If a coordinate $i$ is badly estimated, then it will appear very large in the residual signal and hence will be identified in later rounds, till it is estimated correctly. One can observe that the correct estimation of that round for coordinate
$i$ cancels out (remedies) the mistakes of previous rounds on coordinate $i$. Thus, the identification and estimation procedure, which are interleaved, work complementary to each other. The authors of [GLPS10] were the first that carefully managed to argue that identifying and estimating a constant fraction of heavy hitters per iteration, gives the optimal number of measurements.

More formally, the authors prove the following iterative loop invariant, where $k_{r}=$ $k 3^{-r}, \epsilon_{r}=\epsilon_{r} 2^{-r}$ for $r \in[R]$ with $R=\log k$ : Given $x \in \mathbb{R}^{n}$ there exists a sequence of vectors $\left\{x^{(r)}\right\}_{r \in[R]}$, such that $x^{(r+1)}=x^{(r)}-\widehat{x}^{(r)}$ and

$$
\begin{equation*}
\left\|(x-\widehat{x})_{-k_{r}}\right\|_{2}^{2} \leq\left(1+\epsilon_{r}\right)\left\|x_{-k_{r}}\right\|_{2}^{2} . \tag{10.1}
\end{equation*}
$$

In the end, one can apply the above inequality inductively to show that

$$
\left\|x-\sum_{r=1}^{R} \widehat{x}^{(r)}\right\|_{2}^{2} \leq(1+\epsilon)\left\|x_{-k}\right\|_{2}^{2}
$$

We now proceed by briefly describing the implementations of the identification and the estimation part of [GLPS10]. In the identification part, in which lies the main technical contribution of that work, every coordinate $i \in[n]$ is hashed to $O(k / \epsilon)$ buckets and in each bucket $O(\log (\epsilon n / k))$-measurement scheme based on error-correcting codes is used to identify a heavy hitter; the authors carefully use a random error-correcting code of length $O(\log \log (\epsilon n / k))$, so they afford to iterate over all codewords and employ a more sophisticated approach and use nearest-neighbor decoding. This difference is one of the main technical ideas that allow them to obtain $O\left(\epsilon^{-1} k \log (n / k)\right)$ measurements, beating previous work, but it is also the main reason why they obtain $k \cdot \operatorname{poly}(\log n)$ decoding time ${ }^{2}$ : performing

[^56]nearest neighbor decoding and storing the code per bucket incurs additional poly $(\log n)$ factors. Moreover, for every iteration $r$ they need to repeat the identification scheme $r$ times in order to bring down the failure probability to $2^{-r}$, so that they afford a unionbound over all iterations. This leads to an additional $O\left(\log ^{2} k\right)$ factor in the update time. The estimation step consists of hashing to $O(k / \epsilon)$ buckets and repeating $O(\log (1 / \epsilon))$ times. Since the identification step returns $O(k / \epsilon)$ coordinates, the $O(\log (1 / \epsilon))$ repetitions of the estimation step ensure that at most $k / 3$ coordinates out of the $O(k / \epsilon)$ will not be estimated correctly. This is a desired property, since it allows the algorithm to keep the $2 k$ coordinates with the largest estimates, subtract them from $x$ and iterate.

In the next section, we will lay out our approach which improves the decoding time and the column sparsity of [GLPS10]. The iterative procedure of [GLPS10] lies in the heart of most compressed sensing schemes, so we believe that this new approach could be applied elsewhere in the sparse recovery literature.

### 10.1.5.2 Our approach

As we mentioned before, our approach is totally different from previous work, avoiding the iterative loop that all algorithms before applied. Our algorithm consists of four steps, each one being a different matrix responsible for a different task. The first matrix, with a constant number of rows allows us to approximate the tail of the vector $x$, an approximation that will appear useful in the next step. The second matrix along with its decoding procedure, which should be regarded as the identification step, enables us to find a list $L$ of size $O(k / \epsilon)$
that contains $k$ coordinates ${ }^{3}$, which are sufficient for the $\ell_{2} / \ell_{2}$ guarantee. This matrix has $\left.O\left(\epsilon^{-1} k \cdot \log (\epsilon n / k)\right)\right)$ rows. The third matrix with its decoding procedure, which should be regarded as the pruning step, takes the aforementioned list $L$, and prunes it down to $O(k)$ coordinates, which are sufficient for the $\ell_{2} / \ell_{2}$ guarantee. This matrix again has $O\left(\epsilon^{-1} k\right.$. $\log (1 / \epsilon))$ rows, for a total of $O\left(\epsilon^{-1} k \log (n / k)\right)$ rows. The last matrix is a standard set-query sketch.

## Step 1: Tail Estimation

Lemma 10.1.3 (tail estimation). Let $c_{1} \geq 1$ denote some fixed constant. There is an oblivious construction of matrix $\Phi \in \mathbb{R}^{m \times n}$ with $m=O(\log (1 / \delta))$ and column sparsity $O(\log (1 / \delta))$ such that, given $\Phi x$, there is an algorithm that outputs a value $V \in \mathbb{R}$ in time $O(m)$ such that

$$
\frac{1}{10 k}\left\|x_{-c_{1} \cdot k}\right\|_{2}^{2} \leq V \leq \frac{1}{k}\left\|x_{-k}\right\|_{2}^{2}
$$

holds with probability $1-\delta$.

Our first step towards the way for stronger sparse recovery is the design a routine that estimates the $\ell_{2}$ norm of the tail of a vector $x \in \mathbb{R}^{n}$, which we believe might be interesting in its own right. More generally, our algorithm obtains a value $V$ such that $\frac{1}{10 k}\left\|x_{-c_{1} \cdot k}\right\|_{p}^{p} \leq V \leq$ $\frac{1}{k}\left\|x_{-k}\right\|_{p}^{p}$, using $O(1)$ measurements. Here $c_{1}$ is some absolute constant. To obtain this result we subsample the vector at rate $\Theta(1 / k)$ and then use a $p$-stable distribution to approximate the subsampled vector. While the upper bound is immediate, the Paley-Zygmund inequality

[^57]does not give a sufficient result for the lower bound, so more careful arguments are needed to prove the desired result. We obtain our result by employing a random walk argument.

One additional possible application in sparse recovery applications where a two-stage scheme is allowed, e.g. [DLTY06, $\mathrm{DDT}^{+} 08$ ], would be to first use the above routine to roughly estimate how many heavy coordinates exist, before setting up the measurements. For example, we could first run the above routine for $k=1,2,2^{2}, \ldots, 2^{\log n}$, obtaining values $V_{1}, V_{2}, V_{4}, \ldots, V_{\log n}$, and then use these values to estimate the size of the tail of the vector is, or equivalently approximate the size of the set of the heavy coordinates by a number $k^{\prime}$. We can then run a sparse recovery algorithm with sparsity $k^{\prime}$. Details can be found in Section 10.4.

Step 2: The Identification Step The goal of this step is to output a list $L$ of size $O(k / \epsilon)$ that contains a set of $k$ coordinates that are sufficient to satisfy the $\ell_{2} / \ell_{2}$ guarantee. The column sparsity we are shooting for at this point is $O(\log (\epsilon n / k))$, and the decoding time should be $O(m \log (\epsilon n / k))$.

Lemma 10.1.4 (identification sketch). There exists a randomized construction of a matrix $\Phi \in \mathbb{R}^{m \times n}$, with $m=O\left(\epsilon^{-1} k \cdot \log (\epsilon n / k)\right)$ and column sparsity $O(\log (\epsilon n / k))$, such that given $y=\Phi x$, one can find a set $L$ of size $O(k / \epsilon)$ in $O(m \log (\epsilon n / k))$ time, such that

$$
\exists T \subset S,|T| \leq k:\left\|x-x_{T}\right\|_{2} \leq(1+\epsilon)\left\|x_{-k}\right\|_{2}
$$

holds with probability at least $9 / 10$.

For this routine, we will set up a hierarchical separation of $[n]$ to trees. We will call this separation interval forest. we set $\tau=k / \epsilon$. Then we partition [n] into $\tau$ intervals of
length $q=n / \tau$, and set up an interval tree for each interval in the following way: every interval tree has branching factor

$$
\frac{\log q}{\log \log q}
$$

with the same height (this is consistent with the fact that every tree contains $q$ nodes). At the leaves of every interval tree there are the nodes of the corresponding interval.

Our approach consists is now the following. For each level of the interval forest we hash everyone to $k / \epsilon$ buckets in the following way: if two coordinates $i, i^{\prime}$ are in the same interval (node of the interval forest) they are hashed to the same bucket. The above property can be regarded as "hashing interval to buckets". Moreover, every $x_{i}$ is multiplied by a Gaussian random variable. We repeat this process for $\log \log (\epsilon n / k)$ per level.

The decoding algorithm is the following. First, we obtain a value $V$ using the routine in Step 1, Lemma 10.1.3. Then we proceed in a breadth-first (or depth-first, it will make no difference) search manner and find an estimate for every interval, similarly to [LNNT16, CH09], by taking the median of the $\log \log (\epsilon n / k)$ buckets it participates to. There are two technical things one needs to show: First, we should bound the decoding time, as well as the size of the output list $L$. Second, we need to show that there exists a set $T^{\prime}$ of size at most $k$ that satisfies the guarantee of the Lemma 10.1.4. For the first part, we show that the branching process defined by the execution of the algorithm is bounded due to the $\log \log (\epsilon n / k)$ repetitions per level. For the second part, we show that for every coordinate $i \in H(x, k)$ the probability that $i \in L$ is proportional to $k\left|x_{i}\right|^{2} /\left(\epsilon\left\|x_{-k}\right\|_{2}^{2}\right)$. Then we show that the expected $\ell_{2}^{2}$ mass of coordinates $i \in L \backslash H(x, k)$ is $\epsilon\left\|x_{-k}\right\|_{2}^{2}$. This suffices to give the desired guarantee for Lemma 10.1.4. We provide details in Section 10.2.

## Step 3: The Pruning Step

Lemma 10.1.5 (pruning sketch). Let $c_{2}, c_{3}>1$ denote two fixed constants. There exists a randomized construction of a matrix $\Phi \in \mathbb{R}^{m \times n}$, with $m=O\left(\epsilon^{-1} k \cdot \log (1 / \epsilon)\right)$, with column sparsity $O(\log (1 / \epsilon))$ such that the following holds :

Suppose that one is given a (fixed) set $L \subseteq[n]$ such that

$$
|L|=O(k / \epsilon), \quad \exists T \subset L,|T| \leq k:\left\|x-x_{T}\right\|_{2} \leq(1+\epsilon)\left\|x_{-k}\right\|_{2}
$$

Then one can find a set $S$ of size $c_{2} \cdot k$ in time $O(m)$, such that

$$
\left\|x-x_{S}\right\|_{2} \leq\left(1+c_{3} \cdot \epsilon\right)\left\|x_{-k}\right\|_{2}
$$

holds with probability 9/10.

We will now prune the list $L$ obtained from the previous step, to $O(k)$ coordinates. We are going to use $O\left(\epsilon^{-1} k \cdot \log (1 / \epsilon)\right)$ measurements. We hash every coordinate to $k / \epsilon$ buckets, combining with Gaussians, and repeating $O(\log (1 / \epsilon))$ times. A similar matrix was also used in [GLPS10], but there the functionality and the analysis were very different; moreover, the authors used random signs instead of Gaussians. We will heavily exploit the fact that a standard normal $g$ satisfies $\operatorname{Pr}[|g|<x]=O(x)$. The reasons why we need Gaussians is the following: if we have a lot of large coordinates which are equal and much larger than $\sqrt{\epsilon}\left\|x_{-k}\right\|_{2}$, we want to find all of them, but due to the use of random signs, they might cancel each other in a measurement. Switching to Gaussians is the easiest way of avoiding this undesirable case.

Our algorithm computes, for every $i \in L$, an estimate $\widehat{x}_{i}$ by taking the median of the $O(\log (1 / \epsilon))$ buckets it participates to, and then keeps the largest $O(k)$ coordinates in
magnitude to form a set $S$. We then show that these coordinates satisfy the $\ell_{2} / \ell_{2}$ guarantee. We will say that a coordinate is well-estimated if $\left|\widehat{x}_{i}\right|=\Theta\left(\left|x_{i}\right|\right) \pm \sqrt{\epsilon k^{-1}}\left\|x_{-k}\right\|_{2}$. For the analysis, we define a threshold $\tau=\left\|x_{-k}\right\|_{2} / \sqrt{k}$, and classify coordinates based on whether $\left|x_{i}\right| \geq \tau$ or not.

- In the case $\left|x_{i}\right| \geq \tau$ the expected mass of these coordinates $i \notin S$ is small;
- In the other case the number of coordinates $i$ with $\left|x_{i}\right|<\tau$ are $O(k)$. This allows us to employ an exchange argument, similar to previous work, e.g. [PW11], but more tricky due to the use of Gaussians instead of random signs.

We note that the way we compute the estimates $\widehat{x}_{i}$ is different from previous work: one would expect to divide the content of a bucket that $i$ hashed to by the coefficient assigned to $x_{i}$, in order to get an unbiased estimator, but this will not work. The details can be found in Section 10.3.

In the end, this step gives us a set $S$ suitable for our goal, but does not give good estimations of the coordinates inside that set. For that we need another, standard step.

Step 4: Set Query We estimate every coordinate in $S$ using a set query algorithm of Price [Pri11], obtaining the desired guarantee. This matrix needs only $O(k / \epsilon)$ measurements, and runs in $O(k)$ time, while having constant column sparsity.

Lemma 10.1.6 (set query, [Pri11]). For any $\epsilon \in(0,1 / 10]$. There exists a randomized construction of a matrix $\Phi \in \mathbb{R}^{m \times n}$, with $m=O(k / \epsilon)$ and column sparsity $O(1)$, such that given $y=\Phi x$ and $a$ set $S \subseteq[n]$ of size at most $k$, one can find $a k$-sparse vector $\widehat{x}_{S}$, supported
on $S$ in $O(m)$ time, such that

$$
\left\|\widehat{x}_{S}-x_{S}\right\|_{2}^{2} \leq \epsilon\left\|x_{[n] \backslash S}\right\|_{2}^{2} .
$$

holds with probability at least $1-1 / \operatorname{poly}(k)$.

Our Theorem 10.1.2 follows from the above four steps by feeding the output of each step to the next one. In the end, we rescale $\epsilon$. More specifically, by the identification step we obtain a set $L$ of size $O(k / \epsilon)$ which contains a subset of size $k$ that satisfies the $\ell_{2} / \ell_{2}$ guarantee. Then, the conditions for applying the pruning step are satisfied, and hence we can prune the set $L$ down to $O(k)$ coordinates, which satisfy the $\ell_{2} / \ell_{2}$ guarantee. Then we apply the set-query sketch to obtain estimates of these coordinates.

In what follows we ignore constant terms. The number of measurements in total is

$$
\underbrace{1}_{\text {tail estimation }}+\underbrace{(k / \epsilon) \log (\epsilon n / k)}_{\text {identification step }}+\underbrace{(k / \epsilon) \log (1 / \epsilon)}_{\text {pruning step }}+\underbrace{(k / \epsilon)}_{\text {set query }} .
$$

The decoding time equals

$$
\underbrace{1}_{\text {tail estimation }}+\underbrace{(k / \epsilon) \log (\epsilon n / k) \log (\epsilon n / k)}_{\text {identification step }}+\underbrace{(k / \epsilon) \log (1 / \epsilon)}_{\text {pruning step }}+\underbrace{k}_{\text {set query }}
$$

The column sparsity equals

$$
\underbrace{1}_{\text {tail estimation }}+\underbrace{\log (\epsilon n / k)}_{\text {identification step }}+\underbrace{\log (1 / \epsilon)}_{\text {pruning step }}+\underbrace{1}_{\text {set query }} .
$$

### 10.1.6 Possible applications of our approach to other problems

Exactly $k$-sparse signals. When the vector we have to output has to be $k$-sparse, and not $O(k)$-sparse, the dependence on $\epsilon$ has to be quadratic [PW11]. Our algorithm yields a
state of the art result for this case, too. One can observe that the analysis of the algorithm in the Identification Linear Sketch outputs a $O(k / \epsilon)$-sized set which contains a set $T$ of size $k$ that allows $\ell_{2} / \ell_{2}$ sparse recovery. Performing a CountSketch with $O\left(k / \epsilon^{2}\right)$ columns and $O(\log (k / \epsilon)$ rows, and following a standard analysis, one can obtain a sublinear algorithm with measurement complexity $O\left(\epsilon^{-1} k \log (\epsilon n / k)+\epsilon^{-2} k \log (k / \epsilon)\right)$. This is an improvemnet on both the runtime and measurement complexity over previous work [CCF02].

Block-Sparse Signals. Our algorithm easily extends to block-sparse signals. For a signal of block size $b$ we obtain sample complexity $O(k /(\epsilon b) \log (b n / k)+(k / \epsilon))$, with a running time nearly linear in $k$. This matches the sample complexity of previous super-linear algorithms [BCDH10, CIHB09], which also could not facilitate $\epsilon<1$.

Phaseless Compressed Sensing. In Phaseless Compressed Sensing, one wants to design a matrix $\Phi$ with $m$ rows, such that given $y=\Phi x$, one can find a vector $\widehat{x}$ such that $\min _{\theta \in[0,2 \pi]}\left\|x-e^{i \theta} \widehat{x}\right\|_{2} \leq(1+\epsilon)\left\|x_{-k}\right\|_{2}$. This problem has received a fair amount of attention [OYDS11, LV13, CBJC14, YLPR15, PYLR17, Nak17a, LN18], and the state of the art algorithm has $O(k \log n)$ measurement complexity [Nak17a, LN18]. One of the problems is that the iterative loop approach cannot be used here, since it is heavily based on the linearity of the sketch. However, our identification and pruning step do not use the linearity of the sketch, and work also with phaseless measurements. Previous algorithms such as [Nak17a, LN18] suffered a $k \log n$ factor in the number of measurements already from the first step, but this is avoidablue using our new approach. We hope to see an extension down this avenue that gives $O(k \log (n / k))$ measurements.

One-Bit Compressed Sensing. Another important subfield of Sparse Recovery is one-bit compressed sensing, where one has access only to one-bit measurements, i.e. $y=\operatorname{sign}(A x)$, where the sign function on vectors should be understood as pointwise application of the sign function on each entry. Sublinear algorithms appear in [Nak17b, Nak19], but they both do not obtain the optimal number of measurements in terms of $k$ and $n$, which is $k \log (n / k)$, but rather the slightly suboptimal $k \log n$. One of the most important reasons is that the iterative loop cannot be implemented in such a scenario. It is a natural question whether our new approach can give the optimal number of measurements. The thresholding step, namely the part where we take use $V$ to filter out non-heavy intervals cannot be implemented here, but perhaps there is still a way to make a similar argument. One first approach should be to show that sublinear decoding with optimal measurements is achieved using non-adaptive threshold measurements, such as in [KSW16] and $\left[\mathrm{BFN}^{+} 17\right]$ (note that the latter one uses adaptive measurements though).

Sparse Fourier Transform. The standard approach to discrete sparse Fourier transform, is to implement linear measurements by using Fourier measurements [GGI ${ }^{+} 02$, GMS05, HIKP12a, HIKP12b, Iwe13, IKP14, IK14, Kap16, CKSZ17, Kap17]. The idea is to hash the spectrum to $B$ buckets by carefully multiplying in the time-domain the vector $x$ with a sparse vector $z$. In the frequency domain this corresponds to convolving the spectrum of the $x$ with an approximation of the filter of an indicator function of an interval of length roughly $B$. Due to the Uncertainty Principle, however, one has to exchange measurement complexity and decoding time with the quality of the filter. For example, implementing hashing to buckets using "crude" filters leads to leakage in subsequent buckets, giving additional error terms.

When iterating as usual, these errors accumulate and make identification much harder. The sophisticated approach of [Kap17] manages to design an iterative algorithm, in the same vein with previous algorithms, which takes $O\left(\epsilon^{-1} k \log n\right)$ measurements. It would be interesting to see if the approach we suggest avoids some of the problems created by this iterative loop, and can give simpler and faster sparse Fourier transform schemes. It would be interesting to obtain such a result even using adaptive measurements. The work [CKSZ17] has some interesting ideas in the context of block-sparse vectors that could be relevant.

### 10.2 The Identification Linear Sketch

The goal of this section is to prove the following result,
Theorem 10.2.1 (Restatement of Lemma 10.1.4). Let $C_{L}>1$ be a fixed constant. There exists a randomized construction of a matrix $\Phi \in \mathbb{R}^{m \times n}$, with

$$
m=O\left(\epsilon^{-1} k \log (\epsilon n / k)\right)
$$

with column sparsity $O(\log (\epsilon n / k))$ such that given $y=\Phi x$, one can find in time $O(m \log (n / k))$ a set $L$ of size $C_{L} \cdot k / \epsilon$, such that

$$
\exists T \subset L,|T| \leq k:\left\|x-x_{T}\right\|_{2} \leq(1+\epsilon)\left\|x_{-k}\right\|_{2}
$$

with probability 9/10.

In Section 10.2.1, we provide the definition of sketching matrix $\Phi$ and present the decoding algorithm. We proved some concentration result in Section 10.2.2. We analyzed the running time of algorithm in Section 10.2.3. We proved the guarantees of the algorithm in Section 10.2.4. Finally, we bound the number of measurements in Section 10.2.5.

### 10.2.1 Design of the sketch and decoding algorithm

We are going to use a hierarchical separation of $[n]$ into intervals. We will call this separation an interval forest.

Before discussing the matrix, we need the following definitions. We define
Definition 10.2.1 (size of the each tree in the forest). Let

$$
\tau=k / \epsilon
$$

| Notation | Choice | Statement | Parameter |
| :--- | :--- | :--- | :--- |
| $C_{H}$ | 4 | Definition 10.2 .2 | $H$ |
| $C_{R}$ | 100 | Definition 10.2 .3 | $R$ |
| $C_{B}$ | $10^{5}$ | Definition 10.2 .4 | $B$ |
| $C_{0}$ | $10^{3}$ | Lemma 10.4.3 | Blow up on tail size |
| $C_{L}$ | $10^{4}$ | Lemma 10.2.5 | $L$ |
| $\eta$ | $1 / 9$ | Lemma $10.2 .2,10.2 .3$ | Shrinking factor on $V$ |
| $\zeta$ | $1 / 4000$ | Lemma $10.2 .2,10.2 .3$ | $\zeta \leq \eta / 400$ |

Table 10.2: Summary of constants in Section 10.2, the column "Parameter" indicates which parameter is depending on that constant. Note that constants $C_{H}, C_{R}, C_{B}, C_{0}, \eta$ are used in both algorithm and analysis, but constants $C_{L}$ and $\zeta$ are only being used in analysis. $C_{L}$ is the related to the guarantee of the output of the algorithm.
assuming that $k / \epsilon \leq n / 16$. The size of each tree in the forest is

$$
q=n / \tau=n \epsilon / k
$$

Definition 10.2.2 (degree and height of the interval forest). Let $C_{H}>1$ be a sufficiently large constant such that

$$
(\log q / \log \log q)^{C_{H} \log q / \log \log q} \geq q
$$

Let $D$ denote the degree of the tree, and let $H$ denote the height of the tree. We set $D$ and $H, D=\lceil\log q / \log \log q\rceil$, and $H=\left\lceil C_{H} \log q / \log \log q\right\rceil$.

Definition 10.2.3 (number of repetitions per level). Let $R$ denote the number of repetitions in each level. Let $C_{R}>1$ denote some sufficiently large constant. We set $R=C_{R} \log \log q$.

For $\ell \in\{0,1, \ldots, H\}$ we define $\mathcal{J}_{\ell}$, which is a family of sets. Every set $\mathcal{J}_{\ell}$ is a decomposition of $[n]$ to $\tau D^{\ell}$ intervals of (roughly) the same length. The set $\mathcal{J}_{0}$ is a decomposition
of $[n]$ to $\tau$ intervals of (roughly) the same length length $q$. If needed, we can round $q$ to a power of 2 . This means that

$$
\mathcal{J}_{0}=\left\{I_{0,1}, I_{0,2}, \ldots\right\}
$$

where

$$
I_{0,1}=[1, \tau], I_{0,2}=[\tau+1,2 \tau], \ldots
$$

We can conceptualize these sets as a forest consisting of $\tau$ trees, each of branching factor $D$ and height $H$, where the $\ell$-th level partitions $[n]$ into disjoint $\tau D^{\ell}$ intervals of length $n /\left(\tau \cdot D^{\ell}\right)=q \cdot D^{-\ell}$. For $\ell \in\{0, \ldots, H\}$, interval $I_{\ell, j}$ is decomposed to $D$ disjoint and continuous intervals

$$
I_{\ell+1, j \cdot D+1}, \ldots, I_{\ell+1,(j+1) \cdot D}
$$

of the same length, except possibly the last interval.
We say that an interval $I_{\ell+1, j}$ is a child of interval $I_{\ell, j^{\prime}}$ if $j^{\prime}=\lceil j / D\rceil$.
Definition 10.2.4 (sketching matrix $\Phi$ ). Let $C_{B}>1$ be a sufficiently large constant. Let $B=C_{B} k / \epsilon$. Let matrices $\Phi^{(1)}, \ldots, \Phi^{(H)}$, where every matrix $\Phi^{(\ell)}$ consists of $R$ submatrices $\left\{\Phi_{r}^{(\ell)}\right\}_{r \in[R]}$. For every $\ell \in[H]$ and $r \in[R]$, we pick 2 -wise independent hash functions $h_{\ell, r}:\left[\tau D^{\ell}\right] \rightarrow[B]$.

We define measurement $y_{\ell, r, b}=\left(\Phi_{r}^{(\ell)} x\right)_{b}$ as:

$$
y_{\ell, r, b}=\sum_{j \in h_{\ell, r}^{-1}(b)} \sum_{i \in I_{\ell, j}} g_{i, \ell, r} x_{j}
$$

where $g_{i, \ell, r} \sim \mathcal{N}(0,1)$, i.e. independent standard Gaussians.
We slightly abuse notation and treat $y$ as matrix the mapping to vector should be clear.

Note that $C_{B}$ should be chosen such that $C_{B} \gg C_{0}$, where $C_{0}$ appears in tail estimation in Lemma 10.4.3.

### 10.2.2 Concentration of estimation

In the following lemmata, $\zeta, \eta \in(0,1)$ are absolute constants with $1>\eta>1 / 10>\zeta$ (See a choice for our application in Table 10.2). The exact values of the constants will be chosen below.

The following lemma handles the probability of detecting a heavy interval at level $\ell$.
Lemma 10.2.2 (handling the probability of catching a heavy hitter). Let $\bar{V}$ be the value in Line 8 of Algorithm 10.1. Let $V=\epsilon \bar{V}$ be the value in Line 9 of Algorithm 10.1. Let $j^{\prime} \in T_{\ell-1}$, and let $j$ be one of its children. Let $z_{j}$ be defined as follows,

$$
z_{j}=\operatorname{median}_{r \in[R]}\left|y_{\ell, r, h_{\ell, r}(j)}\right|^{2}
$$

If $\left\|x_{I_{\ell, j}}\right\|_{2}^{2} \geq C_{j} \frac{\epsilon}{k}\left\|x_{-k}\right\|_{2}^{2}$, where $C_{j} \geq 2$, then with probability $1-C_{j}^{-R / 6}$ (over the randomness of $h_{\ell, r}$ and $g_{i, \ell, r}$ for $r \in[R], i \in[n]$ in Definition 10.2.4), we have that

$$
z_{j} \geq \eta V
$$

Proof. Fix $r \in[R]$. Let $b=h_{\ell, r}(j)$, and define $J=\cup_{t \in h_{\ell, r}^{-1}(b)} I_{\ell, t}$. We observe that

$$
\left|y_{\ell, r, b}\right|^{2}=\left\|x_{J}\right\|_{2}^{2} g^{2}, \text { where } g \sim \mathcal{N}(0,1)
$$

```
Algorithm 10.1 Interval-forest Sparse Recovery
    procedure IntervalforestSparseRecovery \((x, n, k, \epsilon) \quad \triangleright\) Theorem 10.2.1
        Choose constants \(C_{H}, C_{R}, \eta, C_{0} \quad \triangleright\) According to Table 10.2
        \(\tau \leftarrow(k / \epsilon)\)
        \(q \leftarrow n / \tau \quad \triangleright\) Definition 10.2.1
        \(H \leftarrow\left\lceil C_{H} \log q / \log \log q\right\rceil \quad \triangleright\) Definition 10.2.2
        \(D \leftarrow\lceil\log q / \log \log q\rceil \quad \triangleright\) Definition 10.2.2
        \(R \leftarrow\left\lceil C_{R} \log \log q\right\rceil \quad \triangleright\) Definition 10.2.3
        \(\bar{V} \leftarrow \operatorname{LpLpTaiLEstimation}\left(x, k, 2, C_{0}, 1 / 100\right) \quad \triangleright\) Algorithm 10.3
        \(V \leftarrow \epsilon \bar{V} \quad \triangleright\) Lemma 10.2.2
        \(T_{0} \leftarrow\left\{I_{0,1}, \ldots, I_{0, \tau}\right\}\)
        for \(\ell=1 \rightarrow H\) do
            \(T_{\ell} \leftarrow \operatorname{RecursiveBTree}\left(\ell, R, D, \eta, T_{\ell-1}, V\right) \quad \triangleright\) Lemma 10.2.5
        end for
        \(L \leftarrow T_{H}\)
        return \(L \quad \triangleright\) Lemma 10.2.6
    end procedure
    procedure RecursiveBTree \((\ell, R, D, \eta, T, V)\)
        \(T^{\prime} \leftarrow \emptyset\)
        for \(t \in T\) do
            Let \(I_{\ell, j_{1}}, I_{\ell, j_{2}}, \ldots, I_{\ell, j_{D}}\) denote the child intervals of \(I_{\ell-1, t}\)
            for \(p \in[D]\) do
                    \(z_{j_{p}} \leftarrow \operatorname{median}_{r \in[R]}\left|y_{\ell, r, h_{\ell, r}\left(j_{p}\right)}\right|^{2} \quad \triangleright\) Definition 10.2.4
                end for
                if \(z_{j_{p}} \geq \eta V\) then
                    \(T^{\prime} \leftarrow T^{\prime} \cup\left\{j_{p}\right\}\)
                end if
        end for
        return \(T^{\prime}\)
    end procedure
```

By property of Gaussian distribution, we have

$$
\operatorname{Pr}\left[\left|y_{\ell, r, b}\right|^{2} \leq\left(\eta / C_{j}\right)\left\|x_{J}\right\|_{2}^{2}\right] \leq \frac{2}{\sqrt{2 \pi}} \sqrt{\frac{\eta}{C_{j}}} \leq \sqrt{\frac{2 \eta}{\pi C_{j}}},
$$

It implies, since $\left\|x_{J}\right\|_{2}^{2} \geq\left\|x_{I_{\ell, j}}\right\|_{2}^{2} \geq C_{j} \frac{\epsilon}{k}\left\|x_{-k}\right\|_{2}^{2}$, that

$$
\operatorname{Pr}\left[\left|y_{\ell, r, b}\right|^{2} \leq \eta \frac{\epsilon}{k}\left\|x_{-k}\right\|_{2}^{2}\right] \leq \sqrt{\frac{2 \eta}{\pi C_{j}}}
$$

Since Lemma 10.4.3, we have $\bar{V} \leq \frac{1}{k}\left\|x_{-k}\right\|_{2}^{2}$. Because $V=\epsilon \bar{V}$,

$$
\operatorname{Pr}\left[\left|y_{\ell, r, b}\right|^{2} \leq \eta V\right] \leq \sqrt{\frac{2 \eta}{\pi C_{j}}}
$$

The $R$ repetitions ensure that the failure probability can be driven down to $C_{j}^{-R / 6}$, because

$$
\begin{aligned}
\operatorname{Pr}\left[z_{j} \leq \eta V\right] & \leq\binom{ R}{R / 2} \cdot\left(\sqrt{\frac{2 \eta}{\pi C_{j}}}\right)^{R / 2} \\
& \leq 2^{R} \cdot(2 \eta / \pi)^{R / 4} \cdot C_{j}^{-R / 4} \\
& \leq\left(2^{R} \cdot(2 \eta / \pi)^{R / 4} \cdot 2^{-R / 12}\right) \cdot C_{j}^{-R / 6} \\
& \leq\left(2^{11} \cdot(2 / 9 \pi)^{3}\right)^{R / 12} \cdot C_{j}^{-R / 6} \\
& \leq C_{j}^{-R / 6}
\end{aligned}
$$

where the first step follows from a union bound, the third step follows from $C_{j} \geq 2$, and the forth step follows from $\eta \leq 1 / 9$.

The following lemma handles the probability of a non-heavy interval being considered "heavy" by the algorithm at level $\ell$.

Lemma 10.2.3 (handling the probability of false positives). Let $V$ be the value in Line 9 of Algorithm 10.1. Let $j^{\prime}$ be an index in $T_{\ell-1}$, and let $j$ be one of its children. If $\left\|x_{I_{\ell, j}}\right\|_{2}^{2} \leq$
$\zeta_{k}^{\epsilon}\left\|x_{-C_{0} k}\right\|_{2}^{2}$, then with probability $1-2^{-R / 3}$ (over the randomness of $h_{\ell, r}$ and $g_{i, \ell, r}$ for $r \in$ $[R], i \in[n]$ in Definition 10.2.4) we have that

$$
z_{j}<\eta V .
$$

Proof. Fix $\ell \in[H]$ and consider the set $H_{\ell}$ that contains the $C_{0} k$ coordinates $j^{\prime \prime}$ with the largest $\left\|x_{I_{\ell, j^{\prime \prime}}}\right\|_{2}^{2}$ values. Define $H_{\ell}^{(j)}=H_{\ell} \backslash\{j\}$. Fix $r \in[R]$ and observe that by a unionbound we get that

$$
\operatorname{Pr}\left[\exists j^{\prime \prime} \in H_{\ell}^{(j)} \mid h_{\ell, r}(j)=h_{\ell, r}\left(j^{\prime \prime}\right)\right] \leq C_{0} k \cdot \frac{1}{C_{B} k / \epsilon}=\frac{C_{0} \epsilon}{C_{B}} \leq \frac{1}{20},
$$

because $C_{B} \geq 20 C_{0}$.
We condition on the event $\forall j^{\prime \prime} \in H_{\ell}^{(j)}: h_{\ell, r}(j) \neq h_{\ell, r}\left(j^{\prime \prime}\right)$. A standard calculation now shows that

$$
\begin{aligned}
\mathbb{E}\left[\left|y_{\ell, r, h_{r, \ell}(j)}\right|^{2}\right] & \leq\left\|x_{I_{\ell, j}}\right\|_{2}^{2}+\frac{1}{C_{B}} \frac{\epsilon}{k}\left\|x_{-C_{0} k}\right\|_{2}^{2} \\
& \leq \zeta \frac{\epsilon}{k}\left\|x_{-C_{0} k}\right\|_{2}^{2}+\frac{1}{C_{B}} \frac{\epsilon}{k}\left\|x_{-C_{0} k}\right\|_{2}^{2} \\
& \leq \frac{2 \zeta \epsilon}{k}\left\|x_{-C_{0} k}\right\|_{2}^{2}
\end{aligned}
$$

where the last step follows from $\frac{1}{C_{B}}<\zeta$.
We now apply Markov's inequality to obtain

$$
\begin{array}{rlrl}
\operatorname{Pr}\left[\left|y_{\ell, r, h_{\ell, r}(j)}\right|^{2} \geq \eta V\right] & \leq \operatorname{Pr}\left[\left|y_{\ell, r, h_{\ell, r}(j)}\right|^{2} \geq \frac{\eta \epsilon}{10 k}\left\|x_{-C_{0} k}\right\|_{2}^{2}\right] & \\
& \leq \frac{\frac{2 \zeta \epsilon}{k}\left\|x_{-C_{0} k}\right\|_{2}^{2}}{\frac{\eta \epsilon}{10 k}\left\|x_{-C_{0} k}\right\|_{2}^{2}} \\
& =\frac{20 \zeta}{\eta} & & \\
& \leq \frac{1}{20}, & \text { by } \zeta \leq \eta / 400 .
\end{array}
$$

By a union bound, the unconditional probability $\operatorname{Pr}\left[\left|y_{\ell, r, h_{\ell, r}(j)}\right|^{2} \geq \eta V\right] \leq \frac{1}{10}$. Finally, we can upper bound the probability that $z_{j}=\operatorname{median}_{r \in[R]}\left|y_{\ell, r, h_{\ell, r}(j)}\right|^{2}$ is greater than $\eta V$,

$$
\begin{aligned}
\operatorname{Pr}\left[z_{j} \geq \eta V\right] & \leq\binom{ R}{R / 2}(1 / 10)^{R / 2} \\
& <2^{R} \cdot 2^{-\frac{3}{2} R} \\
& =2^{-\frac{R}{2}} \\
& <2^{-\frac{R}{3}}
\end{aligned}
$$

### 10.2.3 Analysis of running time

Lemma 10.2.4 (bounds of $D, H$ with respect to $R$ ). Let $D, H$ as in Definition 10.2.2. It holds that $D \leq 2^{\frac{R}{6}-10}$, and $H \leq 2^{\frac{R}{6}-10}$.

Proof. Since $H \geq D$, it suffices to prove the claim only for $H$.

$$
H=C_{H} \frac{\log q}{\log \log q}<C_{H} \log q=C_{H} \log (\epsilon n / k)=C_{H} 2^{\log \log (\epsilon n / k)} \leq 2^{\frac{R}{6}-10},
$$

where the third step follows from $q=\epsilon n / k$, and the last step follows from $\log \log (\epsilon n / k)+$ $\log C_{H} \leq\left(C_{R} / 6\right) \log \log (\epsilon n / k)-10$ and note that $\log \log (\epsilon n / k) \geq 2$ because we assume $k / \epsilon \leq n / 16$.

Lemma 10.2.5 (running time). Let $R$ as in Definition 10.2.3 and D, H as in Definition 10.2.2. Let $T_{H}$ be the set obtained by applying procedure RecursiveBTree $H$ times (lines

11-13) of Algorithm 10.1, and let $C_{L}>1$ be some sufficiently large absolute constant. With probability $1-H \cdot 2^{-R / 6+1}$ we have that:

- $\left|T_{H}\right| \leq C_{L} \cdot k / \epsilon$,
- The running time of BTreeSparseRecovery is

$$
O\left(\epsilon^{-1} k \cdot \log (\epsilon n / k) \cdot D\right)
$$

Proof. Let $C_{L}=2\left(C_{0}+1 / \zeta\right)$.
First, it is easy to see that $\left|T_{0}\right|$ is bounded by

$$
\left|T_{0}\right|=\tau=k / \epsilon<C_{L} k / \epsilon
$$

We claim that if we condition on the event that $\left|T_{\ell-1}\right| \leq C_{L} k / \epsilon$, then with probability $1-2^{-R / 6+1},\left|T_{\ell}\right| \leq C_{L} k / \epsilon$. The proof of both bullets will then follow by a union-bound over all $H$ levels. Indeed, consider the set $Q_{\ell}$ containing the $\left|T_{\ell-1}\right| D \leq C_{L} \cdot(k / \epsilon) \cdot D$ coordinates $j$ that are children of some $j^{\prime} \in T_{\ell-1}$. Define

$$
B_{\ell}=\left\{j \in Q_{\ell} \left\lvert\,\left\|x_{I_{\ell, j}}\right\|_{2}^{2} \leq \zeta \frac{\epsilon}{k}\left\|x_{-C_{0} k}\right\|_{2}^{2}\right.\right\} .
$$

By definition of $B_{\ell}$ and $Q_{\ell}$, we have

$$
\left|B_{\ell}\right| \leq\left|Q_{\ell}\right| \leq C_{L} \cdot(k / \epsilon) \cdot D .
$$

Moreover, Lemma 10.2.3 gives

$$
\forall j \in B_{\ell}, \operatorname{Pr}\left[z_{j} \geq \eta V\right] \leq 2^{-R / 3}
$$

Define random variables $W_{j}$ to be 1 if $z_{j} \geq \eta V$, and 0 otherwise. Then

$$
\mathbb{E}\left[\sum_{j \in B_{\ell}} W_{j}\right] \leq \frac{C_{L} k}{\epsilon} D \cdot 2^{-R / 3}
$$

An application of Markov's inequality gives

$$
\operatorname{Pr}\left[\sum_{j \in B_{\ell}} W_{j} \geq \frac{C_{L} k}{2 \epsilon} D \cdot 2^{-R / 6}\right] \leq 2^{-R / 6+1} .
$$

Conditioning on $\sum_{j \in B_{\ell}} W_{j} \leq \frac{C_{L} k}{2 \epsilon} D \cdot 2^{-R / 6}$ we will upper bound the size of $T_{\ell}$.
First, observe that there exist at most $\left(C_{0} k+k /(\zeta \epsilon)\right) j \in \mathcal{J}_{\ell}$ for which $\left\|x_{I_{\ell, j}}\right\|_{2}^{2}>$ $\zeta \frac{\epsilon}{k}\left\|x_{-C_{0} k}\right\|_{2}^{2}$. This gives

$$
\begin{equation*}
\left|T_{\ell}\right| \leq\left(C_{0} k+\frac{k}{\zeta \epsilon}\right)+\frac{C_{L} k}{2 \epsilon} D 2^{-R / 6} \tag{10.2}
\end{equation*}
$$

If $C_{L} \geq 2\left(C_{0}+1 / \zeta\right)$, then we can upper bound the first term in Eq. (10.2),

$$
C_{0} k+\frac{k}{\zeta \epsilon} \leq\left(C_{0}+\frac{1}{\zeta}\right) \frac{k}{\epsilon} \leq \frac{1}{2} \frac{C_{L} k}{\epsilon} .
$$

For the second term in Eq. (10.2), we can show that

$$
\frac{C_{L} k}{2 \epsilon} \cdot D 2^{-R / 6} \leq \frac{1}{2} \frac{C_{L} k}{\epsilon}
$$

or equivalently $D \leq 2^{R / 6}$, which holds by Lemma 10.2.4.
We have $D$ levels, and at each level $\ell$ we have $\left|T_{\ell}\right|=O(k / \epsilon)$, conditioned on the aforementioned events happening. The children of $T_{\ell}$ is then $O(k / \epsilon \cdot D)$. Since we have $R$ repetitions the total running time per level is $O((k / \epsilon) \cdot D \cdot R)$, and the total running time is

$$
O((k / \epsilon) \cdot D \cdot R \cdot H)=O\left((k / \epsilon) \cdot D \cdot \log \log q \cdot \frac{\log q}{\log \log q}\right)=O((k / \epsilon) \cdot D \cdot \log (\epsilon n / k))
$$

where the first step follows from definition of $R$ and $H$, and the last step follows from definition of $q$.

Therefore, it gives the desired result.

### 10.2.4 Guarantees of the algorithm

Lemma 10.2.6 (guarantees). Let $L=T_{H}$ be the set obtained by applying procedure RECURSiveBTree $R$ times (lines 11-13) of Algorithm 10.1, we have that, with probability 9/10, there exist $T^{\prime} \subseteq L$ of size at most $k$, such that

$$
\left\|x-x_{T^{\prime}}\right\|_{2}^{2} \leq(1+\epsilon)\left\|x_{-k}\right\|_{2}^{2}
$$

Proof. Define

$$
\mathcal{H}=\left\{j \in H(x, k)\left|\exists C_{j} \geq 2,\left|x_{j}\right|^{2} \geq C_{j} \frac{\epsilon}{k}\left\|x_{-k}\right\|_{2}^{2}\right\} .\right.
$$

Moreover, associate every $j \in \mathcal{H}$ with its corresponding $C_{j}=\frac{\left|x_{j}\right|^{2}}{\frac{\epsilon}{k}\left\|x_{-k}\right\|_{2}^{2}}$.
Pick $j \in \mathcal{H}$. Let also $j_{1}, j_{2}, \ldots, j_{H}$, be numbers such that

$$
j \in I_{1, j_{1}}, j \in I_{2, j_{2}}, \ldots, j \in I_{R, j_{H}}
$$

For any $t \in\{1, \ldots, H-1\}$, if $I_{t, j_{t}} \in T_{t}$, then $I_{t+1, j_{t+1}} \in T_{t+1}$ with probability $1-C_{j}^{-R / 6}$, by Lemma 10.2.2. Since $C_{j} \geq 2$, this allows us to take a union bound over all $H$ levels and claim that with probablity $1-H C_{j}^{-R / 6}, j \in T_{R}$. For $j \in \mathcal{H}$ define random variable to $\delta_{j}$ to
be 1 if $j \notin T_{H}$.

$$
\begin{aligned}
\mathbb{E}\left[\delta_{j} x_{j}^{2}\right] & \leq H \cdot C_{j}^{-R / 6} x_{j}^{2} \\
& \leq H \cdot C_{j}^{-R / 6} C_{j} \frac{\epsilon}{k}\left\|x_{-k}\right\|_{2}^{2} \\
& =H \cdot C_{j}^{-R / 6+1} \frac{\epsilon}{k}\left\|x_{-k}\right\|_{2}^{2} \\
& \leq \frac{\epsilon}{80 k}\left\|x_{-k}\right\|_{2}^{2},
\end{aligned}
$$

where the first step follows by definition of $\delta_{j}$, the second step follows by the fact that $j \in \mathcal{H}$, and the last step follows by Lemma 10.2.4. Since $|\mathcal{H}| \leq k$, we have that

$$
\mathbb{E}\left[\sum_{j \in \mathcal{H}} \delta_{j} x_{j}^{2}\right] \leq|\mathcal{H}| \cdot \frac{\epsilon}{80 k}\left\|x_{-k}\right\|_{2}^{2} \leq \frac{\epsilon}{80}\left\|x_{-k}\right\|_{2}^{2}
$$

Then applying Markov's inequality, we have

$$
\operatorname{Pr}\left[\sum_{j \in \mathcal{H}} \delta_{j} x_{j}^{2}>\frac{\epsilon}{2}\left\|x_{-k}\right\|_{2}^{2}\right] \leq \frac{1}{40} .
$$

We condition on the event $\sum_{j \in \mathcal{H}} \delta_{j} x_{j}^{2} \leq \frac{\epsilon}{2}\left\|x_{-k}\right\|_{2}^{2}$. Setting $T^{\prime}=\mathcal{H} \cap T_{H}$ we observe that

$$
\begin{aligned}
\left\|x-x_{T^{\prime}}\right\|_{2}^{2} & =\sum_{j \in \mathcal{H} \cap H(x, k)} \delta_{j} x_{j}^{2}+\left\|x_{H(x, k) \backslash \mathcal{H}}\right\|_{2}^{2}+\left\|x_{[n \backslash \backslash H(x, k)}\right\|_{2}^{2} \\
& \leq \frac{\epsilon}{2}\left\|x_{-k}\right\|_{2}^{2}+k \cdot \frac{2 \epsilon}{k}\left\|x_{-k}\right\|_{2}^{2}+\left\|x_{-k}\right\|_{2}^{2} \\
& \leq(1+3 \epsilon)\left\|x_{-k}\right\|_{2}^{2} .
\end{aligned}
$$

where the second step follows by the bound on $\sum_{j \in T} \delta_{j} x_{j}^{2}$ and the fact that every $j \notin T$ satisfies $\left|x_{j}\right|^{2} \leq(2 \epsilon / k)\left\|x_{-k}\right\|_{2}^{2}$.

Rescaling for $\epsilon$ we get the desired result.

### 10.2.5 Bounding the number of measurements

In this subsection, we prove that

Claim 10.2.7 (\#measurements). The number of measurements is $O\left(\epsilon^{-1} k \cdot \log (\epsilon n / k)\right)$.

Proof. Recall the definition of $\tau$ and $q$,

$$
\tau=k / \epsilon, \quad q=n / \tau .
$$

We thus have the following bound on the number of measurements:
$B \cdot H \cdot R=C_{B}(k / \epsilon) \cdot C_{H} \frac{\log (\epsilon n / k)}{\log \log (\epsilon n / k)} \cdot C_{R} \log \log (\epsilon n / k)=O((k / \epsilon) \log (\epsilon n / k))$.

### 10.3 The Pruning Linear Sketch

The goal of this section is to prove Theorem 10.3.1.
Theorem 10.3.1 (Restatement of Lemma 10.1.5). Let $C_{L}, \alpha, \beta>1$ be three fixed constants. There exists a randomized construction of a matrix $\Phi \in \mathbb{R}^{m \times n}$, with $m=O((k / \epsilon) \cdot \log (1 / \epsilon))$, with column sparsity $O(\log (1 / \epsilon))$ such that the following holds :

Suppose that one is given a set $L \subseteq[n]$ such that

$$
|L|=C_{L} \cdot k / \epsilon, \quad \exists T \subset L,|T| \leq k:\left\|x-x_{T}\right\|_{2} \leq(1+\epsilon)\left\|x_{-k}\right\|_{2}
$$

Then procedure Prune (Algorithm 10.2) can find a set $S$ of size $\beta \cdot k$ in time $O(m)$, such that

$$
\left\|x-x_{S}\right\|_{2} \leq(1+\alpha \cdot \epsilon)\left\|x_{-k}\right\|_{2}
$$

holds with probability 9/10.

In Section 10.3.1, we provide some basic definitions and description of our algorithm. We analyze the coordinates from several perspectives in Section 10.3.2. We prove the correctness of our algorithm in Section 10.3.3 and analyze time, number of measurements column sparsity, success probability of algorithm in Section 10.3.4.

### 10.3.1 Design of the sketching matrix, and helpful definitions

Definition 10.3.1 (sketching matrix $\Phi$ ). Let $C_{R}, C_{B}>1$ be absolute constants. Let $R=$ $C_{R} \log (1 / \epsilon)$. Let $B=C_{B} k / \epsilon$. For $r \in[R]$, we pick 2-wise independent hash function $h_{r}:[n] \rightarrow[B]$, as well as normal random variables $\left\{g_{i, r}\right\}_{i \in[n], r \in[R]}$ and take measurements

$$
y_{r, b}=\sum_{i \in h_{r}^{-1}(b)} x_{i} g_{i, r} .
$$

| Notation | Choice | Statement | Parameter |
| :--- | :--- | :--- | :--- |
| $C_{R}$ | $10^{4}+500 C_{L}$ | Definition 10.3.1 | $R$ |
| $C_{B}$ | $5 \times 10^{5}$ | Definition 10.3.1 | $B$ |
| $C_{g}$ | $4 / 5$ | Fact 10.3.2 | Gaussian variable |
| $C_{L}$ | $10^{4}$ | Theorem 10.3.1 | $L$ |
| $\alpha$ | 5 | Theorem 10.3.1 | Blow up on $\epsilon$ |
| $\beta$ | 100 | Theorem 10.3.1 | Blow up on $k$ |

Table 10.3: Summary of constants in Section 10.3, the column "Parameter" indicates which parameter is depending on that constant. Note that set $L$ is the input of the algorithm in Section 10.3 and the output of the algorithm in Section 10.2.

Given the set $L$, for every $i \in L$ we calculate

$$
z_{i}=\underset{r \in[R]}{\operatorname{median}}\left|y_{r, h_{r}(i)}\right|,
$$

and keep the indices $i$ with the $\beta k$ largest $z_{i}$ values to form a set $S$ of indices, for some absolute constant $\beta$ sufficiently large. We describe this pruning step in Algorithm 10.2. For the analysis, we define the threshold

$$
\begin{equation*}
\tau=\left\|x_{-k}\right\|_{2} / \sqrt{k} \tag{10.3}
\end{equation*}
$$

We will need the following standard fact about the Gaussian distribution. Then we proceed with a series of definitions and lemmata.

Fact 10.3 .2 (property of Gaussian). Suppose $x \sim \mathcal{N}\left(0, \sigma^{2}\right)$ is a Gaussian random variable. For any $t \in(0, \sigma]$ we have

$$
\operatorname{Pr}[x \geq t] \in\left[\frac{1}{2}\left(1-\frac{4}{5} \frac{t}{\sigma}\right), \frac{1}{2}\left(1-\frac{2}{3} \frac{t}{\sigma}\right)\right] .
$$

Similarly, if $x \sim \mathcal{N}\left(\mu, \sigma^{2}\right)$, for any $t \in(0, \sigma]$, we have

$$
\operatorname{Pr}[|x| \geq t] \in\left[1-\frac{4}{5} \frac{t}{\sigma}, 1-\frac{2}{3} \frac{t}{\sigma}\right] .
$$

```
Algorithm 10.2 The Prune Procedure
    procedure \(\operatorname{Prune}(x, n, k, \epsilon, L)\)
                                    \(\triangleright\) Theorem 10.3.1
        \(R \leftarrow C_{R} \log (1 / \epsilon)\)
        \(B \leftarrow C_{B} k / \epsilon\)
        for \(r=1 \rightarrow R\) do
            Sample \(h_{r}:[n] \rightarrow[B] \sim 2\)-wise independent family
            for \(i=1 \rightarrow n\) do
                Sample \(g_{i, r} \sim \mathcal{N}(0,1)\)
            end for
        end for
        for \(r=1 \rightarrow R\) do
            for \(b=1 \rightarrow B\) do
                \(y_{r, b} \leftarrow \sum_{i \in h_{r}^{-1}(b)} x_{i} g_{i, r}\)
            end for
        end for
        for \(i \in L\) do
            \(z_{i} \leftarrow \operatorname{median}_{r \in[R]}\left|y_{r, h_{r}(i)}\right|\)
        end for
        \(S \leftarrow\left\{i \in L: z_{i}\right.\) is in the top \(\beta k\) largest coordinates in vector \(\left.z\right\}\)
        return \(S\)
    end procedure
```

The form we will need is the following:

$$
\underset{g \sim \mathcal{N}(0,1)}{\operatorname{Pr}}[|g| \leq t] \leq \frac{4}{5} t
$$

Thought the analysis, for convenience we will set $C_{g}=4 / 5$. Another form we will need is:

$$
\operatorname{Pr}_{g \sim \mathcal{N}(0,1)}\left[|g| \in\left[\frac{1}{3 C_{g}}, 2\right]\right] \geq 0.63
$$

Proof. The first form is true by simple calculation. The second form is holding due to
numerical values of cdf for normal distribution, $\operatorname{Pr}\left[\left|g_{i, r}\right| \in\left[\frac{1}{3 C_{g}}, 2\right]\right]=2\left(f(2)-f\left(1 / 3 C_{g}\right)\right)=2(f(2)-f(5 / 12)) \geq 2(0.977-0.662)=0.63$, where $f(x)=\int_{-\infty}^{x} \frac{e^{-x^{2} / 2}}{\sqrt{2 \pi}} d x$ is the cdf of normal distribution.

Stochastic dominance is a partial order between random variables and it is a classical concept in decision theory and decision analysis [HR69, Baw75]. We give the simplest definition below and it is sufficient for our application.

Definition 10.3.2 (stochastic domination of Gaussian random variables). Let $\sigma_{1}<\sigma_{2}$ and random variables $X \sim \mathcal{N}\left(0, \sigma_{1}^{2}\right), Y \sim \mathcal{N}\left(0, \sigma_{2}^{2}\right)$. Then we say that $|Y|$ stochastically dominates $|X|$, and it holds that

$$
\operatorname{Pr}[|Y| \geq \lambda] \geq \operatorname{Pr}[|X| \geq \lambda], \quad \forall \lambda \geq 0
$$

We formally define the set $L$ as follows:

Definition 10.3.3 (set $L$, input of the algorithm). Let $C_{L}>1$ be a fixed constant, and let set $L \subseteq[n]$ be defined as:

$$
|L|=C_{L} \cdot k / \epsilon, \quad \exists T \subset|T| \leq k:\left\|x-x_{T}\right\|_{2} \leq(1+\epsilon)\left\|x_{-k}\right\|_{2}
$$

We provide a definition called "badly-estimated coordinate",
Definition 10.3.4 (badly-estimated coordinate). We will say a coordinate $i \in[n]$ is badlyestimated if

$$
z_{i} \notin\left[\frac{1}{3 C_{g}}\left|x_{i}\right|-\frac{1}{100} \frac{\sqrt{\epsilon}}{\sqrt{k}}\left\|x_{-k}\right\|_{2}, 2\left|x_{i}\right|+\frac{1}{100} \frac{\sqrt{\epsilon}}{\sqrt{k}}\left\|x_{-k}\right\|_{2}\right],
$$

Then, we can define "badly-estimated set",

Definition 10.3.5 (badly-estimated set $\mathcal{B}$ ). Let set $L$ be defined as Definition 10.3.3. We say $\mathcal{B} \subseteq L$ is a badly-estimated set if for all $i \in \mathcal{B}, z_{i}$ is a badly estimated coordinate (see Definition 10.3.4).

We define a set of large coordinates in head,

Definition 10.3.6 (large coordinates in head). Let $\tau$ be defined in (10.3). Let $L$ be defined in Definition 10.3.3. Let $C_{g}$ be the constant from Fact 10.3.2. Define set

$$
\mathcal{M}=\left\{i \in L \cap H(x, k):\left|x_{i}\right| \geq 3 C_{g} \tau\right\},
$$

which contains the head coordinates of $x$ that are in $L$ and are larger in magnitude than $3 C_{g} \tau$.

### 10.3.2 Analyzing head and badly-estimated coordinates

Lemma 10.3.3 (expected error from coordinates above $\tau$ ). We have that

$$
\mathbb{E}\left[\sum_{i \in \mathcal{M}} x_{i}^{2} \cdot \mathbb{1}_{z_{i}<\tau}\right] \leq \frac{\epsilon}{100}\left\|x_{-k}\right\|_{2}^{2}
$$

Proof. Fix $i \in \mathcal{M}$. Observe that for $r \in[R]$

$$
\left|y_{r, h_{r}(i)}^{\prime}\right| \sim\left\|x_{h_{r}^{-1}(i)}\right\|_{2}|\mathcal{N}(0,1)| .
$$

Since

$$
\left\|x_{h_{r}^{-1}(i)}\right\|_{2} \geq\left|x_{i}\right|
$$

we have that the random variable $\left|y_{r, h_{r}(i)}^{\prime}\right|$ stochastically dominates the random variable $\left|x_{i}\right| \cdot|\mathcal{N}(0,1)|$.

By Fact 10.3.2, we have that

$$
\operatorname{Pr}\left[\left|y_{r, h_{r}(i)}^{\prime}\right| \leq \tau\right] \leq C_{g} \frac{\tau}{\left|x_{i}\right|}
$$

Because of the $R=C_{R} \log (1 / \epsilon)$ repetitions, a standard argument gives that

$$
\operatorname{Pr}\left[\mathbb{1}_{z_{i}<\tau}=1\right] \leq\left(C_{g} \frac{\tau}{\left|x_{i}\right|}\right)^{C^{\prime} \log (1 / \epsilon)}
$$

for some absolute constant $C^{\prime}>C_{R} / 3$.
We now bound

$$
\begin{aligned}
\mathbb{E}\left[\sum_{i \in \mathcal{M}} x_{i}^{2} \cdot \mathbb{1}_{z_{i}<\tau}\right] & \leq \sum_{i \in \mathcal{M}} x_{i}^{2}\left(C_{g} \frac{\tau}{\left|x_{i}\right|}\right)^{C^{\prime} \log (1 / \epsilon)} \\
& =\sum_{i \in \mathcal{M}} C_{g}^{2} \tau^{2}\left(C_{g} \frac{\tau}{\left|x_{i}\right|}\right)^{C^{\prime} \log (1 / \epsilon)-2} \\
& \leq k \cdot \tau^{2} \cdot \frac{\epsilon}{100} \\
& =\frac{\epsilon}{100}\left\|x_{-k}\right\|_{2}^{2}
\end{aligned}
$$

where the first step follows by the bound on $\mathbb{E}\left[\mathbb{1}_{z_{i}<\tau}\right]=\operatorname{Pr}\left[\mathbb{1}_{z_{i}<\tau}=1\right]$, and the third step by choosing by choosing $C_{R}>1$ to be some sufficiently large constant and the facts $C^{\prime}>C_{R} / 3$ and $\left(C_{g} \tau\right) /\left|x_{i}\right| \leq 1 / 3$.

Lemma 10.3.4 (probability of a fixed coordinate is badly-estimated). A coordinate $i$ is badly-estimated (as in Definition 10.3.4) probability at most

$$
\frac{\epsilon^{3}}{100^{2} C_{L}}
$$

Proof. Fix $r$ and set set $b=h_{r}(i)$. Recall the definition of $y_{r, b}=\sum_{i \in h_{r}^{-1}(b)} x_{i} g_{i, r}$ in Definition 10.3.1. We have that

$$
\left|\left|g_{i, r} x_{i}\right|-\left|\sum_{j \in h_{r}^{-1}(b) \backslash\{i\}} g_{j, r} x_{j}\right|\right| \leq\left|y_{r, b}\right| \leq\left|g_{i, r} x_{i}\right|+\left|\sum_{j \in h_{r}^{-1}(b) \backslash\{i\}} g_{j, r} x_{j}\right|
$$

Now, $\left|g_{i, r} x_{i}\right|$ will be at in $\left[\left(1 / 3 C_{g}\right)\left|x_{i}\right|, 2\left|x_{i}\right|\right]$ with probability at least 0.63 (due to Fact 10.3.2).
Moreover, for any $j \in H(x, k) \backslash\{i\}, h_{r}(j) \neq b$ with probability $1-1 / B=1-$ $\epsilon /\left(C_{B} k\right) \geq 1-1 /\left(C_{B} k\right)$. By a union bound, we get with probability at least $1-1 / C_{B}$, for all $j \in H(x, k) \backslash\{i\}, h_{r}(j) \neq b$. Conditioning on this event, we have,

$$
\mathbb{E}\left[\left(\sum_{j \in h_{r}^{-1}(b) \backslash\{i\}} g_{j, r} x_{j}\right)^{2}\right]=\frac{\epsilon}{C_{B} k}\left\|x_{-k}\right\|_{2}^{2}
$$

We then apply Markov's inequality to get that with probability at least $1-10^{4} / C_{B}$,

$$
\left(\sum_{j \in h_{r}^{-1}(b) \backslash\{i\}} g_{j, r} x_{j}\right)^{2} \leq \frac{\epsilon}{10^{4} k}\left\|x_{-k}\right\|_{2}^{2}
$$

Therefore, by a union bound,

$$
\begin{aligned}
\operatorname{Pr}\left[\left|y_{r, b}\right| \in\left[\frac{1}{3 C_{g}}\left|x_{i}\right|-\frac{1}{100} \frac{\sqrt{\epsilon}}{\sqrt{k}}, 2\left|x_{i}\right|+\frac{1}{100} \frac{\sqrt{\epsilon}}{\sqrt{k}}\right]\right] & \geq 0.63-\frac{1}{C_{B}}-\frac{10^{4}}{C_{B}} \\
& \geq 0.6
\end{aligned}
$$

where the last step follows by $C_{B} \geq 5 \times 10^{5}$.
Note that $z_{i}$ is obtained by taking median of $R$ copies of i.i.d. $\left|y_{r, b}\right|$. For each $r \in[R]$, we define $Z_{r}=1$ if $\left|y_{r, b}\right|$ falls into that region, and 0 otherwise. We have $\mathbb{E}\left[\sum_{r=1}^{R} Z_{r}\right] \geq 0.6 R$. Using Chernoff bound, we have

$$
\begin{aligned}
\operatorname{Pr}\left[\sum_{r=1}^{R} Z_{r}<0.9 \cdot 0.6 R\right] & \leq \operatorname{Pr}\left[\sum_{r=1}^{R} Z_{r}<0.9 \mathbb{E}\left[\sum_{r=1}^{R} Z_{r}\right]\right] \\
& \leq e^{-\frac{1}{3} 0.1^{2} \mathbb{E}\left[\sum_{r=1}^{R} Z_{r}\right]} \\
& \leq e^{-\frac{1}{3} 0.1^{2} \cdot 0.6 R}
\end{aligned}
$$

Thus,

$$
\begin{aligned}
\operatorname{Pr}\left[z_{i} \notin\left[\frac{1}{3 C_{g}}\left|x_{i}\right|-\frac{1}{100} \frac{\sqrt{\epsilon}}{\sqrt{k}}, 2\left|x_{i}\right|+\frac{1}{100} \frac{\sqrt{\epsilon}}{\sqrt{k}}\right]\right] & \leq e^{-\frac{1}{3} 0.1^{2} \cdot 0.6 R} \\
& \leq 2^{-0.002 R} \\
& =2^{-0.002 C_{R} \log (1 / \epsilon)} \\
& \leq 2^{-0.002\left(10000+500 C_{L}\right) \log (1 / \epsilon)} \\
& \leq \frac{\epsilon^{3}}{100^{2} C_{L}},
\end{aligned}
$$

where the third step follows from choice of $C_{R}$.

### 10.3.3 Guarantees of the algorithm

We now proceed with the proof of Theorem 10.3.1.

Proof. By Lemma 10.3.3 and an application of Markov's inequality we have that

$$
\sum_{i \in \mathcal{M}} x_{i}^{2} \cdot \mathbb{1}_{z_{i}<\tau} \leq \epsilon\left\|x_{-k}\right\|_{2}^{2}
$$

with probability $99 / 100$. Let this event be $\mathcal{E}_{1}$.
Moreover, by Lemma 10.3.4,

$$
\mathbb{E}[|\mathcal{B}|] \leq \frac{\epsilon^{3}|L|}{100^{2} C_{L}},
$$

so, by Markov's inequality, we have

$$
\operatorname{Pr}\left[|\mathcal{B}| \leq \frac{\epsilon^{2}|L|}{100 C_{L}}\right] \geq 1-\epsilon / 100 \geq 99 / 100
$$

Let this event be $\mathcal{E}_{2}$.
By taking a union bound, $\mathcal{E}_{1}$ and $\mathcal{E}_{2}$ both hold with probability $98 / 100$. Plugging size of $|L|\left(\leq C_{L} \cdot k / \epsilon\right)$ into equation of event $\mathcal{E}_{2}$, we get

$$
\begin{equation*}
|\mathcal{B}| \leq \frac{\epsilon^{2}|L|}{100 C_{L}} \leq \frac{\epsilon k}{100} . \tag{10.4}
\end{equation*}
$$

It means there are at most $\epsilon k / 100$ coordinates that badly-estimated.
We remind that our goal is to bound

$$
\begin{aligned}
\left\|x-x_{S}\right\|_{2}^{2} & =\left\|x_{\bar{S}}\right\|_{2}^{2} \\
& =\left\|x_{\bar{S} \cap \mathcal{M}}\right\|_{2}^{2}+\left\|x_{\bar{S} \backslash \mathfrak{\mathcal { M }}}\right\|_{2}^{2} \\
& =\left\|x_{\bar{S} \cap \mathcal{M}}\right\|_{2}^{2}+\left\|x_{(\bar{S} \backslash \mathcal{M}) \cap \mathcal{B}}\right\|_{2}^{2}+\left\|x_{(\bar{S} \backslash \mathcal{M}) \backslash \mathcal{B}}\right\|_{2}^{2}
\end{aligned}
$$

1. Bounding $\left\|x_{\mathcal{M} \cap \bar{S}}\right\|_{2}^{2}$. Consider the set

$$
I=\left\{i \in L \backslash \mathcal{M}:\left|x_{i}\right| \geq \tau / 3\right\}
$$

which contains the coordinates in $L$ with magnitude in the range $\left[\frac{1}{3} \tau, 3 C_{g} \tau\right)$. By the definition of $\tau$, clearly, $|I| \leq 3 k+k=4 k$, because we can have at most $k$ such elements in $H(x, k)$,
and at most $3 k$ such elements in the tail $[n] \backslash H(x, k)$. Since the number of badly estimated coordinates is at most $\epsilon k / 100$ and the size of $S$ is $\beta k$ for sufficiently large $\beta$, we can have at most $4 k+\epsilon k / 100<\beta k$ coordinates $i \in L$ which are not in $\mathcal{M}$ and are larger than $\tau$. This means that all coordinates in $\mathcal{M}$ with estimate $z_{i} \geq \tau$ will belong to $S$. This implies that

$$
\mathcal{M} \cap \bar{S}=\left\{i \in \mathcal{M}: z_{i}<\tau\right\},
$$

and hence

$$
\left\|x_{\mathcal{M} \cap \bar{S}}\right\|_{2}^{2}=\sum_{i \in \mathcal{M}} x_{i}^{2} \cdot \mathbb{1}_{z_{i}<\tau} \leq \epsilon\left\|x_{-k}\right\|_{2}^{2}
$$

since we conditioned on event $\mathcal{E}_{1}$.
2. Bounding $\left\|x_{(\bar{S} \backslash \mathcal{M}) \cap \mathcal{B}}\right\|_{2}^{2}$. For every $\left.i \in(\bar{S} \backslash \mathcal{M}) \cap \mathcal{B}\right)$ we have the trivial bound $\left|x_{i}\right| \leq \tau$. Since $(\bar{S} \backslash \mathcal{M}) \cap \mathcal{B}) \subseteq \mathcal{B}$, because the event $\mathcal{E}_{2}$ we get that

$$
\left\|x_{(\bar{S} \backslash \mathcal{M}) \cap \mathcal{B}}\right\|_{2}^{2} \leq|\mathcal{B}| \cdot \tau^{2} \leq \frac{\epsilon k}{100} \cdot \frac{\left\|x_{-k}\right\|_{2}^{2}}{k}=\frac{\epsilon}{100}\left\|x_{-k}\right\|_{2}^{2},
$$

where the second step follows from (10.4) and (10.3).
3. Bounding $\left\|x_{(\bar{S} \backslash \mathcal{M}) \backslash \mathcal{B}}\right\|_{2}^{2}$. Observe that set $(\bar{S} \backslash \mathcal{M}) \backslash \mathcal{B}$ consists of well-estimated coordinates that are less than $\tau$ in magnitude, and their estimates do not belong to the largest $\beta k$ estimates. For convenience, set $Q=(\bar{S} \backslash \mathcal{M}) \backslash \mathcal{B}$, then it is obvious that $Q=\bar{S} \backslash(\mathcal{M} \cup \mathcal{B})$.

We define three sets $H_{1}, H_{2}, H_{3}$ as follows,

$$
H_{1}=Q \cap T, \quad H_{2}=Q \cap \bar{T}, \quad \text { and } \quad H_{3}=(\bar{T} \backslash(\mathcal{M} \cup \mathcal{B})) \backslash \bar{S} .
$$

Using the definition of $Q$, we can rewrite $H_{1}, H_{2}$, and $H_{3}$ as follows

$$
\begin{aligned}
& H_{1}=(\bar{S} \backslash(\mathcal{M} \cup \mathcal{B})) \cap T=\bar{S} \cap \overline{\mathcal{M} \cup \mathcal{B}} \cap T, \\
& H_{2}=(\bar{S} \backslash(\mathcal{M} \cup \mathcal{B})) \cap \bar{T}=\bar{S} \cap \overline{\mathcal{M} \cup \mathcal{B}} \cap \bar{T}, \\
& H_{3}=(\bar{T} \backslash(\mathcal{M} \cup \mathcal{B})) \backslash \bar{S}=S \cap \overline{\mathcal{M} \cup \mathcal{B}} \cap \bar{T} .
\end{aligned}
$$

We can show that

$$
\begin{align*}
H_{2} \cap H_{3} & =(\bar{S} \cap \overline{\mathcal{M} \cup \mathcal{B}} \cap \bar{T}) \cup(S \cap \overline{\mathcal{M} \cup \mathcal{B}} \cap \bar{T}) \\
& =\emptyset \tag{10.5}
\end{align*}
$$

and

$$
\begin{align*}
H_{2} \cup H_{3} & =(\bar{S} \cap \overline{\mathcal{M} \cup \mathcal{B}} \cap \bar{T}) \cap(S \cap \overline{\mathcal{M} \cup \mathcal{B}} \cap \bar{T}) \\
& =\overline{\mathcal{M} \cup \mathcal{B}} \cap \bar{T} \\
& =\bar{T} \backslash(\mathcal{M} \cup \mathcal{B}) . \tag{10.6}
\end{align*}
$$

Then,

$$
\begin{aligned}
\left\|x_{Q}\right\|_{2}^{2} & =\left\|x_{H_{1}}\right\|_{2}^{2}+\left\|x_{H_{2}}\right\|_{2}^{2} \\
& =\left\|x_{H_{1}}\right\|_{2}^{2}+\left(\left\|x_{\bar{T} \backslash(\mathcal{M \cup B})}\right\|_{2}^{2}-\left\|x_{H_{3}}\right\|_{2}^{2}\right) \\
& \leq\left\|x_{H_{1}}\right\|_{2}^{2}+\left\|x_{\bar{T}}\right\|_{2}^{2}-\left\|x_{H_{3}}\right\|_{2}^{2} \\
& \leq\left\|x_{H_{1}}\right\|_{2}^{2}+(1+\epsilon)\left\|x_{-k}\right\|_{2}^{2}-\left\|x_{H_{3}}\right\|_{2}^{2}
\end{aligned}
$$

where first step follows from $H_{1} \cap H_{2}=\emptyset$ and $H_{1} \cup H_{2}=Q$, the second step follows from Eq. (10.5) and (10.6), the third step follows from $\left\|x_{\bar{T} \backslash(\mathcal{M \cup B})}\right\|_{2}^{2} \leq\left\|x_{\bar{T}}\right\|_{2}^{2}$ and the last step follows from $\left\|x_{\bar{T}}\right\|_{2}^{2} \leq(1+\epsilon)\left\|x_{-k}\right\|_{2}^{2}$.

We define $d, E, a, b$ as follows
$d=\left|H_{1}\right|, \quad E=\frac{1}{4} \sqrt{\epsilon / k}\left\|x_{-k}\right\|_{2}, \quad a=\max _{i \in H_{1}}\left|x_{i}\right|, \quad b=\min _{i \in H_{3}}\left|x_{i}\right|$.
Let $i^{*}$ and $j^{*}$ be defined as follows:

$$
\begin{equation*}
i^{*}=\arg \max _{i \in H_{1}}\left|x_{i}\right|, \quad \text { and } \quad j^{*}=\arg \min _{j \in H_{3}}\left|x_{j}\right| . \tag{10.8}
\end{equation*}
$$

Recall the definitions of $H_{1}$ and $H_{3}$, we know $H_{3}$ is a subset of $S$ and $H_{1}$ is a subset of $\bar{S}$. Since the set $S$ contains the largest $\beta k$ coordinates, thus we have

$$
z_{j} \geq z_{i}, \forall i \in H_{1}, j \in H_{3}
$$

It further implies $z_{j^{*}} \geq z_{i^{*}}$.
By Definition 10.3.4, we have

$$
\begin{equation*}
z_{i^{*}} \geq \frac{1}{3 C_{g}}\left|x_{i^{*}}\right|-\frac{1}{100} \frac{\sqrt{\epsilon}}{\sqrt{k}}\left\|x_{-k}\right\|_{2}, \tag{10.9}
\end{equation*}
$$

and

$$
\begin{equation*}
z_{j^{*}} \leq 2\left|x_{j^{*}}\right|+\frac{1}{100} \frac{\sqrt{\epsilon}}{\sqrt{k}}\left\|x_{-k}\right\|_{2} \tag{10.10}
\end{equation*}
$$

Then, we can show that $a \leq 6 C_{g} b+E$ in the following sense:

$$
\begin{array}{rlr}
a & =\left|x_{i^{*}}\right| & \text { by def. of } i^{*}, a,(10.8),(10.7) \\
& \leq 3 C_{g} z_{i^{*}}+\frac{3 C_{g}}{100} \sqrt{\epsilon / k}\left\|x_{-k}\right\|_{2} & \text { by }(10.9) \\
& \leq 3 C_{g} z_{j^{*}}+\frac{3 C_{g}}{200} \sqrt{\epsilon / k}\left\|x_{-k}\right\|_{2} & \text { by } z_{i^{*}} \leq z_{j^{*}} \\
& \leq 6 C_{g}\left|x_{j^{*}}\right|+\frac{6 C_{g}}{200} \sqrt{\epsilon / k}\left\|x_{-k}\right\|_{2} & \text { by }(10.10)  \tag{10.10}\\
& =6 C_{g} b+\frac{6 C_{g}}{200} \sqrt{\epsilon / k}\left\|x_{-k}\right\|_{2} & \text { by def. of } j^{*}, b,(10.8),(10.7) \\
& \leq 6 C_{g} b+E & \text { by def. of } E,(10.7)
\end{array}
$$

Note that $H_{3}=(\bar{T} \backslash(\mathcal{M} \cup \mathcal{B})) \backslash \bar{S}=S \backslash(T \cup \mathcal{M} \cup \mathcal{B})$. Therefore,

$$
\begin{aligned}
\left|H_{3}\right| & \geq|S|-|T|-|\mathcal{M}|-|\mathcal{B}| \\
& \geq \beta k-k-k-k \\
& =(\beta-3) k .
\end{aligned}
$$

Finally, we can have

$$
\begin{array}{rlr}
\left\|x_{H_{1}}\right\|_{2}^{2}-\left\|x_{H_{3}}\right\|_{2}^{2} & \leq d a^{2}-(\beta-3) k b^{2} & \\
& \leq d\left(6 C_{g} b+E\right)^{2}-(\beta-3) k b^{2} & \text { by } a \leq 6 C_{g} b+E \\
& =\left(36 C_{g}^{2} d-(\beta-3) k\right) b^{2}+12 C_{g} b d E+d E^{2} & \\
& \leq\left(36 C_{g}^{2} k-(\beta-3) k\right) b^{2}+12 C_{g} b k E+k E^{2} & \text { by } d \leq k \\
& \leq\left(36 C_{g}^{2} k-\left(\beta-5 C_{g}^{2}\right) k\right) b^{2}+12 C_{g} b k E+k E^{2} & \text { by } C_{g} \geq 4 / 5 \\
& \leq-36 k C_{g}^{2} b^{2}+12 C_{g} b k E+k E^{2} & \text { by } \beta \geq 77 C_{g}^{2} \\
& =-k\left(6 C_{g} b-E\right)^{2}+2 k E^{2} & \\
& \leq 2 k E^{2} & \\
& \leq \epsilon\left\|x_{-k}\right\|_{2}^{2} &
\end{array}
$$

where the last step follows from definition of $E$.
Thus, we have

$$
\left\|x_{Q}\right\|_{2}^{2} \leq(1+2 \epsilon)\left\|x_{-k}\right\|_{2}^{2}
$$

Putting it all together. We have

$$
\begin{aligned}
\left\|x-x_{S}\right\|_{2}^{2} & =\left\|x_{\bar{S} \cap \mathcal{M}}\right\|_{2}^{2}+\left\|x_{(\bar{S} \backslash \mathcal{M}) \cap \mathcal{B}}\right\|_{2}^{2}+\left\|x_{(\bar{S} \backslash \mathcal{M}) \backslash \mathcal{B}}\right\|_{2}^{2} \\
& \leq \epsilon\left\|x_{-k}\right\|_{2}^{2}+\frac{\epsilon}{100}\left\|x_{-k}\right\|_{2}^{2}+(1+2 \epsilon)\left\|x_{-k}\right\|_{2}^{2} \\
& \leq(1+4 \epsilon)\left\|x_{-k}\right\|_{2}^{2}
\end{aligned}
$$

Finally, we can conclude $\alpha=5$ and $\beta=100$.

### 10.3.4 Time, measurements, column sparsity, and probability

In this section, we will bound the decoding time, the number of measurements, column sparsity and success probability of algorithm.

Decoding time. For each $i \in L$, we compute $z_{i}$ to be the median of $R$ values. For this part, we spend $O(|L| \cdot R)=O((k / \epsilon) \cdot \log (1 / \epsilon))$ time. Moreover, calculating the top $\beta k$ estimates in $L$ only takes $O(|L|)$ time. Therefore, the decoding time is $O((k / \epsilon) \cdot \log (1 / \epsilon))$.

The number of measurements. The number of measurements is the bucket size $B$ times the number of repetitions $R$, which is $O(B R)=O((k / \epsilon) \cdot \log (1 / \epsilon))$.

Column sparsity. Each $i \in[n]$ goes to one bucket for each hash function, and we repeat $R$ times, so the column sparsity is $O(R)=O(\log (1 / \epsilon))$.

Success probability. By analysis in Section 10.3.3, the success probability is at least 0.98.

### 10.4 Tail Estimation

In Section 10.4.1, we present a standard result on random walks. In Section 10.4.2, we present some results on $p$-stable distribution. In what follows we asssume that $0<p \leq 2$. We show an algorithm for $\ell_{p}$ tail estimation in Section 10.4.3.

### 10.4.1 Random walks

Theorem 10.4.1. We consider the following random walk. We go right if $B_{i}=1$ and we go left if $B_{i}=0$. The probability of $B_{i}=1$ is at least $9 / 10$ and the probability of $B_{i}=0$ is at most $1 / 10$. With at least some constant probability bounded away from $\frac{1}{2}$, for all the possible length of the random walk, it will never return to the origin.

This is a standard claim, that can be proved in numerous ways, such as martingales etc. For the completeness, we still provide a folklore proof here.

Proof. Let $p>1 / 2$ be the probability of stepping to the right, and let $q=1-p$. For integer $m \geq 1$, let $P_{m}$ be the probability of first hitting 0 in exactly $m$ steps. It is obvious that $P_{m}=0$ if $n$ is even, and $P_{1}=q$. In order to hit 0 for the first time on the third step you must Right-Left-Left, so $P_{3}=p q^{2}$. To hit 0 for the first time in exactly $2 k+1$ steps, you must go right $k$ times and left $k+1$ times, your last step must be to the left, and through the first $2 k$ steps you must always have made at least many right steps as left steps. It is well known that the number o such path is $C_{k}$, which is the $k$-th Catalan number. Thus,

$$
P_{2 k+1}=C_{k} q^{k} q^{k+1}=C_{k} \cdot q(p q)^{k}=\frac{q(p q)^{k}}{k+1}\binom{2 k}{k}
$$

since

$$
C_{k}=\frac{1}{k+1}\binom{2 k}{k}
$$

By [Wil05], the generating function for the Catalan numbers is

$$
c(x)=\sum_{k \geq 0} C_{k} x^{k}=\frac{1-\sqrt{1-4 x}}{2 x}
$$

so the probability that the random walk will hit 0 is

$$
\begin{aligned}
\sum_{k \geq 0} P_{2 k+1} & =q \sum_{k \geq 0} C_{k}(p q)^{k} \\
& =q \cdot c(p q) \\
& =q \cdot \frac{1-\sqrt{1-4 p q}}{2 p q} \\
& =\frac{1-\sqrt{1-4 q(1-q)}}{2 p} \\
& =\frac{1-\sqrt{1-4 q+4 q^{2}}}{2 p} \\
& =\frac{1-(1-2 q)}{2 p} \\
& =q / p \\
& \leq 1 / 9
\end{aligned}
$$

Thus, we complete the proof.

### 10.4.2 $p$-stable distributions

We first provide the definition of $p$-stable distribution. For the more details, we refer the readers to [Ind06].


Figure 10.1: This is a visualization of part of the proof in Claim 10.4.7. We consider an example where there are $l=10$ blocks, $B_{1}=1, B_{2}=1, B_{1}=1, B_{3}=0, B_{4}=0$, $B_{5}=1, B_{6}=1, B_{7}=0, B_{8}=0, B_{9}=1$ and $B_{10}=0$. Recall the two important conditions in the proof of Claim 10.4.7, the first one is $B_{1}=1$ and the second one is, for all $j \in[l], \sum_{j^{\prime}=2}^{j} B_{j^{\prime}}>(j-1) / 2$. The number on the green arrow is $\sum_{j^{\prime}=2}^{j} B_{j^{\prime}}$. It is to see that the example we provided here is satisfying those two conditions. Recall the definition of set $S_{1}$ and $S_{0}$. Here $S_{1}=\{2,3,5,6,9\}$ and $S_{0}=\{4,7,8,10\}$. Then $S_{1}^{\prime}=\{2,3,5,6\}$. The mapping $\pi$ satisfies that $\pi(4)=2, \pi(7)=3, \pi(8)=5$ and $\pi(10)=6$.

Definition 10.4.1 ( $p$-stable distribution). A distribution $\mathcal{D}$ over $\mathbb{R}$ is called $p$-stable, if there exists $p \geq 0$ such that for any $n$ real numbers $a_{1}, a_{2}, \cdots, a_{n}$ and i.i.d. variables $x_{1}, x_{2}, \cdots, x_{n}$ from distribution $\mathcal{D}$, the random variable $\sum_{i=1}^{n} a_{i} x_{i}$ has the same distribution as the variable $\|a\|_{p} y$, where $y$ is a random variable from distribution $\mathcal{D}$.

Theorem 10.4.2 ([Zol86]). For any $p \in(0,2]$, there exists a $p$-stable distribution.

Gaussian distribution defined by the density function $f(x)=\frac{1}{\sqrt{2 \pi}} e^{-x^{2} / 2}$, is 2 -stable. Cauchy distribution defined by density function $f(x)=\frac{1}{\pi} \frac{1}{1+x^{2}}$ is 1 -stable. Let $\mathcal{D}_{p}$ denote the
$p$-stable distribution. For $p=2, \mathcal{D}_{p}$ is $\mathcal{N}(0,1)$ and for $p=1, \mathcal{D}_{p}$ is $\mathcal{C}(0,1)$.

### 10.4.3 $\quad \ell_{p}$-tail estimation algorithm

The goal of this Section is prove Lemma 10.4.3.

```
Algorithm \(10.3 \ell_{p}\)-tail Estimation Algorithm
    procedure LpLpTailEstimation \(\left(x, k, p, C_{0}, \delta\right) \quad \triangleright\) Lemma 10.4.3
    \(\triangleright\) Requires \(C_{0} \geq 1000\)
        \(m \leftarrow O(\log (1 / \delta))\)
        Choose \(g_{i, t}\) to be random variable that sampled i.i.d. from distribution \(\mathcal{D}_{p}, \forall i, t \in\)
    \([n] \times[m]\)
        Choose \(\delta_{i, t}\) to be Bernoulli random variable with \(\mathbb{E}\left[\delta_{i, t}\right]=1 /(100 k), \forall i, t \in[n] \times[m]\)
                            \(\triangleright\) Matrix \(A\) is implicitly constructed based on \(g_{i, t}\) and \(\delta_{i, t}\)
        for \(t \in[m]\) do
            \(y_{t} \leftarrow \sum_{i=1}^{n} \delta_{i, t} \cdot g_{i, t} \cdot x_{i}\)
        end for
        \(V \leftarrow \operatorname{median}_{t \in[m]}\left|y_{t}\right|^{2}\)
        return \(V\)
    end procedure
```

One can try to prove such a claim for $p=2$ with random signs, instead of Gaussians, by applying the Paley-Zygmund inequality to obtain the lower bound. A straightforward calculation indicates that this approach does not give the desired result, hence we need a new argument to deal with the lower bound.

Lemma 10.4.3 (Restatement of Lemma 10.1.3). Let $C_{0} \geq 1000$ denote some fixed constant. There is an oblivious construction of matrix $A \in \mathbb{R}^{m \times n}$ with $m=O(\log (1 / \delta))$ along with a decoding procedure LpLpTailEstimation $\left(x, k, p, C_{0}, \delta\right)$ (Algorithm 10.3) such that, given $A x$, it is possible to output a value $V$ in time $O(m)$ such that

$$
\frac{1}{10 k}\left\|x_{-C_{0} k}\right\|_{p}^{p} \leq V \leq \frac{1}{k}\left\|x_{-k}\right\|_{p}^{p}
$$

holds with probability $1-\delta$.

Proof. Let $m=O(\log (1 / \delta))$. For each $i \in[n], t \in[m]$, we use $g_{i, t}$ to denote a random variable that sample from distribution $\mathcal{D}_{p}$.

For each $i \in[n], t \in[m]$, we use $\delta_{i, t}$ to denote a Bernoulli random variable such that

$$
\delta_{i, t}= \begin{cases}1, & \text { with prob. } \frac{1}{100 k} \\ 0, & \text { otherwise }\end{cases}
$$

Then we have

$$
\mathbb{E}\left[\delta_{i, t}\right]=\frac{1}{100 k} .
$$

For each $t \in[m]$, we define $y_{t}$ as follows

$$
\begin{equation*}
y_{t}=\sum_{i=1}^{n} \delta_{i, t} g_{i, t} x_{i} . \tag{10.11}
\end{equation*}
$$

For each $t \in[m]$, we define $\Delta_{t}$ as follows

$$
\begin{equation*}
\Delta_{t}=\left(\sum_{i=1}^{n} \delta_{i, t}^{p} x_{i}^{p}\right)^{1 / p} \tag{10.12}
\end{equation*}
$$

Using Claim 10.4.4 and Claim 10.4.7

$$
\underset{g, \delta}{\operatorname{Pr}}\left[\left|y_{t}\right|<\alpha \frac{1}{\left(2 C_{0} k\right)^{1 / p}}\left\|x_{-C_{0} k}\right\|_{p}\right] \leq 1 / 5 .
$$

Using Claim 10.4.5 and Claim 10.4.6

$$
\underset{g, \delta}{\operatorname{Pr}}\left[\left|y_{t}\right|>\beta \frac{1}{k^{1 / p}}\left\|x_{-k}\right\|_{p}\right] \leq 1 / 5 .
$$

Finally, we just take the median over $m$ different independent repeats. Since $m=$ $O(\log (1 / \delta))$, thus, we can boost the failure probability to $\delta$.

It is a standard fact, due to $p$-stability, that $y_{t}$ follows the $p$-stable distribution : $\Delta_{t} \cdot \mathcal{D}_{p}$. Since $p$-stable distributions are continuous functions, we have the following two Claims:

Claim 10.4.4 (upper bound on $\left.\left|y_{t}\right|\right)$. Let $y_{t}$ be defined in Eq. (10.11), let $\Delta_{t}$ be defined in Eq. (10.12). There is some sufficiently small constant $\alpha \in(0,1)$ such that

$$
\underset{g}{\operatorname{Pr}}\left[\left|y_{t}\right|<\alpha \cdot \Delta_{t}\right] \leq 1 / 10
$$

Claim 10.4.5 (lower bound on $\left.\left|y_{t}\right|\right)$. Let $y_{t}$ be defined in Eq. (10.11), let $\Delta_{t}$ be defined in Eq. (10.12). There is some sufficiently large constant $\beta>1$ such that

$$
\underset{g}{\operatorname{Pr}}\left[\left|y_{t}\right|>\beta \cdot \Delta_{t}\right] \leq 1 / 10
$$

It remains to prove Claim 10.4.6 and Claim 10.4.7.

Claim 10.4.6 (lower bound on $\Delta_{t}$ ). Let $\Delta_{t}$ be defined in Eq. (10.12). Then we have

$$
\operatorname{Pr}_{\delta}\left[\Delta_{t}>\frac{1}{k^{1 / p}}\left\|x_{-k}\right\|_{p}\right] \leq 1 / 10
$$

Proof. The proof mainly includes three steps,
First, for a fixed coordinate $i \in[n]$, with probability at most $1 /(100 k)$, it got sampled. Taking a union bound over all $k$ largest coordinates. We can show that with probability at least $1-1 / 100$, none of $k$ largest coordinates is sampled. Let $\xi$ be that event.

Second, conditioning on event $\xi$, we can show that

$$
\mathbb{E}\left[\Delta_{t}\right] \leq \frac{1}{(100 k)^{1 / p}}\left\|x_{-k}\right\|_{p}
$$

Third, applying Markov's inequality, we have

$$
\operatorname{Pr}\left[\Delta_{t} \geq a\right] \leq \mathbb{E}\left[\Delta_{t}\right] / a
$$

Choosing $a=\frac{1}{k^{1 / p}}\left\|x_{-k}\right\|_{p}$, we have

$$
\operatorname{Pr}\left[\Delta_{t} \geq \frac{1}{k^{1 / p}}\left\|x_{-k}\right\|_{p}\right] \leq 1 / 10
$$

Claim 10.4.7 (upper bound on $\Delta_{t}$ ). Let $\Delta_{t}$ be defined in Eq. (10.12). For any $C_{0} \geq 1000$, we have

$$
\operatorname{Pr}_{\delta}\left[\Delta_{t}<\frac{1}{\left(2 C_{0} k\right)^{1 / p}}\left\|x_{-C_{0} k}\right\|_{p}\right] \leq 1 / 10
$$

Proof. Without loss of generality, we can assume that all coordinates of $x_{i}$ are sorted, i.e. $x_{1} \geq x_{2} \geq \cdots \geq x_{n}$. Then we split length $n$ vector into $l$ blocks where each block has length $s=C_{0} k$. Note that it is obvious $l \cdot s=n$.

For each $j \in[l]$, we use boolean variable $B_{j}$ to denote that if at least one coordinate in $j$-th block has been sampled. For a fixed block $j \in[l]$, the probability of sampling at least one coordinate from that block is at least

$$
1-\left(1-\frac{1}{100 k}\right)^{s}=1-\left(1-\frac{1}{100 k}\right)^{C_{0} k} \geq 9 / 10
$$

Thus, we know $1 \geq \mathbb{E}\left[B_{j}\right] \geq 9 / 10$.
Warm-up. Note the probability is not allowed to take a union over all the blocks. However, if we conditioned on that each block has been sampled at least one coordinate,
then we have

$$
\begin{aligned}
\Delta_{t}^{p} & =\sum_{i=1}^{n} \delta_{i, t} x_{i}^{p} \\
& \geq \sum_{j=1}^{l-1} x_{j s}^{p} \\
& \geq \sum_{j=1}^{l-1} \frac{1}{s}\left(x_{j s+1}^{p}+x_{j s+2}^{p}+\cdots+x_{j s+s}^{p}\right) \\
& =\frac{1}{s}\left\|x_{s}\right\|_{p}^{p} .
\end{aligned}
$$

Fixed. For simplicity, for each $j \in[l]$, we use set $T_{j}$ to denote $\{(j-1) s+1,(j-$ 1) $s+2, \cdots,(j-1) s+s\}$.

Using random walk Lemma 10.4.1, with probability at least 99/100, we have : for all $j \in\{2, \cdots, l\}$,

$$
\sum_{j^{\prime}=2}^{j} B_{j^{\prime}}>(j-1) / 2 .
$$

We know that with probability at least $99 / 100, B_{1}=1$. Then with probability at least 99/100, we have

$$
B_{1}=1, \text { and } \sum_{j^{\prime}=2}^{j} B_{j^{\prime}}>(j-1) / 2, \forall j \in[l] .
$$

We conditioned on the above event holds. Let set $S_{1} \subset[n]$ denote the set of indices $j$ such that $B_{j}=1$, i.e.,

$$
S_{1}=\left\{j \mid B_{j}=1, j \in[n] \backslash\{1\}\right\} .
$$

Let set $S_{0} \subset[n]$ denote the set of indices $j$ such that $B_{j}=0$, i.e.,

$$
S_{0}=\left\{j \mid B_{j}=0, j \in[n] \backslash\{1\}\right\} .
$$

Due to $\sum_{j^{\prime}=2}^{j} B_{j^{\prime}}>(j-1) / 2, \forall j \in[l]$, then it is easy to see $S_{1}>S_{0}$ and there exists a one-to-one mapping $\pi: S_{0} \rightarrow S_{1}^{\prime}$ where $S_{1}^{\prime} \subseteq S_{1}$ such that for each coordinate $j \in S_{0}$, $\pi(j)<j$. Since we are the coordinates are being sorted already, thus

$$
\begin{aligned}
\sum_{j \in S_{1}}\left\|x_{T_{j}}\right\|_{p}^{p} & =\sum_{j \in S_{1}^{\prime}}\left\|x_{T_{j}}\right\|_{p}^{p} \\
& =\sum_{j \in S_{1}^{\prime}}\left\|x_{T_{\pi-1}(j)}\right\|_{p}^{p} \\
& \geq \sum_{j \in S_{0}}\left\|x_{T_{j}}\right\|_{p}^{p}
\end{aligned}
$$

which implies that

$$
\Delta_{t}^{p}=\sum_{i=1}^{n} \delta_{i}^{p} x_{i}^{p}=\sum_{j \in S_{1}}\left\|x_{T_{j}}\right\|_{p}^{p} \geq \frac{1}{2 s}\left\|x_{-s}\right\|_{p}^{p}
$$

Thus, with probability at least $9 / 10$, we have

$$
\Delta_{t} \geq \frac{1}{(2 s)^{1 / p}}\left\|x_{-s}\right\|_{p}
$$

### 10.5 Putting It All Together

Our full algorithm first applies interval forest sparse recovery algorithm in Algorithm 10.1 with precision parameter $\frac{\epsilon}{10}$ to obtain a set $L$ of size $O(k / \epsilon)$. By Theorem 10.2.1, with probability $\frac{9}{10}, L$ contains a subset $T$ of size at most $k$ so that $\left\|x-x_{T}\right\|_{2} \leq\left(1+\frac{\epsilon}{10}\right)\left\|x_{-k}\right\|_{2}$. Then the algorithm feeds $L$ into the pruning procedure in Algorithm 10.2 also with precision parameter $\frac{\epsilon}{10}$ to get $S$. By Theorem 10.3.1, with probability $\frac{9}{10},|S|=O(k)$ and $\left\|x-x_{S}\right\|_{2} \leq$ $\left(1+\frac{\epsilon}{2}\right)\left\|x-x_{-k}\right\|_{2}$. Finally, the algorithm uses set query data structure in Lemma 10.1.6 to give $\widehat{x}_{S}$ as approximation to $x_{S}$. With probability at least $1-1 / \operatorname{poly}(k),\left\|\widehat{x}_{S}-x_{S}\right\|_{2}^{2} \leq \frac{\epsilon}{2}\left\|x-x_{S}\right\|_{2}^{2}$. Therefore, $\left\|x-\widehat{x}_{S}\right\|_{2}^{2}=\left\|x-x_{S}\right\|_{2}^{2}+\left\|x_{S}-\widehat{x}_{S}\right\|_{2}^{2} \leq\left(1+\frac{\epsilon}{4}\right)\left(1+\frac{\epsilon}{2}\right)^{2}\left\|x-x_{-k}\right\|_{2}^{2} \leq(1+\epsilon)^{2}\left\|x-x_{-k}\right\|_{2}^{2}$. By union bound, the overall failure probability is at most $1 / 4$.

```
Algorithm 10.4 Stronger \(\ell_{2} / \ell_{2}\) Algorithm
    procedure \(\operatorname{Main}(x, n, k, \epsilon) \quad \triangleright\) Theorem 10.1.2
        \(L \leftarrow \operatorname{IntervalForestSparseRecovery}\left(x, n, k, \frac{\epsilon}{10}\right) \quad \triangleright\) Algorithm 10.1
        \(S \leftarrow \operatorname{Prune}\left(x, n, k, \frac{\epsilon}{10}, L\right) \quad \triangleright\) Algorithm 10.2
        \(\widehat{x}_{S} \leftarrow \operatorname{SetQuery}\left(x, n, \frac{\epsilon}{4}, S\right) \quad \triangleright\) Lemma 10.1.6
        return \(\widehat{x}_{S}\)
    end procedure
```


## Chapter 11

## Continuous Fourier Transform I

In recent years, a number of works have studied methods for computing the Fourier transform in sublinear time if the output is sparse. Most of these have focused on the discrete setting, even though in many applications the input signal is continuous and naive discretization significantly worsens the sparsity level.

We present an algorithm for robustly computing sparse Fourier transforms in the continuous setting. Let $x(t)=x^{*}(t)+g(t)$, where $x^{*}$ has a $k$-sparse Fourier transform and $g$ is an arbitrary noise term. Given sample access to $x(t)$ for some duration $T$, we show how to find a $k$-Fourier-sparse reconstruction $x^{\prime}(t)$ with

$$
\frac{1}{T} \int_{0}^{T}\left|x^{\prime}(t)-x(t)\right|^{2} \mathrm{~d} t \lesssim \frac{1}{T} \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t
$$

The sample complexity is linear in $k$ and logarithmic in the signal-to-noise ratio and the frequency resolution. Previous results with similar sample complexities could not tolerate an infinitesimal amount of i.i.d. Gaussian noise, and even algorithms with higher sample complexities increased the noise by a polynomial factor. We also give new results for how precisely the individual frequencies of $x^{*}$ can be recovered.

### 11.1 Introduction

The Fourier transform is ubiquitous in digital signal processing of a diverse set of signals, including sound, image, and video. Much of this is enabled by the Fast Fourier Transform (FFT) [CT65], which computes the $n$-point discrete Fourier transform in $O(n \log n)$ time. But can we do better?

In many situations, much of the reason for using Fourier transforms is because the transformed signal is sparse - i.e., the energy is concentrated in a small set of $k$ locations. In such situations, one could hope for a dependency that depends nearly linearly on $k$ rather than $n$. Moreover, one may be able to find these frequencies while only sampling the signal for some period of time. This idea has led to a number of results on sparse Fourier transforms, including [GGI ${ }^{+} 02$, GMS05, HIKP12a, IK14], that can achieve $O(k \log (n / k) \log n)$ running time and $O(k \log (n / k))$ sample complexity (although not quite both at the same time) in a robust setting.

These works apply to the discrete Fourier transform, but lots of signals including audio or radio originally come from a continuous domain. The standard way to convert a continuous Fourier transform into a discrete one is to apply a window function then subsample. Unfortunately, doing so "smears out" the frequencies, blowing up the sparsity. Thus, one can hope for significant efficiency gains by directly solving the sparse Fourier transform problem in the continuous setting. This has led researchers to adapt techniques from the discrete setting to the continuous both in theory $\left[\mathrm{BCG}^{+} 12\right.$, TBSR13, CF14, DB13] and in practice $\left[\mathrm{SAH}^{+} 13\right]$. However, these results are not robust to noise: if the signal is sampled with a tiny amount of Gaussian noise or decays very slightly over time, no method has been known for computing a sparse Fourier transform in the continuous setting with sample
complexity linear in $k$ and logarithmic in other factors. That is what we present in this paper.

Formally, a vector $x^{*}(t)$ has a $k$-sparse Fourier transform if it can be written as

$$
x^{*}(t)=\sum_{i=1}^{k} v_{i} e^{2 \pi \mathbf{i} f_{i} t}
$$

for some tones $\left\{\left(v_{i}, f_{i}\right)\right\}$. We consider the problem where we can sample some signal

$$
x(t)=x^{*}(t)+g(t)
$$

at any $t$ we choose in some interval $[0, T]$, where $x^{*}(t)$ has a $k$-sparse Fourier transform and $g(t)$ is arbitrary noise. As long as $g$ is "small enough," one would like to recover a good approximation to $x$ (or to $x^{*}$, or to $\left\{\left(v_{i}, f_{i}\right)\right\}$ ) using relatively few samples $t \in[0, T]$ and fast running time. Our algorithm achieves several results of this form, but a simple one is an $\ell_{2} / \ell_{2}$ guarantee: we reconstruct an $x^{\prime}(t)$ with $k$-sparse Fourier transform such that

$$
\frac{1}{T} \int_{0}^{T}\left|x^{\prime}(t)-x(t)\right|^{2} \mathrm{~d} t \lesssim \frac{1}{T} \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t
$$

using a number of samples that is $k$ times logarithmic factors ${ }^{1}$. To the best of our knowledge, this is the first algorithm achieving such a constant factor approximation with a sample complexity sublinear in $T$ and the signal-to-noise ratio.

Our algorithm also gives fairly precise estimates of the individual tones $\left(v_{i}, f_{i}\right)$ of the signal $x^{*}$. To demonstrate what factors are important, it is helpful to think about a concrete setting. Let us consider sound from a simplified model of a piano.

[^58]Thought experiment: piano tuning In a simplified model of a piano, we have keys corresponding to frequencies over some range $[-F, F]$. The noise $g(t)$ comes from ambient noise and the signals not being pure tones (because, for example, the notes might decay slowly over time). For concrete numbers, a modern piano has 88 keys spaced from about 27.5 Hz to $F=4200 \mathrm{~Hz}$. The space between keys ranges from a few Hz to a few hundred Hz , but most chords will have an $\eta=30 \mathrm{~Hz}$ or more gap between the frequencies being played. One typically would like to tune the keys to within about $\pm \nu=1 \mathrm{~Hz}$. And piano music typically has $k$ around 5 .

Now, suppose you would like to build a piano tuner that can listen to a chord and tell you what notes are played and how they are tuned. For such a system, how long must we wait for the tuner to identify the frequencies? How many samples must the tuner take? And how robust is it to the noise?

If you have a constant signal-to-noise ratio, you need to sample for a time $T$ of at least order $1 / \nu=1$ second in order to get 1 Hz precision-frequencies within 1 Hz of each other will behave very similarly over small fractions of a second, which noise can make indistinguishable. You also need at least $\Omega\left(k \log \frac{F}{k \nu}\right) \approx 50$ samples, because the support of the signal contains that many bits of information and you only get a constant number per measurement (at constant SNR). At higher signal-to-noise ratios $\rho$, these results extend to $\Omega\left(\frac{1}{\nu \rho}\right)$ duration and $\Omega\left(k \log _{\rho} \frac{F}{k \nu}\right)$ samples. But as the signal-to-noise ratio gets very high, there is another constraint on the duration: for $T<\frac{1}{\eta} \approx 33$ milliseconds the different frequencies start becoming hard to distinguish, which causes the robustness to degrade exponentially in $k$ [Moi15] (though the lower bound there only directly applies to a somewhat restricted version of our setting).

This suggests the form of a result: with a duration $T>\frac{1}{\eta}$, one can hope to recover the frequencies to within $\frac{1}{\rho T}$ using $O\left(k \log _{\rho} \frac{F T}{k}\right)$ samples. We give an algorithm that is within logarithmic factors of this ideal: with a duration $T>\frac{O(\log (k / \delta))}{\eta}$, we recover the frequencies to within $O\left(\frac{1}{\rho T}\right)$ using $O\left(k \log _{\rho}(F T) \cdot \log (k / \delta) \log k\right)$ samples, where $\rho$ and $1 / \delta$ are (roughly speaking) the minimum and maximum signal-to-noise ratios that you can tolerate, respectively.

Instead of trying to tune the piano by recovering the frequencies precisely, one may simply wish to record the sound for future playback with relatively few samples. Our algorithm works for this as well: the combination $x^{\prime}(t)$ of our recovered frequencies satisfies

$$
\frac{1}{T} \int_{0}^{T}\left|x^{\prime}(t)-x(t)\right|^{2} \mathrm{~d} t \lesssim \frac{1}{T} \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t
$$

Let us now state our main theorems. The first shows how well we can estimate the frequencies $f_{i}$ and their weights $v_{i}$; we refer to this $\left(v_{i}, f_{i}\right)$ pair as a tone.

Theorem 11.1.1 (Tone estimation). Consider any signal $x(t):[0, T] \rightarrow \mathbb{C}$ of the form

$$
x(t)=x^{*}(t)+g(t),
$$

for arbitrary "noise" $g(t)$ and an exactly $k$-sparse $x^{*}=\sum_{i \in[k]} v_{i} e^{2 \pi \mathbf{i} f_{i} t}$ with frequencies $f_{i} \in$ $[-F, F]$ and frequency separation $\eta=\min _{i \neq j}\left|f_{i}-f_{j}\right|$. For some parameter $\delta>0$, define the "noise level"

$$
\mathcal{N}^{2}:=\frac{1}{T} \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t+\delta \sum_{i=1}^{k}\left|v_{i}\right|^{2}
$$

We give an algorithm that takes samples from $x(t)$ over any duration $T>O\left(\frac{\log (k / \delta)}{\eta}\right)$ and returns a set of $k$ tones $\left\{\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\right\}$ that approximates $x^{*}$ with error proportional to $\mathcal{N}$. In
particular, every large tone is recovered: for any $v_{i}$ with $\left|v_{i}\right| \gtrsim \mathcal{N}$, we have for an appropriate permutation of the indices that

$$
\begin{equation*}
\left|f_{i}^{\prime}-f_{i}\right| \lesssim \frac{\mathcal{N}}{T\left|v_{i}\right|} \quad \text { and } \quad\left|v_{i}^{\prime}-v_{i}\right| \lesssim \mathcal{N} \tag{11.1}
\end{equation*}
$$

In fact, we satisfy a stronger guarantee that the total error is bounded:

$$
\begin{equation*}
\sum_{i=1}^{k} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi \mathrm{i} \mathrm{f}_{i}^{\prime} t}-v_{i} e^{2 \pi \mathrm{i} \mathrm{f}_{i} t}\right|^{2} \mathrm{~d} t \lesssim \mathcal{N}^{2} \tag{11.2}
\end{equation*}
$$

The algorithm takes $O\left(k \log (F T) \log \left(\frac{k}{\delta}\right) \log (k)\right)$ samples and $O\left(k \log (F T) \log \left(\frac{F T}{\delta}\right) \log (k)\right)$ running time, and succeeds with probability at least $1-1 / k^{c}$ for an arbitrarily large constant c.

We then show that the above approximation of the individual tones is good enough to estimate the overall signal $x(t)$ to within constant factors:

Theorem 11.1.2 (Signal estimation). In the same setting as Theorem 11.1.1, if the duration is slightly longer at $T>O\left(\frac{\log (1 / \delta)+\log ^{2} k}{\eta}\right)$, the reconstructed signal $x^{\prime}(t)=\sum_{i=1}^{k} v_{i}^{\prime} e^{2 \pi \mathrm{i} \mathrm{f}_{i}^{\prime} t}$ achieves a constant factor approximation to the complete signal $x$ :

$$
\begin{equation*}
\frac{1}{T} \int_{0}^{T}\left|x^{\prime}(t)-x(t)\right|^{2} \mathrm{~d} t \lesssim \mathcal{N}^{2} \tag{11.3}
\end{equation*}
$$

The algorithm takes $O\left(k \log (F T) \log \left(\frac{k}{\delta}\right) \log (k)\right)$ samples and $O\left(k \log (F T) \log \left(\frac{F T}{\delta}\right) \log (k)\right)$ running time, and succeeds with probability at least $1-1 / k^{c}$ for an arbitrarily large constant c.

The above theorems give three different error guarantees, which are all in terms of a "noise level" $\mathcal{N}^{2}$ that is the variance of the noise $g(t)$ plus $\delta$ times the energy of the signal. The algorithm depends logarithmically on $\delta$, so one should think of $\mathcal{N}^{2}$ as being the variance of the noise, e.g. $\sigma^{2}$ if samples have error $N\left(0, \sigma^{2}\right)$.

Error guarantees. Our algorithm does a good job of estimating the signal, but how exactly should we quantify this? Because very few previous results have shown robust recovery in the continuous setting, there is no standard error measure to use. We therefore bound the error in three different ways: the maximum error in the estimation of any tone; the weighted total error in the estimation of all tones; and the difference between the reconstructed signal and the true signal over the sampled interval. The first measure has been studied before, while the other two are to the best of our knowledge new but useful to fully explain the robustness we achieve.

The error guarantee (11.1) says that we achieve good recovery of any tones with magnitude larger than $C \mathcal{N}$ for some constant $C$. Note that such a requirement is necessary: for tones with $\left|v_{i}\right| \leq \mathcal{N}$, one could have $g(t)=-v_{i} e^{2 \pi \mathrm{i} f_{i} t}$, completely removing the tone $\left(v_{i}, f_{i}\right)$ from the observed $x(t)$ and making it impossible to find. For the tones of sufficiently large magnitude, we find them to within $\frac{\mathcal{N}}{T\left|v_{i}\right|}$. This is always less than $1 / T$, and converges to 0 as the noise level decreases. This is known as superresolution-one can achieve very high frequency resolution in sparse, nearly noiseless settings. Moreover, by Lemma 11.3.15 our "superresolution" precision $\left|f_{i}^{\prime}-f_{i}\right| \lesssim \frac{\mathcal{N}}{T\left|v_{i}\right|}$ is optimal.

While the guarantee of (11.1) is simple and optimal given its form, it is somewhat unsatisfying. It shows that the maximum error over all $k$ tones is $\mathcal{N}$, while one can hope to bound the total error over all $k$ tones by $\mathcal{N}$. This is precisely what Equation (11.2) does. The guarantee (11.1) is the precision necessary to recover the tone to within $O(\mathcal{N})$ average error in time, that is (11.1) is equivalent to

$$
\frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi \mathrm{i} f_{i}^{\prime} t}-v_{i} e^{2 \pi \mathrm{i} f_{i} t}\right|^{2} \mathrm{~d} t \lesssim \mathcal{N}^{2} \quad \forall i \in[k]
$$

In (11.2), we show that this bound holds even if we sum the left hand side over all $i \in[k]$, so the average error is a factor $k$ better than would be implied by (11.1). It also means that the total mass of all tones that are not recovered to within $\pm 1 / T$ is $O(\mathcal{N})$, not just that every tone larger than $O(\mathcal{N})$ is recovered to within $\pm 1 / T$.

The stronger bound (11.2) can be converted to the guarantee (11.3), which is analogous to the $\ell_{2} / \ell_{2}$ recovery guarantee standard in compressive sensing. It bounds the error of our estimate in terms of the input noise, on average over the sampled duration. The standard form of the $\ell_{2} / \ell_{2}$ guarantee in the discrete sparse Fourier transform setting [GGI+ 02 , GMS05, HIKP12a, IKP14, IK14] compares the energy in frequency domain rather than time domain. This cannot be achieved directly in the continuous setting, since frequencies infinitesimally close to each other are indistinguishable over a bounded time interval $T$. But if the signal is periodic with period $T$ (the case where a discrete sparse Fourier transform applies directly), then (11.3) is equivalent to the standard guarantee by Parseval's identity. So (11.3) seems like the right generalization of $\ell_{2} / \ell_{2}$ recovery to our setting.

Other factors. Our algorithm succeeds with high probability in $k$, which could of course be amplified by repetition (but increasing the sample complexity and running time). Our running time, at $O(k \log (F T) \log (F T / \delta) \log k)$, is (after translating between the discrete and continuous setting) basically a $\log k$ factor larger than the fastest known algorithm for discrete sparse Fourier transforms ([HIKP12a]). But since that result only succeeds with constant probability, and repetition would also increase that by a $\log k$ factor, our running time is actually equivalent to the fastest known results that succeed with high probability in $k$.

Our sample complexity is $O(\log (k / \delta) \log k)$ worse than the presumptive optimal, which is known to be achievable in the discrete setting ([IK14]). However, the techniques that [IK14] uses to avoid the $\log (k / \delta)$ factor seem hard to adapt to the continuous setting without losing some of the robustness/precision guarantees.

Another useful property of our method, not explicitly stated in the theorem, is that the sampling method is relatively simple: it chooses $O(\log (F T) \log k)$ different random arithmetic sequences of points, where each sequence has $O(k \log (k / \delta))$ points spread out over a constant fraction of the time domain $T$. Thus one could implement this in hardware using a relatively small number of samplers, each of which performs regular sampling at a rate of $\frac{k \log (k / \delta)}{T}$. This is in contrast to the Nyquist rate for non-sparse signals of $2 F-$ in the piano example, each sampler has a rate on the order of 50 Hz rather than 8000 Hz .

The superresolution frequency is optimal because two signals of magnitude $v_{i}$ and frequency separation $\nu<1 / T$ will differ by $O\left(\nu^{2} T^{2}\left|v_{i}\right|^{2}\right)$ over the duration $T$, so for $\nu$ below our threshold the difference is just $\mathcal{N}^{2}$. Hence if the observed signal $x(t)$ looks identical to $\left(v_{i}, f_{i}\right)$, it might actually be $\left(v_{i}, f_{i}^{\prime}\right)$ with noise equaling the difference between the two.

The only previous result known in the form of (11.1) was [Moi15], which lost a $\operatorname{poly}\left(k, \frac{\max \left|v_{i}\right|}{\min \left|v_{i}\right|}, \eta\right)$ factor in noise tolerance and also did not optimize for sample complexity.

### 11.1.1 Comparison to Naive Methods

This section compares our result to some naive ways one could try to solve this problem by applying algorithms not designed for a continuous, sparse Fourier transform setting. The next section will compare our result to algorithms that are designed for such a setting.

Nyquist Sampling. The traditional theory of band-limited signals from discrete samples says that, from samples taken at the Nyquist rate $2 F$, one can reconstruct an $F$-band-limited signal exactly. The Whittaker-Shannon interpolation formula then says that

$$
\begin{equation*}
x^{*}(t)=\sum_{i=-\infty}^{\infty} x^{*}(2 F i) \operatorname{sinc}(2 F t-i) \tag{11.4}
\end{equation*}
$$

where $\operatorname{sinc}(t)$ is the normalized sinc function $\frac{\sin (\pi t)}{\pi t}$. This is for the band-limited "pure" signal $x^{*}$, but one could then get a relationship for samples of the actual signal $x(t)$. This has no direct implications for learning the tones (e.g. our (11.1) or (11.2)), but for learning the signal (our (11.3)) there is also an issue. Even in the absence of noise and for $k=1$, this method will have error polynomially rather than exponentially small in the number of samples.

That is, if there is no noise the method has zero error given infinitely many samples. But we only receive samples over the interval $[0, T]$, leading to error. Consider the trivial setting of $x(t)=1$. The partial sum of (11.4) at a given $t$ will be missing terms for $i>2 F t$ and $i<0$, which (for a random $t$ in $[0, T / 2]$ ) have magnitude at most $1 /(F t)$. The terms alternate in sign, so the sum has error approximately $1 /(F t)^{2}$. This means that the error over the first $1 / F$ time is a constant, leading to average error of $\frac{1}{F T}$. This is with an algorithm that uses FT samples and time. By contrast, our algorithm in the noiseless setting has error exponentially small in the samples and running time.

Discrete Sparse Fourier Transforms. An option related to the previous would be to discretize very finely, then apply a discrete sparse Fourier transform algorithm to keep the sample complexity and runtime small. The trouble here is that sparse Fourier transforms
require sparsity, and this process decreases the sparsity. In particular, this process supposes that the signal is periodic with period $T$, so one can analyze this process as first converting the signal to one, equivalent over $[0, T]$, but with frequency spectrum only containing integer multiples of $1 / T$. This is done by convolving each frequency $f_{i}$ with a sinc function (corresponding to windowing to $[0, T]$ ) then restricting to multiples of $1 / T$ (corresponding to aliasing). The result is that a one-sparse signal $e^{2 \pi \mathbf{i} f_{i} t}$ is viewed as having Fourier spectrum

$$
\widehat{x^{\prime \prime}}[j]=\operatorname{sinc}\left(f_{i} T-j\right)
$$

for $j \in \mathbb{Z}$. When $f_{i}$ is not a multiple of $1 / T$, this means the signal is not a perfectly sparse signal. And this is true regardless of the discretization level, which only affects the subsequent error from aliasing $j \in \mathbb{Z}$ down to $\mathbb{Z}_{n}$. To have error proportional to $\delta\left\|\widehat{x}^{*}\right\|_{2}$, one would need to run such methods for a sparsity level of $k / \delta$. Thus, as with Nyquist sampling, the sample and runtime will be polynomial, rather than logarithmic, in $\delta$.

The above discussion refers to methods for learning the signal (our (11.3)). In terms of learning the tones, one could run the algorithm for sparsity $O(k)$ so that $\delta$ is a small constant, which would let one learn roughly where the peaks are and get most of the frequencies to the nearest $1 / T$. This would give a similar bound to our (11.1), but without the superresolution effect as the noise becomes small. On the plus side, the duration could be just $O(1 / \eta)$ which is sufficient for the different peaks to be distinguishable - rather than $O\left(\frac{\log k}{\eta}\right)$ as our method would require, and the time and sample complexities could save a $\log k$ factor (if one did not want to recover all the tones, just most of them).

Essentially, this algorithm tries to round each frequency to the nearest multiple of $1 / T$, which introduces noise that is a constant fraction of the signal. If the signal-to-noise
ratio is already low, this does not increase the noise level by that much so such an algorithm will work reasonably well. If the signal-to-noise ratio is fairly high, however, then the added noise leads to much worse performance. Getting a constant factor approximation to the whole signal is only nontrivial for high SNR, so such a method does badly in that setting. For approximating the tones, it is comparable to our method in the low SNR setting but does not improve much as SNR increases.

### 11.1.2 Previous Work In Similar Settings

There have been several works that recover continuous frequencies from samples in the time domain. Some of these are in our setting where the samples can be taken at arbitrary positions over $[0, T]$ and others are in the discrete-time (DTFT) setting where the samples must be taken at multiples of the Nyquist frequency $\frac{1}{2 F}$.

The results of [TBSR13, CF14, YX15, DB13] show that a convex program can solve the problem in the DTFT setting using $O(k \log k \log (F T))$ samples if the duration is $T>$ $O\left(\frac{1}{\eta}\right)$, in the setting where $g(t)=0$ and the coefficients of $x^{*}$ have random phases. The sample complexity can be one log factor better than ours, which one would expect for the noiseless setting. None of these results show robustness to noise, and some additionally require a running time polynomial in $F T$ rather than $k$.

The result of $\left[\mathrm{BCG}^{+} 12\right]$ is in a similar setting to our paper, using techniques of the same lineage. It achieves very similar sample complexity and running time to our algorithm, and a guarantee similar in spirit to (11.1) with some notion of robustness. However, the robustness is weaker than ours in significant ways. They consider the noise $g(t)$ in frequency space (i.e. $\widehat{g}(f)$ ), then require that $\widehat{g}(f)$ is zero at any frequency within $\eta$ of the signal
frequencies $f_{i}$, and bound the error in terms of $\mathcal{N}^{\prime}=\|\widehat{g}\|_{1} / k$ instead of $\|g\|_{2}$. This fails to cover simple examples of noise, including i.i.d. Gaussian noise $g(t) \sim N\left(0, \sigma^{2}\right)$ and the noise one would get from slow decay of the signal over time (e.g. $x(t)=x^{*}(t) e^{-\frac{t}{100 T}}$.). Both types of noise violate both assumptions on the noise: $\widehat{g}(f)$ will be nonzero arbitrarily close to each $f_{i}$ and $\|\widehat{g}\|_{1}$ will be unbounded. Their result also requires a longer duration than our algorithm and has worse precision for any fixed duration.

The result of [Moi15] studies noise tolerance in the DTFT setting, ignoring sample complexity and running time. It shows that the matrix pencil method [HS90], using $F T$ samples, achieves a guarantee of the form (11.1), except that the bounds are an additional $\operatorname{poly}(F T, k, \delta)$ factor larger. Furthermore, it shows a sharp characterization of the minimal $T$ for which this is possible by any algorithm: $T=(1 \pm o(1)) \frac{2}{\eta}$ is necessary and sufficient. It is an interesting question whether the lower bound generalizes to our non-DTFT setting, where the samples are not necessarily taken from an even grid.

Lastly, $\left[\mathrm{SAH}^{+} 13\right]$ tries to apply sparse Fourier transforms to a domain with continuous signals. They first apply a discrete sparse Fourier transform then use hill-climbing to optimize their solution into a decent set of continuous frequencies. They have interesting empirical results but no theoretical ones.

### 11.2 Algorithm Overview

At a high level, our algorithm is an adaptation of the framework used by [HIKP12a] to the continuous setting. However, getting our result requires a number of subtle changes to the algorithm. This section will describe the most significant ones. We assume some familiarity with previous work in the area [GGI+ 02 , CCF02, GMS05, GLPS10, HIKP12a].

First we describe a high-level overview of the structure. The algorithm proceeds in $\log k$ stages, where each stage attempts to recover each tone with a large constant probability (e.g. 9/10). In each stage, we choose a parameter $\sigma \approx \frac{T}{k \log (k / \delta)}$ that we think of as "hashing" the frequencies into random positions. For this $\sigma$, we will choose about $\log (F T)$ different random "start times" $t_{0}$ and sample an arithmetic sequence starting at $t_{0}$, i.e. observe

$$
x\left(t_{0}\right), x\left(t_{0}+\sigma\right), x\left(t_{0}+2 \sigma\right), \ldots, x\left(t_{0}+(k \log (k / \delta)) \sigma\right)
$$

We then scale these observations by a "window function," which has specific properties but among other things scales down the values near the ends of the sequence, giving a smoother transition between the time before and after we start/end sampling. We alias this down to $B=O(k)$ terms (i.e. add together terms $1, B+1,2 B+1, \ldots$ to get a $B$-dimensional vector) and take the $B$-dimensional DFT. This gives a set of $B$ values $\widehat{u}_{i}$. The observation made in previous papers is that $\widehat{u}$ is effectively a hashing of the tones of $\widehat{x}$ into $B$ buckets, where $\sigma$ defines a permutation on the frequencies that affects whether two different tones land in the same bucket, and $\widehat{u}_{j}$ approximately equals the sum of all the tones that land in bucket $j$, each scaled by a phase shift depending on $t_{0}$.

Because of this phase shift, for each choice of $t_{0}$ the value of $\widehat{u}_{j}$ is effectively a sample from the Fourier transform of a signal that contains only the tones of $\widehat{x}^{*}$ that land in bucket
$j$, with zeros elsewhere. And since there are $k$ tones and $O(k)$ buckets, most tones are alone in their bucket. Therefore this sampling strategy reduces the original problem of $k$-sparse recovery to one of 1-sparse recovery - we simply choose $t_{0}$ according to some strategy that lets us achieve 1 -sparse recovery, and recover a tone for each bin.

One-sparse recovery. The algorithm for one-sparse recovery in [HIKP12a] is a good choice for adaptation to the continuous setting. It narrows down to the frequency in a locality-aware way, maintaining an interval of frequencies that decreases in size at each stage (in contrast to the method in [GMS05], which starts from the least significant bit rather than most significant bit).

If a frequency is perturbed slightly in time (e.g., by multiplying by a very slow decay over time) this will blur the frequency slightly into a narrow band. The one-sparse recovery algorithm of [HIKP12a] will proceed normally until it gets to the narrow scale, at which point it will behave semi-arbitrarily and return something near that band. This gives a desired level of robustness - the error in the recovered frequency will be proportional to the perturbation.

Still, to achieve our result we need a few changes to the one-sparse algorithm. One is related to the duration $T$ : in the very last stage of the algorithm, when the interval is stretched at the maximal amount, we can only afford one "fold" rather than the typical $O(\log n)$. The only cost to this is in failure probability, and doing it for one stage is fine - but showing this requires a different proof. Another difference is that we need the final interval to have precision $\frac{1}{T \rho}$ if the signal-to-noise ratio is $\rho$ - the previous analysis showed $\frac{1}{T \sqrt{\rho}}$ and needed to be told $\rho$, but (as we shall see) to achieve an $\ell_{2} / \ell_{2}$ guarantee we need the optimal
$\rho$-dependence and for the algorithm to be oblivious to the value of $\rho$. Doing so requires a modification to the algorithm and slightly more clever analysis.
$k$-sparse recovery. The changes to the $k$-sparse recovery structure are broader. First, to make the algorithm simpler we drop the [GLPS10]-style recursion with smaller $k$, and just repeat an $O(k)$-size hashing $O(\log k)$ times. This loses a $\log k$ factor in time and sample complexity, but because of the other changes it is not easy to avoid, and at the same time improves our success probability.

The most significant changes come because we can no longer measure the noise in frequency space or rely on the hash function to randomize the energy that collides with a given heavy hitter. Because we only look at a bounded time window $T$, Parseval's identity does not hold and the energy of the noise in frequency space may be unrelated to its observed energy. Moreover, if the noise consists of frequencies infinitesimally close to a true frequency, then because $\sigma$ is bounded the true frequency will always hash to the same bin as the noise. These two issues are what drive the restrictions on noise in the previous work $\left[\mathrm{BCG}^{+} 12\right]$ - assuming the noise is bounded in $\ell_{1}$ norm in frequency domain and is zero in a neighborhood of the true frequencies fixes both issues. But we want a guarantee in terms of the average $\ell_{2}$ noise level $\mathcal{N}^{2}$ in time domain over the observed duration. If the noise level is $\mathcal{N}^{2}$, because we cannot hash the noise independently of the signal, we can only hope to guarantee reliable recovery of tones with magnitude larger than $\mathcal{N}^{2}$. This is in contrast to the $\mathcal{N}^{2} / k$ that is possible in the discrete setting, and would naively lose a factor of $k$ in the $\ell_{2} / \ell_{2}$ approximation.

The insight here is that, even though the noise is not distributed randomly across bins, the total amount of noise is still bounded. If a heavy hitter of magnitude $v^{2}$ is not
recovered due to noise, that requires $\Omega\left(v^{2}\right)$ noise mass in the bin that is not in any other bin. Thus the total amount of signal mass not recovered due to noise is $O\left(\mathcal{N}^{2}\right)$, which allows for $\ell_{2} / \ell_{2}$ recovery.

This difference is why our algorithm only gets a constant factor approximation rather than the $1+\epsilon$ guarantee that hashing techniques for sparse recovery can achieve in other settings. These techniques hash into $B=O(k / \epsilon)$ bins so the average noise per bin is $O\left(\frac{\epsilon}{k} \mathcal{N}^{2}\right)$. In our setting, where the noise is not hashed independently of the signal, this would give no benefit.

Another difference arises in the choice of the parameter $\sigma$, which is the separation between samples in the arithmetic sequence used for a single hashing, and gives the permutation during hashing. In the discrete setting, one chooses $\sigma$ uniformly over $n$, which in our setting would correspond to a scale of $\sigma \approx \frac{1}{\eta}$. Since the arithmetic sequences have $O(k \log (k / \delta))$ samples, the duration would then become at least $\frac{k \log (k / \delta)}{\eta}$ (which is why $\left[\mathrm{BCG}^{+} 12\right]$ has this duration). What we observe is that $\sigma$ can actually be chosen at the scale of $\frac{1}{k \eta}$, giving the desired $O\left(\frac{\log (k / \delta)}{\eta}\right)$ duration. This causes frequencies at the minimum separation $\eta$ to always land in bins that are a constant separation apart. This is sufficient because we use [HIKP12a]-style window functions with strong isolation properties (and, in fact, [HIKP12a] could have chosen $\sigma \approx n / B)$; it would be an issue if we were using the window functions of [GMS05, IK14] that have smaller supports but less isolation.

Getting an $\ell_{2}$ bound Lastly, converting the guarantee (11.2) into (11.3) is a nontrivial task that is trivial in the discrete setting. In the discrete setting, it follows immediately from the different frequencies being orthogonal to each other. In our setting, we use that
the recovered frequencies should themselves have $\Omega(\eta)$ separation, and that well-separated frequencies are nearly orthogonal over long enough time scales $T \gg 1 / \eta$.

This bears some similarity to issues that arise in sparse recovery with overcomplete dictionaries. It would be interesting to see whether further connections can be made between the problems.

### 11.3 Proof outline

In this section we present the key lemmas along the path to producing the algorithm. The full proof are presented later.

Notation. First we define the notation necessary to understand the lemmas. The full notation as used in the proofs appears in Section 11.4.

The algorithm proceeds in stages, each of which hashes the frequencies to $B$ bins. The hash function depends on two parameters $\sigma$ and $b$, and so we define it as $h_{\sigma, b}(f):[-F, F] \rightarrow$ $[B]$.

A tone with a given frequency $f$ can have two "bad events" $E_{\text {coll }}(f)$ or $E_{\text {off }}(f)$ hold for a given hashing. These correspond to colliding with another frequency of $x^{*}$ or landing within an $\alpha$ fraction of the edge, respectively; they each will occur with small constant probability.

For a given hashing, we will choose a number of different offsets $a$ that let us perform recovery of the tones that have neither bad event in this stage.

We use $f \lesssim g$ to denote that there exists a constant $C$ such that $f \leq C g$, and $f \approx g$ to denote $f \lesssim g \lesssim f$.

Key Lemmas First, we need to be able to compare the distance between two pure tone signals in time domain to their differences in parameters. The relation is as follows:

Lemma 11.3.1. Let $(v, f)$ and $\left(v^{\prime}, f^{\prime}\right)$ denote any two tones, i.e., (magnitude, frequency)
pairs. Then for

$$
\operatorname{dist}\left((v, f),\left(v^{\prime}, f^{\prime}\right)\right)^{2}:=\frac{1}{T} \int_{0}^{T}\left|v e^{2 \pi f t \mathbf{i}}-v^{\prime} e^{2 \pi f^{\prime} t \mathbf{i}}\right|^{2} \mathrm{~d} t
$$

we have

$$
\operatorname{dist}\left((v, f),\left(v^{\prime}, f^{\prime}\right)\right)^{2} \bar{\sim}\left(|v|^{2}+\left|v^{\prime}\right|^{2}\right) \cdot \min \left(1, T^{2}\left|f-f^{\prime}\right|^{2}\right)+\left|v-v^{\prime}\right|^{2}
$$

and

$$
\operatorname{dist}\left((v, f),\left(v^{\prime}, f^{\prime}\right)\right) \approx|v| \cdot \min \left(1, T\left|f-f^{\prime}\right|\right)+\left|v-v^{\prime}\right|
$$

The basic building block for our algorithm is a function HashToBins, which is very similar to one of the same name in [HIKP12a].

The key property of HAshToBins is that, if neither "bad" event holds for a frequency $f$ (i.e. it does not collide or land near the boundary of the bin), then for the bin $j=h_{\sigma, b}(f)$ we have that $\left|\widehat{u}_{j}\right| \approx\left|\widehat{x^{*}}(f)\right|$ with a phase depending on $a$.

How good is the approximation? In the discrete setting, one can show that each tone has error about $\mathcal{N}^{2} / B$ in expectation. Here, because the hash function cannot randomize the noise, we instead show that the total error over all tones is about $\mathcal{N}^{2}$ :

Lemma 11.3.2. Let $\sigma \in\left[\frac{1}{B \eta}, \frac{2}{B \eta}\right]$ uniformly at random, then $b \in\left[0, \frac{[F / \eta\rceil}{\sigma B}\right], a \in\left[0, \frac{c T}{\sigma}\right]$ be sampled uniformly at random for some constant $c>0$. Let the other parameters be arbitrary in $\widehat{u}=\operatorname{HashToBins}\left(x, P_{\sigma, a, b}, B, \delta, \alpha\right)$, and consider

$$
H=\left\{f \in \operatorname{supp}\left(\widehat{x}^{*}\right) \mid \text { neither } E_{\text {coll }}(f) \text { nor } E_{\text {off }}(f) \text { holds }\right\}
$$

and $I=[B] \backslash h_{\sigma, b}\left(\operatorname{supp}\left(\widehat{x}^{*}\right)\right)$ to be the bins that have no frequencies hashed to them. Then

$$
\underset{\sigma, b, a}{\mathbb{E}}\left[\sum_{f \in H}\left|\widehat{u}_{h_{\sigma, b}(f)}-\widehat{x}^{*}(f) e^{a \sigma 2 \pi f \mathbf{i}}\right|^{2}+\sum_{j \in I} \widehat{u}_{j}^{2}\right] \lesssim \mathcal{N}^{2}
$$

We prove Lemma 11.3.2 by considering the cases of $x^{*}=0$ and $g=0$ separately; linearity then gives the result. Both follow from properties of our window functions.

Lemma 11.3.3. If $x^{*}(t)=0, \forall t \in[0, T]$, then

$$
\underset{\sigma, a, b}{\mathbb{E}}\left[\sum_{j=1}^{B}\left|\widehat{u}_{j}\right|^{2}\right] \lesssim \frac{1}{T} \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t
$$

Lemma 11.3.4. If $g(t)=0, \forall t \in[0, T]$. Let $H$ denote a set of frequencies, $H=\{f \in$ $\operatorname{supp}\left(\widehat{x}^{*}\right) \mid$ neither $E_{\text {coll }}(f)$ nor $E_{\text {off }}(f)$ holds $\}$. Then,

$$
\underset{\sigma, a, b}{\mathbb{E}}\left[\sum_{f \in H}\left|\widehat{u}_{h_{\sigma, b}(f)}-\widehat{x^{*}}(f) e^{a \sigma 2 \pi f \mathbf{i}}\right|^{2}\right] \leq \delta\left\|\widehat{x^{*}}\right\|_{1}^{2}
$$

Lemma 11.3.2 is essentially what we need for 1 -sparse recovery. We first show a lemma about the inner call, which narrows the frequency from a range of size $\Delta l$ to one of size $\frac{\Delta l}{\rho s t}$ for some parameters $\rho s t$. This gives improved performance (superresolution) when the signal-to-noise ratio $\rho$ within the bucket is high. The parameter $s$ and $t$ provide a tradeoff between success probability, performance, running time, and duration.

Lemma 11.3.5. Consider any $B, \delta, \alpha$. Algorithm HashToBins takes $O(B \log (k / \delta))$ samples and runs in $O\left(\frac{B}{\alpha} \log (k / \delta)+B \log B\right)$ time.

Lemma 11.3.6. Given $\sigma$ and $b$, consider any frequency $f$ for which neither $E_{\text {coll }}(f)$ nor $E_{\text {off }}(f)$ holds, and let $j=h_{\sigma, b}(f)$. Let $\mu^{2}(f)=\mathbb{E}_{a}\left[\left|\widehat{u_{j}}-\widehat{x^{*}}(f) e^{a \sigma 2 \pi f \mathbf{i}}\right|^{2}\right]$ and $\rho^{2}=\left|\widehat{x^{*}}(f)\right|^{2} / \mu^{2}(f)$. For sufficiently large $\rho$, and $\forall 0<s<1, t \geq 4$, consider any run of LOcATEInNER with $f \in\left[l_{j}-\frac{\Delta l}{2}, l_{j}+\frac{\Delta l}{2}\right]$. It takes $O\left(R_{l o c}\right)$ random $(\gamma, \beta) \in\left[\frac{1}{2}, 1\right] \times\left[\frac{s t}{4 \sigma \Delta l}, \frac{s t}{2 \sigma \Delta l}\right]$ samples over duration $\beta \sigma=\Theta\left(\frac{s t}{\Delta l}\right)$, runs in $O\left(s t R_{l o c}\right)$ time, to learn $f$ within a region that has length $\Theta\left(\frac{\Delta l}{t}\right)$ with failure probability at most $\left(\frac{4}{s \rho}\right)^{R_{l o c}}+t \cdot(60 s)^{R_{l o c} / 2}$.

By repeating this inner loop, we can recover the tones in almost every bin that does not have the "bad" events happen, so we recover a large fraction of the heavy hitters in each stage.

Lemma 11.3.7. Algorithm LOCATEKSignal takes $O\left(k \log _{C}(F T) \log (k / \delta)\right)$ samples over $O\left(\frac{\log (k / \delta)}{\eta}\right)$ duration, runs in $O\left(k \log _{C}(F T) \log (F T / \delta)\right)$ time, and outputs a set $L \subset[-F, F]$ of $O(k)$ frequencies with minimum separation $\Omega(\eta)$.

Given $\sigma$ and $b$, consider any frequency $f$ for which neither of $E_{\text {coll }}(f)$ or $E_{\text {off }}(f)$ hold. Let $j=h_{\sigma, b}(f), \mu^{2}(f)=\mathbb{E}_{a}\left[\left|\widehat{u}_{j}-\widehat{x^{*}}(f) e^{a \sigma 2 \pi f \mathbf{i}}\right|^{2}\right]$, and $\rho^{2}=\left|\widehat{x^{*}}(f)\right|^{2} / \mu^{2}(f)$. If $\rho>C$, then with an arbitrarily large constant probability there exists an $f^{\prime} \in L$ with

$$
\left|f-f^{\prime}\right| \lesssim \frac{1}{T \rho}
$$

Combining this with estimation of the magnitudes of recovered frequencies, we can show that the total error over all bins without "bad" events - that is, bins with either one well placed frequency or zero frequencies - is small. At this point we give no guarantee for the (relatively few) bins with bad events; the recovered values there may be arbitrarily large. Lemma 11.3.8. Algorithm OneStage takes $O\left(k \log _{C}(F T) \log (k / \delta)\right)$ samples over $O\left(\frac{\log (k / \delta)}{\eta}\right)$ duration, runs in $O\left(k\left(\log _{C}(F T) \log (F T / \delta)\right)\right)$ time, and outputs a set of $\left\{\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\right\}$ of size $O(k)$ with $\min _{i \neq j}\left|f_{i}^{\prime}-f_{j}^{\prime}\right| \gtrsim \eta$. Moreover, one can imagine a subset $S \subseteq[k]$ of "successful" recoveries, where $\operatorname{Pr}[i \in S] \geq \frac{9}{10} \forall i \in[k]$ and for which there exists an injective function $\pi:[k] \rightarrow[O(k)]$ so that

$$
\underset{\sigma, b}{\mathbb{E}}\left[\sum_{i \in S} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} \mathbf{t i}}\right|^{2} \mathrm{~d} t\right] \lesssim C^{2} \mathcal{N}^{2}
$$

with $1-1 / k^{c}$ probability for an arbitrarily large constant $c$.

We can repeat the procedure for $O(\log k)$ stages and merge the results, getting a list of $O(k)$ tones that includes $k$ tones that match up well to the true tones. However, we give no guarantee for the rest of the recovered tones at this point - as far as the analysis is concerned, mistakes from bins with collisions may cause arbitrarily large spurious tones.

Lemma 11.3.9. Repeating algorithm OneStage $O(\log k)$ times, MergedStages returns a set $\left\{\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\right\}$ of size $O(k)$ with $\min _{i \neq j}\left|f_{i}^{\prime}-f_{j}^{\prime}\right| \gtrsim \eta$ that can be indexed by $\pi$ such that

$$
\sum_{i=1}^{k} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} \mathbf{t} \mathbf{i}}\right|^{2} \mathrm{~d} t \lesssim C^{2} \mathcal{N}^{2}
$$

with probability $1-1 / k^{c}$ for an arbitrarily large constant $c$.
To address the issue of spurious tones, we run the above algorithm twice and only take the tones that are recovered in both stages. We show that the resulting $O(k)$ tones are together a good approximation to the vector.

Lemma 11.3.10. If we run Mergedstages twice and take the tones $\left\{\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\right\}$ from the first result that have $f_{i}^{\prime}$ within $c \eta$ for small $c$ of some frequency in the second result, we get a set of $k^{\prime \prime}=O(k)$ tones that can be indexed by some permutation $\pi$ such that

$$
\begin{equation*}
\sum_{i=1}^{k} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} t \mathbf{i}}\right|^{2} \mathrm{~d} t+\sum_{i=k+1}^{k^{\prime \prime}}\left|v_{i}^{\prime}\right|^{2} \lesssim C^{2} \mathcal{N}^{2} \tag{11.5}
\end{equation*}
$$

Simply picking out the largest $k$ recovered tones then gives the result (11.2).
Theorem 11.3.11. Algorithm ContinuousFourierSparseRecovery returns a set $\left\{\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\right\}$ of size $k$ with $\min _{i \neq j}\left|f_{i}^{\prime}-f_{j}^{\prime}\right| \gtrsim \eta$ for which

$$
\sum_{i=1}^{k} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} \mathbf{i} \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} t \mathbf{i}}\right|^{2} \mathrm{~d} t \lesssim C^{2} \mathcal{N}^{2}
$$

with probability $1-1 / k^{c}$ for an arbitrarily large constant $c$.

By only considering the term in the sum corresponding to tone $i$ and applying Lemma 11.3.1, we get result (11.1):

Corollary 11.3.12. With probability $1-1 / k^{c}$ for an arbitrarily large constant $c$, we recover a set of tones $\left\{\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\right\}$ such that, for any $v_{i}$ with $\left|v_{i}\right| \gtrsim \mathcal{N}$, we have for an appropriate permutation of the indices that

$$
\begin{equation*}
\left|f_{i}^{\prime}-f_{i}\right| \lesssim \frac{\mathcal{N}}{T\left|v_{i}\right|} \quad \text { and } \quad\left|v_{i}^{\prime}-v_{i}\right| \lesssim \mathcal{N} \tag{11.6}
\end{equation*}
$$

We then show that (11.2) implies (11.3) for sufficiently long durations $T$. A long duration helps because it decreases the correlation between $\eta$-separated frequencies.

Lemma 11.3.13. Let $\left\{\left(v_{i}, f_{i}\right)\right\}$ and $\left\{\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\right\}$ be two sets of $k$ tones for which $\min _{i \neq j}\left|f_{i}-f_{j}\right| \geq$ $\eta$ and $\min _{i \neq j}\left|f_{i}^{\prime}-f_{j}^{\prime}\right| \gtrsim \eta$ for some $\eta>0$. Suppose that $T>O\left(\frac{\log ^{2} k}{\eta}\right)$. Then these sets can be indexed such that

$$
\begin{equation*}
\frac{1}{T} \int_{0}^{T}\left|\sum_{i=1}^{k}\left(v_{i}^{\prime} e^{2 \pi \mathbf{i} f_{i}^{\prime} t}-v_{i} e^{2 \pi \mathbf{i} f_{i} t}\right)\right|^{2} \mathrm{~d} t \lesssim \sum_{i=1}^{k} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi \mathbf{i} f_{i}^{\prime} t}-v_{i} e^{2 \pi \mathrm{i} \mathrm{f}_{i} t}\right|^{2} \mathrm{~d} t \tag{11.7}
\end{equation*}
$$

Combining Theorem 11.3.11 and Lemma 11.3.13 immediately implies
Theorem 11.3.14. Suppose we sample for a duration $T>O\left(\frac{\log (1 / \delta)+\log ^{2} k}{\eta}\right)$. Then the reconstructed signal $x^{\prime}(t)=\sum_{i=1}^{k} v_{i}^{\prime} e^{2 \pi \mathrm{i} f_{i}^{\prime} t}$ achieves a constant factor approximation to the complete signal $x$ :

$$
\begin{equation*}
\frac{1}{T} \int_{0}^{T}\left|x^{\prime}(t)-x(t)\right|^{2} \mathrm{~d} t \lesssim C^{2} \mathcal{N}^{2} \tag{11.8}
\end{equation*}
$$

The algorithm takes $O\left(k \log \frac{F}{\eta} \log \left(\frac{k}{\delta}\right) \log (k)\right)$ samples, runs in $O\left(k \log \frac{F}{\eta} \log \left(\frac{F T}{\delta}\right) \log (k)\right)$ time, and succeeds with probability at least $1-1 / k^{c}$ for an arbitrarily large constant $c$.

That finishes the proof of our main theorem. We also show that our "superresolution" precision from (11.1) is optimal, which is a simple corollary of Lemma 11.3.1.

Lemma 11.3.15. There exists a constant $c>0$ such that, for a given sample duration $T$, one cannot recover the frequency $f$ to within

$$
c \frac{\mathcal{N}}{T\left|\widehat{x}^{*}(f)\right|}
$$

with $3 / 4$ probability, for all $\delta>0$, even if $k=1$.

### 11.4 Notation and Definitions: Permutation, Hashing, Filters

This section gives definitions about the permutation, hashing, and filters that are used throughout the proofs. Let $[n]$ denote the set $\{1,2, \cdots, n-1, n\}$. $\mathbb{R}$ denotes the real numbers, $\mathbb{Z}$ denotes the integer numbers and $\mathbb{C}$ denotes the complex numbers. The convolution of two continuous functions $f$ and $g$ is written as $f * g$,

$$
(f * g)(t):=\int_{-\infty}^{+\infty} f(\tau) g(t-\tau) \mathrm{d} \tau
$$

and the discrete convolution of $f$ and $g$ is given by,

$$
(f * g)[n]:=\sum_{m=-\infty}^{+\infty} f[m] g[n-m]
$$

Let $\mathbf{i}$ denote $\sqrt{-1}$, and $e^{\mathbf{i} \theta}=\cos (\theta)+\mathbf{i} \sin (\theta)$. For any complex number $z \in \mathbb{C}$, we have $z=a+\mathbf{i} b$, where $a, b \in \mathbb{R}$. Define $\bar{z}=a-\mathbf{i} b,|z|^{2}=z \bar{z}=a^{2}+b^{2}$ and let $\phi(z)$ be the phase of $z$. Let $\operatorname{supp}(f)$ denote the support of function/vector $f$, and $\|f\|_{0}=|\operatorname{supp}(f)|$. For any $p \in[1, \infty]$, the $\ell_{p}$ norm of a vector of $x$ is $\|x\|_{p}=\left(\sum_{i}\left|x_{i}\right|^{p}\right)^{\frac{1}{p}}$, defined to be $\max _{i}\left|x_{i}\right|$ for $p=\infty$. Let $k$ denote the sparsity of frequency domain. All the frequencies $\left\{f_{1}, f_{2}, \cdots, f_{k}\right\}$ are from $[-F, F]$. Let $B=O(k)$ denote the number of hash bins in our algorithm.

We translate the "permutation" $P_{\sigma, a, b}$ of [HIKP12a] from the DFT setting to the DTFT setting.

Definition 11.4.1. $\left(P_{\sigma, a, b} x^{*}\right)(t)=x^{*}(\sigma(t-a)) e^{-2 \pi \mathbf{i} \sigma b t}$.
Lemma 11.4.1. $\widehat{P_{\sigma, a, b} x^{*}}(\sigma(f-b))=e^{-2 \pi f \sigma a \mathrm{i}} \widehat{x}^{*}(f)$

Proof. The time domain representation of the given Fourier definition would be

$$
\begin{aligned}
x^{\prime}(t) & =\sum_{f \in \operatorname{supp}\left(\widehat{x}^{*}\right)} e^{-2 \pi \mathbf{i} \sigma a f} \widehat{x}^{*}(f) e^{2 \pi \mathbf{i} \sigma(f-b) t} \\
& =\sum_{f \in \operatorname{supp}\left(\widehat{x}^{*}\right)} e^{-2 \pi \mathbf{i} \sigma b t} \widehat{x}^{*}(f) e^{2 \pi \mathbf{i} \sigma(t-a) f} \\
& =e^{-2 \pi \mathbf{i} \sigma b t} x^{*}(\sigma(t-a))
\end{aligned}
$$

which matches, so the formula is right.

We also extend the flat window function for the DFT setting [HIKP12a], [HIKP12b] to the DTFT setting:

Definition 11.4.2. Let $M=O\left(B \log \frac{k}{\delta}\right)$. We say that $\left(G, \widehat{G^{\prime}}\right)=\left(G_{B, \delta, \alpha},{\widehat{G^{\prime}}}_{B, \delta, \alpha}\right) \in \mathbb{R}^{M} \times$ $\mathbb{R}^{[-F, F]}$ is a flat window function with parameters $B \geq 1, \delta>0$, and $\alpha>0$. For simplicity, let's say $B$ is a function of $\alpha$. Define $|\operatorname{supp}(G)|=M$ and $\widehat{G^{\prime}}$ satisfies

- $G_{i}=\frac{\sin \left(i \frac{1}{B}\right)}{i} \cdot e^{-\frac{i^{2}}{2 \sigma^{2}}}$, where $\sigma=\Theta(B \sqrt{\log (k / \delta)})$.
- $\widehat{G}(f)=\sum_{i=1}^{M} G_{i} e^{f \cdot \frac{i}{M} 2 \pi \mathrm{i}}$.
- $\operatorname{supp}\left(\widehat{G^{\prime}}\right) \subset\left[-\frac{2 \pi}{2 B}, \frac{2 \pi}{2 B}\right]$.
- $\widehat{G^{\prime}}(f)=1$ for all $f \in\left[-\frac{(1-\alpha) 2 \pi}{2 B}, \frac{(1-\alpha) 2 \pi}{2 B}\right]$.
- $\widehat{G^{\prime}}(f)=0$ for all $|f| \geq \frac{2 \pi}{2 B}$.
- $\widehat{G^{\prime}}(f) \in[0,1]$ for all $f$.
- $\left\|\widehat{G^{\prime}}-\widehat{G}\right\|_{\infty}^{2}<\delta / k$.

Claim 11.4.2. $\sum_{i=1}^{M} G_{i}^{2} \approx \frac{1}{B}$, where $M=O(B \log k / \delta)$.
Proof. By definition of $G_{i}$, we have

$$
\sum_{i=1}^{M} G_{i}^{2}=2 \sum_{i=1}^{M / 2} \frac{\sin ^{2}\left(i \frac{1}{B}\right)}{(i)^{2}}\left(e^{-\frac{(i)^{2}}{2 \sigma^{2}}}\right)^{2}
$$

There exists some constant $c \in[0,2 \pi)$, such that $\frac{\sin (i / B)}{i} \approx \frac{1}{B}$ if $i / B<c \pi$.

$$
\begin{aligned}
\sum_{i=1}^{M} G_{i}^{2} & =2 \sum_{i=1}^{\lfloor B c \pi\rfloor} \frac{\sin ^{2}\left(i \frac{1}{B}\right)}{(i)^{2}}\left(e^{\left.-\frac{(i)^{2}}{2 \sigma^{2}}\right)^{2}+2 \sum_{i=\lceil B c \pi\rceil}^{M / 2} \frac{\sin ^{2}\left(i \frac{1}{B}\right)}{(i)^{2}}\left(e^{-\frac{(i)^{2}}{2 \sigma^{2}}}\right)^{2}}\right. \\
& \leq 2 \sum_{i=1}^{\lfloor B c \pi\rfloor} \frac{\sin ^{2}\left(i \frac{1}{B}\right)}{(i)^{2}} \cdot 1+2 \sum_{i=\lceil B c \pi\rceil}^{M / 2} \frac{\sin ^{2}\left(i \frac{1}{B}\right)}{(i)^{2}} \cdot 1 \\
& \lesssim \sum_{i=1}^{\lfloor B c \pi\rfloor} \frac{1}{B^{2}}+\sum_{i=\lceil B c \pi\rceil}^{M / 2} \frac{1}{i^{2}} \\
& \lesssim \frac{1}{B} .
\end{aligned}
$$

Thus, we show an upper bound. It remains to prove the lower bound.

$$
\begin{aligned}
\sum_{i=1}^{M} G_{i}^{2} & \geq 2 \sum_{i=1}^{\lfloor B c \pi\rfloor} \frac{\sin ^{2}\left(i \frac{1}{B}\right)}{(i)^{2}}\left(e^{\left.-\frac{(i)^{2}}{2 \sigma^{2}}\right)^{2}}\right. \\
& \geq 2 \sum_{i=1}^{\lfloor B c \pi\rfloor} \frac{\sin ^{2}\left(i \frac{1}{B}\right)}{(i)^{2}}\left(e^{-\frac{(B c \pi)^{2}}{2 \sigma^{2}}}\right)^{2} \\
& \gtrsim 2 \sum_{i=1}^{\lfloor B c \pi\rfloor} \frac{1}{B^{2}}\left(e^{-\frac{(B c \pi)^{2}}{2 \sigma^{2}}}\right)^{2} \\
& \gtrsim \frac{1}{B}\left(e^{-\frac{(B c \pi)^{2}}{2 \sigma^{2}}}\right)^{2} \\
& \gtrsim \frac{1}{B}\left(e^{-c_{0}}\right)^{2}
\end{aligned}
$$

The last inequality follows by there exists some universal constant $c_{0}>0$ such that $-\frac{1}{\log (k / \delta)} \gtrsim$ $-c_{0}$. Thus, we show $\sum_{i=1}^{M} G_{i}^{2} \gtrsim \frac{1}{B}$.


Figure 11.1: Filter $\widehat{G^{\prime}}(f)$

To analyze the details of our algorithm, we explain some lower-level definitions and claims first. Here we give the definition of three notations that related to hash function.

Definition 11.4.3. $\pi_{\sigma, b}(f)=2 \pi \sigma(f-b)(\bmod 2 \pi)$. We maps frequency to a circle $[0,2 \pi)$, since our observation of sample is the phase of some complex number, which also belongs to $[0,2 \pi)$.

Definition 11.4.4. $h_{\sigma, b}(f)=\operatorname{round}\left(\pi_{\sigma, b}(f) \cdot \frac{B}{2 \pi}\right) . h_{\sigma, b}(f)$ is a "hash function" that hashes frequency $f$ into one of the $B$ bins. The motivation is, it is very likely that each bin only has 1 heavy hitters if we choose large enough $B$. Then, for each bin, we can run a 1 -sparse algorithm to recover the frequency.

Definition 11.4.5. $o_{\sigma, b}(f)=\pi_{\sigma, b}(f)-\frac{2 \pi}{B} \cdot h_{\sigma, b}(f)$. Offset $o_{\sigma, b}(f)$ denotes the distance from $\pi_{\sigma, b}(f)$ to the center of the corresponding bin that frequency $f$ was hashed into.

Then we define some events that might happen after applying hash function to the entire frequency domain.

Definition 11.4.6. "Collision" event $E_{\text {coll }}(f)$ : holds iff $h_{\sigma, b}(f) \in h_{\sigma, b}\left(\operatorname{supp}\left(\widehat{x^{*}}\right) \backslash\{f\}\right)$. The "collision" event happening means there exists some other frequency $f^{\prime}$ such that both $f$ and $f^{\prime}$ are hashed into the same bin. Once two frequencies are colliding in one bin, the algorithm will not be able to recover them.

Definition 11.4.7. "Large offset" event $E_{\text {off }}(f)$ : holds iff $\left|o_{\sigma, b}(f)\right| \geq(1-\alpha) \frac{2 \pi}{2 B}$. The event holds if frequency $f$ is not within factor $1-\alpha$ of the radius close to the center of that hash bin. It causes the frequency to be in the intermediate regime of filter and not recoverable, see Figure 11.1.

Definition 11.4.8. We sample $\sigma$ uniformly at random from $\left[\frac{1}{B \eta}, \frac{2}{B \eta}\right]$. Conditioning on $\sigma$ is chosen first, we sample $b$ uniformly at random from $\left[0, \frac{[F / \eta]}{B \sigma}\right]$. Then we sample $\gamma$ uniformly at random from $\left[\frac{1}{2}, 1\right]$ and $\beta$ uniformly at random from $[\widehat{\beta}, 2 \widehat{\beta}]$, where $\widehat{\beta}$ is dynamically changing during our algorithm(The details of setting $\widehat{\beta}$ are explained in Lemma 11.3.6). For $P_{\sigma, \gamma, b}$ and $P_{\sigma, \gamma+\beta, b}$, we take the following two sets of samples over time domain,

$$
\begin{aligned}
& x(\sigma(1-\gamma)), x(\sigma(2-\gamma)), x(\sigma(3-\gamma)), \cdots, x(\sigma(B \log (k / \delta)-\gamma)) \\
& x(\sigma(1-\gamma-\beta)), x(\sigma(2-\gamma-\beta)), x(\sigma(3-\gamma-\beta)), \cdots, x(\sigma(B \log (k / \delta)-\gamma-\beta))
\end{aligned}
$$

Conditioning on drawing $\sigma, b$ from some distribution, we are able to show that the probability of "Collision" and "Large offset" event holding are small.

Lemma 11.4.3. For any $\widetilde{T}$, and $0 \leq \widetilde{\epsilon}, \widetilde{\delta} \leq \widetilde{T}$, if we sample $\widetilde{\sigma}$ uniformly at random from $[A, 2 A]$, then

$$
\begin{equation*}
\frac{2 \widetilde{\epsilon}}{\widetilde{T}}-\frac{2 \widetilde{\epsilon}}{A} \leq \operatorname{Pr}[\widetilde{\sigma} \quad(\bmod \widetilde{T}) \in[\widetilde{\delta}-\widetilde{\epsilon}, \widetilde{\delta}+\widetilde{\epsilon}]] \leq \frac{2 \widetilde{\epsilon}}{\widetilde{T}}+\frac{4 \widetilde{\epsilon}}{A} \tag{11.9}
\end{equation*}
$$

Proof. Since we sample $\widetilde{\sigma}$ uniformly at random from $[A, 2 A]$, let $I$ denote a set of candidate integers, then the smallest one is $\lfloor A\rfloor$ and the largest one is $\lceil 2 A\rceil$. Thus, the original probability equation is equivalent to

$$
\begin{equation*}
\operatorname{Pr}[\widetilde{\sigma} \in[s \cdot \widetilde{T}+\widetilde{\delta}-\widetilde{\epsilon}, s \cdot \widetilde{T}+\widetilde{\delta}+\widetilde{\epsilon}] \exists s \in I] \tag{11.10}
\end{equation*}
$$

where $I=\{\lfloor A / \widetilde{T}\rfloor, \cdots,\lceil 2 A / \widetilde{T}\rceil\}$.
Consider any $s \in I$, the probability of $\widetilde{\sigma}$ belonging to the interval $[s \cdot \widetilde{T}+\widetilde{\delta}-\widetilde{\epsilon}, s \cdot \widetilde{T}+\widetilde{\delta}+\widetilde{\epsilon}]$ is

$$
\operatorname{Pr}[\widetilde{\sigma} \in[s \cdot \widetilde{T}+\widetilde{\delta}-\widetilde{\epsilon}, s \cdot \widetilde{T}+\widetilde{\delta}+\widetilde{\epsilon}]]=\frac{2 \widetilde{\epsilon}}{A}
$$

Taking the summation over all $s \in I$, we obtain

$$
\begin{aligned}
& \operatorname{Pr}[\widetilde{\sigma} \in[s \cdot \widetilde{T}+\widetilde{\delta}-\widetilde{\epsilon}, s \cdot \widetilde{T}+\widetilde{\delta}+\widetilde{\epsilon}] \exists s \in I] \\
= & \sum_{s \in I} \operatorname{Pr}[\widetilde{\sigma} \in[s \cdot \widetilde{T}+\widetilde{\delta}-\widetilde{\epsilon}, s \cdot \widetilde{T}+\widetilde{\delta}+\widetilde{\epsilon}]] \\
= & \sum_{s \in I} \frac{2 \widetilde{\epsilon}}{A} \\
= & \frac{2 \widetilde{\epsilon}|I|}{A}
\end{aligned}
$$

It remains to bound $\frac{2 \widetilde{\mid} I \mid}{A}$. Since $|I|=\lceil 2 A / \widetilde{T}\rceil-\lfloor A / \widetilde{T}\rfloor+1$, then we have an upper bound for $I$,

$$
|I| \leq A / \widetilde{T}+2
$$

On the other side, we have an lower bound,

$$
|I| \geq A / \widetilde{T}-1
$$

Plugging upper bound of $I$ into $\frac{2 \widetilde{\epsilon}|I|}{A}$,

$$
\frac{2 \widetilde{\epsilon}|I|}{A} \leq \frac{2 \widetilde{\epsilon}}{A}(A / \widetilde{T}+2)=\frac{2 \widetilde{\epsilon}}{\widetilde{T}}+\frac{4 \widetilde{\epsilon}}{A}
$$

Using the lower bound of $|I|$, we have

$$
\frac{2 \widetilde{\epsilon}|I|}{A} \geq \frac{2 \widetilde{\epsilon}}{A}(A / \widetilde{T}-1)=\frac{2 \widetilde{\epsilon}}{\widetilde{T}}-\frac{2 \widetilde{\epsilon}}{A}
$$

Thus, we complete the proof.

The following corollary will be used many times in this paper. The proof directly follows by Lemma 11.4.3.

Corollary 11.4.4. For any $\widetilde{T}, \Delta f$, and $0 \leq \widetilde{\epsilon}, \widetilde{\delta} \leq \widetilde{T}$, if we sample $\widetilde{\sigma}$ uniformly at random from $[A, 2 A]$, then

$$
\begin{equation*}
\frac{2 \widetilde{\epsilon}}{\widetilde{T}}-\frac{2 \widetilde{\epsilon}}{A \Delta f} \leq \operatorname{Pr}[\widetilde{\sigma} \Delta f \quad(\bmod \widetilde{T}) \in[\widetilde{\delta}-\widetilde{\epsilon}, \widetilde{\delta}+\widetilde{\epsilon}]] \leq \frac{2 \widetilde{\epsilon}}{\widetilde{T}}+\frac{4 \widetilde{\epsilon}}{A \Delta f} \tag{11.11}
\end{equation*}
$$

Proof. Since $\widetilde{\sigma}$ is sampled uniformly at random from $[A, 2 A]$, then $\widetilde{\sigma} \Delta f$ is sampled uniformly at random from $[A \Delta f, 2 A \Delta f]$. Now applying Lemma 11.4.3 by only replacing $A \Delta f$ by $A$.

Claim 11.4.5. Let $\sigma$ be sampled uniformly at random from $\left[\frac{1}{B \eta}, \frac{2}{B \eta}\right]$ and $\min _{i \neq j}\left|f_{i}-f_{j}\right|>\eta$. $\forall i, j \in[k]$, if $i \neq j$, then $\operatorname{Pr}\left[h_{\sigma, b}\left(f_{i}\right)=h_{\sigma, b}\left(f_{j}\right)\right] \lesssim \frac{1}{B}$.

Proof. To simplify the proof, define $\Delta f=\left|f_{i}-f_{j}\right|$. We consider two cases: (I) $\eta \leq\left|f_{i}-f_{j}\right|<$ $\frac{(B-1) \eta}{2}$, (II) $\frac{(B-1) \eta}{2} \leq\left|f_{i}-f_{j}\right|$.
(I) If $\Delta f=\eta$, then $2 \pi \sigma \Delta f$ is at least $\frac{2 \pi}{\eta B} \cdot \eta=\frac{2 \pi}{B}$, which means two frequencies have to go to different bins after hashing. If $\Delta f=\frac{(B-1) \eta}{2}$, then $2 \pi \sigma \Delta f$ is at most $\frac{4 \pi}{B \eta} \cdot \frac{(B-1) \eta}{2}=$
$(1-1 / B) 2 \pi$. In order to make two frequencies collide, $2 \pi \sigma \Delta f$ should belong to [(1$1 / B) 2 \pi,(1+1 / B) 2 \pi)$. Since for any $\Delta f \in\left[\eta, \frac{(B-1) \eta}{2}\right)$, we have $2 \pi \sigma \Delta f \in\left[\frac{1}{B} 2 \pi,(1-1 / B) 2 \pi\right)$, which does not intersect interval $[(1-1 / B) 2 \pi,(1+1 / B) 2 \pi)$. Thus,

$$
\underset{\sigma, b}{\operatorname{Pr}}\left[h_{\sigma, b}\left(f_{i}\right)=h_{\sigma, b}\left(f_{j}\right)\right]=0 .
$$

(II) We apply Corollary 11.4 .4 by setting $\widetilde{T}=2 \pi, \widetilde{\sigma}=2 \pi \sigma, \widetilde{\delta}=0, \widetilde{\epsilon}=\frac{2 \pi}{2 B}, A=2 \pi \frac{1}{B \eta}$. Then we have

$$
\underset{\sigma, b}{\operatorname{Pr}}\left[h_{\sigma, b}\left(f_{i}\right)=h_{\sigma, b}\left(f_{j}\right)\right]=\underset{\sigma, b}{\operatorname{Pr}}\left[2 \pi \sigma \Delta f \in\left[s \cdot 2 \pi-\frac{2 \pi}{2 B}, s \cdot 2 \pi+\frac{2 \pi}{2 B}\right] \exists s \in I\right]
$$

where

$$
I=\left\{\left\lfloor\frac{1}{B \eta} \Delta f\right\rfloor, \cdots,\left\lceil\frac{2}{B \eta} \Delta f\right\rceil\right\}
$$

By upper bound of Corollary 11.4.4, Equation (11.12) is at most

$$
\frac{1}{\frac{2 \pi}{B \eta} \Delta f} \cdot \frac{2 \pi}{B} \cdot\left(\frac{1}{B \eta} \Delta f+2\right)=\frac{1}{B}+\frac{2 \eta}{\Delta f} \leq \frac{1}{B}+\frac{4}{B-1} \lesssim \frac{1}{B},
$$

where the first inequality follows by $\frac{(B-1) \eta}{2} \leq\left|f_{i}-f_{j}\right|$ which is the assumption of part (II).

Claim 11.4.6. $\forall f, \forall 0<\alpha<1, \underset{\sigma, b}{\operatorname{Pr}}\left[\left|o_{\sigma, b}(f)\right| \leq(1-\alpha) \frac{2 \pi}{2 B}\right] \geq 1-O(\alpha)$.

Proof. Since we draw $\sigma$ uniformly at random from $\left[\frac{1}{B \eta}, \frac{2}{B \eta}\right]$, then $2 \pi \sigma(f-b) \in\left[\frac{2 \pi}{B \eta}(f-\right.$ $\left.b), \frac{4 \pi}{B \eta}(f-b)\right]$ uniformly at random. The probability is equal to

$$
\underset{\sigma, b}{\operatorname{Pr}}\left[2 \pi \sigma(f-b) \in\left[s \cdot \frac{2 \pi}{B}-(1-\alpha) \frac{2 \pi}{2 B}, s \cdot \frac{2 \pi}{B}+(1-\alpha) \frac{2 \pi}{2 B}\right] \exists s \in I\right],
$$

where

$$
I=\left\{\left\lfloor\frac{2 \pi}{B \eta}(f-b) \frac{B}{2 \pi}-\frac{1-\alpha}{2}\right\rfloor, \cdots,\left\lceil\frac{4 \pi}{B \eta}(f-b) \frac{B}{2 \pi}+\frac{1-\alpha}{2}\right\rceil\right\} .
$$

We apply Corollary 11.4 .4 by setting $\widetilde{T}=\frac{2 \pi}{B}, \widetilde{\sigma}=2 \pi \sigma, \widetilde{\delta}=0, \widetilde{\epsilon}=(1-\alpha) \frac{2 \pi}{2 B}, A=2 \pi \frac{1}{B \eta}$, $\Delta f=|f-b|$.

By lower bound of Corollary 11.4.4, we have

$$
\operatorname{Pr}\left[\left|o_{\sigma, b}(f)\right| \leq(1-\alpha) \frac{2 \pi}{2 B}\right] \geq(1-\alpha)-(1-\alpha) \cdot \frac{\eta}{|f-b|}
$$

Since $\alpha \cdot \frac{\eta}{|f-b|}>0$, then

$$
\operatorname{Pr}\left[\left|o_{\sigma, b}(f)\right| \leq(1-\alpha) \frac{2 \pi}{2 B}\right] \geq(1-\alpha)-\frac{\eta}{|f-b|}
$$

Recall that we sample $\sigma$ uniformly at random from $\left[\frac{1}{B \eta}, \frac{2}{B \eta}\right]$ and sample $b$ uniformly at random from $\left[0, \frac{[F / \eta\rceil}{B \sigma}\right]$. Since $f \in[-F, F]$ and $b$ is uniformly chosen from range $\left(0, \frac{[F / \eta\rceil}{B \sigma}\right]$, thus for any $C>0, \underset{b}{\operatorname{Pr}}[|f-b| \leq C \eta] \lesssim \frac{C \eta}{F}$. Replacing $C$ by $1 / \alpha$, we have $\operatorname{Pr}\left[\frac{\eta}{|f-b|} \leq \alpha\right] \geq$ $1-\Theta\left(\frac{\eta}{\alpha F}\right)$. Compared to $\frac{\eta}{F}, \alpha$ is just a constant. Thus, we finish the proof.

### 11.5 Proofs of basic hashing-related lemmas

Lemma 11.3.1. Let $(v, f)$ and $\left(v^{\prime}, f^{\prime}\right)$ denote any two tones, i.e., (magnitude, frequency) pairs. Then for

$$
\operatorname{dist}\left((v, f),\left(v^{\prime}, f^{\prime}\right)\right)^{2}:=\frac{1}{T} \int_{0}^{T}\left|v e^{2 \pi f t \mathrm{i}}-v^{\prime} e^{2 \pi f^{\prime} t i}\right|^{2} \mathrm{~d} t
$$

we have

$$
\operatorname{dist}\left((v, f),\left(v^{\prime}, f^{\prime}\right)\right)^{2} \approx\left(|v|^{2}+\left|v^{\prime}\right|^{2}\right) \cdot \min \left(1, T^{2}\left|f-f^{\prime}\right|^{2}\right)+\left|v-v^{\prime}\right|^{2}
$$

and

$$
\operatorname{dist}\left((v, f),\left(v^{\prime}, f^{\prime}\right)\right) \bar{\sim}|v| \cdot \min \left(1, T\left|f-f^{\prime}\right|\right)+\left|v-v^{\prime}\right| .
$$

Proof. Define $\nu=\left|f-f^{\prime}\right|$. First, let's show the first upper bound. We have that

$$
\begin{aligned}
\text { LHS } & =\frac{1}{T} \int_{0}^{T}\left|v e^{2 \pi f t i}-v^{\prime} e^{2 \pi f^{\prime} t \mathrm{i}}\right|^{2} \mathrm{~d} t \\
& \leq 2 \cdot \frac{1}{T} \int_{0}^{T}\left|v e^{2 \pi f t \mathrm{i}}-v e^{2 \pi f^{\prime} t \mathrm{t}}\right|^{2}+\left|v e^{2 \pi f^{\prime} t \mathrm{i}}-v^{\prime} e^{2 \pi f^{\prime} t \mathrm{i}}\right|^{2} \mathrm{~d} t \\
& =2|v|^{2} \cdot \frac{1}{T} \int_{0}^{T}\left|e^{2 \pi \nu t \mathrm{i}}-1\right|^{2} \mathrm{~d} t+2\left|v-v^{\prime}\right|^{2} \\
& \leq 2|v|^{2} \cdot \frac{1}{T} \int_{0}^{T} \min (2,2 \pi \nu t)^{2} \mathrm{~d} t+2\left|v-v^{\prime}\right|^{2} \\
& \leq 2|v|^{2} \cdot \min \left(4, \frac{4 \pi^{2}}{3} \nu^{2} T^{2}\right)+2\left|v-v^{\prime}\right|^{2},
\end{aligned}
$$

as desired.
Now consider the lower bound. First we show this in the setting where $|v|=\left|v^{\prime}\right|$.

Suppose $v^{\prime}=v e^{-\theta \mathbf{i}}$. Then we want to bound

$$
\begin{aligned}
\mathrm{LHS} & =\frac{1}{T} \int_{0}^{T}\left|v e^{2 \pi f t \mathbf{i}}-v e^{\left(2 \pi f^{\prime} t-\theta\right) \mathbf{i}}\right|^{2} \mathrm{~d} t \\
& =|v|^{2} \frac{1}{T} \int_{0}^{T}\left|e^{(2 \pi \nu t+\theta) \mathbf{i}}-1\right|^{2} \mathrm{~d} t,
\end{aligned}
$$

as being at least $\Omega\left(|v|^{2}\left(\min \left(1, \nu^{2} T^{2}\right)+\theta^{2}\right)\right)$. In the case that $\nu T<1 / 10$, then

$$
\left|e^{(2 \pi \nu t+\theta) \mathbf{i}}-1\right| \gtrsim|2 \pi \nu t+\theta|,
$$

and

$$
\underset{t}{\mathbb{E}}\left[(2 \pi \nu t+\theta)^{2}\right] \geq\left(\theta-\frac{2 \pi \nu T}{2}\right)^{2}+\mathbb{E}\left[(2 \pi \nu(t-T / 2))^{2}\right] \gtrsim \nu^{2} T^{2}+\left(\theta-\frac{2 \pi \nu T}{2}\right)^{2} \gtrsim \nu^{2} T^{2}+\theta^{2}
$$

On the other hand, if $\nu T>1 / 10$, then $2 \pi \nu t-\theta$ is $\Omega(1)$ for at least a constant fraction of the $t$, giving that

$$
\left|e^{(2 \pi \nu t+\theta) \mathbf{i}}-1\right| \gtrsim 1
$$

Hence the lower bound holds whenever $|v|=\left|v^{\prime}\right|$.
Finally, consider the lower bound for $|v| \neq\left|v^{\prime}\right|$. Without loss of generality assume $\left|v^{\prime}\right| \geq|v|$, and define $v^{*}=\frac{|v|}{\left|v^{\prime}\right|} v^{\prime \prime}$. For any two angles $\theta, \theta^{\prime}$ we have that

$$
\left|v e^{\theta \mathbf{i}}-v^{\prime} e^{\theta^{\prime} \mathbf{i}}\right|^{2} \geq\left|v e^{\theta \mathbf{i}}-v^{*} e^{\theta^{\prime} \mathbf{i}}\right|^{2}+\left|v^{*}-v^{\prime}\right|^{2}
$$

because the angle $\angle v v^{*} v^{\prime}$ is obtuse. Therefore

$$
\begin{aligned}
\text { LHS } & \geq \frac{1}{T} \int_{0}^{T}\left|v e^{2 \pi f t \mathbf{i}}-v^{*} e^{2 \pi f^{\prime} \mathbf{t}}\right|^{2}+\left|v^{*}-v^{\prime}\right|^{2} \mathrm{~d} t \\
& \gtrsim|v|^{2} \min \left(1, \nu^{2} T^{2}\right)+\left|v-v^{*}\right|^{2}+\left|v^{*}-v^{\prime}\right|^{2} \\
& \gtrsim|v|^{2} \min \left(1, \nu^{2} T^{2}\right)+\left|v-v^{\prime}\right|^{2} .
\end{aligned}
$$

Now, if $\left|v^{\prime}\right|^{2} \leq 2|v|^{2}$, this gives the desired bound. But otherwise, it also gives the desired bound because $\left|v-v^{\prime}\right|^{2} \gtrsim\left|v^{\prime}\right|^{2}$. So we get the bound in all settings.

The second equation follows from the first, using that $(a+b)^{2} \sim a^{2}+b^{2}$ for positive $a, b$ to show

$$
\operatorname{dist}\left((v, f),\left(v^{\prime}, f^{\prime}\right)\right) \approx\left(|v|+\left|v^{\prime}\right|\right) \cdot \min \left(1, T\left|f-f^{\prime}\right|\right)+\left|v-v^{\prime}\right|
$$

We can then replace $|v|+\left|v^{\prime}\right|$ with $|v|$ because either they are equivalent up to constants or $\left|v-v^{\prime}\right|$ is within a constant factor of $|v|+\left|v^{\prime}\right|$.

Lemma 11.3.3. If $x^{*}(t)=0, \forall t \in[0, T]$, then

$$
\underset{\sigma, a, b}{\mathbb{E}}\left[\sum_{j=1}^{B}\left|\widehat{u}_{j}\right|^{2}\right] \lesssim \frac{1}{T} \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t
$$

Proof. Since $\widehat{u}_{j}=\operatorname{FFT}\left(u_{j}\right)$, then $\sum_{j=1}^{B}\left|\widehat{u}_{j}\right|^{2}=B \sum_{j=1}^{B}\left|u_{j}\right|^{2}$. Recall that $\left(P_{\sigma, a, b} x\right)(t)=$ $x(\sigma(t-a)) e^{2 \pi \sigma b t \mathbf{i}}, u_{j}=\sum_{i=1}^{\log (k / \delta)} y_{j+B i}$ and $y_{j}=G_{j} \cdot\left(P_{\sigma, a, b} g\right)_{j}=G_{j} \cdot g(\sigma(j-a)) e^{2 \pi \mathbf{i} \sigma b j}$. Then

$$
\begin{aligned}
\underset{\sigma, a, b}{\mathbb{E}}\left[\sum_{j=1}^{B}\left|u_{j}\right|^{2}\right] & =\underset{\sigma, a, b}{\mathbb{E}}\left[\sum_{j=1}^{B}\left|\sum_{i=1}^{\log (k / \delta)} y_{j+B i}\right|^{2}\right] \\
& =\underset{\sigma, a}{\mathbb{E}}\left[\sum_{j=1}^{B} \underset{b}{\mathbb{E}}\left[\left.\sum_{i=1}^{\log (k / \delta)} y_{j+B i}\right|^{2}\right]\right. \\
& =\underset{\sigma, a}{\mathbb{E}}\left[\sum_{j=1}^{B} \underset{b}{\mathbb{E}}\left[\left(\sum_{i=1}^{\log (k / \delta)} y_{j+B i}\right)\left(\sum_{i^{\prime}=1}^{\log (k / \delta)} \overline{y_{j+B i^{\prime}}}\right)\right]\right] \\
& =\underset{\sigma, a}{\mathbb{E}}\left[\sum_{j=1}^{B} \underset{b}{\mathbb{E}}\left[\sum_{i=1}^{\log (k / \delta)} y_{j+B i} \overline{y_{j+B i}}+\sum_{i \neq i^{\prime}}^{\log (k / \delta)} y_{j+B i} \overline{y_{j+B i^{\prime}}}\right]\right] .
\end{aligned}
$$

For any $(i, j) \in[\log (k / \delta)] \times[B]$, let $S_{i, j}=G_{j+B i} g(\sigma(j+B i-a))=y_{j+B i} e^{-2 \pi \mathrm{i} \sigma b(j+B i)}$, then

$$
\underset{\sigma, a, b}{\mathbb{E}}\left[\sum_{j=1}^{B}\left|u_{j}\right|^{2}\right]=\underset{\sigma, a}{\mathbb{E}}[\sum_{j=1}^{B} \underset{b}{\mathbb{E}}[\underbrace{\sum_{i=1}^{\log (k / \delta)}\left|S_{i, j}\right|^{2}}_{C_{1}}+\underbrace{\sum_{i \neq i^{\prime}}^{\log (k / \delta)} S_{i, j} \overline{S_{i^{\prime}, j}} e^{2 \pi \mathbf{i} \sigma b B\left(i-i^{\prime}\right)}}_{C_{2}}]] .
$$

Consider the expectation of $C_{2}$ :

$$
\begin{array}{rlr}
\underset{b}{\mathbb{E}}\left[C_{2}\right] & =\underset{b}{\mathbb{E}}\left[\sum_{i \neq i^{\prime}}^{\log (k / \delta)} S_{i, j} \overline{S_{i^{\prime}, j}} e^{2 \pi \mathbf{i} \sigma b B\left(i-i^{\prime}\right)}\right] \\
& =\sum_{i \neq i^{\prime}}^{\log (k / \delta)} S_{i, j} \overline{S_{i^{\prime}, j}} \frac{\mathbb{E}}{b}\left[e^{2 \pi \mathbf{i} \sigma b B\left(i-i^{\prime}\right)}\right] & \\
& =0 & \text { by Definition 11.4.8 }
\end{array}
$$

Note that term $C_{1}$ is independent of $b$ which means $\underset{b}{\mathbb{E}} C_{1}=C_{1}$. Thus, we can remove the expectation over $b$. Then,

$$
\begin{aligned}
\underset{\sigma, a, b}{\mathbb{E}}\left[\sum_{j=1}^{B}\left|u_{j}\right|^{2}\right] & =\underset{\sigma, a}{\mathbb{E}}\left[\sum_{j=1}^{B} \sum_{i=1}^{\log (k / \delta)}\left|G_{j+B i}\right|^{2} \cdot|g(\sigma(j+B i-a))|^{2}\right] \\
& =\underset{\sigma, a}{\mathbb{E}}\left[\sum_{i=1}^{B \log (k / \delta)}\left|G_{i}\right|^{2} \cdot|g(\sigma(i-a))|^{2}\right]
\end{aligned}
$$

Now, the idea is to replace the expectation term $\mathbb{E}_{a}$ by an integral term $\int_{a \in A}(*) \mathrm{d} a$. Then, replace it by another integral term $\int_{0}^{T}(\star) \mathrm{d} t$. Let $A$ denote a set of intervals that we will sample $a$ from. It is easy to verify that $|A| \lesssim T / \sigma$, since $(\sigma(i-a))$ is sampled from $[0, T]$. If we choose $T$ to be a constant factor larger than $\sigma|\operatorname{supp}(G)|$, then we also have $|A| \gtrsim T / \sigma$.

$$
\begin{aligned}
\underset{\sigma}{\mathbb{E}}\left[\underset{a}{\mathbb{E}}\left[\sum_{i=1}^{B \log (k / \delta)}\left|G_{i}\right|^{2} \cdot|g(\sigma(i-a))|^{2}\right]\right] & =\underset{\sigma}{\mathbb{E}}\left[\frac{1}{|A|} \int_{a \in A} \sum_{i=1}^{B \log (k / \delta)}\left|G_{i}\right|^{2} \cdot|g(\sigma(i-a))|^{2} \mathrm{~d} a\right] \\
& =\underset{\sigma}{\mathbb{E}}\left[\sum_{i=1}^{B \log (k / \delta)}\left|G_{i}\right|^{2} \cdot \frac{1}{|A|} \int_{a \in A}|g(\sigma(i-a))|^{2} \mathrm{~d} a\right] \\
& =\underset{\sigma}{\mathbb{E}}\left[\sum_{i=1}^{B \log (k / \delta)}\left|G_{i}\right|^{2} \cdot \frac{1}{\sigma|A|} \int_{a \in A}|g(\sigma(i-a))|^{2} \mathrm{~d} \sigma a\right] \\
& \lesssim \underset{\sigma}{\mathbb{E}}\left[\|G\|_{2}^{2} \cdot \frac{1}{T} \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t\right] .
\end{aligned}
$$

By Claim 11.4.2, we know that $\|G\|_{2}^{2} \approx \frac{1}{B}$. Combining $\sum_{j=1}^{B}\left|\widehat{u}_{j}\right|^{2}=B \sum_{j=1}^{B}\left|u_{j}\right|^{2}$ and $\|G\|_{2}^{2} \approx \frac{1}{B}$ gives the desired result.

Lemma 11.3.4. If $g(t)=0, \forall t \in[0, T]$. Let $H$ denote a set of frequencies, $H=\{f \in$ $\operatorname{supp}\left(\widehat{x}^{*}\right) \mid$ neither $E_{\text {coll }}(f)$ nor $E_{\text {off }}(f)$ holds $\}$. Then,

$$
\underset{\sigma, a, b}{\mathbb{E}}\left[\sum_{f \in H}\left|\widehat{u}_{h_{\sigma, b}(f)}-\widehat{x^{*}}(f) e^{a \sigma 2 \pi f \mathbf{i}}\right|^{2}\right] \leq \delta\left\|\widehat{x^{*}}\right\|_{1}^{2}
$$

Proof. For simplicity, let $G=G_{B, \delta, \alpha}$ and $\widehat{G^{\prime}}=\widehat{G^{\prime}}{ }_{B, \delta, \alpha}$. we have

$$
\begin{aligned}
\widehat{y} & =G \widehat{G \cdot \widehat{P_{\sigma, a, b}} x} \\
& =\widehat{G} * \widehat{P_{\sigma, a, b} x} \\
& =\widehat{G^{\prime}} * \widehat{P_{\sigma, a, b} x}+\left(\widehat{G}-\widehat{G^{\prime}}\right) * \widehat{P_{\sigma, a, b} x}
\end{aligned}
$$

The $\ell_{\infty}$ norm of second term can be bounded :

$$
\left\|\left(\widehat{G}-\widehat{G^{\prime}}\right) * \widehat{P_{\sigma, a, b} x}\right\|_{\infty} \leq\left\|\widehat{G}-\widehat{G^{\prime}}\right\|_{\infty}\left\|\widehat{P_{\sigma, a, b} x}\right\|_{1} \leq \sqrt{\delta / k}\left\|\widehat{x^{*}}\right\|_{1}
$$

Thus, consider the $j$ th term of $\widehat{u}$,

$$
\begin{aligned}
\widehat{u}_{j} & =\widehat{y}_{j F / B} \\
& =\sum_{|l|<F /(2 B)} \widehat{G^{\prime}}{ }_{-l}\left(\widehat{P_{\sigma, a, b} x}\right)_{j F / B+l} \pm \sqrt{\delta / k}\left\|\widehat{x^{*}}\right\|_{1} \\
& =\sum_{\left|\pi_{\sigma, b}(f)-j F / B\right|<F /(2 B)} \widehat{G_{j}^{\prime}}{ }_{j F / B-\pi_{\sigma, b}(f)} \widehat{P_{\sigma, a, b} x} \pi_{\pi_{\sigma, b}(f)} \pm \sqrt{\delta / k}\left\|\widehat{x^{*}}\right\|_{1} \\
& =\sum_{h_{\sigma, b}(f)=j} \widehat{G^{\prime}}{ }_{-o_{\sigma, b}(f)} \widehat{x^{*}}(f) e^{2 \pi f \sigma a \mathbf{i}} \pm \sqrt{\delta / k}\left\|\widehat{x^{*}}\right\|_{1} .
\end{aligned}
$$

If neither $E_{\text {coll }}(f)$ nor $E_{\text {off }}(f)$ happens, then we know that frequency $f$ is the only heavy hitter hashed into bin $j$ and $\widehat{G^{\prime}}-o_{\sigma, b}(f)=1$ for frequency $f$. Thus,

$$
\underset{\sigma, a, b}{\mathbb{E}}\left[\left|\widehat{u}_{h_{\sigma, b}(f)}-\widehat{x^{*}}(f) e^{a \sigma 2 \pi f \mathbf{i}}\right|^{2}\right] \leq \delta / k\left\|\widehat{x^{*}}\right\|_{1}^{2} .
$$

Since the above equation holds for all $f \in H$, we get

$$
\sum_{f \in H} \underset{\sigma, a, b}{\mathbb{E}}\left[\left|\widehat{u}_{h_{\sigma, b}(f)}-\widehat{x^{*}}(f) e^{a \sigma 2 \pi f \mathbf{i}}\right|^{2}\right] \leq k \delta / k\left\|\widehat{x^{*}}\right\|_{1}^{2}=\delta\left\|\widehat{x^{*}}\right\|_{1}^{2}
$$

### 11.6 Proofs for one stage of recovery

Binary search of one-sparse algorithm We first explain a simple, clean, but not optimal one-sparse algorithm, then we try to optimize the algorithm step by step. Let "heavy" frequency $f \in[-F, F]$, we can split the frequency interval into two regions: left region $[-F, 0)$ and right region $[0, F]$. We can observe $\theta \equiv 2 \pi \beta f(\bmod 2 \pi)$ by checking the phase difference between using $P_{1, a, b}$ and $P_{1, a+\beta, b}$, where $\beta$ is uniformly at random sampled from some suitable range $[\widehat{\beta}, 2 \widehat{\beta}]$. For each observation $\theta$, we can guess $m$ different possibilities for $f$, say $\theta_{1}, \theta_{2}, \cdots, \theta_{m}$, if $\theta_{i}$ belong to the left region, we add a vote to that, otherwise we add a vote to the right region. After taking enough samples, we choose the region that has the largest vote. This decision will let us narrow down the searching range of the true frequency by half with some good probability. Suppose we decide to choose the right region $[0, F]$, then we can just repeat the above binary search over $[0, F]$ again to get into a region that has size $F / 2$. Repeating it $D$ times, we can learn the frequency with a region that has size at most $2 F / 2^{D}$. But the binary search is not the best approach, actually, we can do much better with using $t$-ary search.
$k$-sparse To locate those $k$ heavy signals in the frequency domain, we need to consider the "bins" computed by HashToBins with $P_{\sigma, a, b}$. One of the main difference from previous work [HIKP12a] is, instead of permuting the discrete coordinates according to $P_{\sigma, a, b}$ and partitioning the coordinates into $B=O(k)$ bins, we permute the continuous frequency domain and partition the frequency domain into $B=O(k)$ bins. With a large constant probability, we can obtain for each heavy signal $f$ that neither $E_{\text {coll }}$ nor $E_{\text {off }}$ happens. After splitting those $k$ frequencies into different bins, then we can run the one-sparse algorithm
for all the bins simultaneously.
$t$-ary search To argue the final succeed probability of our algorithm, we need to take the union bound for each array/region in each round and also take the union bound over each round. There are several benefits of changing binary search to $t$-ary search, (1) to reach the same accuracy, the number of rounds $D$ for $t$-ary search is smaller than the number of rounds for binary search; (2) our searching procedure is a "noisy" searching problem, for the noiseless version of the searching problem, we do not need to take care of the union bound argument. Having a parameter for the number of arrays/regions is important to optimize the entire procedure.

Recall the traditional binary search problem(noiseless version), given a list of sorted numbers $a[1,2 \cdots, n]$ in increasing order. We want to determine if some number $x$ belongs to $a[1,2, \cdots, n]$. When we compare some $a[i]$ with $x$, we will know the true answer.

But the noisy binary search problem is slightly harder. $a[1,2, \cdots, n]$ is still a list of sorted numbers in increasing order, we want to determine if some number $x$ belongs to $a[1,2, \cdots, n]$. But when we compare the $a[i]$ with $x$, we will know the true answer with $9 / 10$ probability, and get the false answer with $1 / 10$ probability. In this case, we can not finish the task by following the procedure of traditional binary search algorithm, e.g. making the decision by just comparing $a[i]$ with $x$ once. The reason is after taking the union bound of $\log n$ rounds, the failure probability can be arbitrarily large. One idea to fix this issue is independently comparing $a[i]$ with $x$ multiple times, e.g. $\log \log n$ times. Then we can amplify the succeed probability of each round from $9 / 10$ to $1-\frac{1}{\operatorname{poly}(\log n)}$, after taking the union bound over $\log n$ rounds, we still have $1-\frac{1}{\operatorname{poly}(\log n)}$ succeed probability.

Our problem is still more complicated than the above noisy binary search problem. In each round of algorithm LOCATEINNER, we do a $t$-array search by splitting the candidate frequency region( that has length $\Delta l$ ) into $t$ consecutive regions, $Q_{1}, Q_{2}, \cdots, Q_{t}$, each of them has the equal size $\frac{\Delta l}{t}$. By using the hash values of $P_{\sigma, a, b}$ and $P_{\sigma, a+\beta, b}$ (Line 26-27 in Algorithm 11.2), we can have an observation over $[0,2 \pi)$. For each such observation over $[0,2 \pi)$, it was in fact scaled by $2 \pi \sigma \beta$ and rounded over $[0,2 \pi)$. Thus the corresponding frequency location of this observation might belong to $m=\Theta(\sigma \beta \Delta l)$ different possible regions. Since we do not know which one, we just add a vote to all of the possible regions. To understand $t$-ary search, let's consider this example. Suppose we split the frequency into 20 regions, the true frequency belongs to region 9 and each observation is correct with probability $4 / 5$. For each observation, we will add a vote to a batch of roughly evenly spaced regions.

1. For the observation 1 , we add a vote to region $1,5,9,13,17$.
2. For the observation 2, we add a vote to region 3, 9, 15 .
3. For the observation 3, we add a vote to region 2, 9, 16 .
4. For the observation 4 , we add a vote to region $4,9,14,19$.
5. For the observation 5, we add a vote to region 2, 6, 10, 14, 20.
where the first four observations are the correct observations and the last one is wrong. Then the true region will have more than half of the $R_{l o c}=5$ votes with some good probability.

Details of adding vote In previous description, to let people have a better understanding of $t$-ary search, we simplify the step of adding vote. In fact, adding a vote to each of the $m$
candidate region is not enough. The right thing to do is, not only adding a vote to those $m$ regions, but also adding a vote to a constant number $c_{n}$ of neighbors of each possible region (e.g. two neighbors nearby that region, line 33 in Algorithm 11.2). The reason for doing this is, if the frequency $f$ is located very close to the boundary of two regions, then neither of them will get more than $R_{\text {loc }} / 2$ votes. After we already have $R_{\text {loc }}$ independent observations, we just choose any region that contains more than $R_{\text {loc }} / 2$ votes and enlarge it by the same constant factor $c_{n}$ to be the next candidate frequency region for $t$-array search, and plugging the new parameters into algorithm LOCATEINNER and running it again.

The slightly different last round Recall that, in the previous description of each round of $t$-ary search, we're able to decrease the searching range of frequency geometrically. To achieve this progress, we have to increase the sample duration of $\beta$ is geometrically. At some point, the sample duration will reach $T$. Suppose the sample duration of $\beta$ is geometrically increasing during the first $D-1$ rounds (Line 13-15 in Algorithm 11.2) and it becomes $T$ after $D-1$ rounds. If we want to use the same duration $T$ to perform one more round again, can we get some benefit for learning the frequency by doing some tricks? (1) Suppose in all the first $D-1$ rounds, we use $R_{l o c}(D-1)$ samples. Then we can use $R_{l o c}(D-1)$ for the last round, since it does not increase the sample complexity. This way allows us to learn frequency within $\frac{1}{C T}$. But can we do better than that? (2) Using more samples is a nice observation, another right thing to do is changing the algorithm from reporting region to reporting frequency. In the previous $D-1$ rounds, as the description of $t$-ary search, we just report the region that has more than $R_{\text {loc }} / 2$ votes. But, we can take the median over all the values that were assigned to any region that has votes more than $R_{\text {loc }} / 2$. This way
can actually allow us to learn frequency location more accurately ( $\sim \frac{1}{\rho T}$ ) without increasing the resolution for $\beta$ ! Another question people might ask is, if we take median in the last round, why not just do it in every round? The answer is, in the first $D-1$ rounds, our goal is just narrowing down the search range of frequency location and we do not need to report a frequency location. Another reason is reporting a candidate region needs less samples and has higher succeed probability, but taking the median is more expensive than reporting a candidate region. We cannot pay for taking the median in every round.

To prove main Lemma 11.3.6 for one stage recovery, we introduce Lemma 11.6.1. Before explaining the proof, we give some definitions. Consider a frequency $f$, and define $j=h_{\sigma, b}(f)$ to be the bin that frequency $f$ was hashed into. Define $\theta=f-b(\bmod F)$. Define

$$
\widehat{u}=\operatorname{HashToBins}\left(x, P_{\sigma, \gamma, b}, B, \delta, \alpha\right) \quad \text { and } \quad \widehat{u}^{\prime}=\operatorname{HashToBins}\left(x, P_{\sigma, \gamma+\beta, b}, B, \delta, \alpha\right) .
$$

Note that "bin" and "region" are representing different things in this paper. "bin" is related to hash function $h_{\sigma, b}(f)$. "region" is only used in the algorithm of one stage recovery in this section. Define "true" region to be the region that contains frequency $f$. Define "wrong" region to be the region that is not within a constant number $c_{n}$ of neighbors of the "true" region. Part (I) of Lemma 11.6.1 shows that for each observation generated by a batch of samples drawn from time domain, the counter $v_{j, q^{\prime}}$ corresponding to the true region will increase by one, with some "good" probability. On the other side, Part (II) of Lemma 11.6.1 shows that the counter $v_{j, q}$ corresponding to wrong region will not increase by one, with some "good" probability. Note that, [HIKP12a] proved the case when $c_{n}=6$ under discrete setting, we translate it into continuous setting, here.

Lemma 11.6.1. Given $\sigma$ and b. Assume $f \in \operatorname{region}\left(j, q^{\prime}\right)$ and $\underset{\gamma}{\mathbb{E}}\left[\left|\widehat{u}_{j}-e^{2 \pi \gamma \sigma \theta \mathrm{i}} \widehat{x}(\theta)\right|^{2}\right] \leq$ $\frac{1}{\rho^{2}}|\widehat{x}(\theta)|^{2}$. $\forall 0<s<1$, for each two samples $(\gamma, 0)$ and $(\gamma, \beta)$, where $\gamma$ is sampled from $\left[\frac{1}{2}, 1\right]$ uniformly at random and $\beta$ is sampled from $\left[\frac{s t}{4 \sigma \Delta l}, \frac{s t}{2 \sigma \Delta l}\right]$ uniformly at random, we have
(I) for the $q^{\prime}$, with probability at least $1-\left(\frac{2}{\rho s}\right)^{2}, v_{j, q^{\prime}}$ will increase by one.
(II) for any $q$ such that $\left|q-q^{\prime}\right|>3$, with probability at least $1-15 s$, $v_{j, q}$ will not increase.

Proof. Part (I) of Lemma 11.6.1. We have,

$$
\underset{\gamma}{\mathbb{E}}\left[\left|\widehat{u}_{j}-e^{2 \pi \gamma \sigma \theta \mathrm{i}} \widehat{x}(\theta)\right|^{2}\right] \leq \frac{1}{\rho^{2}}|\widehat{x}(\theta)|^{2}
$$

By Chebyshev's Inequality, we have $\forall g>0$, with probability $1-g$ we have

$$
\begin{aligned}
\left|\widehat{u}_{j}-e^{2 \pi \gamma \sigma \theta} \widehat{x}(\theta)\right| & \leq \frac{1}{\rho} \sqrt{\frac{1}{g}}|\widehat{x}(\theta)| \\
\left\|\phi\left(\widehat{u}_{j}\right)-(\phi(\widehat{x}(\theta))-2 \pi \gamma \sigma \theta)\right\|_{\circ} & \leq \sin ^{-1}\left(\frac{1}{\rho} \sqrt{\frac{1}{g}}\right)
\end{aligned}
$$

where $\|x-y\|_{\bigcirc}=\min _{z \in \mathbb{Z}}|x-y+2 \pi z|$ denote the "circular distance" between $x$ and $y$. Similarly, replacing $\gamma$ by $\gamma+\beta$, with probability $1-g$, we also have that

$$
\left\|\phi\left(\widehat{u}_{j}^{\prime}\right)-(\phi(\widehat{x}(\theta))-2 \pi(\gamma+\beta) \sigma \theta)\right\|_{\circ} \leq \sin ^{-1}\left(\frac{1}{\rho} \sqrt{\frac{1}{g}}\right)
$$

Define $c_{j}=\phi\left(\widehat{u}_{j} / \widehat{u}_{j}^{\prime}\right)$. Combining the above two results, with probability $1-2 g$ we have

$$
\begin{aligned}
\left\|c_{j}-2 \pi \beta \sigma \theta\right\|_{\bigcirc} & =\left\|\phi\left(\widehat{u}_{j}\right)-\phi\left(\widehat{u}_{j}^{\prime}\right)-2 \pi \beta \sigma \theta\right\|_{\bigcirc} \\
& \left.\left.=\|\left(\phi\left(\widehat{u}_{j}\right)-(\phi(\widehat{x}(\theta))-2 \pi \gamma \sigma \theta)\right)\right)-\left(\phi\left(\widehat{u}_{j}^{\prime}\right)-(\phi(\widehat{x}(\theta))-2 \pi(\gamma+\beta) \sigma \theta)\right)\right) \|_{\bigcirc} \\
& \leq\left\|\phi\left(\widehat{u}_{j}\right)-(\phi(\widehat{x}(\theta))-2 \pi \gamma \sigma \theta)\right\|_{\bigcirc}+\left\|\phi\left(\widehat{u}_{j}^{\prime}\right)-(\phi(\widehat{x}(\theta))-2 \pi(\gamma+\beta) \sigma \theta)\right\|_{\bigcirc} \\
& \leq 2 \sin ^{-1}\left(\frac{1}{\rho} \sqrt{\frac{1}{g}}\right) .
\end{aligned}
$$

Here, we want to set $g=\left(\frac{4}{s \pi \rho}\right)^{2}$, thus, with probability at least $1-\left(\frac{2}{s \rho}\right)^{2}$

$$
\begin{equation*}
\left\|c_{j}-2 \pi \beta \sigma \theta\right\|_{\bigcirc}<s \pi / 2 \tag{11.14}
\end{equation*}
$$

The above equation shows that $c_{j}$ is a good estimate for $2 \pi \beta \sigma \theta$ with good probability. We will now show that this means the true region $Q_{q^{\prime}}$ gets a vote with large probability.

For each $q^{\prime}$ with $f \in\left[l_{j}-\frac{\Delta l}{2}+\frac{q^{\prime}-1}{t} \Delta l, l_{j}-\frac{\Delta l}{2}+\frac{q^{\prime}}{t} \Delta l\right] \subset[-F, F]$, we have that $m_{j, q^{\prime}}=l_{j}-\frac{\Delta l}{2}+\frac{q^{\prime}-0.5}{t} \Delta l$ and $\theta_{j, q^{\prime}}=m_{j, q^{\prime}}-b(\bmod F)$ satisfies

$$
\left|f-m_{j, q^{\prime}}\right| \leq \frac{\Delta l}{2 t} \text { and }\left|\theta-\theta_{j, q^{\prime}}\right| \leq \frac{\Delta l}{2 t}
$$

Since we sample $\beta$ uniformly at random from $\left[\frac{s t}{4 \sigma \Delta l}, \frac{s t}{2 \sigma \Delta l}\right]$, then $\beta \leq \frac{s t}{2 \sigma \Delta l}$, which implies that $2 \pi \beta \sigma \frac{\Delta l}{2 t} \leq \frac{s \pi}{2}$. Thus, we can show the observation $c_{j}$ is close to the true region in the following sense,

$$
\begin{aligned}
& \left\|c_{j}-2 \pi \beta \sigma \theta_{j, q^{\prime}}\right\|_{\bigcirc} \\
\leq & \left\|c_{j}-2 \pi \beta \sigma \theta\right\|_{\bigcirc}+\left\|2 \pi \beta \sigma \theta-2 \pi \beta \sigma \theta_{j, q^{\prime}}\right\|_{\bigcirc} \text { by triangle inequality } \\
< & \frac{s \pi}{2}+2 \pi\left\|\beta \sigma \theta-\beta \sigma \theta_{j, q^{\prime}}\right\|_{\bigcirc} \text { by Equation (11.14) } \\
\leq & \frac{s \pi}{2}+2 \pi \beta \sigma \frac{\Delta l}{2 t} \\
= & \frac{s \pi}{2}+\frac{s \pi}{2} \\
\leq & s \pi
\end{aligned}
$$

Thus, $v_{j, q^{\prime}}$ will increase in each round with probability at least $1-\left(\frac{2}{\rho s}\right)^{2}$.

Proof. Part (II) of Lemma 11.6.1.
Consider $q$ with $\left|q-q^{\prime}\right|>3$. Then $\left|f-m_{j, q}\right| \geq \frac{7 \Delta l}{2 t}$, and (assuming $\beta \geq \frac{s t}{4 \sigma \Delta l}$ ) we have

$$
\begin{equation*}
2 \pi \beta \sigma\left|f-m_{j, q}\right| \geq 2 \pi \frac{s t}{4 \sigma \Delta l} \sigma\left|f-m_{j, q}\right|=\frac{s \pi t}{2 \Delta l}\left|f-m_{j, q}\right| \geq \frac{7 s \pi}{4}>\frac{3 s \pi}{2} \tag{11.15}
\end{equation*}
$$

There are two cases: $\left|f-m_{j, q}\right| \leq \frac{\Delta l}{s t}$ and $\left|f-m_{j, q}\right|>\frac{\Delta l}{s t}$.
First, if $\left|f-m_{j, q}\right| \leq \frac{\Delta l}{s t}$. In this case, from the definition of $\beta$ it follows that

$$
\begin{equation*}
2 \pi \beta \sigma\left|f-m_{j, q}\right| \leq \frac{s \pi t}{\sigma \Delta l} \sigma\left|f-m_{j, q}\right| \leq \pi \tag{11.16}
\end{equation*}
$$

Combining equations (11.15) and (11.16) implies that

$$
\operatorname{Pr}\left[2 \pi \beta \sigma\left(f-m_{j, q}\right) \quad \bmod 2 \pi \in\left[-\frac{3 s}{4} 2 \pi, \frac{3 s}{4} 2 \pi\right]\right]=0 .
$$

Second, if $\left|f-m_{j, q}\right|>\frac{\Delta l}{s t}$. We show this claim is true: $\operatorname{Pr}\left[2 \pi \beta \sigma\left(f-m_{j, q}\right)(\bmod 2 \pi) \in\right.$ $\left.\left[-\frac{3 s}{4} 2 \pi, \frac{3 s}{4} 2 \pi\right]\right] \lesssim s$. To prove it, we apply Corollary 11.4 . by setting $\widetilde{T}=2 \pi, \widetilde{\sigma}=2 \pi \sigma \beta$, $\widetilde{\delta}=0, \widetilde{\epsilon}=\frac{3 s}{4} 2 \pi, A=2 \pi \sigma \widehat{\beta}, \Delta f=\left|f-m_{j, q}\right|$. By upper bound of Corollary 11.4.4, the probability is at most

$$
\frac{2 \widetilde{\epsilon}}{\widetilde{T}}+\frac{4 \widetilde{\epsilon}}{A \Delta f}=\frac{3 s}{2}+\frac{3 s}{\sigma \widehat{\beta} \Delta f} \leq \frac{3 s}{2}+\frac{3 s}{\sigma \frac{s t}{4 \sigma \Delta l} \frac{\Delta l}{s t}}<15 s
$$

Then in either case, with probability at least $1-15 \mathrm{~s}$, we have

$$
\left\|2 \pi \beta \sigma m_{j, q}-2 \pi \beta \sigma f\right\|_{\bigcirc}>\frac{3 s}{4} 2 \pi .
$$

which implies that $v_{j, q}$ will not increase.

Lemma 11.3.6. Given $\sigma$ and $b$, consider any frequency $f$ for which neither $E_{\text {coll }}(f)$ nor $E_{\text {off }}(f)$ holds, and let $j=h_{\sigma, b}(f)$. Let $\mu^{2}(f)=\mathbb{E}_{a}\left[\left|\widehat{u_{j}}-\widehat{x^{*}}(f) e^{a \sigma 2 \pi f i}\right|^{2}\right]$ and $\rho^{2}=\left|\widehat{x^{*}}(f)\right|^{2} / \mu^{2}(f)$. For sufficiently large $\rho$, and $\forall 0<s<1, t \geq 4$, consider any run of LocateInner with $f \in\left[l_{j}-\frac{\Delta l}{2}, l_{j}+\frac{\Delta l}{2}\right]$. It takes $O\left(R_{l o c}\right)$ random $(\gamma, \beta) \in\left[\frac{1}{2}, 1\right] \times\left[\frac{s t}{4 \sigma \Delta l}, \frac{s t}{2 \sigma \Delta l}\right]$ samples over duration $\beta \sigma=\Theta\left(\frac{s t}{\Delta l}\right)$, runs in $O\left(s t R_{l o c}\right)$ time, to learn $f$ within a region that has length $\Theta\left(\frac{\Delta l}{t}\right)$ with failure probability at most $\left(\frac{4}{s \rho}\right)^{R_{l o c}}+t \cdot(60 s)^{R_{\text {loc }} / 2}$.


Figure 11.2: For an arbitrary frequency interval $\left[l_{j}-\frac{\Delta l}{2}, l_{j}+\frac{\Delta l}{2}\right]$, we scale it by $2 \pi \sigma \beta$ to get a longer interval $\left[2 \pi \sigma \beta\left(l_{j}+\frac{\Delta l}{2}\right), 2 \pi \sigma \beta\left(l_{j}-\frac{\Delta l}{2}\right)\right]$. Then, we wrap the longer interval on a circle $[0,2 \pi)$. The number of folds after wrapping is $\lceil\sigma \beta \Delta l\rceil$. For any random sample, the observation $c_{j}$ is close to the true answer within $1 / \rho$ with some "good" probability.

Proof. Let $t$ denote the number of regions, and $\left[l_{j}-\frac{\Delta l}{2}, l_{j}+\frac{\Delta l}{2}\right]$ be the interval that contains frequency $f$. Let $Q_{q}$ denote a region that is $\left[l_{j}-\frac{\Delta l}{2}+(q-1) \frac{\Delta l}{t}, l_{j}-\frac{\Delta l}{2}+q \frac{\Delta l}{t}\right]$. Let $\theta=f-b$ $(\bmod F)$. Recall that we sample $\sigma$ uniformly at random from $\left[\frac{1}{B \eta}, \frac{2}{B \eta}\right]$. Then we sample $\gamma$ uniformly at random from $\left[\frac{1}{2}, 1\right]$ and sample $\beta$ uniformly at random from $\left[\frac{s t}{4 \sigma \Delta l}, \frac{s t}{2 \sigma \Delta l}\right]$. Define $c_{j}=\phi\left(\widehat{u}_{j} / \widehat{u}_{j}^{\prime}\right)$. Let $m$ denote the number of folds, which is equal to $\lceil\sigma \beta \Delta l\rceil$. Let $v_{j, q}$ denote the vote of region $(j, q)$.

We hope to show that in any round $r$, each observed $c_{j}$ is close to $2 \pi \sigma \beta \theta$ with good probability. On the other hand, for each observed $c_{j}$, we need to assign it to some regions and increase the vote of the corresponding region. The straightforward way is just checking all possible $t$ regions, which takes $O(t)$ time. In fact, there are only $\Theta(m)$ regions close enough to the observation $c_{j}$, where $m=\Theta\left(\frac{2 \pi \sigma \beta \Delta l}{2 \pi}\right)=\Theta(s t)$. The reason is $2 \pi \sigma \beta$ will scale
the original length $\Delta l$ interval to a new interval that has length $2 \pi \sigma \beta \Delta l$. This new interval can only wrap around circle $[0,2 \pi)$ at most $\lceil\sigma \beta \Delta l\rceil$ times. The running time is $O\left(s t R_{l o c}\right)$, since we take $R_{l o c}$ independent observations.

For each observed $c_{j}$ : Part (I) of Lemma 11.6 .1 says, we assign it to the true region with some good probability; Part (II) of Lemma 11.6.1 says, we do not assign it to the wrong region with some good probability. Thus taking $R_{l o c}$ independent $c_{j}$, we can analyze the failure probability of this algorithm based on these three cases.
(I) What's the probability of true fold fails?

$$
\begin{aligned}
& \operatorname{Pr}[\text { True fold fails }] \\
= & \operatorname{Pr}\left[\text { True region fails } \geq R_{l o c} / 2 \text { times }\right] \\
\leq & 2 \cdot\binom{R_{l o c}}{R_{l o c} / 2} \cdot(\operatorname{Pr}[\text { True region fails once }])^{R_{l o c} / 2} \\
\leq & 2 \cdot\binom{R_{l o c}}{R_{l o c} / 2} \cdot\left(\frac{2}{s \rho}\right)^{2 R_{l o c} / 2} \quad \text { by Lemma 11.6.1 } \\
\leq & \left(\frac{4}{s \rho}\right)^{R_{l o c}}
\end{aligned}
$$

(II) What if the region that is "near"(within $c_{n}$ neighbors) true region becomes true?

Any of those region gets a vote only if true region also gets a vote. Since our algorithm choosing any region that has more than $R_{\text {loc }} / 2$ votes and enlarging the region size by containing $c_{n}$ nearby neighbors of that chosen region, then the new larger region must contain the "real" true region.
(III) What if the region that is "far away"(not within $c_{n}$ neighbors) from true region becomes true?

By Part (II) of Lemma 11.6.1, the probability of one such wrong region gets a vote is at most 15 s . Thus, one of the wrong region gets more than $R_{\text {loc }} / 2$ votes is at most $(60 s)^{R_{\text {loc }} / 2}$. By taking the union bound over all $t$ regions, we have the probability of existing one wrong region getting more than $R_{l o c} / 2$ vote is at most $t \cdot(60 s)^{R_{l o c} / 2}$.

Thus, if we first find any region $Q_{q}$ that has more than $R_{\text {loc }} / 2$ votes, and report a slightly larger region $\left[l_{j}-\frac{\Delta l}{2}+(q-1) \frac{\Delta l}{t}-\frac{c_{n}}{2} \frac{\Delta l}{t}, l_{j}-\frac{\Delta l}{2}+q \frac{\Delta l}{t}+\frac{c_{n}}{2} \frac{\Delta l}{t}\right]$, it is very likely this large region contains the frequency $f$. Finally, the failure probability of this algorithm is at $\operatorname{most} \Theta\left(\left(\frac{4}{s \rho}\right)^{R_{l o c}}+t \cdot(60 s)^{R_{l o c} / 2}\right)$.

Lemma 11.6.2. Taking the median of values belong to any region getting at least $\frac{1}{2} R_{\text {loc }}$ votes, then we can learn frequency $f$ within $\Theta\left(\frac{\Delta l}{\rho s t}\right)$ with probability $1-\exp \left(-\Omega\left(R_{\text {loc }}\right)\right)$.

Proof. Let region $\left(j, q^{\prime}\right)$ be the region that getting at least $\frac{1}{2} R_{\text {loc }}$ votes. Let $R=\left|\operatorname{region}\left(j, q^{\prime}\right)\right|$ denote the number of observations/votes assigned to region $\left(j, q^{\prime}\right)$. Since this region getting at least $\frac{1}{2} R_{\text {loc }}$ votes, then $\frac{1}{2} R_{\text {loc }} \leq R \leq R_{\text {loc }}$. Using Equation (11.14) in Lemma 11.6.1, $\forall g>0$, we have

$$
\left\|c_{j}-2 \pi \beta \sigma \theta\right\|_{\bigcirc} \leq 2 \cdot \sin ^{-1}\left(\frac{1}{\rho} \sqrt{\frac{1}{g}}\right)
$$

holds with probability $1-2 g$. Choosing $g=\Theta(1)$, we have with constant success probability $p>\frac{1}{2}$,

$$
\left\|c_{j}-2 \pi \beta \sigma \theta\right\|_{\circ} \lesssim \frac{1}{\rho}
$$

holds.
Taking the median over all the observations that belong to region $\left(j, q^{\prime}\right)$ gives

$$
\begin{align*}
& \operatorname{Pr}\left[\left|\operatorname{median}_{r \in \operatorname{region}\left(j, q^{\prime}\right)} c_{j}^{r}-2 \pi \beta^{r} \sigma \theta\right| \gtrsim \frac{1}{\rho}\right] \\
& <\sum_{i=R / 2}^{R_{\text {loc }}}\binom{R_{\text {loc }}}{i}(1-p)^{i} p^{R_{\text {loc }}-i} \\
& =\sum_{i=R_{\text {loc }} / 2}^{R_{\text {loc }}}\binom{R_{\text {loc }}}{i}(1-p)^{i} p^{R_{\text {loc }}-i}+\sum_{i=R / 2}^{R_{\text {loc }} / 2}\binom{R_{\text {loc }}}{i}(1-p)^{i} p^{R_{\text {loc }}-i} \\
& <\sum_{i=R_{l o c} / 2}^{R_{l o c}}\binom{R_{l o c}}{R_{l o c} / 2}(1-p)^{i}+\sum_{i=R / 2}^{R_{l o c} / 2}\binom{R_{l o c}}{R_{l o c} / 2}(1-p)^{i} \\
& <2\binom{R_{\text {loc }}}{R_{\text {loc }} / 2}(1-p)^{R / 2} \\
& \leq 2(2 e)^{R_{\text {loc }} / 2}(1-p)^{R_{\text {loc }} / 4} \\
& \leq e^{-c R_{l o c}}, \tag{11.17}
\end{align*}
$$

where the the second inequality follows by $\binom{R}{i} \leq\binom{ R}{R / 2}$ and $p^{R_{\text {loc }}-i}<1$, $\forall i$; the fourth inequality follows by $\binom{n}{k} \leq(n e / k)^{k}$; the last inequality follows by choosing some $p$ such that $\frac{1}{2} \log _{2 e} \frac{1}{1-p}-1>2 c$ where $c>0$ is some constant. Equation (11.17) implies that

$$
\operatorname{Pr}\left(\left|\operatorname{median}_{r \in \operatorname{region}\left(j, q^{\prime}\right)} \theta^{r}-\theta\right|<\frac{1}{\rho 2 \pi \sigma \widehat{\beta}}\right)>1-\exp \left(-\Omega\left(R_{l o c}\right)\right),
$$

where $\forall r \in\left[R_{l o c}\right], \beta^{r}$ is sampled uniformly at random from $[\widehat{\beta}, 2 \widehat{\beta}]=\left[\frac{s t}{4 \sigma \Delta l}, \frac{s t}{2 \sigma \Delta l}\right]$. Thus, we can learn $f$ within $\Theta\left(\frac{1}{\rho 2 \pi \sigma \widehat{\beta}}\right)=\Theta\left(\frac{\Delta l}{\rho s t}\right)$.

Lemma 11.3.7. Algorithm LocateKSignal takes $O\left(k \log _{C}(F T) \log (k / \delta)\right)$ samples over $O\left(\frac{\log (k / \delta)}{\eta}\right)$ duration, runs in $O\left(k \log _{C}(F T) \log (F T / \delta)\right)$ time, and outputs a set $L \subset[-F, F]$ of $O(k)$ frequencies with minimum separation $\Omega(\eta)$.


Figure 11.3: The number of "blue" regions is equal to the number of folds $m$. The number of total regions is $t$. For each observed $c_{j}$, instead of checking all the $t$ regions, we only assign vote to the these "blue" regions. Since only these $m$ "blue" regions can be the candidate region that contains frequency $f$.

Given $\sigma$ and $b$, consider any frequency $f$ for which neither of $E_{\text {coll }}(f)$ or $E_{\text {off }}(f)$ hold. Let $j=h_{\sigma, b}(f), \mu^{2}(f)=\mathbb{E}_{a}\left[\left|\widehat{u}_{j}-\widehat{x^{*}}(f) e^{a \sigma 2 \pi f \mathbf{i}}\right|^{2}\right]$, and $\rho^{2}=\left|\widehat{x^{*}}(f)\right|^{2} / \mu^{2}(f)$. If $\rho>C$, then with an arbitrarily large constant probability there exists an $f^{\prime} \in L$ with

$$
\left|f-f^{\prime}\right| \lesssim \frac{1}{T \rho}
$$

Proof. Algorithm LocateKSignal rerun procedure LocateInner $D$ times. For the first $D-1$ rounds, the sampling range for $\beta$ is increased by $t$ every time. For the last round, the sampling range for $\beta$ is not increasing any more. On the other hand, the sampling range for $\beta$ for $D-1$ round and the last round are the same.

Recall that, we sample $\sigma$ uniformly at random from $\left[\frac{1}{B \eta}, \frac{2}{B \eta}\right]$. Then we sample $\gamma$ uniformly at random from $\left[\frac{1}{2}, 1\right]$ and $\beta$ uniformly at random from $\left[\frac{s t}{4 \sigma \Delta l}, \frac{s t}{2 \sigma \Delta l}\right] . c_{n}$ is some constant for the number of neighbor regions nearby "true" region. In the first $D-1$ rounds, we set $s=1 / \sqrt{C}, \Delta l=F /\left(t^{\prime}\right)_{0}^{i-1}, \forall i \in[D-1], t \approx \log (F T)$, and $t^{\prime}=\frac{t}{c_{n}+1}$. For the last round, we set $s \approx 1 / C, \Delta l \bar{\sim} s t / T$ and $t \bar{\sim} \log (F T) / s . C$ is known as "approximation" factor.

Finally, we need to choose depth $D=\log _{t^{\prime}}(F T / s t) \approx \log _{t^{\prime}}(F T)$ and set $R_{\text {loc }}=O\left(\log _{C}(t C)\right)$ for all the rounds. Define $D R=\sum_{i=1}^{D} R_{l o c}^{i}$, which is $(D-1) R_{l o c}+R_{l o c}^{D}=O\left(\log _{C}(F T)\right)$.

After setting all the parameters for the first $D-1$ rounds and the last round, we explain some intuitions and motivations for setting the last round in a different way of the first $D-1$ rounds. For the first $D-1$ rounds, it is not acceptable to have constant failure probability for each round, since we need to take the union bound over $D-1$ rounds. But, for the last round, it is acceptable to allow just constant failure probability, since it is just a single round. That's the reason for setting $C$ in a different way.

We have the following reason for choosing the number of regions $(=t)$ in the last round larger than that of first $D-1$ rounds. For the first $D-1$ rounds, we do not need to learn frequency within $\approx \frac{1}{T \rho}$. It is enough to know which region does frequency belong to, although the diameter of the region is large at the beginning. The algorithm is making progress round by round, since the diameter of each region is geometrically decreasing while $\widehat{\beta}$ is geometrically increasing. For the last round, by Lemma 11.6.2, we can learn $f$ within $\Theta\left(\frac{\Delta l}{\rho s t}\right)$. Since after the last round, we hope to learn frequency within $\frac{1}{T \rho}$, thus we need to choose some $s, t$ and $\Delta l$ such that $\frac{1}{T} \bar{\sim} \frac{\Delta l}{s t}$. To get more accuracy result in the last round, we'd like to choose a larger $t$. But there is no reason to increase $\widehat{\beta}$ again, since the $\widehat{\beta}$ of the ( $D-1$ )th rounds can tolerance the $t$ we choose at the last round.

To show the constant succeed probability of this Lemma, we separately consider about the failure probability of the first $D-1$ rounds and the last round. By the union bound,
the probability of existing one of the first $D-1$ rounds is failing is,

$$
\begin{aligned}
& (D-1)\left(\left(\frac{4}{s \rho}\right)^{R_{l o c}}+t \cdot(60 s)^{R_{l o c} / 2}\right) \\
\leq & D\left(\left(\frac{4}{s \rho}\right)^{R_{l o c}}+t \cdot(60 s)^{R_{l o c} / 2}\right) \\
\leq & D\left(\left(\frac{4}{s C}\right)^{R_{l o c}}+t \cdot(60 s)^{R_{l o c} / 2}\right) \text { by } \rho>C>1 \\
= & D\left(\left(\frac{4}{\sqrt{C}}\right)^{R_{l o c}}+t \cdot\left(\frac{60}{\sqrt{C}}\right)^{R_{l o c} / 2}\right) \text { by setting } s \bar{\sim} 1 / \sqrt{C} \\
\leq & D \cdot \frac{1}{(C t)^{c}} \text { by setting } R_{l o c}=O\left(\log _{C}(t C)\right),
\end{aligned}
$$

where $c$ is some arbitrarily large constant. Using $t>D$, we can show that failure happening in any of the first $D-1$ rounds is small. Then, we still need to show that the probability of the last round is failing is also small,

$$
\begin{aligned}
& \left(\frac{4}{s \rho}\right)^{R_{l o c}}+t \cdot(60 s)^{R_{l o c} / 2}+e^{-\Theta\left(R_{l o c}\right)} \\
\leq & \left(\Theta\left(\frac{C}{\rho}\right)\right)^{R_{l o c}}+t \cdot\left(\Theta\left(\frac{1}{C}\right)\right)^{R_{l o c} / 2}+e^{-\Theta\left(R_{l o c}\right)} \text { by setting } s \bar{\sim} 1 / C \\
\leq & \frac{1}{c_{1}}+\frac{1}{(t C)^{c_{2}}}+e^{-\Theta\left(R_{l o c}\right)} \text { by setting } R_{l o c}=O\left(\log _{C}(t C)\right) \\
\leq & \frac{1}{c_{1}}+\frac{1}{(t C)^{c_{2}}}+\frac{1}{c_{3}}
\end{aligned}
$$

where in the first line, the first two terms are from Lemma 11.3.6 and the third term comes from Lemma 11.6.2; in the last line $c_{1}, c_{2}$ and $c_{3}$ are some arbitrarily large constants.

The expected running time includes the following part: Running HashtoBins algorithm $O(D R)$ times, each run takes $O\left(\frac{B}{\alpha} \log \frac{k}{\delta}+B \log B\right)$. Updating the counter $v$, which
takes $O(D R \cdot B t)$ time. The total running time should be

$$
\begin{aligned}
& O\left(D R\left(\frac{B}{\alpha} \log \frac{k}{\delta}+B \log B\right)+\left(D R_{l o c} B t\right)\right) \\
= & O(D R B \log (k / \delta \cdot B \cdot F T)) \\
= & O\left(B \log _{C}(F T) \log \left(\frac{k}{\delta} B F T\right)\right) \\
= & O\left(B \log _{C}(F T) \log (F T / \delta)\right) \quad \text { by } F T \gg F \frac{1}{\eta} \gg k .
\end{aligned}
$$

The total number of samples is

$$
O\left(D R \cdot B \log \left(\frac{k}{\delta}\right)\right)=O\left(B \log _{C}(F T) \log (k / \delta)\right)
$$

The sample duration of Algorithm LocateKSignal is $O\left(\frac{\log \frac{k}{\delta}}{\eta}\right)$.
In conclusion, we can show that for frequency where neither $E_{\text {coll }}$ nor $E_{\text {off }}$ holds, we recover an $f^{\prime}$ with $\left|f-f^{\prime}\right| \lesssim \frac{1}{T \rho}$ as long as $\rho>C$, with an arbitrarily large constant probability.

Lemma 11.3.8. Algorithm OneStage takes $O\left(k \log _{C}(F T) \log (k / \delta)\right)$ samples over $O\left(\frac{\log (k / \delta)}{\eta}\right)$ duration, runs in $O\left(k\left(\log _{C}(F T) \log (F T / \delta)\right)\right)$ time, and outputs a set of $\left\{\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\right\}$ of size $O(k)$ with $\min _{i \neq j}\left|f_{i}^{\prime}-f_{j}^{\prime}\right| \gtrsim \eta$. Moreover, one can imagine a subset $S \subseteq[k]$ of "successful" recoveries, where $\operatorname{Pr}[i \in S] \geq \frac{9}{10} \forall i \in[k]$ and for which there exists an injective function $\pi:[k] \rightarrow[O(k)]$ so that

$$
\underset{\sigma, b}{\mathbb{E}}\left[\sum_{i \in S} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} \mathbf{t i}}\right|^{2} \mathrm{~d} t\right] \lesssim C^{2} \mathcal{N}^{2}
$$

with $1-1 / k^{c}$ probability for an arbitrarily large constant $c$.

Proof. Let $H$ denote the set of frequencies $f$ for which neither $E_{\text {coll }}(f)$ nor $E_{\text {off }}(f)$ holds. For each such $f$, let $v=\widehat{x^{*}}(f)$ and denote

$$
\begin{equation*}
\mu^{2}(f)=\underset{a}{\mathbb{E}}\left[\left|\widehat{u}_{j}-v e^{a \sigma 2 \pi f \mathbf{i}}\right|^{2}\right] \tag{11.18}
\end{equation*}
$$

and $\rho^{2}(f)=|v|^{2} / \mu^{2}(f)$, as in Lemma 11.3.7. We have that, if $\rho^{2}(f)>C^{2}$, then with an arbitrarily large constant probability one of the recovered $f^{\prime} \in L$ has $\left|f^{\prime}-f\right| \lesssim \frac{1}{T \rho}$. If this happens, then OnEStage will estimate $v$ using $v^{\prime}=\widehat{u}_{j} e^{-a \sigma 2 \pi f^{\prime} i}$. By triangle inequality,

$$
\begin{equation*}
\left|v^{\prime}-v\right|^{2} \lesssim|v|^{2}\left|e^{a 2 \pi \sigma\left(f^{\prime}-f\right) \mathbf{i}}-1\right|^{2}+\left|\widehat{u}_{j}-v e^{a \sigma 2 \pi f \mathbf{i}}\right|^{2} \tag{11.19}
\end{equation*}
$$

Since $a \sigma \leq T$ and $\left|f^{\prime}-f\right| \lesssim \frac{1}{T \rho}$, then the first term of RHS of Equation (11.19) have

$$
|v|^{2}\left|e^{a 2 \pi \sigma\left(f^{\prime}-f\right) \mathbf{i}}-1\right|^{2} \lesssim|v|^{2}\left|a \sigma\left(f^{\prime}-f\right)\right|^{2}
$$

For the second term of RHS of Equation (11.19). Using Equation (11.18), we have

$$
\left|\widehat{u}_{j}-v e^{a \sigma 2 \pi f \mathbf{i}}\right|^{2} \lesssim \mu^{2}(f),
$$

with arbitrarily large constant probability. Combining the bounds for those two terms gives

$$
\left|v^{\prime}-v\right|^{2} \lesssim|v|^{2}\left|a \sigma\left(f^{\prime}-f\right)\right|^{2}+\mu^{2}(f)
$$

with arbitrarily large constant probability. Since $a \sigma \leq T$, the first term is $|v|^{2} / \rho^{2}=\mu^{2}$, for

$$
\left|v^{\prime}-v\right|^{2} \lesssim \mu^{2}(f)
$$

On the other hand, if $\rho^{2}(f)<C^{2}$, then $|v|=\rho(f) \mu(f) \lesssim C \mu(f)$ so regardless of the frequency $f^{\prime}$ recovered, the estimate $v^{\prime}$ will have

$$
\left|v^{\prime}-v\right|^{2} \lesssim C^{2} \mu^{2}(f)
$$

with arbitrarily large constant probability.
Combining with Lemma 11.3.1, we get for any $f \in H$ that the recovered $f^{\prime}, v^{\prime}$ will have

$$
\frac{1}{T} \int_{0}^{T}\left|v^{\prime} e^{2 \pi f^{\prime} t \mathbf{i}}-v e^{2 \pi f t \mathbf{i}}\right|^{2} \mathrm{~d} t \lesssim C^{2} \mu^{2}(f)
$$

with arbitrarily large constant probability. Let $S \subset H$ be the set of frequencies for which this happens. We can choose our permutation $\pi$ to match frequencies in $S$ to their nearest approximation. By Lemma 11.3.2, this means that

$$
\underset{\sigma, b}{\mathbb{E}}\left[\sum_{i \in S} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} \mathbf{t} \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} \mathbf{t} \mathbf{i}}\right|^{2} \mathrm{~d} t\right] \lesssim C^{2} \mathcal{N}^{2}
$$

as desired.

### 11.7 Proofs for combining multiple stages

We first prove that the median of a bunch of estimates of a frequency has small error if most of the estimates have small error.

Lemma 11.7.1. Let $\left(v_{i}, f_{i}\right)$ be a set of tones for $i \in S$. Define $v^{\prime}$ and $f^{\prime}$ to be the (coordinatewise) median of the $\left(v_{i}, f_{i}\right)$. Then for any $\left(v^{*}, f^{*}\right)$ we have

$$
\frac{1}{T} \int_{0}^{T}\left|v^{*} e^{2 \pi f^{*} t \mathbf{i}}-v^{\prime} e^{2 \pi f^{\prime} t \mathbf{i}}\right|^{2} \mathrm{~d} t \lesssim \underset{i}{\operatorname{median}} \frac{1}{T} \int_{0}^{T}\left|v^{*} e^{2 \pi f^{*} t \mathbf{i}}-v_{i} e^{2 \pi f_{i} \mathbf{i} \mathbf{i}}\right|^{2} \mathrm{~d} t
$$

Proof. By Lemma 11.3.1 we have that

$$
\frac{1}{T} \int_{0}^{T}\left|v^{*} e^{2 \pi f^{*} t \mathbf{i}}-v^{\prime} e^{2 \pi f^{\prime} t \mathbf{i}}\right|^{2} \mathrm{~d} t \approx\left|v^{*}\right|^{2} \min \left(1, T^{2}\left|f^{*}-f^{\prime}\right|^{2}\right)+\left|v^{*}-v^{\prime}\right|^{2}
$$

Using that $v^{\prime}$ is taken as a two dimensional median, it suffices to show: if $x^{(1)}, x^{(2)}, \ldots \in \mathbb{R}^{3}$ then $x^{\prime}=\underset{i}{\operatorname{median}} x^{(i)}$ has

$$
\begin{equation*}
\left\|x^{\prime}\right\|_{2}^{2} \lesssim \underset{i}{\operatorname{median}}\left\|x^{(i)}\right\|_{2}^{2} . \tag{11.20}
\end{equation*}
$$

This follows because in each of the three coordinates $j$, we have

$$
\left(x_{j}^{\prime}\right)^{2}=\left(\underset{i}{\operatorname{median}} x_{j}^{(i)}\right)^{2} \leq \underset{i}{\operatorname{median}}\left(x_{j}^{(i)}\right)^{2} \leq \underset{i}{\operatorname{median}}\left\|x^{(i)}\right\|_{2}^{2} .
$$

so summing over the three coordinates gives (11.20), as desired.
Therefore, for two dimensional median and one dimension median, we have

$$
\begin{equation*}
\left|v^{*}-v^{\prime}\right|^{2} \lesssim \operatorname{median}_{i}\left|v^{*}-v_{i}\right|^{2} \tag{11.21}
\end{equation*}
$$

and

$$
\left|f^{*}-f^{\prime}\right|^{2} \lesssim \operatorname{median}_{i}\left|f^{*}-f_{i}\right|^{2}
$$



Figure 11.4: We demonstrate the algorithm for merging various stages $(R=O(\log k))$ on 2-dimensional data. Note that the true data should be 3-dimensional, since for each tone $\left(v_{i}, f_{i}\right), v_{i} \in \mathbb{C}$ and $f_{i} \in \mathbb{R}$. The $x$-axis represents the frequency and the $y$-axis represents the real part of magnitude.

Moreover,

$$
\begin{align*}
\left|v^{*}\right|^{2} \cdot \min \left(1, T^{2}\left|f^{*}-f^{\prime}\right|^{2}\right) & \lesssim\left|v^{*}\right|^{2} \cdot \min \left(1, T^{2} \underset{i}{\operatorname{median}}\left|f^{*}-f_{i}\right|^{2}\right) \\
& =\left|v^{*}\right|^{2} \cdot \min \left(1, \operatorname{median}_{i}^{2}\left|f^{*}-f_{i}\right|^{2}\right) \\
& =\left|v^{*}\right|^{2} \cdot \operatorname{median} \min \left(1, T^{2}\left|f^{*}-f_{i}\right|^{2}\right) \tag{11.22}
\end{align*}
$$

Combining Equation (11.21) and (11.22), we have

$$
\begin{aligned}
& \left|v^{*}-v^{\prime}\right|^{2}+\left|v^{*}\right|^{2} \cdot \min \left(1, T^{2}\left|f^{*}-f^{\prime}\right|^{2}\right) \\
\lesssim & \operatorname{median}_{i}\left|v^{*}-v_{i}\right|^{2}+\underset{i}{\operatorname{median}}\left|v^{*}\right|^{2} \cdot \min \left(1, T^{2}\left|f^{*}-f_{i}\right|^{2}\right) \\
= & \underset{i}{\operatorname{median}}\left|v^{*}-v_{i}\right|^{2}+\left|v^{*}\right|^{2} \cdot \min \left(1, T^{2}\left|f^{*}-f_{i}\right|^{2}\right) .
\end{aligned}
$$

Thus, we complete the proof.

Lemma 11.3.9. Repeating algorithm OneStage $O(\log k)$ times, MergedStages returns a set $\left\{\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\right\}$ of size $O(k)$ with $\min _{i \neq j}\left|f_{i}^{\prime}-f_{j}^{\prime}\right| \gtrsim \eta$ that can be indexed by $\pi$ such that

$$
\sum_{i=1}^{k} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} \mathbf{t} \mathbf{i}}\right|^{2} \mathrm{~d} t \lesssim C^{2} \mathcal{N}^{2}
$$

with probability $1-1 / k^{c}$ for an arbitrarily large constant $c$.

Proof. Our only goal here is to recover the actual tones well, not to worry about spurious tones.

Suppose the number of stages we perform is $R=O(\log k)$. The algorithm for merging the various stages is to scan over a $c \eta$ size region for small $c$, and take the median (in both frequency and magnitude) over $3 c \eta$ region around that $c \eta$ if there are at least $\frac{6}{10} R$ results in that $c \eta$ region. If so, the algorithm will jump to the first right point that is at least $\eta$ far away from current region and look for the next $c \eta$ region. Because the minimum separation between frequencies for a given stage is $\Omega(\eta)$, this will have minimum separation $\eta$ in the output, and because there are $O(k)$ tones output at each stage so will this method. What remains is to show that the total error is small.

We say a stage is "good" if the term inside the expectation of Lemma 11.3.8 is less than 10 times its expectation, as happens with $9 / 10$ probability:

$$
\operatorname{Pr}\left[\sum_{i \in S} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} \mathbf{t} \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} t \mathbf{i}}\right|^{2} \mathrm{~d} t \leq 10 \underset{\sigma, b}{\mathbb{E}}\left[\sum_{i \in S} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} t \mathbf{i}}\right|^{2} \mathrm{~d} t\right]\right] \geq 9 / 10 .
$$

We say a frequency $f_{i}$ is "successful" in a given stage if the stage is good and $i$ lies in $S$ for that stage. Therefore a frequency is successful in each stage with at least $8 / 10$ probability. For sufficiently large $R=O(\log k)$, this will cause all $k$ frequencies to be successful more than $\frac{7}{10} R$ times with high probability. Suppose this happens.

Let $\mu^{2}\left(f_{i}\right)$ denote the error of $\left(v_{i}, f_{i}\right)$. By Lemma 11.3.8, the total error over all good stages and every successful recovery of a tone in a good stage is $O\left(C^{2} \mathcal{N}^{2} R\right)$. We define $\mu^{2}(f)$ to be the $6 / 10 R$ worst amount of error in the recovery of $f$ over all stages. Because there are $R / 10$ worse successful recoveries for each $f$, we have that $\sum_{i} \mu^{2}\left(f_{i}\right) \lesssim C^{2} \mathcal{N}^{2}$.

We will match each $f_{i}$ with a recovered frequency $f_{i}^{\prime}$ with cost at most $\mu^{2}\left(f_{i}\right)$. If $\mu^{2}\left(f_{i}\right) \gtrsim\left|v_{i}\right|^{2}$, then we can set an arbitrary $f_{i}^{\prime}$ with $v_{i}^{\prime}=0$. Otherwise, more that $\frac{6}{10} R$ successful recoveries of $f_{i}$ also yield $f^{\prime}$ that are within $O\left(\frac{1}{T}\right) \ll c \eta$ of $f_{i}$. Thus the algorithm for merging tones will find enough tones to report something. What it reports will be the median of at most $R$ values, $6 / 10$ of which have less error than $\mu^{2}(f)$. Therefore by Lemma 11.7.1 the reported frequency and magnitude will have error $O\left(\mu^{2}(f)\right)$. This suffices to get the result.

Lemma 11.3.10. If we run MergedStages twice and take the tones $\left\{\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\right\}$ from the first result that have $f_{i}^{\prime}$ within $c \eta$ for small $c$ of some frequency in the second result, we get a set of $k^{\prime \prime}=O(k)$ tones that can be indexed by some permutation $\pi$ such that

$$
\begin{equation*}
\sum_{i=1}^{k} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} \mathbf{t} \mathbf{i}}\right|^{2} \mathrm{~d} t+\sum_{i=k+1}^{k^{\prime \prime}}\left|v_{i}^{\prime}\right|^{2} \lesssim C^{2} \mathcal{N}^{2} \tag{11.5}
\end{equation*}
$$

Proof. By Lemma 11.3.9 and Lemma 11.3.1, the first run gives us a set of $k^{\prime}=O(k)$ pairs $\left\{\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\right\}$ such that they may be indexed by using permutation $\pi$ such that the first $k$ are a good approximation to $\left\{\left(v_{i}, f_{i}\right)\right\}$ in the sense that

$$
\sum_{i=1}^{k}\left(\left|v_{\pi(i)}\right|^{2}+\left|v_{i}^{\prime}\right|^{2}\right) \cdot \min \left(1, T^{2}\left|f_{\pi(i)}-f_{i}^{\prime}\right|^{2}\right)+\left|v_{\pi(i)}-v_{i}^{\prime}\right|^{2} \lesssim C^{2} \mathcal{N}^{2}
$$

We may as well index to match tones to their nearest match in frequency (because the separation for both $f_{i}^{\prime}$ and $f_{i}$ is at least $\left.\Omega(\eta)>1 / T\right)$. That is, we may index such that $\left|f_{j}^{\prime}-f_{i}\right| \gtrsim \eta$ for any $j \neq i$.

Now, for indices where $\left|f_{i}^{\prime}-f_{\pi(i)}\right|>1 / T$, one would do better by setting the corresponding $v_{i}$ to zero. In particular, suppose you knew the true frequencies $f_{\pi(i)}$ and only took
the $\left(v_{i}^{\prime}, f_{i}^{\prime}\right)$ where $f_{i}^{\prime}$ is close to $f_{\pi(i)}$. That is, there exists some permutation $\pi$, for the subset $S \subseteq[k]$ containing $\left\{i:\left|f_{\pi(i)}-f_{i}^{\prime}\right| \leq c / T\right\}$ for any $c \gtrsim 1$, we have

$$
\begin{equation*}
\sum_{i \in S}\left(\left(\left|v_{\pi(i)}\right|^{2}+\left|v_{i}^{\prime}\right|^{2}\right) \cdot \min \left(1, T^{2}\left|f_{\pi(i)}-f_{i}^{\prime}\right|^{2}\right)+\left|v_{\pi(i)}-v_{i}^{\prime}\right|^{2}\right)+\sum_{i \in[k] \backslash S}\left(\left|v_{\pi(i)}\right|^{2}+\left|v_{i}^{\prime}\right|^{2}\right) \lesssim C^{2} \mathcal{N}^{2} \tag{11.23}
\end{equation*}
$$

The problem is that we do not know the set $S$, so we can not throw out the other frequencies. However, we can fake it by running the algorithm again on the signal. In particular, we apply Lemma 11.3.9 again to sparse recovery of the signal defined by

$$
\left\{\left(v_{i}, f_{i}\right) \mid i \in[k]\right\} \cup\left\{\left(0, f_{i}^{\prime}\right) \mid i \in\left\{k+1, \ldots, k^{\prime}\right\}\right\} .
$$

That is, we pretend the "signal" has terms at the other $f_{i}^{\prime}$ for $k<i \leq k$ ', but with magnitude zero. This is an identical signal in time domain, so it's really just an analytical tool; for the analysis, it has $k^{\prime}=O(k)$ sparsity and $\Omega(\eta)$ separation, so Lemma 11.3.9 applies again and gets a set of $k^{\prime \prime}=O(k)$ pairs $\left\{\left(v_{j}^{\prime \prime}, f_{j}^{\prime \prime}\right)\right\}$ such that, for the subset $S^{*} \subset\left[k^{\prime}\right]$ containing the indices $i$ where $\left|f_{j}^{\prime \prime}-f_{i}^{\prime}\right| \leq c / T$ for some $j \in\left[k^{\prime \prime}\right]$. In other words, there exists some permutation $\tau$ such that for the subset $S^{*} \subset\left[k^{\prime}\right]$ containing $\left\{i:\left|f_{\tau(i)}^{\prime \prime}-f_{i}^{\prime}\right| \leq c / T\right\}$. We have
by analogy to (11.23) that

$$
\begin{aligned}
C^{2} \mathcal{N}^{2} & \gtrsim \sum_{i \in S^{*} \cap[k]}\left(\left(\left|v_{\pi(i)}\right|^{2}+\left|v_{\tau(i)}^{\prime \prime}\right|^{2}\right) \cdot \min \left(1, T^{2}\left|f_{\pi(i)}-f_{\tau(i)}^{\prime \prime}\right|^{2}\right)+\left|v_{\pi(i)}-v_{\tau(i)}^{\prime \prime}\right|^{2}\right) \\
& +\sum_{i \in S^{*} \backslash[k]}\left(\left(|0|^{2}+\left|v_{\tau(i)}^{\prime \prime}\right|^{2}\right) \cdot \min \left(1, T^{2}\left|f_{i}^{\prime}-f_{\tau(i)}^{\prime \prime}\right|^{2}\right)+\left|v_{\tau(i)}^{\prime \prime}-0\right|^{2}\right) \\
& +\sum_{i \in[k] \backslash S^{*}}\left(\left|v_{\pi(i)}\right|^{2}+\left|v_{\tau(i)}^{\prime \prime}\right|^{2}\right)+\sum_{i \in\left[k^{\prime}\right] \backslash S^{*} \backslash[k]}\left(\left|v_{\tau(i)}^{\prime \prime}\right|^{2}\right) \\
& \gtrsim \sum_{i \in S^{*} \cap[k]}\left(\left(\left|v_{\pi(i)}\right|^{2}+\left|v_{\tau(i)}^{\prime \prime}\right|^{2}\right) \cdot \min \left(1, T^{2}\left|f_{\pi(i)}-f_{\tau(i)}^{\prime \prime}\right|^{2}\right)+\left|v_{\pi(i)}-v_{\tau(i)}^{\prime \prime}\right|^{2}\right) \\
& +\sum_{i \in S^{*} \backslash[k]}\left|v_{\tau(i)}^{\prime \prime}\right|^{2}+\sum_{i \in[k] \backslash S^{*}}\left|v_{\pi(i)}\right|^{2},
\end{aligned}
$$

where $\pi$ and $\tau$ are two permutations and $\left|S^{*}\right|=k^{*}=O(k)$. This last term is precisely the desired total error for the set of tones $\left\{\left(v_{\tau(i)}^{\prime \prime}, f_{\tau(i)}^{\prime \prime}\right): i \in S^{*}\right\}$, giving the result.

To prove Theorem 11.3.11, we still need the following "local" Lemma.

Lemma 11.7.2. For any three tones $\left(v_{\pi(i)}, f_{\pi(i)}\right),\left(v_{i}^{*}, f_{i}^{*}\right)$ and $\left(v_{i}^{\prime}, f_{i}^{\prime}\right)$, if $\left|v_{i}^{\prime}\right| \geq\left|v_{i}^{*}\right|$ then,

$$
\begin{aligned}
& \left(\left|v_{i}^{\prime}\right|^{2}+\left|v_{\pi(i)}\right|^{2}\right) \cdot \min \left(1, T^{2}\left|f_{i}^{\prime}-f_{\pi(i)}\right|^{2}\right)+\left|v_{i}^{\prime}-v_{\pi(i)}\right|^{2} \\
\lesssim & \left(\left|v_{i}^{*}\right|^{2}+\left|v_{\pi(i)}\right|^{2}\right) \cdot \min \left(1, T^{2}\left|f_{i}^{*}-f_{\pi(i)}\right|^{2}\right)+\left|v_{i}^{*}-v_{\pi(i)}\right|^{2}+\left|v_{i}^{\prime}\right|^{2}
\end{aligned}
$$

Proof. First, we can show an upper bound for LHS. Using inequality $\min \left(1, T^{2}\left|f_{i}^{\prime}-f_{\pi(i)}\right|^{2}\right) \leq$ 1 ,

$$
\operatorname{LHS} \leq\left|v_{i}^{\prime}\right|^{2}+\left|v_{\pi(i)}\right|^{2}+\left|v_{i}^{\prime}-v_{\pi(i)}\right|^{2}
$$

By triangle inequality,

$$
\left|v_{i}^{\prime}-v_{\pi(i)}\right|^{2} \leq 2\left|v_{i}^{\prime}\right|^{2}+2\left|v_{\pi(i)}\right|^{2}
$$

Thus, we obtain,

$$
\text { LHS } \lesssim\left|v_{i}^{\prime}\right|^{2}+\left|v_{\pi(i)}\right|^{2}
$$

Second, we can show a lower bound for RHS. Since first term of RHS is nonnegative, then

$$
\mathrm{RHS} \geq\left|v_{i}^{*}-v_{\pi(i)}\right|^{2}+\left|v_{i}^{\prime}\right|^{2} .
$$

Using $\left|v_{i}^{\prime}\right| \geq\left|v_{i}^{*}\right|$, we have

$$
\text { RHS } \gtrsim\left|v_{i}^{\prime}\right|^{2}+2\left|v_{i}^{*}\right|^{2}+2\left|v_{i}^{*}-v_{\pi(i)}\right|^{2} .
$$

By triangle inequality,

$$
2\left|v_{i}^{*}\right|^{2}+2\left|v_{i}^{*}-v_{\pi(i)}\right|^{2} \geq\left|v_{\pi(i)}\right|^{2} .
$$

Then, we prove the lower bound for RHS,

$$
\text { RHS } \gtrsim\left|v_{i}^{\prime}\right|^{2}+\left|v_{\pi(i)}\right|^{2} .
$$

Combining the lower bound of RHS and the upper bound of LHS completes the proof.

By plugging Lemma 11.3.1 into Lemma 11.7.2, we have

Corollary 11.7.3. For any three tones $\left(v_{\pi(i)}, f_{\pi(i)}\right)$, $\left(v_{i}^{*}, f_{i}^{*}\right)$ and $\left(v_{i}^{\prime}, f_{i}^{\prime}\right)$, if $\left|v_{i}^{\prime}\right| \geq\left|v_{i}^{*}\right|$, then

$$
\frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} t \mathbf{i}}\right|^{2} \mathrm{~d} t \lesssim \frac{1}{T} \int_{0}^{T}\left|v_{i}^{*} e^{2 \pi f_{i}^{*} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} \mathbf{t} \mathbf{i}}\right|^{2} \mathrm{~d} t+\left|v_{i}^{\prime}\right|^{2}
$$

We use Corollary 11.7.3 to present the proof of Theorem 11.3.11,

Theorem 11.3.11. Algorithm ContinuousFourierSparseRecovery returns a set $\left\{\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\right\}$ of size $k$ with $\min _{i \neq j}\left|f_{i}^{\prime}-f_{j}^{\prime}\right| \gtrsim \eta$ for which

$$
\sum_{i=1}^{k} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} \mathbf{i} \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} t \mathbf{i}}\right|^{2} \mathrm{~d} t \lesssim C^{2} \mathcal{N}^{2}
$$

with probability $1-1 / k^{c}$ for an arbitrarily large constant $c$.

Proof. Let $\left\{\left(v_{i}^{*}, f_{i}^{*}\right)\right\}_{i=1,2, \cdots, k^{\prime \prime}}$ denote the set of tones returned by Lemma 11.3.10, where $k^{\prime \prime}=O(k)$ and $k^{\prime \prime}>k$. For any $i \in[k]$, tone $\left(v_{i}^{*}, f_{i}^{*}\right)$ was mapped to tone $\left(v_{\pi(i)}, f_{\pi(i)}\right)$ in Lemma 11.3.10. Let $\left\{\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\right\}_{i=1,2, \cdots, k}$ denote a subset of $\left\{\left(v_{i}^{*}, f_{i}^{*}\right)\right\}_{i=1,2, \cdots, k^{\prime \prime}}$ that satisfies the following two conditions (1) for any $i \in[k],\left|v_{i}^{\prime}\right|$ is one of the top-k largest magnitude tones; (2) for any $i \in[k],\left|v_{i}^{\prime}\right| \geq\left|v_{i}^{*}\right|$. By Lemma 11.3.10, we also know that $\min _{i \neq j}\left|f_{i}^{\prime}-f_{j}^{\prime}\right| \gtrsim \eta$.

For any $i \in[k]$, we consider these three tones $\left(v_{\pi(i)}, f_{\pi(i)}\right),\left(v_{i}^{*}, f_{i}^{*}\right),\left(v_{i}^{\prime}, f_{i}^{\prime}\right)$. If $\left(v_{i}^{*}, f_{i}^{*}\right) \neq$ $\left(v_{i}^{\prime}, f_{i}^{\prime}\right)$, then applying Corollary 11.7.3 we have

$$
\frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} \mathbf{t} \mathbf{i}}\right|^{2} \mathrm{~d} t \lesssim \frac{1}{T} \int_{0}^{T}\left|v_{i}^{*} e^{2 \pi f_{i}^{*} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} t \mathbf{i}}\right|^{2} \mathrm{~d} t+\left|v_{i}^{\prime}\right|^{2}
$$

Otherwise $\left(v_{i}^{*}, f_{i}^{*}\right)=\left(v_{i}^{\prime}, f_{i}^{\prime}\right)$, we also have

$$
\frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} t \mathbf{i}}\right|^{2} \mathrm{~d} t=\frac{1}{T} \int_{0}^{T}\left|v_{i}^{*} e^{2 \pi f_{i}^{*} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} t \mathbf{i}}\right|^{2} \mathrm{~d} t
$$

which means there exists some universal constant $\alpha$ such that $\forall i \in[k]$, if $\left(v_{i}^{*}, f_{i}^{*}\right)=\left(v_{i}^{\prime}, f_{i}^{\prime}\right)$ then

$$
\frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} \mathbf{t} \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} \mathbf{t i}}\right|^{2} \mathrm{~d} t \leq \alpha \cdot\left(\frac{1}{T} \int_{0}^{T}\left|v_{i}^{*} e^{2 \pi f_{i}^{*} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} t \mathbf{i}}\right|^{2} \mathrm{~d} t+\left|v_{i}^{\prime}\right|^{2}\right)
$$

holds. Otherwise $\left(v_{i}^{*}, f_{i}^{*}\right) \neq\left(v_{i}^{\prime}, f_{i}^{\prime}\right)$, then

$$
\frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} \mathbf{i} \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} t \mathbf{i}}\right|^{2} \mathrm{~d} t=\alpha \frac{1}{T} \int_{0}^{T}\left|v_{i}^{*} e^{2 \pi f_{i}^{*} \mathbf{t} \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} \mathbf{t}}\right|^{2} \mathrm{~d} t
$$

holds. Let $S$ denote the set of indices $i$ such that $\left(v_{i}^{*}, f_{i}^{*}\right) \neq\left(v_{i}^{\prime}, f_{i}^{\prime}\right)$. Taking the summation from $i=1$ to $i=k$,

$$
\begin{aligned}
& \sum_{i=1}^{k} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} \mathbf{t} \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} t \mathbf{i}}\right|^{2} \mathrm{~d} t \\
\leq & \sum_{i \in S} \alpha \cdot\left(\frac{1}{T} \int_{0}^{T}\left|v_{i}^{*} e^{2 \pi f_{i}^{*} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} \mathbf{t} \mathbf{i}}\right|^{2} \mathrm{~d} t+\left|v_{i}^{\prime}\right|^{2}\right) \\
+ & \sum_{i \in[k] \backslash S} \alpha \cdot\left(\frac{1}{T} \int_{0}^{T}\left|v_{i}^{*} e^{2 \pi f_{i}^{*} \mathbf{i} \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} \mathbf{t}}\right|^{2} \mathrm{~d} t\right),
\end{aligned}
$$

furthermore, we have

$$
\begin{aligned}
\sum_{i=1}^{k} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} \mathbf{i} \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} \mathbf{i}}\right|^{2} \mathrm{~d} t & \leq \alpha \sum_{i=1}^{k}\left(\frac{1}{T} \int_{0}^{T}\left|v_{i}^{*} e^{2 \pi f_{i}^{*} \mathbf{t} \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} t \mathbf{i}}\right|^{2} \mathrm{~d} t\right) \\
& +\alpha \sum_{i \in S}\left|v_{i}^{\prime}\right|^{2}
\end{aligned}
$$

To finish the proof, we need to show that $\sum_{i \in S}\left|v_{i}^{\prime}\right|^{2} \leq \sum_{i=k+1}^{k^{\prime \prime}}\left|v_{i}^{*}\right|^{2}$. The point is, for any $i \in S$, we know that $\left(v_{i}^{\prime}, f_{i}^{\prime}\right) \neq\left(v_{i}^{*}, f_{i}^{*}\right)$ which implies that $\left(v_{i}^{\prime}, f_{i}^{\prime}\right) \notin\left\{\left(v_{i}^{*}, f_{i}^{*}\right)\right\}_{i=1}^{k}$. Thus,

$$
\begin{aligned}
& \sum_{i=1}^{k} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} t \mathbf{i}}\right|^{2} \mathrm{~d} t \\
\leq & \alpha \sum_{i=1}^{k}\left(\frac{1}{T} \int_{0}^{T}\left|v_{i}^{*} e^{2 \pi f_{i}^{*} \mathbf{t} \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} \mathbf{t}}\right|^{2} \mathrm{~d} t\right)+\alpha \sum_{i=k+1}^{k^{\prime \prime}}\left|v_{i}^{*}\right|^{2} \\
\lesssim & \sum_{i=1}^{k}\left(\frac{1}{T} \int_{0}^{T}\left|v_{i}^{*} e^{2 \pi f_{i}^{*} t \mathbf{i}}-v_{\pi(i)} e^{2 \pi f_{\pi(i)} t \mathbf{i}}\right|^{2} \mathrm{~d} t\right)+\sum_{i=k+1}^{k^{\prime \prime}}\left|v_{i}^{*}\right|^{2} \\
\lesssim & C^{2} \mathcal{N}^{2},
\end{aligned}
$$

where $k^{\prime \prime}=O(k)$ is defined in Lemma 11.3 .10 and the last inequality follows by using Equation (11.5) in Lemma 11.3.10.

### 11.8 Proofs for converting (11.2) into (11.3)

In this section, we show that as long as the sample duration $T$ is sufficiently large, it is possible to convert Equation (11.2) to Equation (11.3). First, we show an auxiliary lemma, Lemma 11.8.3, which bounds an integral that will appear in the analysis.

We will show that

$$
\int_{-\infty}^{+\infty} \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right) \cdot \min \left(T, \frac{1}{\left|f_{j}-\theta\right|}\right) \mathrm{d} \theta \lesssim \frac{\log \left(T\left|f_{i}-f_{j}\right|\right)}{\left|f_{i}-f_{j}\right|}
$$

for $f_{j}-f_{i} \geq 2 / T$. We split this into two pieces.
Claim 11.8.1. Given two frequencies $f_{i}, f_{j}$ and $f_{j}-f_{i} \geq \frac{2}{T}$, we have

$$
\int_{f_{i}-\frac{1}{T}}^{f_{i}} \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right) \cdot \frac{1}{\left|f_{j}-\theta\right|} \mathrm{d} \theta \lesssim \frac{1}{f_{j}-f_{i}}
$$

Proof. By $f_{i}-\frac{1}{T}<\theta<f_{i}$, we have

$$
\mathrm{LHS}=\int_{f_{i}-\frac{1}{T}}^{f_{i}} T \cdot \frac{1}{f_{j}-\theta} \mathrm{d} \theta .
$$

Since $\frac{1}{f_{j}-\theta} \bar{\sim} \frac{1}{f_{j}-f_{i}}$ for all $\theta \in\left[f_{i}-\frac{1}{T}, f_{i}\right]$,

$$
\mathrm{LHS} \lesssim \int_{f_{i}-\frac{1}{T}}^{f_{i}} \frac{T}{f_{j}-f_{i}} \mathrm{~d} \theta=\frac{1}{f_{j}-f_{i}}
$$

Claim 11.8.2. Given two frequencies $f_{i}, f_{j}$ and $f_{j}-f_{i} \geq \frac{2}{T}$, we have

$$
\int_{-\infty}^{f_{i}-\frac{1}{T}} \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right) \cdot \frac{1}{\left|f_{j}-\theta\right|} \mathrm{d} \theta \lesssim \frac{\log \left(T\left|f_{j}-f_{i}\right|\right)}{f_{j}-f_{i}}
$$

Proof. By $\theta<f_{i}-\frac{1}{T}<f_{j}$, we have that

$$
\begin{aligned}
\text { LHS } & =\int_{-\infty}^{f_{i}-\frac{1}{T}} \frac{1}{f_{i}-\theta} \cdot \frac{1}{f_{j}-\theta} \mathrm{d} \theta \\
& =\frac{1}{f_{j}-f_{i}} \int_{-\infty}^{f_{i}-\frac{1}{T}} \frac{f_{j}-f_{i}}{\left(f_{i}-\theta\right)\left(f_{j}-\theta\right)} \mathrm{d} \theta \\
& =\frac{1}{f_{j}-f_{i}} \int_{-\infty}^{f_{i}-\frac{1}{T}} \frac{1}{f_{i}-\theta}-\frac{1}{f_{j}-\theta} \mathrm{d} \theta \\
& =-\frac{1}{f_{j}-f_{i}} \log \frac{f_{i}-f_{i}+\frac{1}{T}}{f_{j}-f_{i}+\frac{1}{T}} \\
& =-\frac{1}{f_{j}-f_{i}} \log \frac{1}{T\left(f_{j}-f_{i}\right)+1} \\
& \lesssim \frac{\log \left(T\left(f_{j}-f_{i}\right)\right)}{f_{j}-f_{i}} .
\end{aligned}
$$

Lemma 11.8.3. Given two frequencies $f_{i}, f_{j}$ and $f_{j}-f_{i} \geq \frac{2}{T}$, we have

$$
\int_{-\infty}^{+\infty} \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right) \cdot \min \left(T, \frac{1}{\left|f_{j}-\theta\right|}\right) \mathrm{d} \theta \lesssim \frac{\log \left(T\left|f_{i}-f_{j}\right|\right)}{\left|f_{i}-f_{j}\right|}
$$

Proof. By symmetry, we have

$$
\operatorname{LHS}=2 \int_{-\infty}^{\frac{f_{i}+f_{j}}{2}} \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right) \cdot \min \left(T, \frac{1}{\left|f_{j}-\theta\right|}\right) \mathrm{d} \theta .
$$

Since $T>\frac{1}{f_{j}-\theta}$ when $\theta<\frac{f_{i}+f_{j}}{2}$,

$$
\operatorname{LHS} \leq 2 \int_{-\infty}^{\frac{f_{i}+f_{j}}{2}} \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right) \cdot \frac{1}{\left|f_{j}-\theta\right|} \mathrm{d} \theta
$$



Figure 11.5: $F(\theta)$ is a sinc function and has derivate function $F^{\prime}(\theta)$.

We also observe that $\frac{1}{\left|f_{j}-\theta\right|} \bar{\sim} \frac{1}{\left|f_{j}-f_{i}\right|}$ for all $\theta \in\left[f_{i}-\frac{f_{j}-f_{i}}{2}, f_{i}+\frac{f_{j}-f_{i}}{2}\right]$,

$$
\int_{f_{i}-\frac{f_{j}-f_{i}}{2}}^{\frac{f_{j}+f_{i}}{2}} \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right) \cdot \frac{1}{\left|f_{j}-\theta\right|} \mathrm{d} \theta \approx \int_{f_{i}-\frac{f_{j}-f_{i}}{2}}^{f_{i}} \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right) \cdot \frac{1}{\left|f_{j}-\theta\right|} \mathrm{d} \theta
$$

Thus, we get

$$
\text { LHS } \lesssim \int_{-\infty}^{f_{i}} \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right) \cdot \frac{1}{\left|f_{j}-\theta\right|} \mathrm{d} \theta
$$

Plugging Claim 11.8.1 and 11.8.2 into the above formula completes the proof.

Lemma 11.8.4. For any $i$, let $a_{i}(t)=v_{i} e^{2 \pi f_{i} t \mathbf{i}}-v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} \mathbf{i} \mathbf{i}}$, then for $i \neq j$,

$$
\begin{equation*}
\frac{1}{T} \int_{0}^{T} a_{i}(t) \overline{a_{j}(t)} \mathrm{d} t \lesssim \frac{\log \left(\Delta f_{i, j} T\right)}{\Delta f_{i, j} T} \cdot\left(\frac{1}{T} \int_{0}^{T}\left|a_{i}(t)\right|^{2} \mathrm{~d} t \cdot \frac{1}{T} \int_{0}^{T}\left|a_{j}(t)\right|^{2} \mathrm{~d} t\right)^{\frac{1}{2}} \tag{11.24}
\end{equation*}
$$

where $\Delta f_{i, j}=\min \left(\left|f_{i}-f_{j}\right|,\left|f_{i}-f_{j}^{\prime}\right|,\left|f_{i}^{\prime}-f_{j}\right|,\left|f_{i}-f_{j}^{\prime}\right|\right)$.

Proof. Let $\nu_{i}=\left|f_{i}-f_{i}^{\prime}\right|$ and $\nu_{j}=\left|f_{j}-f_{j}^{\prime}\right|$. Define $\left\|a_{i}\right\|=\left(\frac{1}{T} \int_{0}^{T}\left|a_{i}(t)\right|^{2} \mathrm{~d} t\right)^{1 / 2}$. We define $f(t)$ and $F(\theta)$ to be a rectangle and sinc function respectively:

$$
\begin{gathered}
f(t)= \begin{cases}1 & \text { if } 0 \leq t \leq T \\
0 & \text { otherwise }\end{cases} \\
F(\theta)=\frac{\sin (2 \pi \theta T)}{2 \pi \theta}
\end{gathered}
$$

where $F=\widehat{f}$.
$F(\theta)$ has the derivative,

$$
F^{\prime}(\theta)=\frac{2 \pi \theta T \cos (2 \pi \theta T)-\sin (2 \pi \theta T)}{2 \pi \theta^{2}}
$$

which means that

$$
\left|F^{\prime}(\theta)\right| \lesssim \begin{cases}T^{2} & \text { if } \theta \leq 1 / T \\ T /|\theta| & \text { otherwise }\end{cases}
$$

Let $y_{i}(t)=a_{i}(t) \cdot f(t)$, then

$$
\begin{aligned}
\widehat{y}_{i}(\theta) & =\widehat{a}_{i}(\theta) * \widehat{f}(\theta) \\
& =\widehat{a}_{i}(\theta) * F(\theta) \text { by } F=\widehat{f} \\
& =v_{i} F\left(f_{i}-\theta\right)-v_{i}^{\prime} F\left(f_{i}^{\prime}-\theta\right) \\
& =\left(v_{i}-v_{i}^{\prime}\right) F\left(f_{i}-\theta\right)+v_{i}^{\prime}\left(f_{i}-f_{i}^{\prime}\right) \cdot F^{\prime}(x-\theta) \text { some } x \in\left[f_{i}, f_{i}^{\prime}\right]
\end{aligned}
$$

We split into two cases. First, if $\nu_{i} \leq \frac{1}{T}$, then

$$
\begin{align*}
\left|\widehat{y}_{i}(\theta)\right| & \lesssim\left(\left|v_{i}-v_{i}^{\prime}\right|+\nu_{i} T\left|v_{i}^{\prime}\right|\right) \cdot \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right) \\
& \lesssim\left\|a_{i}\right\| \cdot \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right) \quad \text { by Lemma 11.3.1, } \tag{11.25}
\end{align*}
$$



Figure 11.6: $\int_{-\infty}^{+\infty} \min \left(T, \frac{1}{\left|\theta-f_{i}\right|}\right) \cdot \min \left(T, \frac{1}{\left|\theta-f_{j}\right|}\right) \mathrm{d} \theta$
where the first step holds for both $f_{i}>f_{i}^{\prime}$ and $f_{i} \leq f_{i}^{\prime}$ since the triangle inequality.
Therefore

$$
\begin{aligned}
\frac{1}{T} \int_{0}^{T} a_{i}(t) \overline{a_{j}(t)} \mathrm{d} t & =\frac{1}{T} \int_{-\infty}^{\infty} y_{i}(t) \overline{y_{j}(t)} \mathrm{d} t \\
& =\frac{1}{T} \int_{-\infty}^{\infty} \widehat{y_{i}}(\theta) \overline{\widehat{y}_{j}(\theta)} \mathrm{d} \theta \\
& \lesssim \frac{1}{T}\left\|a_{i}\right\|\left\|a_{j}\right\| \underbrace{\int_{-\infty}^{+\infty} \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right) \cdot \min \left(T, \frac{1}{\left|f_{j}-\theta\right|}\right) \mathrm{d} \theta}_{C},
\end{aligned}
$$

where the first step follows from $y_{i}(t)=a_{i}(t) \cdot f(t)$ and $f(t)=1$ if $t \in[0, T]$, the second step follows from the property of Fourier transform, and the last step follows from Equation (11.25).

Using Lemma 11.8.3, we have following bound for term $C$,

$$
\int_{-\infty}^{+\infty} \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right) \cdot \min \left(T, \frac{1}{\left|f_{j}-\theta\right|}\right) \mathrm{d} \theta \lesssim \frac{\log T\left|f_{j}-f_{i}\right|}{\left|f_{j}-f_{i}\right|}
$$

This gives the result for $\nu_{i} \leq \frac{1}{T}$. In the alternate case, we have $\nu_{i}>\frac{1}{T}$, then

$$
\begin{aligned}
\left|\widehat{y}_{i}(\theta)\right| & \lesssim v_{i} \cdot \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right)+v_{i}^{\prime} \cdot \min \left(T, \frac{1}{\left|f_{i}^{\prime}-\theta\right|}\right) \\
& \lesssim\left\|a_{i}\right\| \cdot\left(\min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right)+\min \left(T, \frac{1}{\left|f_{i}^{\prime}-\theta\right|}\right)\right) \quad \text { by Lemma 11.3.1. }
\end{aligned}
$$

By similar reason for Equation (11.26), we have

$$
\begin{aligned}
& \frac{1}{T} \int_{0}^{T} a_{i}(t) \overline{a_{j}(t)} \mathrm{d} t \\
\lesssim & \frac{1}{T}\left\|a_{i}\right\|\left\|a_{j}\right\| \underbrace{\int_{-\infty}^{+\infty} \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right) \cdot \min \left(T, \frac{1}{\left|f_{j}-\theta\right|}\right) \mathrm{d} \theta}_{C_{1}} \\
+ & \frac{1}{T}\left\|a_{i}\right\|\left\|a_{j}\right\| \underbrace{\int_{-\infty}^{+\infty} \min \left(T, \frac{1}{\left|f_{i}-\theta\right|}\right) \cdot \min \left(T, \frac{1}{\left|f_{j}^{\prime}-\theta\right|}\right) \mathrm{d} \theta}_{C_{2}} \\
+ & \frac{1}{T}\left\|a_{i}\right\|\left\|a_{j}\right\| \underbrace{\int_{-\infty}^{+\infty} \min \left(T, \frac{1}{\left|f_{i}^{\prime}-\theta\right|}\right) \cdot \min \left(T, \frac{1}{\left|f_{j}-\theta\right|}\right) \mathrm{d} \theta}_{C_{4}} \\
+ & \frac{1}{T}\left\|a_{i}\right\|\left\|a_{j}\right\| \underbrace{\int_{-\infty}^{+\infty} \min \left(T, \frac{1}{\left|f_{i}^{\prime}-\theta\right|}\right) \cdot \min \left(T, \frac{1}{\left|f_{j}^{\prime}-\theta\right|}\right) \mathrm{d} \theta} .
\end{aligned}
$$

Applying Lemma 11.8.3 on the term $C_{1}, C_{2}, C_{3}$ and $C_{4}$ respectively,

$$
\frac{1}{T} \int_{0}^{T} a_{i}(t) \overline{a_{j}(t)} \lesssim\left\|a_{i}\right\|\left\|a_{j}\right\| \frac{\log \left(T \Delta f_{i, j}\right)}{\Delta f_{i, j}}
$$

where $\Delta f_{i, j}=\min \left(\left|f_{i}-f_{j}\right|,\left|f_{i}-f_{j}^{\prime}\right|,\left|f_{i}^{\prime}-f_{j}\right|,\left|f_{i}^{\prime}-f_{j}^{\prime}\right|\right)$.

Lemma 11.8.5. Let $\left\{\left(v_{i}, f_{i}\right)\right\}$ and $\left\{\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\right\}$ be two sets of $k$ tones for which $\min _{i \neq j}\left|f_{i}-f_{j}\right| \geq$ $\eta$ and $\min _{i \neq j}\left|f_{i}^{\prime}-f_{j}^{\prime}\right| \geq \eta$ for some $\eta>0$. Suppose that $T>C / \eta$ for a sufficiently large constant $C$. Then these sets can be indexed such that

$$
\begin{equation*}
\frac{1}{T} \int_{0}^{T}\left|\sum_{i=1}^{k}\left(v_{i}^{\prime} e^{2 \pi \mathbf{i} f_{i}^{\prime} t}-v_{i} e^{2 \pi \mathbf{i} f_{i} t}\right)\right|^{2} \mathrm{~d} t \leq \alpha \cdot \sum_{i=1}^{k} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi \mathbf{i} f_{i}^{\prime} t}-v_{i} e^{2 \pi \mathbf{i} f_{i} t}\right|^{2} \mathrm{~d} t \tag{11.27}
\end{equation*}
$$

where

$$
\alpha=\left(1+O\left(\frac{\log (k \eta T) \log (k)}{\eta T}\right)\right) .
$$

Proof. For simplicity, let $a_{i}(t)=v_{i} e^{2 \pi f_{i} t \mathbf{i}}-v_{i}^{\prime} e^{2 \pi f_{i}^{\prime} t \mathbf{i}}$. Let's express the square of summations by diagonal term and off-diagonal term, and then bound them separately.

$$
\begin{align*}
& \int_{0}^{T}\left|\sum_{i=1}^{k} a_{i}(t)\right|^{2} \mathrm{~d} t \\
= & \int_{0}^{T}\left(\sum_{i=1}^{k} a_{i}(t)\right)\left(\sum_{i=1}^{k} \overline{a_{i}(t)}\right) \mathrm{d} t \\
= & \int_{0}^{T} \sum_{i=1}^{k} \underbrace{a_{i}(t) \overline{a_{i}(t)}}_{\text {diagonal }}+\sum_{i \neq j}^{k} \underbrace{a_{i}(t) \overline{a_{j}(t)}}_{\text {off-diagonal }} \mathrm{d} t . \tag{11.28}
\end{align*}
$$

Using the result of Lemma 11.8.4 and $a^{2}+b^{2} \geq 2 a b$, we can upper bound the offdiagonal term,

$$
\begin{array}{rlr} 
& \int_{0}^{T} a_{i}(t) \overline{a_{j}(t)} \mathrm{d} t & \\
\lesssim & \frac{\log \left(\Delta f_{i, j} T\right)}{\Delta f_{i, j} T} \cdot\left(\int_{0}^{T}\left|a_{i}(t)\right|^{2} \mathrm{~d} t \int_{0}^{T}\left|a_{j}(t)\right|^{2} \mathrm{~d} t\right)^{1 / 2} & \text { by Lemma 11.8.4 } \\
\lesssim & \frac{\log \left(\Delta f_{i, j} T\right)}{\Delta f_{i, j} T} \cdot\left(\int_{0}^{T}\left|a_{i}(t)\right|^{2} \mathrm{~d} t+\int_{0}^{T}\left|a_{j}(t)\right|^{2} \mathrm{~d} t\right) . & \\
\text { by } 2 a b \leq a^{2}+b^{2}
\end{array}
$$

where $\Delta f_{i, j}=\min \left(\left|f_{i}-f_{j}\right|,\left|f_{i}-f_{j}^{\prime}\right|,\left|f_{i}^{\prime}-f_{j}\right|,\left|f_{i}^{\prime}-f_{j}^{\prime}\right|\right)$. If we index such that any $f_{i}$ is matched with any $f_{j}^{\prime}$ with $\left|f_{i}-f_{j}^{\prime}\right|<\eta / 3$ - which is possible, since at most one such $f_{j}^{\prime}$ will exist by the separation among the $f_{j}^{\prime}$, and that $f_{j}^{\prime}$ will be within $\eta / 3$ of at most on $f_{i}$ - then we have $\Delta f_{i, j} \gtrsim\left|f_{i}-f_{j}\right|$. If we order the $f_{i}$ in increasing order, then in fact $\Delta f_{i, j} \gtrsim \eta|i-j|$.

If $T>C / \eta$ for a sufficiently large constant $C$, this means that $\Delta f_{i, j} T \gtrsim|i-j| \eta T \geq e$. Since $\frac{\log x}{x}$ is decreasing on the region, this implies

$$
\frac{\log \left(\Delta f_{i, j} T\right)}{\Delta f_{i, j} T} \lesssim \frac{\log (|i-j| \eta T)}{|i-j| \eta T} .
$$

Thus, we have

$$
\begin{equation*}
\int_{0}^{T} a_{i}(t) \overline{a_{j}(t)} \mathrm{d} t \lesssim \frac{\log (|i-j| \eta T)}{|i-j| \eta T} \cdot\left(\int_{0}^{T}\left|a_{i}(t)\right|^{2} \mathrm{~d} t+\int_{0}^{T}\left|a_{j}(t)\right|^{2} \mathrm{~d} t\right) \tag{11.29}
\end{equation*}
$$

Finally, we have

$$
\begin{align*}
& \int_{0}^{T}\left|\sum_{i=1}^{k} a_{i}(t)\right|^{2} \mathrm{~d} t-\sum_{i=1}^{k} \int_{0}^{T}\left|a_{i}(t)\right|^{2} \mathrm{~d} t \\
= & \sum_{i \neq j}^{k} \int_{0}^{T} a_{i}(t) \overline{a_{j}(t)} \mathrm{d} t  \tag{11.28}\\
\lesssim & \sum_{i \neq j}^{k} \frac{\log (|i-j| \eta T)}{|i-j| \eta T} \cdot\left(\int_{0}^{T}\left|a_{i}(t)\right|^{2} \mathrm{~d} t+\int_{0}^{T}\left|a_{j}(t)\right|^{2} \mathrm{~d} t\right)  \tag{11.29}\\
\leq & \text { by Equation (11.28) } \\
= & \text { by Equation (11.29) } \\
= & \frac{\log (k \eta T)}{\eta T} \sum_{i=1}^{k} \sum_{j \neq i}^{k} \frac{1}{|i-j|} \cdot\left(\int_{0}^{T}\left|a_{i}(t)\right|^{2} \mathrm{~d} t+\int_{0}^{T}\left|a_{j}(t)\right|^{2} \mathrm{~d} t\right) \\
\lesssim & \frac{\log (k \eta T)}{\eta T} \sum_{i=1}^{k} \int_{j \neq i}^{k} \frac{1}{|i-j|} \int_{0}^{T}\left|a_{i}(t)\right|^{2} \mathrm{~d} t \\
\lesssim & \\
\lesssim & \frac{\left.\log (k \eta T)\right|^{2} \mathrm{~d} t \sum_{j \neq i}^{k} \frac{1}{\eta i-j \mid}}{\eta T}(k) \\
& \sum_{i=1}^{k} \int_{0}^{T}\left|a_{i}(t)\right|^{2} \mathrm{~d} t \quad \text { by } \sum_{i=1}^{k} \frac{1}{i} \sim \log (k)
\end{align*}
$$

Thus, we complete the proof.
Lemma 11.3.13. Let $\left\{\left(v_{i}, f_{i}\right)\right\}$ and $\left\{\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\right\}$ be two sets of $k$ tones for which $\min _{i \neq j}\left|f_{i}-f_{j}\right| \geq$ $\eta$ and $\min _{i \neq j}\left|f_{i}^{\prime}-f_{j}^{\prime}\right| \gtrsim \eta$ for some $\eta>0$. Suppose that $T>O\left(\frac{\log ^{2} k}{\eta}\right)$. Then these sets can be indexed such that

$$
\begin{equation*}
\frac{1}{T} \int_{0}^{T}\left|\sum_{i=1}^{k}\left(v_{i}^{\prime} e^{2 \pi \mathbf{i} \mathrm{f}_{i}^{\prime} t}-v_{i} e^{2 \pi \mathbf{i} \mathrm{f}_{i} t}\right)\right|^{2} \mathrm{~d} t \lesssim \sum_{i=1}^{k} \frac{1}{T} \int_{0}^{T}\left|v_{i}^{\prime} e^{2 \pi \mathrm{i} \mathrm{f}_{i}^{\prime} t}-v_{i} e^{2 \pi \mathrm{i} \mathrm{f}_{i} t}\right|^{2} \mathrm{~d} t . \tag{11.30}
\end{equation*}
$$

Proof. Directly follows by Lemma 11.8.5.

### 11.9 Lower Bound

Lemma 11.3.15. There exists a constant $c>0$ such that, for a given sample duration $T$, one cannot recover the frequency $f$ to within

$$
c \frac{\mathcal{N}}{T\left|\widehat{x}^{*}(f)\right|}
$$

with $3 / 4$ probability, for all $\delta>0$, even if $k=1$.

Proof. Suppose this were possible, and consider two one-sparse signals $y$ and $y^{\prime}$ containing tones $(v, f)$ and $\left(v, f^{\prime}\right)$, respectively. By Lemma 11.3.1,

$$
\int_{0}^{T}|y(t)-y(t)|^{2} \mathrm{~d} t \lesssim|v|^{2} T^{2}\left|f-f^{\prime}\right|^{2}
$$

Consider recovery of the signal $x(t)=y(t)$, and suppose it outputs some frequency $f^{*}$. This must simultaneously be a good recovery for the decomposition $\left(x^{*}, g\right)=(y, 0)$ and $\left(x^{*}, g\right)=\left(y^{\prime}, y-y^{\prime}\right)$. These have noise levels $\mathcal{N}^{2}$ bounded by $\delta|v|^{2}$ and $\delta|v|^{2}+O\left(|v|^{2} T^{2}\left|f-f^{\prime}\right|^{2}\right)$, respectively. By the assumption of good recovery, and the triangle inequality, we require

$$
c \cdot \frac{2 \sqrt{\delta|v|^{2}}+\sqrt{O\left(|v|^{2} T^{2}\left|f-f^{\prime}\right|^{2}\right)}}{T v} \gtrsim\left|f-f^{\prime}\right|
$$

or

$$
c \cdot O\left(\left(\frac{\delta}{T\left|f-f^{\prime}\right|}\right)^{1 / 2}+1\right) \geq 1
$$

Because $\delta$ may be chosen arbitrarily small, we can choose a small constant $c$ such that this is a contradiction.

```
Algorithm 11.1 Continuous Fourier Sparse Recovery - Part 1
    procedure ContinuousFourierSparseRecovery \((x, k, \delta, \alpha, C, F, T, \eta)\) - The
    Main Algorithm
        \(F\) is the upper bound of frequency, \(T\) is the sample duration, \(C\) is the approximation
    factor.
        \(\delta, \alpha\) are the parameters associated with Hash function.
        \(c=1 / 10\) and \(b=8 / 10\)
        \(R_{1} \leftarrow \operatorname{NoisyKSparseCFFT}(x, k, \delta, \alpha, C, F)\)
        \(S_{1} \leftarrow \operatorname{MergedStages}\left(R_{1}, O(k \log k), \eta, c, b\right)\)
        \(R_{2} \leftarrow \operatorname{NoisyKSparseCFFT}(x, O(k), \delta, \alpha, C, F, T)\)
        \(S_{2} \leftarrow \operatorname{MERGEDStages}\left(R_{2}, O(k \log k), \Omega(\eta), c, b\right)\)
        \(S \leftarrow S_{1} \cap S_{2}\), which means only keeping the tones that \(S_{1}\) agrees with \(S_{2}\) by Lemma
    11.3.10
        \(S^{*} \leftarrow \operatorname{Prune}(S, k)\), which means only keeping the top- \(k\) largest magnitude tones.
        return \(S^{*}\)
    end procedure
```

```
Algorithm 11.2 Continuous Fourier Sparse Recovery - Part 2
    procedure NoisyKSparseCFFT( \(x, k, \delta, \alpha, C, F, T\) )
        Let \(B=k / \epsilon\).
        for \(c=1 \rightarrow \log (k)\) do
            Choose \(\sigma\) uniformly at random from \(\left[\frac{1}{B \eta}, \frac{2}{B \eta}\right]\).
            Choose \(b\) uniformly at random from \(\left[0, \frac{2 \pi\lceil F / \eta\rceil}{\sigma B}\right]\).
            \(R_{c} \leftarrow \operatorname{OnEStage}(x, B, \delta, \alpha, \sigma, b, C, F, T)\)
        end for
        return \(\left(R_{1}, R_{2}, \cdots, R_{\log (k)}\right)\).
    end procedure
    procedure MergedStages \((R, m, \eta, c, b)\)
        \(R\) is a list of \(m\) tones \(\left(v_{i}^{\prime}, f_{i}^{\prime}\right)\)
        \(c\) is some constant \(<1\).
        \(b\) is some constant \(<1\).
        Sort list \(R\) based on \(f_{i}^{\prime}\).
        Building the 1D range search Tree based on \(m\) points by regarding each frequency
    \(f_{i}^{\prime}\) as a 1 D point on a line where \(x_{i}=f_{i}^{\prime}\).
        \(S \leftarrow \emptyset, i \leftarrow 0\)
        while \(i<m\) do
            if Tree. \(\operatorname{Count}\left(f_{i}^{\prime}, f_{i}^{\prime}+c \eta\right) \geq b \log k\) then
                \(f \leftarrow \operatorname{median}\left\{f_{j}^{\prime} \mid f_{j}^{\prime} \in\left[f_{i}^{\prime}-c \eta, f_{i}^{\prime}+2 c \eta\right]\right\}\)
                \(v \leftarrow \operatorname{median}\left\{v_{j}^{\prime} \mid f_{j}^{\prime} \in\left[f_{i}^{\prime}-c \eta, f_{i}^{\prime}+2 c \eta\right]\right\}\)
                \(S \leftarrow S \cup(f, v)\)
                \(i \leftarrow \operatorname{Tree} . \operatorname{Search}\left(f_{i}^{\prime}+2 c \eta+\eta / 2\right)\), which means walk to the first point that
    is on the right of \(f_{i}^{\prime}+2 c \eta+\eta / 2\)
            else
                \(i \leftarrow i+1\)
            end if
        end while
        return \(S\)
    end procedure
```

```
Algorithm 11.3 Continuous Fourier Sparse Recovery - Part 3
    procedure \(\operatorname{HashToBins}\left(x, P_{\sigma, a, b}, B, \delta, \alpha\right)\)
        Compute \(\widehat{y}_{j F / B}\) for \(j \in[B]\), where \(y=G_{B, \alpha, \delta} \cdot\left(P_{\sigma, a, b} x\right)\)
        return \(\widehat{u}\) given by \(\widehat{u}_{j}=\widehat{y}_{j F / B}\)
    end procedure
    procedure OneStage \((x, B, \delta, \alpha, \sigma, b, C, F, T)\)
        \(L \leftarrow \operatorname{LocateKSignal}(x, B, \delta, \alpha, \sigma, b, C, F, T)\)
        Choose \(a \in[0,1]\) uniformly at random.
        \(\widehat{u} \leftarrow \operatorname{Hash} \operatorname{ToBins}\left(x, P_{\sigma, a, b}, B, \delta, \alpha\right)\)
        return \(\left\{\left(\widehat{u}_{h_{\sigma, b}\left(f^{\prime}\right)} e^{-2 \pi \sigma a f^{\prime} \mathbf{i}}, f^{\prime}\right)\right.\) for \(f^{\prime} \in L\) if not \(\left.E_{\text {off }}\left(f^{\prime}\right)\right\}\).
    end procedure
    procedure LocateKSignal \((x, B, \delta, \alpha, \sigma, b, C, F, T)\)
        Set \(t \approx \log (F T), t^{\prime}=t /\left(c_{n}+1\right), D \approx \log _{t^{\prime}}(F T), R_{\text {loc }} \approx \log _{C}(t C), l^{(1)}=F / 2\).
        for \(i \in[D-1]\) do
            \(\Delta l \approx F /\left(t^{\prime}\right)^{i-1}, s=\frac{1}{\sqrt{C}}, \widehat{\beta}=\frac{t s}{2 \sigma \Delta l}\)
            \(l^{(i+1)} \leftarrow \operatorname{LOCATEINNER}\left(x, B, \delta, \alpha, \sigma, b, \widehat{\beta}, l^{(i)}, \Delta l, t, R_{l o c}\right.\), false \()\).
        end for
            Set \(s=1 / C, t \approx \log (F T) / s, \Delta l \approx s t / T, \widehat{\beta}=\frac{t s}{2 \sigma \Delta l}, R_{l o c} \approx \log _{C}(t C)\)
            \(l^{(*)} \leftarrow \operatorname{LocatEInNER}\left(x, B, \delta, \alpha, \sigma, b, \widehat{\beta}, l^{(D)}, \Delta l, t, R_{l o c}\right.\), true \()\).
            return \(l^{(*)}\).
    end procedure
```

```
Algorithm 11.4 Continuous Fourier Sparse Recovery - Part 4
    procedure LocateInner \(\left(x, B, \delta, \sigma, b, \widehat{\beta}, l, \Delta l, t, R_{l o c}\right.\), last \() \quad \triangleright\) Lemma 11.3.6
        Let \(v_{j, q}=0\) for \((j, q) \in[B] \times[t]\).
        for \(r \in\left[R_{\text {loc }}\right]\) do
            Choose \(\gamma \in\left[\frac{1}{2}, 1\right]\) uniformly at random.
            Choose \(\beta \in\left[\frac{1}{2} \widehat{\beta}, 1 \widehat{\beta}\right]\) uniformly at random.
            \(\widehat{u} \leftarrow \operatorname{HashToBins}\left(x, P_{\sigma, \gamma, b}, B, \delta, \alpha\right)\).
            \(\widehat{u}^{\prime} \leftarrow \operatorname{Hash} T o B i n s\left(x, P_{\sigma, \gamma+\beta, b}, B, \delta, \alpha\right)\).
            for \(j \in[B]\) do
                for \(i \in[m]\) do
                    \(\theta_{j, i}^{r}=\frac{1}{2 \pi \sigma \beta}\left(\phi\left(\widehat{u}_{j} / \widehat{u_{j}^{\prime}}{ }_{j}\right)+2 \pi s_{i}\right), s_{i} \in\left[\sigma \beta\left(l_{j}-\Delta l / 2\right), \sigma \beta\left(l_{j}+\Delta l / 2\right)\right] \cap \mathbb{Z}_{+}\)
                    \(f_{j, i}^{r}=\theta_{j, i}^{r}+b(\bmod F)\)
                suppose \(f_{j, i}^{r}\) belongs to region \((j, q)\),
                add a vote to both region \((j, q)\) and two neighbors nearby that region, e.g.
    region \((j, q-1)\) and region \((j, q+1)\)
                end for
            end for
        end for
        for \(j \in[B]\) do
            \(q_{j}^{*} \leftarrow\left\{q \left\lvert\, v_{j, q}>\frac{R_{l o c}}{2}\right.\right\}\)
            if last \(=\) true then
                \(l_{j}^{*} \leftarrow \operatorname{median}\left\{f_{j, i}^{r} \mid f_{j, i}^{r} \in \operatorname{region}\left(j, q_{j}^{*}\right), i \in[f], r \in\left[R_{l o c}\right]\right\}\)
            else
                \(l_{j}^{*} \leftarrow\) center of region \(\left(j, q_{j}^{*}\right)\)
            end if
        end for
        return \(l^{*}\)
    end procedure
```


## Chapter 12

## Continuous Fourier Transform II

We consider the problem of estimating a Fourier-sparse signal from noisy samples, where the sampling is done over some interval $[0, T]$ and the frequencies can be "off-grid". Previous methods for this problem required the gap between frequencies to be above $1 / T$, the threshold required to robustly identify individual frequencies. We show the frequency gap is not necessary to estimate the signal as a whole: for arbitrary $k$-Fourier-sparse signals under $\ell_{2}$ bounded noise, we show how to estimate the signal with a constant factor growth of the noise and sample complexity polynomial in $k$ and logarithmic in the bandwidth and signal-to-noise ratio.

As a special case, we get an algorithm to interpolate degree $d$ polynomials from noisy measurements, using $O(d)$ samples and increasing the noise by a constant factor in $\ell_{2}$.

### 12.1 Introduction

In an interpolation problem, one can observe $x(t)=x^{*}(t)+g(t)$, where $x^{*}(t)$ is a structured signal and $g(t)$ denotes noise, at points $t_{i}$ of one's choice in some interval $[0, T]$. The goal is to recover an estimate $\widetilde{x}$ of $x^{*}$ (or of $x$ ). Because we can sample over a particular interval, we would like our approximation to be good on that interval, so for any function $y(t)$ we define

$$
\|y\|_{T}^{2}=\frac{1}{T} \int_{0}^{T}|y(t)|^{2} \mathrm{~d} t
$$

to be the $\ell_{2}$ error on the sample interval. For some parameters $C$ and $\delta$, we would then like to get

$$
\begin{equation*}
\left\|\widetilde{x}-x^{*}\right\|_{T} \leq C\|g\|_{T}+\delta\left\|x^{*}\right\|_{T} \tag{12.1}
\end{equation*}
$$

while minimizing the number of samples and running time. Typically, we would like $C$ to be $O(1)$ and to have $\delta$ be very small (either zero, or exponentially small). Note that, if we do not care about changing $C$ by $O(1)$, then by the triangle inequality it doesn't matter whether we want to estimate $x^{*}$ or $x$ (i.e. we could replace the LHS of (12.1) by $\|\widetilde{x}-x\|_{T}$ ).

Of course, to solve an interpolation problem one also needs $x^{*}$ to have structure. One common form of structure is that $x^{*}$ have a sparse Fourier representation. We say that a function $x^{*}$ is $k$-Fourier-sparse if it can be expressed as a sum of $k$ complex exponentials:

$$
x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathbf{i} f_{j} t}
$$

for some $v_{j} \in \mathbb{C}$ and $f_{j} \in[-F, F]$, where $F$ is the "bandlimit". Given $F, T$, and $k$, how many samples must we take for the interpolation (12.1)?

If we ignore sparsity and just use the bandlimit, then Nyquist sampling and ShannonWhittaker interpolation uses $F T+1 / \delta$ samples to achieve (12.1). Alternatively, in the absence of noise, $x^{*}$ can be found from $O(k)$ samples by a variety of methods, including Prony's method from 1795 or Reed-Solomon syndrome decoding [Mas69], but these methods are not robust to noise.

If the signal is periodic with period $T$-i.e., the frequencies are multiples of $1 / T$-then we can use sparse discrete Fourier transform methods, which take $O\left(k \log ^{c}(F T / \delta)\right)$ time and samples (e.g. [GGI ${ }^{+} 02$, HIKP12a, IKP14]). If the frequencies are not multiples of $1 / T$ (are "off the grid"), then the discrete approximation is only $k / \delta$ sparse, making the interpolation less efficient; and even this requires that the frequencies be well separated.

A variety of algorithms have been designed to recover off-grid frequencies directly, but they require the minimum gap among the frequencies to be above some threshold. With frequency gap at least $1 / T$, we can achieve a $k^{c}$ approximation factor using $O(F T)$ samples [Moi15], and with gap above $O\left(\log ^{2} k\right) / T$ we can get a constant approximation using $O\left(k \log ^{c}(F T / \delta)\right)$ samples and time [PS15].

Having a dependence on the frequency gap is natural. If two frequencies are very close together-significantly below $1 / T$-then the corresponding complex exponentials will be close on $[0, T]$, and hard to distinguish in the presence of noise. In fact, from a lower bound in [Moi15], below $1 / T$ frequency gap one cannot recover the frequencies in the presence of noise as small as $2^{-\Omega(k)}$. The lower bound proceeds by constructing two signals using significantly different frequencies that are exponentially close over $[0, T]$.

But if two signals are so close, do we need to distinguish them? Such a lower bound
doesn't apply to the interpolation problem, it just says that you can't solve it by finding the frequencies. Our question becomes: can we benefit from Fourier sparsity in a regime where we can't recover the individual frequencies?

We answer in the affirmative, giving an algorithm for the interpolation using

$$
\operatorname{poly}(k \log (F T / \delta))
$$

samples. Our main theorem is the following:
Theorem 12.1.1. Let $x(t)=x^{*}(t)+g(t)$, where $x^{*}$ is $k$-Fourier-sparse signal with frequencies in $[-F, F]$. Given samples of $x$ over $[0, T]$ we can output $\widetilde{x}(t)$ such that with probability at least $1-2^{-\Omega(k)}$,

$$
\left\|\widetilde{x}-x^{*}\right\|_{T} \lesssim\|g\|_{T}+\delta\left\|x^{*}\right\|_{T} .
$$

Our algorithm uses poly $(k, \log (1 / \delta)) \cdot \log (F T)$ samples and poly $(k, \log (1 / \delta)) \cdot \log ^{2}(F T)$ time. The output $\widetilde{x}$ is poly $(k, \log (1 / \delta))$-Fourier-sparse signal.

Relative to previous work, this result avoids the need for a frequency gap, but loses a polynomial factor in the sample complexity and time. We lose polynomial factors in a number of places; some of these are for ease of exposition, but others are challenging to avoid.

Degree $d$ polynomials are the special case of $d$-Fourier-sparse functions in the limit of $f_{j} \rightarrow 0$, by a Taylor expansion. This is a regime with no frequency gap, so previous sparse Fourier results would not apply but Theorem 12.1.1 shows that poly $(d \log (1 / \delta))$ samples suffices. In fact, in this special case we can get a better polynomial bound:

Theorem 12.1.2. For any degree $d$ polynomial $P(t)$ and an arbitrary function $g(t)$, Procedure RobustPolynomialLearning in Algorithm 12.5 takes $O(d)$ samples from $x(t)=$ $P(t)+g(t)$ over $[0, T]$ and reports a degree $d$ polynomial $Q(t)$ in time $O\left(d^{\omega}\right)$ such that, with probability at least 99/100,

$$
\|P(t)-Q(t)\|_{T}^{2} \lesssim\|g(t)\|_{T}^{2} .
$$

where $\omega<2.373$ is matrix multiplication exponent [Str69],[CW87],[Wil12].

We also show how to reduce the failure probability to an arbitrary $p>0$ with $O(\log (1 / p))$ independent repetitions, in Theorem 12.4.5.

Although we have not seen such a result stated in the literature, our method is quite similar to one used in [CDL13]. Since $d$ samples are necessary to interpolate a polynomial without noise, the result is within constant factors of optimal.

One could apply Theorem 12.1.2 to approximate other functions that are well approximated by polynomials or piecewise polynomials. For example, a Gaussian of standard deviation at least $\sigma$ can be approximated by a polynomial of degree $O\left(\left(\frac{T}{\sigma}\right)^{2}+\log (1 / \delta)\right)$; hence the same bound applies as the sample complexity of improper interpolation of a positive mixture of Gaussians.

### 12.1.1 Related work

Sparse discrete Fourier transforms. There is a large literature on sparse discrete Fourier transforms. Results generally are divided into two categories: one category of results that carefully choose measurements that allow for sublinear recovery time, including [GGI+ 02, GMS05, HIKP12b, Iwe13, HIKP12a, IK14, IKP14, Kap16]. The other category of results expect randomly chosen measurements and show that a generic recovery algorithm such as $\ell_{1}$ minimization will work with high probability; these results often focus on proving the Restricted Isometry Property [CRT06b, RV08, Bou14, HR16]. At the moment, the first category of results have better theoretical sample complexity and running time, while results in the second category have better failure probabilities and empirical performance. Our result falls in the first category. The best results here can achieve $O(k \log n)$ samples [IK14], $O\left(k \log ^{2} n\right)$ time [HIKP12b], or within $\log \log n$ factors of both [IKP14].

For signals that are not periodic, the discrete Fourier transform will not be sparse: it takes $k / \delta$ frequencies to capture a $1-\delta$ fraction of the energy. To get a better dependence on $\delta$, one has to consider frequencies "off the grid", i.e. that are not multiples of $1 / T$.

Off the grid. Finding the frequencies of a signal with sparse Fourier transform off the grid has been a question of extensive study. The first algorithm was by Prony in 1795, which worked in the noiseless setting. This was refined by classical algorithms like MUSIC [Sch81] and ESPRIT [RPK86], which empirically work better with noise. Matrix pencil [BM86] is a method for computing the maximum likelihood signal under Gaussian noise and evenly spaced samples. The question remained how accurate the maximum likelihood estimate is; [Moi15] showed that it has an $O\left(k^{c}\right)$ approximation factor if the frequency gap is at least
$1 / T$.
Now, the above results all use $F T$ samples, which is analogous to $n$ in the discrete setting. This can be decreased down till $O(k)$ by only looking at a subset of time, i.e. decreasing $T$; but doing so increases the frequency gap needed for decent robustness results.

A variety of works have studied how to adapt sparse Fourier techniques from the discrete setting to get sublinear sample complexity; they all rely on the minimum separation among the frequencies to be at least $c / T$ for $c \geq 1$. [TBSR13] showed that a convex program can recover the frequencies exactly in the noiseless setting, for $c \geq 4$. This was improved in [CF14] to $c \geq 2$ for complex signals and $c \geq 1.87$ for real signals. [CF14] also gave a result for $c \geq 2$ that was stable to noise, but this required the signal frequencies to be placed on a finely spaced grid. [YX15] gave a different convex relaxation that empirically requires smaller $c$ in the noiseless setting. [DB13] used model-based compressed sensing when $c=\Omega(1)$, again without theoretical noise stability. Note that, in the noiseless setting, exact recovery can be achieved without any frequency separation using Prony's method or Berlekamp-Massey syndrome decoding [Mas69]; the benefit of the above results is that a convex program might be robust to noise, even if it has not been proven to be so.

In the noisy setting, [FL12] gave an extension of Orthogonal Matching Pursuit (OMP) that can recover signals when $c=\Omega(k)$, with an approximation factor $O(k)$, and a few other assumptions. Similarly, $\left[\mathrm{BCG}^{+} 12\right]$ gave a method that required $c=\Omega(k)$ and was robust to certain kinds of noise. [HK15] got the threshold down to $c=O(1)$, in multiple dimensions, but with approximation factor $O\left(F T k^{O(1)}\right)$.
[TBR15] shows that, under Gaussian noise and with separation $c \geq 4$, a semidefinite
program can optimally estimate $x^{*}\left(t_{i}\right)$ at evenly spaced sample points $t_{i}$ from observations $x^{*}\left(t_{i}\right)+g\left(t_{i}\right)$. This is somewhat analogous to our setting, the differences being that (a) we want to estimate the signal over the entire interval, not just the sampled points, (b) our noise $g$ is adversarial, so we cannot hope to reduce it -if $g$ is also $k$-Fourier-sparse, we cannot distinguish $x^{*}$ and $g$, and of course (c) we want to avoid requiring frequency separation.

In [PS15], we gave the first algorithm with $O(1)$ approximation factor, finding the frequencies when $c \gtrsim \log (1 / \delta)$, and the signal when $c \gtrsim \log (1 / \delta)+\log ^{2} k$.

Now, all of the above results algorithms are designed to recover the frequencies; some of the ones in the noisy setting then show that this yields a good approximation to the overall signal (in the noiseless setting this is trivial). Such an approach necessitates $c \geq 1$ : [Moi15] gave a lower bound, showing that any algorithm finding the frequencies with approximation factor $2^{o(k)}$ must require $c \geq 1$.

Thus, in the current literature, we go from not knowing how to get any approximation for $c<1$, to getting a polynomial approximation at $c=1$ and a constant approximation at $c \gtrsim \log ^{2} k$. In this work, we show how to get a constant factor approximation to the signal regardless of $c$.

Polynomial interpolation. Our result is a generalization of robust polynomial interpolation, and in Theorem 12.1.2 we construct an optimal method for polynomial interpolation as a first step toward interpolating Fourier-sparse signals.

Our result here can be seen as essentially an extension of a technique shown in [CDL13]. The focus of [CDL13] is on the setting where sample points $x_{i}$ are chosen independently, so
$\Theta(d \log d)$ samples are necessary. One of their examples, however, shows essentially the same thing as our Corollary 12.4.2. From this, getting our theorem is not difficult.

The recent work [GZ16] looks at robust polynomial interpolation in a different noise model, featuring $\ell_{\infty}$ bounded noise with some outliers. In this setting they can get a stronger $\ell_{\infty}$ guarantee on the output than is possible in our setting.

Nyquist sampling. The classical method for learning bandlimited signals uses Nyquist sampling-i.e., samples at rate $1 / F$, for $F T$ points-and interpolates them using ShannonNyquist interpolation. This doesn't require any frequency gap, but also doesn't benefit from sparsity like sparse Fourier transform-based techniques. As discussed in [PS15], on the signal $x(t)=1$ it takes $F T+O(1 / \delta)$ samples to get $\delta$ error on average. Our dependence is logarithmic on both those terms.

### 12.1.2 Our techniques

Previous results on sparse Fourier transforms with robust recovery all required a frequency gap. So consider the opposite situation, where all the frequencies converge to zero and the coefficients are adjusted to keep the overall energy fixed. If we take a Taylor expansion of each complex exponential, then the signal will converge to a degree $k$ polynomial. So robust polynomial interpolation is a necessary subproblem for our algorithm.

Polynomial interpolation. Let $P(x)$ be a degree $d$ polynomial, and suppose that we can query $f(x)=P(x)+g(x)$ over the interval $[-1,1]$, where $g$ represents adversarial noise. We would like to query $f$ at $O(d)$ points and output a degree $d$ polynomial $Q(x)$ such that $\|P-Q\| \lesssim\|g\|$, where we define $\|h\|^{2}:=\int_{-1}^{1}|h(x)|^{2} \mathrm{~d} x$.

One way to do this would be to sample points $S \subset[-1,1]$ uniformly, then output the degree $d$ polynomial $Q$ with the smallest empirical error

$$
\|P+g-Q\|_{S}^{2}:=\frac{1}{|S|} \sum_{x \in S}|(P+g-Q)(x)|^{2}
$$

on the observed points. If $\|R\|_{S} \approx\|R\|$ for all degree $d$ polynomials $R$, in particular for $P-Q$, then since usually $\|g\|_{S} \lesssim\|g\|$ by Markov's inequality, the result follows.

This has two problems: first, uniform sampling is poor because polynomials like Chebyshev polynomials can have most of their energy within $O\left(1 / d^{2}\right)$ of the edges of the interval. This necessitates $\Omega\left(d^{2}\right)$ uniform samples before $\|R\|_{S} \approx\|R\|$ with good probability on a single polynomial. Second, the easiest method to extend from approximating one polynomial to approximating all polynomials uses a union bound over a net exponential in $d$, which would give an $O\left(d^{3}\right)$ bound.

To fix this, we need to bias our sampling toward the edges of the interval and we need our sampling to not be iid. We partition $[-1,1]$ into $O(d)$ intervals $I_{1}, \ldots, I_{n}$ so that the interval containing each $x$ has width at most $O\left(\sqrt{1-x^{2}}\right)$, except for the $O\left(1 / d^{2}\right)$ size regions at the edges. For any degree $d$ polynomial $R$ and any choice of $n$ points $x_{i} \in I_{i}$, the appropriately weighted empirical energy is close to $\|R\|$. This takes care of both issues with uniform sampling. If the points are chosen uniformly at random from within their intervals, then $\|g\|$ is probably bounded as well, and the empirically closest degree $d$ polynomial $Q$ will satisfy our requirements.

This result is shown in Section 12.4.

Clusters. Many previous sparse Fourier transform algorithms start with a one-sparse recovery algorithm, then show how to separate frequencies to get a $k$-sparse algorithm by reducing to the one-sparse case. Without a frequency gap, we cannot hope to reduce to the one-sparse case; instead, we reduce to individual clusters of nearby frequencies.

Essentially the problem is that one cannot determine all of the high-energy frequencies of a function $x$ only by sampling it on a bounded interval, as some of the frequencies might cancel each other out on this interval. We also cannot afford to work merely with the frequencies of the truncation of $x$ to the interval $[0, T]$, as the truncation operation will spread the frequencies of $x$ over too wide a range. To fix this problem, we must do something in between the two. In particular, we instead study $x \cdot H$ for a judiciously chosen function $H$. We want $H$ to approximate the indicator function of the interval $[0, T]$ and have small Fourier-support, $\operatorname{supp}(\widehat{H}) \subset\left[-k^{c} / T, k^{c} / T\right]$. By using some non-trivial lemmas about the growth rate of $x^{*}$, we can show that the difference between $x \cdot H$ on $\mathbb{R}$ and the truncation of
$x$ to $[0, T]$ has small $L^{2}$ mass, so that we can use the former as a substitute for the latter.
On the other hand, the Fourier transform of $x \cdot H$ is the convolution $\widehat{x} * \widehat{H}$, which has most of its mass within poly $(k) / T$ of the frequencies of $x^{*}$. Although it is impossible to determine the individual frequencies of $x^{*}$, we can hope to identify $O(k)$ intervals each of length $\operatorname{poly}(k) / T$ so that all but a small fraction of the energy of $\widehat{x}$ is contained within these intervals.

Note that many of these intervals will represent not individual frequencies of $x^{*}$, but small clusters of such frequencies. Furthermore, some frequencies of $x^{*}$ might not show up in these intervals either because they are too small, or because they cancel out other frequencies when convolved with $\widehat{H}$.

One-cluster recovery. Given our notion of clusters, we start looking at Fourier-sparse interpolation in the special case of one-cluster recovery. This is a generalization of onesparse recovery where we can have multiple frequencies, but they all lie in $[f-\Delta, f+\Delta]$ for some base frequency $f$ and bandwidth $\Delta=k^{c} / T$. Because all the frequencies are close to each other, values $x(a)$ and $x(a+\beta)$ will tend to have ratio close to $e^{2 \pi \mathrm{i} f \beta}$ when $\beta$ is small enough. We find that $\beta<\frac{1}{\Delta \sqrt{T \Delta}}$ is sufficient, which lets us figure out a frequency $\widetilde{f}$ with $|\tilde{f}-f| \leq \Delta \sqrt{T \Delta}=k^{O(1)} / T$.

Once we have the frequency $\tilde{f}$, we can consider $x^{\prime}(t)=x(t) e^{-2 \pi \mathbf{i} \tilde{f}}$. This signal is $k$-Fourier-sparse with frequencies bounded by $k^{O(1)} / T$. By taking a Taylor approximation to each complex exponential ${ }^{1}$, can show $x^{*}$ is $\delta$-close to $P(t) e^{2 \pi \mathbf{i} \tilde{f}}$ for a degree $d=O\left(k^{c}+\right.$

[^59]$k \log (1 / \delta))$ polynomial $P$. Thus we could apply our polynomial interpolation algorithm to recover the signal.
$k$-cluster frequency estimation. Reminiscent of algorithms such as [HIKP12a, PS15], we choose random variables $\sigma \approx T / k^{c}, a \in[0,1]$, and $b \in[0,1 / \sigma]$ and look at $v \in \mathbb{C}^{k^{c}}$ given by
$$
v_{i}=(x \cdot H)(\sigma(i-a)) e^{-2 \pi \mathbf{i} \sigma b i} G(i)
$$
where $G$ is a filter function. That is, $G$ has compact $\operatorname{support}\left(\operatorname{supp}(G) \subset\left[-k^{c}, k^{c}\right]\right)$, and $\widehat{G}$ approximates an interval of length $\Theta\left(\frac{2 \pi}{k}\right)$. In other words, $G$ is the same as $\widehat{H}$ with different parameters: an interval convolved with itself $k^{c}$ times, multiplied by a sinc function.

We alias $v$ down to $O(k)$ dimensions and take the discrete Fourier transform, getting $\widehat{u}$. It has been implicit in previous work - and we make it explicit-that $\widehat{u}_{j}$ is equal to $z_{\sigma a}$ for a vector $z$ defined by

$$
\widehat{z}=(\widehat{x} * \widehat{H}) \cdot \widehat{G}_{\sigma, b}^{(j)}
$$

where $\widehat{G}_{\sigma, b}^{(j)}$ is a particular permutation of $\widehat{G}$. In particular, $\widehat{G}_{\sigma, b}^{(j)}$ has period $1 / \sigma$, and approximates an interval of size $\frac{1}{\sigma B}$ within each period.

In previous work, when $\sigma$ and $b$ were chosen randomly, each individual frequency would have a good chance of being the only frequency preserved in $\widehat{z}$, and we could apply one-sparse recovery by choosing a variety of $a$. Without a frequency gap we can't quite say that: we pick $1 / \sigma \gg \Delta$ so that the entire cluster usually lands in the same bin, but then nearby clusters can also often land in the same bin. Fortunately, it is still usually true
with frequency gap $\delta 2^{-k}$, and coefficients at most $2^{k} / \delta$.
that only nearby clusters will collide. Since our 1-cluster algorithm works when the signal frequencies are nearby, we apply it to find a frequency approximation within $\frac{\sqrt{T / \sigma}}{\sigma}=k^{O(1)} / T$ of the cluster.

The above algorithm recovers each individual frequency with constant probability. By repeating it $O(\log k)$ times, with high probability we find a list $L$ of $O(k)$ frequencies within $k^{O(1)} / T$ of each significant cluster.
$k$-sparse recovery. Because different clusters aren't anywhere close to orthogonal, we can't simply approximate each cluster separately and add them up. Instead, given the list $L$ of candidate frequencies, we consider the $O(k d)$-dimensional space of functions

$$
\widetilde{x}(t):=\sum_{\tilde{f} \in L} \sum_{i=0}^{d} \alpha_{\tilde{f}, i} t^{i} e^{2 \pi \mathrm{i} \tilde{\mathrm{f}} t}
$$

where $d=O\left(k^{O(1)}+\log (1 / \delta)\right)$. We then take a bunch of random samples of $x$, and choose the $\widetilde{x}(t)$ minimizing the empirical error using linear regression. This regression can be made slightly faster using oblivious subspace embeddings [CW13], [NN13a], [Woo14b],[CNW15].

Our argument to show this works is analogous to the naive method we considered for polynomial recovery. Similarly to the one-cluster setting, using Taylor approximations and Gram determinants, we can show that this space includes a sufficiently close approximation to $x$. Since polynomials are the limit of sparse Fourier as frequencies tend to zero, these functions are arbitrarily close to $O(k d)$-Fourier-sparse functions. Hence we know that the maximum of $|\widetilde{x}(t)|$ is at most a poly $(k d)$ factor larger than its average over $[0, T]$. Using a net argument, this shows poly $(k d)$ samples are sufficient to find a good approximation to the nearest function in our space.

Growth rate of Fourier-sparse signals. We need that $\frac{1}{\sqrt{T}}\left\|x^{*} \cdot H\right\|_{2} \approx\left\|x^{*}\right\|_{T}$, where $H$ approximates the interval $1_{[0, T]}$. Because $H$ has support size $k^{c} / T$, it has a transition region of size $T / k^{c^{\prime}}$ at the edges, and it decays as $(t / T)^{-c^{c^{\prime \prime}}}$ for $t \gg T$. The difference between $\frac{1}{\sqrt{T}}\left\|x^{*} \cdot H\right\|_{2}$ and $\left\|x^{*}\right\|_{T}$ involves two main components: mass in the transition region that is lost, and mass outside the sampling interval that is gained. To show the approximation, we need that $\left|x^{*}(t)\right| \lesssim \widetilde{O}\left(k^{2}\right)\left\|x^{*}\right\|_{T}$ within the interval and $\left|x^{*}(t)\right| \lesssim(k t / T)^{O(k)}\left\|x^{*}\right\|_{T}$ outside.

We outline the bound of $\max _{t \in[0, T]}\left|x^{*}(t)\right|$ in terms of its average $\left\|x^{*}\right\|_{T}$ to bound $\left|x^{*}(t)\right|$ within the interval. Notice that we can assume $\left|x^{*}(0)\right|=\max _{t \in[0, T]}\left|x^{*}(t)\right|:$ if $t^{*}=\underset{t \in[0, T]}{\arg \max }\left|x^{*}(t)\right|^{2}$ is not 0 or $T$, we can rescale the two intervals $\left[0, t^{*}\right]$ and $\left[t^{*}, T\right]$ to $[0, T]$ separately. Then we show that for any $t^{\prime}$, there exist $m=\widetilde{O}\left(k^{2}\right)$ and constants $C_{1}, \cdots, C_{m}$ such that $x^{*}(0)=$ $\sum_{j \in[m]} C_{j} \cdot x^{*}\left(j \cdot t^{\prime}\right)$. Then we take the integration of $t^{\prime}$ over $[0, T / m]$ to bound $\left|x^{*}(0)\right|^{2}$ by its average. For any outside $t>T$, we follow this approach to show $x^{*}(t)=\sum_{j \in[k]} C_{j} \cdot x^{*}\left(t_{j}\right)$ where $t_{j} \in[0, T]$ and $\left|C_{j}\right| \leq \operatorname{poly}(k) \cdot(k t / T)^{O(k)}$ for each $j \in[k]$. These results are shown in Section 12.5.

### 12.1.3 Organization

This paper is organized as follows. We provide a brief overview about signal recovery in Section 12.2. We introduce some notations and tools in Section 12.3. Then we show our main Theorem 12.1.2 about polynomial interpolation in Section 12.4. For signals with $k$-sparse Fourier transform, we show two bounds on their growth rate in Section 12.5 and describe the hash functions and filter functions in Section 12.6. We provide the algorithm for frequency estimation and its proof in Section 12.7. In Section 12.8, we describe the algorithm for one-cluster recovery. In Section 12.9, we show the proof of Theorem 12.1.1. We defer several technical proofs in Section 12.10. Section 12.11 gives a summary of several well-known facts are existing in literature. We provide the analysis of hash functions and filter functions in Section 12.12.

### 12.2 Proof Sketch

We first consider one-cluster recovery centered at zero, i.e., $x^{*}(t)=\sum_{j=1}^{k} v_{j} \cdot e^{2 \pi \mathbf{i} f_{j} t}$ where every $f_{j}$ is in $[-\Delta, \Delta]$ for some small $\Delta>0$. The road map is to replace $x^{*}$ by a low degree polynomial $P$ such that $\left\|x^{*}(t)-P(t)\right\|_{T}^{2} \lesssim \delta\left\|x^{*}\right\|_{T}^{2}$ then recover a polynomial $Q$ to approximate $P$ through the observation $x(t)=P(t)+g^{\prime}(t)$ where $g^{\prime}(t)=g(t)+\left(x^{*}(t)-P(t)\right)$.

A natural way to replace $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathbf{i} f_{j} t}$ by a low degree polynomial $P(t)$ is the Taylor expansion. To bound the error after taking the low degree terms in the expansion by $\delta\left\|x^{*}\right\|_{T}$, we show the existence of $x^{\prime}(t)=\sum_{j=1}^{k} v_{j}^{\prime} e^{2 \pi \mathrm{i} f_{j}^{\prime} t}$ approximating $x^{*}$ on $[0, T]$ with an extra property-any coefficient $v_{j}^{\prime}$ in $x^{\prime}(t)$ has an upper bound in terms of $\left\|x^{\prime}\right\|_{T}^{2}=\frac{1}{T} \int_{0}^{T}\left|x^{\prime}(t)\right|^{2} \mathrm{~d} t$. We prove the existence of $x^{\prime}(t)$ via two more steps, both of which rely on the estimation of some Gram matrix constituted by these $k$ signals.

The first step is to show the existence of a $k$-Fourier-sparse signal $x^{\prime}(t)$ with frequency gap $\eta \geq \frac{\exp (-\operatorname{poly}(k)) \cdot \delta}{T}$ that is sufficiently close to $x^{*}(t)$.
Lemma 12.2.1. There is a universal constant $C_{1}>0$ such that, for any $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$ and any $\delta>0$, there always exist $\eta \geq \frac{\delta}{T} \cdot k^{-C_{1} k^{2}}$ and $x^{\prime}(t)=\sum_{j=1}^{k} v_{j}^{\prime} e^{2 \pi \mathrm{i} f_{j}^{\prime} t}$ satisfying

$$
\left\|x^{\prime}(t)-x^{*}(t)\right\|_{T} \leq \delta\left\|x^{*}(t)\right\|_{T}
$$

with $\min _{i \neq j}\left|f_{i}^{\prime}-f_{j}^{\prime}\right| \geq \eta$ and $\max _{j \in[k]}\left\{\left|f_{j}^{\prime}-f_{j}\right|\right\} \leq k \eta$.
We outline our approach and defer the proof to Section 12.8. We focus on the replacement of one frequency $f_{k}$ in $x^{*}=\sum_{j \in[k]} v_{j} e^{2 \pi \mathbf{i} f_{j} t}$ by a new frequency $f_{k+1} \neq f_{k}$ and its error. The idea is to consider every signal $e^{2 \pi \mathrm{i} f_{j} t}$ as a vector and prove that for any vector
$x^{*}$ in the linear subspace $\operatorname{span}\left\{e^{2 \pi \mathbf{i} f_{j} t} \mid j \in[k]\right\}$, there exists a vector in the linear subspace $\operatorname{span}\left\{e^{2 \pi \mathbf{i} f_{k+1} t}, e^{2 \pi \mathbf{i} f_{j} t} \mid j \in[k-1]\right\}$ with distance at most $\exp \left(k^{2}\right) \cdot\left(\left|f_{k}-f_{k+1}\right| T\right) \cdot\left\|x^{*}\right\|_{T}$ to $x^{*}$.

The second step is to lower bound $\left\|x^{\prime}\right\|_{T}^{2}$ by its coefficients through the frequency gap $\eta$ in $x^{\prime}$.
Lemma 12.2.2. There exists a universal constant $c>0$ such that for any $x(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$ with frequency gap $\eta=\min _{i \neq j}\left|f_{i}-f_{j}\right|$,

$$
\|x(t)\|_{T}^{2} \geq k^{-c k^{2}} \min \left((\eta T)^{2 k}, 1\right) \sum_{j=1}^{k}\left|v_{j}\right|^{2}
$$

Combining Lemma 12.2.1 and Lemma 12.2.2, we bound $\left|v_{j}^{\prime}\right|$ by $\exp (\operatorname{poly}(k)) \cdot \delta^{-O(k)}$. $\left\|x^{\prime}\right\|_{T}$ for any coefficient $v_{j}^{\prime}$ in $x^{\prime}$. Now we apply the Taylor expansion on $x^{\prime}(t)$ and keep the first $d=O\left(\Delta T+\operatorname{poly}(k)+k \log \frac{1}{\delta}\right)$ terms of every signal $v_{j}^{\prime} \cdot e^{2 \pi \mathrm{i} f_{j}^{\prime} t}$ in the expansion to obtain a polynomial $P(t)$ of degree at most $d$. To bound the distance between $P(t)$ and $x^{\prime}(t)$, we observe that the error of every point $t \in[0, T]$ is at most $\left(\frac{2 \pi \Delta \cdot T}{d}\right)^{d} \sum_{j}\left|v_{j}^{\prime}\right|$, which can be upper bounded by $\delta\left\|x^{\prime}(t)\right\|_{T}$ via the above connection. We summarize all discussion above as follows.

Lemma 12.2.3. For any $\Delta>0$ and any $\delta>0$, let $x^{*}(t)=\sum_{j \in[k]} v_{j} e^{2 \pi \mathbf{i} f_{j} t}$ where $\left|f_{j}\right| \leq \Delta$ for each $j \in[k]$. There exists a polynomial $P(t)$ of degree at most

$$
d=O\left(T \Delta+k^{3} \log k+k \log 1 / \delta\right)
$$

such that

$$
\left\|P(t)-x^{*}(t)\right\|_{T}^{2} \leq \delta\left\|x^{*}\right\|_{T}^{2}
$$

To recover $x^{*}(t)$, we observe $x(t)$ as a degree $d$ polynomial $P(t)$ with noise. We use properties of the Legendre polynomials to design a method of random sampling such that we only need $O(d)$ random samples to find a polynomial $Q(t)$ approximating $P(t)$.

Theorem 12.1.2. For any degree $d$ polynomial $P(t)$ and an arbitrary function $g(t)$, Procedure RobustPolynomialLearning in Algorithm 12.5 takes $O(d)$ samples from $x(t)=$ $P(t)+g(t)$ over $[0, T]$ and reports a degree $d$ polynomial $Q(t)$ in time $O\left(d^{\omega}\right)$ such that, with probability at least 99/100,

$$
\|P(t)-Q(t)\|_{T}^{2} \lesssim\|g(t)\|_{T}^{2}
$$

where $\omega<2.373$ is matrix multiplication exponent [Str69],[CW87],[Wil12].

We can either report the polynomial $Q(t)$ or transfer $Q(t)$ to a signal with $d$-sparse Fourier transform. We defer the technical proofs and the formal statements to Section 12.8 and discuss the recovery of $k$ clusters from now on.

As mentioned before, we apply the filter function $(H(t), \widehat{H}(f))$ on $x^{*}$ such that $\widehat{x^{*} \cdot H}$ has at most $k$ clusters given $\widehat{x^{*}}$ with $k$-sparse Fourier transform. First, we show that all frequencies in the "heavy" clusters of $\widehat{x^{*} \cdot H}$ constitute a good approximation of $x^{*}$ in Section 12.9 .

Definition 12.2.4. Given $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} \mathrm{f}_{j} t}$, any $\mathcal{N}>0$, and a filter function $(H, \widehat{H})$ with bounded support in frequency domain. Let $L_{j}$ denote the interval of $\operatorname{supp}\left(e^{2 \pi \mathrm{if}_{j} t} \cdot H\right)$ for each $j \in[k]$.

Define an equivalence relation $\sim$ on the frequencies $f_{i}$ by the transitive closure of the relation $f_{i} \sim f_{j}$ if $L_{i} \cap L_{j} \neq \emptyset$. Let $S_{1}, \ldots, S_{n}$ be the equivalence classes under this relation.

Define $C_{i}=\underset{f \in S_{i}}{\cup} L_{i}$ for each $i \in[n]$. We say $C_{i}$ is a "heavy" cluster iff $\int_{C_{i}}\left|\widehat{H \cdot x^{*}}(f)\right|^{2} \mathrm{~d} f \geq$ $T \cdot \mathcal{N}^{2} / k$.
Claim 12.2.5. Given $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i}_{j} t}$ and any $\mathcal{N}>0$, let $H$ be the filter function defined in Section 12.12 .1 and $C_{1}, \cdots, C_{l}$ be the heavy clusters from Definition 12.2.4. For

$$
S=\left\{j \in[k] \mid f_{j} \in C_{1} \cup \cdots C_{l}\right\}
$$

we have $x^{(S)}(t)=\sum_{j \in S} v_{j} e^{2 \pi \mathbf{i} f_{j} t}$ approximating $x^{*}$ within distance $\left\|x^{(S)}(t)-x^{*}(t)\right\|_{T}^{2} \lesssim \mathcal{N}^{2}$.
Hence it is enough to recover $x^{(S)}$ for the recovery of $x^{*}$. Let $\Delta_{h}$ denote the bandwidth of $\widehat{H}$. In Section 12.7, we choose $\Delta>k \cdot \Delta_{h}$ such that for any $j \in S, \int_{f_{j}-\Delta}^{f_{j}+\Delta}\left|\widehat{H \cdot x^{*}}(f)\right|^{2} \mathrm{~d} f \geq$ $T \cdot \mathcal{N}^{2} / k$ from the fact $\left|C_{i}\right| \leq k \cdot \Delta_{h}$. Then we prove Theorem 12.2.6 in Section 12.7, which finds $O(k)$ frequencies to cover all heavy clusters of $\widehat{x^{*} \cdot H}$.
Theorem 12.2.6. Let $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathbf{i} f_{j} t}$ and $x(t)=x^{*}(t)+g(t)$ be our observable signal where $\|g(t)\|_{T}^{2} \leq c\left\|x^{*}(t)\right\|_{T}^{2}$ for a sufficiently small constant $c$. Then Procedure FrequenCYRECOVERYKCLUSTER returns a set $L$ of $O(k)$ frequencies that covers all heavy clusters of $x^{*}$, which uses poly $(k, \log (1 / \delta)) \log (F T)$ samples and $\operatorname{poly}(k, \log (1 / \delta)) \log ^{2}(F T)$ time. In particular, for $\Delta=\operatorname{poly}(k, \log (1 / \delta)) / T$ and $\mathcal{N}^{2}:=\|g(t)\|_{T}^{2}+\delta\left\|x^{*}(t)\right\|_{T}^{2}$, with probability $1-2^{-\Omega(k)}$, for any $f^{*}$ with

$$
\begin{equation*}
\int_{f^{*}-\Delta}^{f^{*}+\Delta}|\widehat{x \cdot H}(f)|^{2} \mathrm{~d} f \geq T \mathcal{N}^{2} / k \tag{12.2}
\end{equation*}
$$

there exists an $\tilde{f} \in L$ satisfying

$$
\left|f^{*}-\widetilde{f}\right| \lesssim \Delta \sqrt{\Delta T}
$$

Let $L=\left\{\widetilde{f}_{1}, \cdots, \widetilde{f_{l}}\right\}$ be the list of frequencies from the output of Procedure FrequencyRecoveryKCluster in Theorem 12.2.6. The guarantee is that, for any $f_{j}$ in
$x^{(S)}$, there exists some $p_{j} \in[l]$ such that $\left|\widetilde{f}_{p_{j}}-f_{j}\right| \lesssim \Delta \sqrt{\Delta T}$ for $\Delta=\operatorname{poly}(k, \log (1 / \delta)) / T$. Hence we rewrite $x^{(S)}(t)=\sum_{i \in[l]} e^{2 \pi \mathbf{i} \widetilde{f}_{i} t}\left(\sum_{j \in S: p_{j}=i} e^{2 \pi \mathbf{i}\left(f_{j}-\widetilde{f}_{i}\right) t}\right)$. For each $i \in[l]$, we apply Lemma 12.2 .3 of one-cluster recovery on $\sum_{j \in S: p_{j}=i} e^{2 \pi \mathbf{i}\left(f_{j}-\widetilde{f}_{i}\right) t}$ to approximate it by a degree $d$ polynomial $P_{i}(t)$.

Now we consider $x(t)=\sum_{i \in[l]} e^{2 \pi \mathrm{i} \tilde{f}_{i} t} \cdot P_{i}(t)+g^{\prime \prime}(t)$ where $\left\|g^{\prime \prime}(t)\right\|_{T} \lesssim\|g(t)\|_{T}+$ $\delta\left\|x^{*}(t)\right\|_{T}$. To recover $\sum_{i \in[l]} e^{2 \pi \mathbf{i} \tilde{f}_{i} t} \cdot P_{i}(t)$, we treat it as a vector in the linear subspace

$$
V=\operatorname{span}\left\{e^{2 \pi \mathbf{i} \tilde{\mathrm{i}}_{i} t} \cdot t^{j} \mid j \in\{0, \cdots, d\}, i \in[l]\right\}
$$

with dimension at most $l(d+1)$ and find a vector in this linear subspace approximating it.
We show that for any $v \in V$, the average of poly $(k d)$ random samples on $v$ is enough to estimate $\|v\|_{T}^{2}$. In particular, any vector in this linear subspace satisfies that the maximum of it in $[0, T]$ has an upper bound in terms of its average in $[0, T]$. Then we apply the Chernoff bound to prove that poly $(k d)$ random samples are enough for the estimation of one vector $v \in V$.
Claim 12.2.7. For any $\vec{u} \in \operatorname{span}\left\{e^{2 \pi \tilde{\mathbf{i}}_{i} t} \cdot t^{j} \mid j \in\{0, \cdots, d\}, i \in[l]\right\}$, there exists some universal constants $C_{1} \leq 4$ and $C_{2} \leq 3$ such that

$$
\max _{t \in[0, T]}\left\{|\vec{u}(t)|^{2}\right\} \lesssim(l d)^{C_{1}} \log ^{C_{2}}(l d) \cdot\|\vec{u}\|_{T}^{2}
$$

At last we use an $\epsilon$-net to argue that poly $(k d)$ random samples from $[0, T]$ are enough to interpolate $x(t)$ by a vector $v \in V$. Because the dimension of this linear subspace is at most $l(d+1)=O(k d)$, there exists an $\epsilon$-net in this linear subspace for unit vectors with size at most $\exp (k d)$. Combining the Chernoff bound on all vectors in the $\epsilon$-net and Claim 12.2.7, we know that poly $(k d)$ samples are sufficient to estimate $\|v\|_{T}^{2}$ for any vector $v \in V$.

In Section 12.9, we show that a vector $v \in V$ minimizing the distance on poly $(k d)$ random samples is a good approximation for $\sum_{i \in[l]} e^{2 \pi \tilde{\mathrm{f}}_{i} t} \cdot P_{i}(t)$, which is a good approximation for $x^{*}(t)$ from all discussion above.

Theorem 12.1.1. Let $x(t)=x^{*}(t)+g(t)$, where $x^{*}$ is $k$-Fourier-sparse signal with frequencies in $[-F, F]$. Given samples of $x$ over $[0, T]$ we can output $\widetilde{x}(t)$ such that with probability at least $1-2^{-\Omega(k)}$,

$$
\left\|\widetilde{x}-x^{*}\right\|_{T} \lesssim\|g\|_{T}+\delta\left\|x^{*}\right\|_{T} .
$$

Our algorithm uses poly $(k, \log (1 / \delta)) \cdot \log (F T)$ samples and $\operatorname{poly}(k, \log (1 / \delta)) \cdot \log ^{2}(F T)$ time. The output $\widetilde{x}$ is $\operatorname{poly}(k, \log (1 / \delta))$-Fourier-sparse signal.

### 12.3 Preliminaries

We first provide some notations in Section 12.3.1 and basic Fourier facts in Section 12.3.2. Then we review some probability inequalities in Section 12.3.3. At last, we introduce Legendre polynomials in Section 12.3.4 and review some basic properties of Gram matrix and its determinant in Section 12.3.5.

### 12.3.1 Notation

For any function $f$, we define $\widetilde{O}(f)$ to be $f \cdot \log ^{O(1)}(f)$. We use $[n]$ to denote $\{1,2, \cdots, n\}$. Let $\mathbf{i}$ denote $\sqrt{-1}$. For any Complex number $z=a+\mathbf{i} b \in C$, where $a, b \in \mathbb{R}$. We define $\bar{z}$ to be $a-\mathbf{i} b$ and $|z|=\sqrt{a^{2}+b^{2}}$ such that $|z|^{2}=z \bar{z}$. For any function $f(t): \mathbb{R} \rightarrow \mathbb{C}$, we use $\operatorname{supp}(f)$ to denote the support of $f$.

For convenience, we define the sinc function and the Gaussian distribution Gaussian ${ }_{\mu, \sigma}$ on $\mathbb{R}$ with expectation $\mu$ and variance $\sigma^{2}$ as follows:

$$
\operatorname{sinc}(t)=\frac{\sin (\pi t)}{\pi t}, \quad \operatorname{Gaussian}_{\mu, \sigma}(t)=\frac{1}{\sigma \sqrt{2 \pi}} e^{-\frac{(t-\mu)^{2}}{2 \sigma^{2}}}
$$

For a fixed $T>0$, we define the inner product of two functions $x, y:[0, T] \rightarrow \mathbb{C}$ as

$$
\langle x, y\rangle_{T}=\frac{1}{T} \int_{0}^{T} x(t) \bar{y}(t) \mathrm{d} t
$$

We define the $\|\cdot\|_{T}$ norm as

$$
\|x(t)\|_{T}=\sqrt{\langle x(t), x(t)\rangle_{T}}=\sqrt{\frac{1}{T} \int_{0}^{T}|x(t)|^{2} \mathrm{~d} t}
$$

### 12.3.2 Facts about the Fourier transform

In this work, we always use $x(t)$ to denote a signal from $\mathbb{R} \rightarrow \mathbb{C}$. The Fourier transform $\widehat{x}(f)$ of an integrable function $x: \mathbb{R} \rightarrow \mathbb{C}$ is defined as

$$
\widehat{x}(f)=\int_{-\infty}^{+\infty} x(t) e^{-2 \pi \mathbf{i} f t} \mathrm{~d} t \text {, for any real number } f
$$

Similarly, $x(t)$ is determined from $\widehat{x}(f)$ by the inverse transform:

$$
x(t)=\int_{-\infty}^{+\infty} \widehat{x}(f) e^{2 \pi \mathbf{i} f t} \mathrm{~d} f, \text { for any real number } t
$$

Let CFT denote the continuous Fourier transform, DTFT denote the discrete-time Fourier transform, DFT denote the discrete Fourier transform, and FFT denote the fast Fourier transform.

For any signal $x(t)$ and $n \in \mathbb{N}_{+}$, we define $x^{* n}(t)=\underbrace{x(t) * \cdots * x(t)}_{n}$ and $\widehat{x}^{n}(f)=$ $\underbrace{\widehat{x}(f) \cdots \widehat{x}(f)}_{n}$.

Fact 12.3.1. Let $\delta_{\Delta}(f)$ denote the Dirac delta at $\Delta$. Then

$$
\widehat{\delta}_{\Delta}(t)=\int_{-\infty}^{+\infty} \delta_{\Delta}(f) e^{2 \pi \mathrm{i} f t} \mathrm{~d} f=e^{2 \pi \mathrm{i} t \Delta} .
$$

Fact 12.3.2. For any $s>0$, let $\operatorname{Comb}_{s}(t)=\sum_{j \in \mathbb{Z}} \delta_{j s}(t)$. Then the Fourier transform of $\operatorname{Comb}_{s}(t)$ is

$$
\widehat{\operatorname{Comb}}_{s}(f)=\frac{1}{s} \operatorname{Comb}_{1 / s}(f) .
$$

The following fact says the the Fourier transform of a rectangle function is a sinc function.


Figure 12.1: A picture of a $\mathrm{Comb}_{s}$, rect $_{s}, \operatorname{sinc}_{s}$, Gaussian $_{\mu, \sigma}$.

Fact 12.3.3. We use

$$
\operatorname{rect}_{s}(t)= \begin{cases}1 & \text { if }|t| \leq \frac{s}{2} \\ 0 & \text { otherwise }\end{cases}
$$

Then the Fourier transform of $\operatorname{rect}_{s}(t)$ is $\widehat{\operatorname{rect}}_{s}(f)=\frac{\sin (\pi f s)}{\pi f s}=\operatorname{sinc}(f s)$.

The Fourier transform of a Gaussian function is another Gaussian function.
Fact 12.3.4. For $\operatorname{Gaussian}_{\mu, \sigma}(t)=\frac{1}{\sigma \sqrt{2 \pi}} e^{-\frac{(t-\mu)^{2}}{2 \sigma^{2}}}$. Then the Fourier transform is

$$
\widehat{\operatorname{Gaussia}}_{\mu, \sigma}(f)=e^{-2 \pi \mathbf{i} f u} \frac{1}{\sigma \sqrt{2 \pi}} \operatorname{Gaussian}_{0, \sigma^{\prime}}(f) \text { for } \sigma^{\prime}=1 /(2 \pi \sigma)
$$

Proof. From the definition of the Fourier transform,

$$
\begin{aligned}
\widehat{\operatorname{Gaussian}}_{\mu, \sigma}(f) & =\int_{-\infty}^{+\infty} \frac{1}{\sigma \sqrt{2 \pi}} e^{-\frac{(t-\mu)^{2}}{2 \sigma^{2}}} e^{-2 \pi \mathbf{i} \mathrm{i} t} \mathrm{~d} t \\
& =e^{-2 \pi \mathbf{i} f u} \int_{-\infty}^{+\infty} \frac{1}{\sigma \sqrt{2 \pi}} e^{-\frac{t^{2}}{2 \sigma^{2}}} e^{-2 \pi \mathbf{i} f t} \mathrm{~d} t \\
& =e^{-2 \pi \mathbf{i} \mathrm{i} u} \int_{-\infty}^{+\infty} \frac{1}{\sigma \sqrt{2 \pi}} e^{-\frac{\left(t+2 \pi \mathbf{i} \boldsymbol{\sigma}^{2} f\right)^{2}}{2 \sigma^{2}}-2 \pi^{2} f^{2} \sigma^{2}} \mathrm{~d} t \\
& =e^{-2 \pi \mathbf{i} \mathrm{i} u} e^{-\frac{f^{2}}{2 \sigma^{\prime 2}}}
\end{aligned}
$$

where $\sigma^{\prime}=1 /(2 \sigma \pi)$, which is $e^{-2 \pi \mathbf{i} f u} \cdot \sigma^{\prime} \sqrt{2 \pi} \cdot \operatorname{Gaussian}_{0, \sigma^{\prime}}(f)$.

### 12.3.3 Tools and inequalities

From the Chernoff Bound (Lemma A.1.1), we show that if the maximum of a signal is bounded by $d$ times its energy over some fixed interval, then taking more than $d$ samples (each sample is drawn i.i.d. over that interval) suffices to approximate the energy of the signal on the interval with high probability.

Lemma 12.3.5. Given any function $x(t): \mathbb{R} \rightarrow \mathbb{C}$ with $\max _{t \in[0, T]}|x(t)|^{2} \leq d\|x(t)\|_{T}^{2}$. Let $S$ denote a set of points from 0 to $T$. If each point of $S$ is chosen uniformly at random from $[0, T]$, we have

$$
\left.\left.\operatorname{Pr}\left[\left.\left|\frac{1}{|S|} \sum_{i \in S}\right| x\left(t_{i}\right)\right|^{2}-\|x(t)\|_{T}^{2}\right] \right\rvert\, \geq \epsilon\|x(t)\|_{T}^{2}\right] \leq e^{-\Omega\left(\epsilon^{2}|S| / d\right)}
$$

We provide a proof in Section 12.10.5.
Because $d \cdot \frac{1}{2 d}+\frac{1}{2} \cdot\left(1-\frac{1}{2 d}\right) \leq 1$, we have the following inequality when the maximum of $|x(t)|^{2}$ is at most $d$ times its average.

Lemma 12.3.6. Given any function $x(t): \mathbb{R} \rightarrow \mathbb{C}$ with $\max _{t \in[0, T]}|x(t)|^{2} \leq d\|x(t)\|_{T}^{2}$. Let $S$ denote a set of points from 0 to $T$. For any point $a$ is sampled uniformly at random from $[0, T]$, we have,

$$
\operatorname{Pr}_{a \sim[0, T]}\left[|x(a)|^{2} \geq \frac{1}{2}\|x(t)\|_{T}^{2}\right] \geq \frac{1}{2 d}
$$

### 12.3.4 Legendre polynomials

We provide an brief introduction to Legendre polynomials (please see [Dun10] for a complete introduction). For convenience, we fix $\|f(t)\|_{T}^{2}=\frac{1}{2} \int_{-1}^{1}|f(t)|^{2} \mathrm{~d} t$ in this section.

Definition 12.3.1. Let $L_{n}(x)$ denote the Legendre polynomials of degree $n$, the solution to Legendre's differential equation:

$$
\begin{equation*}
\frac{\mathrm{d}}{\mathrm{~d} x}\left[\left(1-x^{2}\right) \frac{\mathrm{d}}{\mathrm{~d} x} L_{n}(x)\right]+n(n+1) L_{n}(x)=0 \tag{12.3}
\end{equation*}
$$

We will the following two facts about the Legendre polynomials in this work.
Fact 12.3.7. $L_{n}(1)=1$ for any $n \geq 0$ in the Legendre polynomials.
Fact 12.3.8. The Legendre polynomials constitute an orthogonal basis with respect to the inner product on interval $[-1,1]$ :

$$
\int_{-1}^{1} L_{m}(x) L_{n}(x) \mathrm{d} x=\frac{2}{2 n+1} \delta_{m n}
$$

where $\delta_{m n}$ denotes the Kronecker delta, i.e., it equals to 1 if $m=n$ and to 0 otherwise.

For any polynomial $P(x)$ of degree at most $d$ with complex coefficients, there exists a set of coefficients from the above properties such that

$$
P(x)=\sum_{i=0}^{d} \alpha_{i} \cdot L_{i}(x), \text { where } \alpha_{i} \in \mathbb{C}, \forall i \in\{0,1,2, \cdots, d\}
$$

Lemma 12.3.9. For any polynomial $P(t)$ of degree at most $d$ from $R$ to $\mathbb{C}$, for any interval $[S, T]$,

$$
\max _{t \in[S, T]}|P(t)|^{2} \leq(d+1)^{2} \cdot \frac{1}{T-S} \int_{S}^{T}|P(t)|^{2} \mathrm{~d} x
$$

We provide a proof in Section 12.10.6.

### 12.3.5 Gram matrix and its determinant

We provide an brief introduction to Gramian matrices (please see [Haz01] for a complete introduction). We use $\langle x, y\rangle$ to denote the inner product between vector $x$ and vector $y$.

Let $\vec{v}_{1}, \cdots, \vec{v}_{n}$ be $n$ vectors in an inner product space and $\operatorname{span}\left\{\vec{v}_{1}, \cdots, \vec{v}_{n}\right\}$ be the linear subspace spanned by these $n$ vectors with coefficients in $\mathbb{C}$, i.e.,

$$
\left\{\sum_{i \in[n]} \alpha_{i} \vec{v}_{i} \mid \forall i \in[n], \alpha_{i} \in \mathbb{C}\right\}
$$

The Gram matrix $\operatorname{Gram}_{n}$ of $\vec{v}_{1}, \cdots, \vec{v}_{n}$ is an $n \times n$ matrix defined as $\operatorname{Gram}_{n}(i, j)=\left\langle\vec{v}_{i}, \vec{v}_{j}\right\rangle$ for any $i \in[n]$ and $j \in[n]$.

Fact 12.3.10. $\operatorname{det}\left(\operatorname{Gram}_{n}\right)$ is the square of the volume of the parallelotope formed by $\vec{v}_{1}, \cdots, \vec{v}_{n}$.

Let $\operatorname{Gram}_{n-1}$ be the Gram matrix of $\vec{v}_{1}, \cdots, \vec{v}_{n-1}$. Let $\vec{v}_{n}^{\|}$be the projection of $v_{n}$ onto the linear subspace span $\left\{\vec{v}_{1}, \cdots, \vec{v}_{n-1}\right\}$ and $\vec{v}_{n}^{\perp}=\vec{v}_{n}-\vec{v}_{n}^{\|}$. We use $\|\vec{v}\|$ to denote the length of $\vec{v}$ in the inner product space, which is $\sqrt{\langle\vec{v}, \vec{v}\rangle}$.

## Claim 12.3.11.

$$
\left\|\vec{v}_{n}^{\perp}\right\|^{2}=\frac{\operatorname{det}\left(\operatorname{Gram}_{n-1}\right)}{\operatorname{det}\left(\operatorname{Gram}_{n}\right)}
$$

Proof.

$$
\operatorname{det}\left(\operatorname{Gram}_{n}\right)=\operatorname{volume}{ }^{2}\left(\vec{v}_{1}, \cdots, \vec{v}_{n}\right)=\operatorname{volume}{ }^{2}\left(\vec{v}_{1}, \cdots, \vec{v}_{n-1}\right) \cdot\left\|\vec{v}_{n}^{\perp}\right\|^{2}=\operatorname{det}\left(\operatorname{Gram}_{n}\right) \cdot\left\|\vec{v}_{n}^{\perp}\right\|^{2} .
$$

### 12.4 Robust Polynomial Interpolation Algorithm

In Section 12.4.1, we show how to learn a low degree polynomial by using linear number of samples, running polynomial time, and achieving constant success probability. In Section 12.4.2, we show to how boost the success probability by rerunning previous algorithm several times.

### 12.4.1 Constant success probability

We show how to learn a degree- $d$ polynomial $P$ with $n=O(d)$ samples and prove Theorem 12.1.2 in this section. For convenience, we first fix the interval to be $[-1,1]$ and use $\|f\|_{[-1,1]}^{2}=\frac{1}{2} \int_{-1}^{1}|f(t)|^{2} \mathrm{~d} t$.

Lemma 12.4.1. Let $d \in \mathbb{N}$ and $\epsilon \in \mathbb{R}^{+}$, there exists an efficient algorithm to compute $a$ partition of $[-1,1]$ to $n=O(d / \epsilon)$ intervals $I_{1}, \cdots, I_{n}$ such that for any degree $d$ polynomial $P(t): \mathbb{R} \rightarrow \mathbb{C}$ and any $n$ points $x_{1}, \cdots, x_{n}$ in the intervals $I_{1}, \cdots, I_{n}$ respectively, the function $Q(t)$ defined by

$$
Q(t)=P\left(x_{j}\right) \quad \text { if } \quad t \in I_{j}
$$

approximates $P$ by

$$
\begin{equation*}
\|Q-P\|_{[-1,1]} \leq \epsilon\|P\|_{[-1,1]} . \tag{12.4}
\end{equation*}
$$

One direct corollary from the above lemma is that observing $n=O(d / \epsilon)$ points each from $I_{1}, \cdots, I_{n}$ provides a good approximation for all degree $d$ polynomials. For any set $S=\left\{t_{1}, \cdots, t_{m}\right\}$ where each $t_{i} \in[-1,1]$ and a distribution with support $\left\{w_{1}, \cdots, w_{m}\right\}$ on $S$ where $\sum_{i=1}^{m} w_{i}=1$ and $w_{i} \geq 0$ for each $i \in[m]$, we define $\|x\|_{S, w}=\left(\sum_{i=1}^{m} w_{i} \cdot\left|x\left(t_{i}\right)\right|^{2}\right)^{1 / 2}$.

Corollary 12.4.2. Let $I_{1}, \cdots, I_{n}$ be the intervals in the above lemma and $w_{j}=\left|I_{j}\right| / 2$ for each $j \in[n]$. For any $x_{1}, \cdots, x_{n}$ in the intervals $I_{1}, \cdots, I_{n}$ respectively, we consider $S=$ $\left\{x_{1}, \cdots, x_{n}\right\}$ with the distribution $w_{1}, \cdots, w_{n}$. Then for any degree d polynomial $P$, we have

$$
\|P\|_{S, w} \in\left[(1-\epsilon)\|P\|_{[-1,1]},(1+\epsilon)\|P\|_{[-1,1]}\right]
$$

We first state the main technical lemma and finish the proof of the above lemma (we defer the proof of Lemma 12.4.3 to Section 12.10.3).

Lemma 12.4.3. For any degree $d$ polynomial $P(t): \mathbb{R} \rightarrow \mathbb{C}$ with derivative $P^{\prime}(t)$, we have,

$$
\begin{equation*}
\int_{-1}^{1}\left(1-t^{2}\right)\left|P^{\prime}(t)\right|^{2} \mathrm{~d} t \leq 2 d^{2} \int_{-1}^{1}|P(t)|^{2} \mathrm{~d} t \tag{12.5}
\end{equation*}
$$

Proof of Lemma 12.4.1. We set $m=10 d / \epsilon$ and show a partition of $[-1,1]$ into $n \leq 20 m$ intervals. We define $g(t)=\frac{\sqrt{1-t^{2}}}{m}$ and $y_{0}=0$. Then we choose $y_{i}=y_{i-1}+g\left(y_{i-1}\right)$ for $i \in \mathbb{N}^{+}$. Let $l$ be the first index of $y$ such that $y_{l} \geq 1-\frac{9}{m^{2}}$. We show $l \lesssim m$.

Let $j_{k}$ be the first index in the sequence such that $y_{j_{k}} \geq 1-2^{-k}$. Notice that

$$
j_{2} \leq \frac{3 / 4}{\frac{\sqrt{1-(3 / 4)^{2}}}{m}} \leq 1.5 m
$$

and

$$
y_{i}-y_{i-1}=g\left(y_{i-1}\right)=\frac{\sqrt{1-y_{i-1}^{2}}}{m} \geq \frac{\sqrt{1-y_{i-1}}}{m} .
$$

Then for all $k>2$, we have

$$
j_{k}-j_{k-1} \leq \frac{2^{-k}}{\frac{\sqrt{1-y_{\left(j_{k}-1\right)}}}{m}} \leq 2^{-k / 2} m
$$

Therefore $j_{k} \leq\left(1.5+\left(2^{-3 / 2}+\cdots 2^{-k / 2}\right)\right) m$ and $l \leq 10 m$.
Because $y_{l-1} \leq 1-\frac{9}{m^{2}}$, for any $j \in[l]$ and any $x \in\left[y_{i-1}, y_{i}\right]$, we have the following property:

$$
\begin{equation*}
\frac{1-x^{2}}{m^{2}} \geq \frac{1}{2} \cdot \frac{\left(1-y_{i-1}^{2}\right)}{m^{2}}=\left(y_{i}-y_{i-1}\right)^{2} / 2 \tag{12.6}
\end{equation*}
$$

Now we set $n$ and partition $[-1,1]$ into $I_{1}, \cdots, I_{n}$ as follows:

1. $n=2(l+1)$.
2. For $j \in[l], I_{2 j-1}=\left[y_{j-1}, y_{j}\right]$ and $I_{2 j}=\left[-y_{j},-y_{j-1}\right]$.
3. $I_{2 l+1}=\left[y_{l}, 1\right]$ and $I_{2 l+2}=\left[-1,-y_{l}\right]$.

For any $x_{1}, \cdots, x_{n}$ where $x_{j} \in I_{j}$ for each $j \in[n]$, we rewrite the LHS of (12.4) as follows:

$$
\begin{equation*}
\underbrace{\sum_{j=1}^{n-2} \int_{I_{j}}\left|P\left(x_{j}\right)-P(t)\right|^{2} \mathrm{~d} t}_{A}+\underbrace{\int_{I_{n-1}}\left|P\left(x_{n-1}\right)-P(t)\right|^{2} \mathrm{~d} t+\int_{I_{n}}\left|P\left(x_{n}\right)-P(t)\right|^{2} \mathrm{~d} t}_{B} \tag{12.7}
\end{equation*}
$$

For A in Equation (12.7), from the Cauchy-Schwarz inequality, we have

$$
\sum_{j=1}^{n-2} \int_{I_{j}}\left|P\left(x_{j}\right)-P(t)\right|^{2} \mathrm{~d} t=\sum_{j=1}^{n-2} \int_{I_{j}}\left|\int_{x_{j}}^{t} P^{\prime}(y) \mathrm{d} y\right|^{2} \mathrm{~d} t \leq \sum_{j=1}^{n-2} \int_{I_{j}}\left|t-x_{j}\right| \int_{x_{j}}^{t}\left|P^{\prime}(y)\right|^{2} \mathrm{~d} y \mathrm{~d} t
$$

Then we swap $\mathrm{d} t$ with $\mathrm{d} y$ and use Equation (12.6):

$$
\sum_{j=1}^{n-2} \int_{I_{j}}\left|P^{\prime}(y)\right|^{2} \int_{t \notin\left(x_{j}, y\right)}\left|t-x_{j}\right| \mathrm{d} t \mathrm{~d} y \leq \sum_{j=1}^{n-2} \int_{I_{j}}\left|P^{\prime}(t)\right|^{2} \cdot\left|I_{j}\right|^{2} \mathrm{~d} t \leq \sum_{j=1}^{n-2} \int_{I_{j}}\left|P^{\prime}(t)\right|^{2} \frac{2\left(1-t^{2}\right)}{m^{2}} \mathrm{~d} t
$$

We use Lemma 12.4.3 to simplify it by

$$
\sum_{j=1}^{n-2} \int_{I_{j}}\left|P\left(x_{j}\right)-P(t)\right|^{2} \mathrm{~d} t \leq \int_{-1}^{1}\left|P^{\prime}(t)\right|^{2} \frac{2\left(1-t^{2}\right)}{m^{2}} \mathrm{~d} t \leq \frac{2 d^{2}}{m^{2}} \int_{-1}^{1}|P(t)|^{2} \mathrm{~d} t
$$

For B in Equation (12.7), notice that $\left|I_{n-1}\right|=\left|I_{n}\right|=1-y_{l} \leq 9 m^{-2}$ and for $j \in$ $\{n-1, n\}$

$$
\left|P(t)-P\left(x_{j}\right)\right|^{2} \leq 4 \max _{t \in[-1,1]}|P(t)|^{2} \leq 4(d+1)^{2}\|P\|_{[-1,1]}^{2}
$$

from the properties of degree- $d$ polynomials, i.e., Lemma 12.3.9. Therefore B in Equation (12.7) is upper bounded by $2 \cdot 4(d+1)^{2}\left(9 m^{-2}\right)\|P(t)\|_{[-1,1]}^{2}$.

From all discussion above, $\|Q(t)-P(t)\|_{[-1,1]}^{2} \leq \frac{99 d^{2}}{m^{2}} \leq \epsilon^{2}$.

Now we use the above lemma to provide a faster learning algorithm for polynomials on interval $[-1,1]$ with noise instead of using the $\epsilon$-nets argument. Algorithm RobustPolynomialLearningFixedinterval works as follows:

1. Let $\epsilon=1 / 20$ and $I_{1}, \cdots, I_{n}$ be the intervals for $d$ and $\epsilon$ in Lemma 12.4.1.
2. Random choose $x_{j} \in I_{j}$ for every $j \in[n]$ and define $S=\left\{x_{1}, \cdots, x_{n}\right\}$ with weight $w_{1}=\frac{\left|I_{1}\right|}{2}, \cdots, w_{n}=\frac{\left|I_{n}\right|}{2}$.
3. Find the degree $d$ polynomial $Q(t)$ that minimizes $\|P(t)-Q(t)\|_{S, w}$ using Fact 12.11.1.

Lemma 12.4.4. For any degree d polynomial $P(t)$ and an arbitrary function $g(t)$, Algorithm RobustPolynomialLearningFixedInterval takes $O(d)$ samples from $x(t)=$ $P(t)+g(t)$ over $[-1,1]$ and reports a degree d polynomial $Q(t)$ in time $O\left(d^{\omega}\right)$ such that, with probability at least 99/100,

$$
\|P(t)-Q(t)\|_{[-1,1]}^{2} \lesssim\|g(t)\|_{[-1,1]}^{2} .
$$

Proof. Notice that $n=O(d / \epsilon)=O(d)$ and the running time depends on solving a linear regression problem( Fact 12.11.1 ), which takes $O\left(d^{\omega}\right)$ time. It is enough to bound the
distance between $P$ and $Q$ :

$$
\begin{aligned}
& \|P-Q\|_{[-1,1]} \\
\leq & 1.09\|P-Q\|_{S, w} \\
=1.09\|x-g-Q\|_{S, w} & \text { by Corollary } 12.4 .2 \\
\leq 1.09\|g\|_{S, w}+1.09\|x-Q\|_{S, w} & \text { by } x=P+g \\
\leq 1.09\|g\|_{S, w}+1.09\|x-P\|_{S, w} & \text { by triangle inequality } \\
\leq & 2.2\|g\|_{S, w}
\end{aligned}
$$

Because $\underset{S}{\mathbb{E}}\left[\|g\|_{S, w}^{2}\right]=\|g\|_{[-1,1]}^{2}$, we know that $\|P-Q\|_{[-1,1]} \leq 2200\|g\|_{[-1,1]}$ with probability $\geq .999$ by using Markov's inequality.

For any function $f:[0, T] \rightarrow \mathbb{C}$, let $\widetilde{f}(t)=f\left(\frac{2 t-T}{T}\right)$. Then $\|\widetilde{f}\|_{[-1,1]}=\|f\|_{T}$ from the definition. Hence we can switch any interval $[0, T]$ to $[-1,1]$ and use Lemma 12.4.4.

Theorem 12.1.2. For any degree $d$ polynomial $P(t)$ and an arbitrary function $g(t)$, Procedure RobustPolynomialLearning in Algorithm 12.5 takes $O(d)$ samples from $x(t)=$ $P(t)+g(t)$ over $[0, T]$ and reports a degree $d$ polynomial $Q(t)$ in time $O\left(d^{\omega}\right)$ such that, with probability at least 99/100,

$$
\|P(t)-Q(t)\|_{T}^{2} \lesssim\|g(t)\|_{T}^{2}
$$

where $\omega<2.373$ is matrix multiplication exponent [Str69],[CW87],[Wil12].

### 12.4.2 Boosting success probability

Notice that the success probability of Theorem 12.1.2 is only constant, and the proof technique of obtaining that result cannot be modified to $1-1 / \operatorname{poly}(d)$ or $1-2^{-\Omega(d)}$ success probability due to using Markov's inequality. However, we can use that algorithm as a black box, and rerun it $O(\log (1 / p))$ (for any $p>0$ ) times on fresh samples. Using the careful median analysis from [MP14] gives

Theorem 12.4.5. For any degree $d$ polynomial $P(t)$, an arbitrary function $g(t)$, and any $p>0$, Procedure RobustPolynomialLearning ${ }^{+}$in Algorithm 12.5 takes $O(d \log (1 / p))$ samples from $x(t)=P(t)+g(t)$ over $[0, T]$ and reports a degree $d$ polynomial $Q(t)$ in time $O\left(d^{\omega} \log (1 / p)\right)$ such that, with probability at least $1-p$,

$$
\|P(t)-Q(t)\|_{T}^{2} \lesssim\|g(t)\|_{T}^{2}
$$

where $\omega<2.373$ is matrix multiplication exponent.

Proof. We run algorithm RobustPolynomialLearning $R$ rounds with $O(d)$ independent and fresh samples per round. We will obtain $R$ degree- $d$ polynomials $Q_{1}(t), Q_{2}(t), \cdots, Q_{R}(t)$. We say a polynomial $Q_{i}(t)$ is good if $\left\|Q_{i}(t)-P(t)\right\|_{T}^{2} \lesssim\|g(t)\|_{T}^{2}$. Using the Chernoff bound, with probability at least $1-2^{-\Omega(R)}$, at least a $3 / 4$ fraction of the polynomials are "good". We output polynomial $Q(t)=Q_{j^{*}}(t)$ such that

$$
\begin{equation*}
j^{*}=\underset{j \in[R]}{\arg \min }\left(\operatorname{median}\left\{\left\|Q_{j}(t)-Q_{1}(t)\right\|_{T}^{2},\left\|Q_{j}(t)-Q_{2}(t)\right\|_{T}^{2}, \cdots,\left\|Q_{j}(t)-Q_{R}(t)\right\|_{T}^{2}\right\}\right) \tag{12.8}
\end{equation*}
$$

The Equation (12.8) can be solved in following straightforward way. For $i \neq j$, it takes $O(d)$ time to compute $\left\|Q_{j}(t)-Q_{i}(t)\right\|_{T}^{2}$. Because of the number of pairs is $O\left(R^{2}\right)$, thus it takes
$O\left(R^{2} d\right)$ time write down a $R \times R$ matrix. For each column, we run linear time 1-median algorithm. This step takes $O\left(R^{2}\right)$ time. At the end, $j^{*}$ is index of the column that has the smallest median value. Thus, polynomial $Q(t)=Q_{j^{*}}(t) 1$ the 0 with probability at least $1-p$ by choosing $R=O(\log (1 / p))$. The running time is not optimized yet.

To improve the dependence on $R$ for running time, we replace the step of solving Equation (12.8) by an approach that is similar to [MP14]. We choose a new set of samples $S$, say $S=\left\{t_{1}, t_{2}, \cdots, t_{n}\right\}$ and $n=O(d)$. Using Fact 12.11.2, we can compute $Q_{i}\left(t_{j}\right)$ for all $i, j \in[R] \times[n]$ in $O(R d \operatorname{poly}(\log (d)))$ time. Define

$$
\begin{equation*}
\widetilde{Q}_{j}=\underset{i \in[R]}{\operatorname{median}} Q_{i}\left(t_{j}\right), \forall j \in[n] . \tag{12.9}
\end{equation*}
$$

Our algorithm will output a degree- $d$ polynomial $Q$ which is the optimal solution of this problem, $\min _{\text {degree-d } Q^{\prime \prime}}\left\|Q^{\prime}-\widetilde{Q}\right\|_{S, w} \cdot{ }^{2}$ In the rest of the proof, we will show that $\|Q-P\|_{T} \lesssim\|g\|_{T}$ with probability at least $1-2^{-\Omega(R)}$.

Notice that Equation (12.9) implies that $\widetilde{Q}_{j}-P\left(t_{j}\right)=\underset{i \in[R]}{\operatorname{median}}\left(Q_{i}\left(t_{j}\right)-P\left(t_{j}\right)\right)$. Fix a coordinate $j$ and applying the proof argument of Lemma 6.1 in [MP14], we have

$$
\left(\widetilde{Q}_{j}-P\left(t_{j}\right)\right)^{2} \lesssim \operatorname{mean}_{\text {good } i}\left(Q_{i}\left(t_{j}\right)-P\left(t_{j}\right)\right)^{2}
$$

Taking the weighted summation over all the coordinates $j$, we have

$$
\|\widetilde{Q}-P\|_{S, w}^{2} \lesssim \operatorname{mean}_{\text {good } i}\left\|Q_{i}-P\right\|_{S, w}^{2}
$$

Using Corollary 12.4.2, for each good $i$,

$$
\left\|Q_{i}-P\right\|_{S, w}^{2} \lesssim\left\|Q_{i}-P\right\|_{T}^{2}
$$

[^60]Combining the above two inequalities gives

$$
\begin{equation*}
\|\widetilde{Q}-P\|_{S, w}^{2} \lesssim \operatorname{mean}_{\text {good } i}\left\|Q_{i}-P\right\|_{T}^{2} \lesssim\|g\|_{T}^{2} \tag{12.10}
\end{equation*}
$$

Because $Q$ is the optimal solution for $\widetilde{Q}$, then

$$
\begin{equation*}
\|\widetilde{Q}-Q\|_{S, w}^{2} \leq\|\widetilde{Q}-P\|_{S, w}^{2} \lesssim\|g\|_{T}^{2} \tag{12.11}
\end{equation*}
$$

Using Corollary 12.4.2 and for any good $i, i^{\prime},\left\|Q_{i}-Q_{i^{\prime}}\right\|_{T} \lesssim\|g\|_{T}$, we can replace $P$ by $Q_{i^{\prime}}$ in the Equation (12.10). Thus, for any $Q_{i^{\prime}}$ where $i^{\prime}$ is good,

$$
\begin{equation*}
\left\|\widetilde{Q}-Q_{i^{\prime}}\right\|_{S, w}^{2} \lesssim\|g\|_{T}^{2} \tag{12.12}
\end{equation*}
$$

For any good $i^{\prime}$,

$$
\begin{aligned}
& \left\|Q_{i^{\prime}}-Q\right\|_{T} \\
\lesssim & \left\|Q_{i^{\prime}}-Q\right\|_{S, w} \\
\leq & \text { by Corollary } 12.4 .2 \\
\lesssim & \|g\|_{i^{\prime}}-\widetilde{Q}\left\|_{S, w}+\right\| \widetilde{Q}-Q \|_{S, w}
\end{aligned} \quad \text { by triangle inequality }
$$

Thus, our algorithm takes $O(d R)$ samples from $x(t)=P(t)+g(t)$ over $[0, T]$ and reports a polynomial $Q(t)$ in time $O\left(R d^{\omega}\right)$ such that, with probability at least $1-2^{-\Omega(R)}, \| P(t)-$ $Q(t)\left\|_{T}^{2} \lesssim\right\| g(t) \|_{T}^{2}$. Choosing $R=O(\log (1 / p))$ completes the proof.

### 12.5 Bounding the Magnitude of a Fourier-sparse Signal in Terms of Its Average Norm

The main results in this section are two upper bounds, Lemma 12.5.1 on $\max _{t \in[0, T]}|x(t)|^{2}$ and Lemma 12.5.5 on $|x(t)|^{2}$ for $t>T$, in terms of the typical signal value $\|x\|_{T}^{2}=$ $\frac{1}{T} \int_{0}^{T}|x(t)|^{2} \mathrm{~d} t$. We prove Lemma 12.5.1 in Section 12.5 .1 and Lemma 12.5.5 in Section 12.5.2

### 12.5.1 Bounding the maximum inside the interval

The goal of this section is to prove Lemma 12.5.1.
Lemma 12.5.1. For any $k$-Fourier-sparse signal $x(t): \mathbb{R} \rightarrow \mathbb{C}$ and any duration $T$, we have

$$
\max _{t \in[0, T]}|x(t)|^{2} \lesssim k^{4} \log ^{3} k \cdot\|x\|_{T}^{2}
$$

Proof. Without loss of generality, we fix $T=1$. Then $\|x\|_{T}^{2}=\int_{0}^{1}|x(t)|^{2} \mathrm{~d} t$. Because $\|x\|_{T}^{2}$ is the average over the interval $[0, T]$, if $t^{*}=\underset{t \in[0, T]}{\arg \max }|x(t)|^{2}$ is not 0 or $T=1$, we can rescale the two intervals $\left[0, t^{*}\right]$ and $\left[t^{*}, T\right]$ to $[0,1]$ and prove the desired property separately. Hence we assume $|x(0)|^{2}=\max _{t \in[0, T]}|x(t)|^{2}$ in this proof.

Claim 12.5.2. For any $k$, there exists $m=O\left(k^{2} \log k\right)$ such that for any $k$-Fourier-sparse signal $x(t)$, any $t_{0} \geq 0$ and $\tau>0$, there always exist $C_{1}, \cdots, C_{m} \in \mathbb{C}$ such that the following properties hold,

$$
\begin{array}{cl}
\text { Property I } & \left|C_{j}\right| \leq 11 \text { for all } j \in[m] \\
\text { Property II } & x\left(t_{0}\right)=\sum_{j \in[m]} C_{j} \cdot x\left(t_{0}+j \cdot \tau\right)
\end{array}
$$

We first use this claim to finish the proof of Lemma 12.5.1. We choose $t_{0}=0$ such that $\forall \tau>0$, there always exist $C_{1}, \cdots, C_{m} \in \mathbb{C}$, and

$$
x(0)=\sum_{j \in[m]} C_{j} \cdot x(j \cdot \tau) .
$$

By the Cauchy-Schwarz inequality, it implies that for any $\tau$,

$$
\begin{align*}
|x(0)|^{2} & \leq m \sum_{j \in[m]}\left|C_{j}\right|^{2}|x(j \cdot \tau)|^{2} \\
& \lesssim m \sum_{j \in[m]}|x(j \cdot \tau)|^{2} \tag{12.13}
\end{align*}
$$

At last, we obtain

$$
\begin{aligned}
|x(0)|^{2} & =m \int_{0}^{1 / m}|x(0)|^{2} \mathrm{~d} \tau \\
& \lesssim m \cdot \int_{0}^{1 / m}\left(m \sum_{j=1}^{m}|x(j \cdot \tau)|^{2}\right) \mathrm{d} \tau \\
& =m^{2} \cdot \sum_{j=1}^{m} \int_{0}^{1 / m}|x(j \cdot \tau)|^{2} \mathrm{~d} \tau \\
& =m^{2} \cdot \sum_{j=1}^{m} \frac{1}{j} \int_{0}^{j / m}|x(\tau)|^{2} \mathrm{~d} \tau \\
& \leq m^{2} \cdot \sum_{j=1}^{m} \frac{1}{j} \cdot \int_{0}^{1}|x(\tau)|^{2} \mathrm{~d} \tau \\
& \lesssim m^{2} \log m \cdot\|x\|_{T}^{2}
\end{aligned}
$$

where the first inequality follows by Equation (12.13), the second inequality follows by $j / m \leq$ 1 and the last step follows by $\sum_{i=1}^{m} \frac{1}{i}=O(\log m)$. From $m=O\left(k^{2} \log k\right)$, we obtain $|x(0)|^{2}=$ $O\left(k^{4} \log ^{3} k\|x\|_{T}^{2}\right)$.

To prove Claim 12.5.2, we use the following lemmas about polynomials. We defer their proofs to Section 12.10.2.

Lemma 12.5.3. Let $Q(z)$ be a degree $k$ polynomial, all of whose roots are complex numbers with absolute value 1. For any integer $n$, let $r_{n, k}(z)=\sum_{l=0}^{k-1} r_{n, k}^{(l)} \cdot z^{l}$ denote the residual polynomial of

$$
r_{n, k}(z) \equiv z^{n} \quad(\bmod Q(z))
$$

Then, each coefficient of $r_{n, k}$ is bounded: $\left|r_{n, k}^{(l)}\right| \leq 2^{k} n^{k-1}$ for any $l$.

Lemma 12.5.4. For any $k \in \mathbb{Z}$ and any $z_{1}, \cdots, z_{k}$ on the unit circle of $\mathbb{C}$, there always exists a degree $m=O\left(k^{2} \log k\right)$ polynomial $P(z)=\sum_{j=0}^{m} c_{j} z^{j}$ with the following properties:

$$
\text { Property I } \quad P\left(z_{i}\right)=0, \forall i \in\{1, \cdots, k\}
$$

Property II $\quad c_{0}=1$,
Property III $\quad\left|c_{j}\right| \leq 11, \forall j \in\{1, \cdots, m\}$.

Proof of Claim 12.5.2. For $x(t)=\sum_{i=1}^{k} v_{i} e^{2 \pi \mathrm{i} \mathrm{f}_{i} t}$, we fix $t_{0}$ and $\tau$ then rewrite $x\left(t_{0}+j \cdot \tau\right)$ as a polynomial of $b_{i}=v_{i} \cdot e^{2 \pi \mathbf{i} f_{i} t_{0}}$ and $z_{i}=e^{2 \pi \mathbf{i} f_{i} \tau}$ for each $i \in[k]$.

$$
\begin{aligned}
x\left(t_{0}+j \cdot \tau\right) & =\sum_{i=1}^{k} v_{i} e^{2 \pi \mathbf{i} f_{i} \cdot\left(t_{0}+j \cdot \tau\right)} \\
& =\sum_{i=1}^{k} v_{i} e^{2 \pi \mathbf{i} f_{i} t_{0}} \cdot e^{2 \pi \mathbf{i} f_{i} \cdot j \tau} \\
& =\sum_{i=1}^{k} b_{i} \cdot z_{i}^{j}
\end{aligned}
$$

Given $k$ and $z_{1}, \cdots, z_{k}$, let $P(z)=\sum_{j=0}^{m} c_{j} z^{j}$ be the degree $m$ polynomial in Lemma 12.5.4.

$$
\begin{align*}
\sum_{j=0}^{m} c_{j} x\left(t_{0}+j \tau\right) & =\sum_{j=0}^{m} c_{j} \sum_{i=1}^{k} b_{i} \cdot z_{i}^{j} \\
& =\sum_{i=1}^{k} b_{i} \sum_{j=0}^{m} c_{j} \cdot z_{i}^{j} \\
& =\sum_{i=1}^{k} b_{i} P\left(z_{i}\right) \\
& =0 \tag{12.14}
\end{align*}
$$

where the last step follows by Property I of $P(z)$ in Lemma 12.5.4. From the Property II and III of $P(z)$, we obtain $x\left(t_{0}\right)=-\sum_{j=1}^{m} c_{j} x\left(t_{0}+j \tau\right)$.

### 12.5.2 Bounding growth outside the interval

Here we show signals with sparse Fourier transform cannot grow too quickly outside the interval.

Lemma 12.5.5. Let $x(t)$ be a $k$-Fourier-sparse signal. For any $T>0$ and any $t>T$,

$$
|x(t)|^{2} \leq k^{7} \cdot(2 k t / T)^{2.5 k} \cdot\|x\|_{T}^{2} .
$$

Proof. For any $t>T$, let $t=t_{0}+n \cdot \tau$ such that $t_{0} \in[0, T / k], \tau \in[0, T / k]$ and $n \leq \frac{2 k t}{T}$. We define $b_{i}=v_{i} e^{2 \pi \mathbf{i} f_{i} t_{0}}$, and $z_{i}=e^{2 \pi \mathbf{i} f_{i} \tau}$ such that $x\left(t_{0}+n \cdot \tau\right)=\sum_{j \in[k]} b_{j} z_{j}^{n}$.

By Lemma 12.5.3, we have for any $z_{1}, z_{2}, \cdots, z_{k}$ and any $n$,

$$
z^{n} \equiv \sum_{i=0}^{k-1} a_{i} z^{i} \quad\left(\bmod \prod_{i=1}^{k}\left(z-z_{i}\right)\right)
$$

where $\left|a_{i}\right| \leq 2^{k} \cdot n^{k}, \forall i \in\{0,1, \cdots, k-1\}$. Thus, we obtain

$$
x\left(t_{0}+n \tau\right)=\sum_{j=1}^{k} b_{j} z_{j}^{n}=\sum_{j=1}^{k} b_{j}\left(\sum_{i=0}^{k-1} a_{i} z_{j}^{i}\right) .
$$

From the fact that $x\left(t_{0}+i \cdot \tau\right)=\sum_{j \in[k]} b_{j} z_{j}^{i}$, we simplify it to be

$$
x\left(t_{0}+n \tau\right)=\sum_{i=0}^{k-1} a_{i} \sum_{j=1}^{k} b_{j} z_{j}^{i}=\sum_{i=0}^{k-1} a_{i} x\left(t_{0}+i \cdot \tau\right) .
$$

Because $\left(t_{0}+i \cdot \tau\right) \in[0, T]$ for any $i=0, \cdots, k-1$, we have

$$
\left|x\left(t_{0}+i \tau\right)\right|^{2} \leq \max _{t \in[0, T]}|x(t)|^{2} \lesssim k^{4} \log ^{3} k\|x\|_{T}^{2}
$$

from Lemma 12.5.1.

Hence

$$
\begin{aligned}
\left|x\left(t_{0}+n \cdot \tau\right)\right|^{2} & \leq k \sum_{i=0}^{k-1}\left|a_{i}\right|^{2} \cdot\left|x\left(t_{0}+i \cdot \tau\right)\right|^{2} \\
& \leq k \sum_{i=0}^{k-1} n^{2.2 k} \cdot \max _{t \in[0, T]}|x(t)|^{2} \\
& \leq k^{7} \cdot(2 k t / T)^{2.2 k}\|x\|_{T}^{2} .
\end{aligned}
$$

Thus, we complete the proof.

### 12.6 Hash Functions and Filter Functions

### 12.6.1 Permutation function and hash function

We first review the permutation function $P_{\sigma, a, b}$ and the hash function $h_{\sigma, b}$ in [PS15], which translates discrete settings to the continuous setting.

Definition 12.6.1. For any signal $x(t): \mathbb{R} \rightarrow \mathbb{C}$ and $a, b, \sigma \in \mathbb{R}$, let $\left(P_{\sigma, a, b} x\right)(t)=x(\sigma(t-$ a)) $e^{-2 \pi \mathbf{i} \sigma b t}$.

Lemma 12.6.1. $\widehat{P_{\sigma, a, b} x}(\sigma(f-b))=\frac{1}{\sigma} e^{-2 \pi \mathbf{i} \sigma a f} \widehat{x}(f)$ and $\widehat{P_{\sigma, a, b} x}(f)=\frac{1}{\sigma} e^{-2 \pi \mathbf{i} \sigma a(f / \sigma+b)} \widehat{x}(f / \sigma+b)$
For completeness, we provide a proof of Lemma 12.6.1 in Section 12.10.4.

Definition 12.6.2. [PS15] Let $\pi_{\sigma, b}(f)=2 \pi \sigma(f-b)(\bmod 2 \pi)$ and $h_{\sigma, b}(f)=\operatorname{round}\left(\pi_{\sigma, b}(f)\right.$. $\frac{B}{2 \pi}$ ) be the hash function that maps frequency $f \in[-F, F]$ into bins $\{0, \cdots, B-1\}$.

Claim 12.6.2. [PS15] For any $\Delta>0$, let $\sigma$ be a sample uniformly at random from $\left[\frac{1}{B \Delta}, \frac{2}{B \Delta}\right]$.
(I) If $\Delta \leq\left|f^{+}-f^{-}\right|<\frac{(B-1) \Delta}{2}$, then $\operatorname{Pr}\left[h_{\sigma, b}\left(f^{+}\right)=h_{\sigma, b}\left(f^{-}\right)\right]=0$
(II) If $\frac{(B-1) \Delta}{2} \leq\left|f^{+}-f^{-}\right|$, then $\operatorname{Pr}\left[h_{\sigma, b}\left(f^{+}\right)=h_{\sigma, b}\left(f^{-}\right)\right] \lesssim \frac{1}{B}$

From previous work [HIKP12b, HIKP12a, PS15], uniformly sampling from $[A, 2 A]$ for some large $A \geq \widetilde{T}$ provides an almost uniform sample on $[0, \widetilde{T}]$ when taken modulo over $\widetilde{T}$.

Lemma 12.6.3. For any $\widetilde{T}$, and $0 \leq \widetilde{\epsilon}, \widetilde{\delta} \leq \widetilde{T}$, if we sample $\widetilde{\sigma}$ uniformly at random from $[A, 2 A]$, then

$$
\begin{equation*}
\frac{2 \widetilde{\epsilon}}{\widetilde{T}}-\frac{2 \widetilde{\epsilon}}{A} \leq \operatorname{Pr}[\widetilde{\sigma} \quad(\bmod \widetilde{T}) \in[\widetilde{\delta}-\widetilde{\epsilon}, \widetilde{\delta}+\widetilde{\epsilon}]] \leq \frac{2 \widetilde{\epsilon}}{\widetilde{T}}+\frac{4 \widetilde{\epsilon}}{A} \tag{12.15}
\end{equation*}
$$

### 12.6.2 Filter function

We state the properties of filter function $(H(t), \widehat{H}(f))$ and $(G(t), \widehat{G}(f))$, the details of proofs are presented in Section 12.12.1 and 12.12.2.

Lemma 12.6.4. Given $s_{0}, s_{1}, 0<s_{3}<1, \ell>1,0<\delta<1$, where $\ell=\Theta(k \log (k / \delta))$. The filter function $(H(t), \widehat{H}(f))$ has the following properties,

Property I: $\quad H(t) \in[1-\delta, 1]$, when $|t| \leq\left(\frac{1}{2}-\frac{2}{s_{1}}\right) s_{3}$.
Property II : $H(t) \in[0,1]$, when $\left(\frac{1}{2}-\frac{2}{s_{1}}\right) s_{3} \leq|t| \leq \frac{1}{2} s_{3}$.
Property III : $H(t) \leq s_{0} \cdot\left(s_{1}\left(\frac{|t|}{s_{3}}-\frac{1}{2}\right)+2\right)^{-\ell}, \forall|t|>\frac{1}{2} s_{3}$.
Property IV : $\operatorname{supp}(\hat{H}(f)) \subseteq\left[-\frac{s_{1} \ell}{2 s_{3}}, \frac{s_{1} \ell}{2 s_{3}}\right]$.

For any exact $k$-Fourier-sparse signal $x^{*}(t)$, we shift the interval from $[0, T]$ to $[-1 / 2,1 / 2]$ and consider $x^{*}(t)$ for $t \in[-1 / 2,1 / 2]$ to be our observation, which is also $x^{*}(t) \cdot \operatorname{rect}_{1}(t)$.

$$
\begin{array}{ll}
\text { Property V }: & \int_{-\infty}^{+\infty}\left|x^{*}(t) \cdot H(t) \cdot\left(1-\operatorname{rect}_{1}(t)\right)\right|^{2} \mathrm{~d} t<\delta \int_{-\infty}^{+\infty}\left|x^{*}(t) \cdot \operatorname{rect}_{1}(t)\right|^{2} \mathrm{~d} t . \\
\text { Property VI : } & \int_{-\infty}^{+\infty}\left|x^{*}(t) \cdot H(t) \cdot \operatorname{rect}_{1}(t)\right|^{2} \mathrm{~d} t \in[1-\epsilon, 1] \cdot \int_{-\infty}^{+\infty}\left|x^{*}(t) \cdot \operatorname{rect}_{1}(t)\right|^{2} \mathrm{~d} t .
\end{array}
$$

for arbitrarily small constant $\epsilon$.

Lemma 12.6.5. Given $B>1, \delta>0, \alpha>0$, we set $l=\Omega(\log (\delta / k))$. The filter function


Figure 12.2: The filter function $(H(t), \widehat{H}(f))$ with a $k$-Fourier-sparse signal. The property I, II and III are presented in the bottom one, the property IV is presented in the top one.
$(G(t), \widehat{G}(f))[B, \delta, \alpha, l]$ satisfies the following properties,
Property I: $\quad \widehat{G}(f) \in[1-\delta / k, 1]$, if $|f| \leq(1-\alpha) \frac{2 \pi}{2 B}$.
Property II : $\quad \widehat{G}(f) \in[0,1]$, if $(1-\alpha) \frac{2 \pi}{2 B} \leq|f| \leq \frac{2 \pi}{2 B}$.
Property III : $\widehat{G}(f) \in[-\delta / k, \delta / k]$, if $|f|>\frac{2 \pi}{2 B}$.
Property IV : $\operatorname{supp}(G(t)) \subset\left[\frac{l}{2} \cdot \frac{-B}{\pi \alpha}, \frac{l}{2} \cdot \frac{B}{\pi \alpha}\right]$.
Property V: $\max _{t}|G(t)| \lesssim \operatorname{poly}(B, l)$.


Figure 12.3: $G$ and $\widehat{G}$. [PS15]

### 12.6.3 HashToBins

We first define two functions $G_{\sigma, b}^{(j)}(t)$ and $\widehat{G}_{\sigma, b}^{(j)}(f)$, then show the result returned by Procedure HashToBins in Algorithm 12.6 satisfying some nice properties. The details of proofs are presented in Section 12.12.4.

Definition 12.6.6. $\forall \sigma>0, b$ and $j \in[B]$. Define,

$$
\begin{aligned}
G_{\sigma, b}^{(j)}(t) & =\frac{1}{\sigma} G(t / \sigma) e^{2 \pi \mathrm{i} t(j / B-\sigma b) / \sigma} \\
\widehat{G}_{\sigma, b}^{(j)}(f) & =\widehat{G}^{\mathrm{dis}}\left(\frac{j}{B}-\sigma f-\sigma b\right)=\sum_{i \in \mathbb{Z}} \widehat{G}\left(i+\frac{j}{B}-\sigma f-\sigma b\right)
\end{aligned}
$$

Lemma 12.6.7. Let $u \in \mathbb{C}^{B}$ be the result of HASHTOBins under permutation $P_{\sigma, a, b}$, and let $j \in[B]$. Define

$$
\widehat{z}=\widehat{x \cdot H} \cdot \widehat{G}_{\sigma, b}^{(j)}
$$

so

$$
z=(x \cdot H) * G_{\sigma, b}^{(j)} .
$$

Let vector $\widehat{u} \in \mathbb{C}^{B}$ denote the $B$-dimensional DFT of $u$, then $\forall j \in[B]$,

$$
\widehat{u}[j]=z_{\sigma a} .
$$

### 12.7 Frequency Recovery

The goal of this section is to prove Theorem 12.2.6, which is able to recover the frequencies of a signal $x^{*}$ has $k$-sparse Fourier transform under noise.

Theorem 12.2.6. Let $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$ and $x(t)=x^{*}(t)+g(t)$ be our observable signal where $\|g(t)\|_{T}^{2} \leq c\left\|x^{*}(t)\right\|_{T}^{2}$ for a sufficiently small constant $c$. Then Procedure FrequenCYRECOVERYKCLUSTER returns a set $L$ of $O(k)$ frequencies that covers all heavy clusters of $x^{*}$, which uses poly $(k, \log (1 / \delta)) \log (F T)$ samples and poly $(k, \log (1 / \delta)) \log ^{2}(F T)$ time. In particular, for $\Delta=\operatorname{poly}(k, \log (1 / \delta)) / T$ and $\mathcal{N}^{2}:=\|g(t)\|_{T}^{2}+\delta\left\|x^{*}(t)\right\|_{T}^{2}$, with probability $1-2^{-\Omega(k)}$, for any $f^{*}$ with

$$
\begin{equation*}
\int_{f^{*}-\Delta}^{f^{*}+\Delta}|\widehat{x \cdot H}(f)|^{2} \mathrm{~d} f \geq T \mathcal{N}^{2} / k \tag{12.16}
\end{equation*}
$$

there exists an $\tilde{f} \in L$ satisfying

$$
\left|f^{*}-\widetilde{f}\right| \lesssim \Delta \sqrt{\Delta T}
$$

### 12.7.1 Overview

We give an overview of proving Theorem 12.2.6. Instead of starting with $k$-cluster recovery, we first show how to achieve one-cluster recovery.

One-cluster recovery. we start with $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$ where there exists $f_{0}$ and $\Delta$ such that $f_{j}$ is in $\left[f_{0}-\Delta, f_{0}+\Delta\right]$ for each $j \in[k]$ and consider its properties for frequency recovery.

Definition 12.7.1 $((\epsilon, \Delta)$-one-cluster signal). We say that a signal $z(t)$ is an $(\epsilon, \Delta)$-onecluster signal around $f_{0}$ iff $z(t)$ and $\widehat{z}(f)$ satisfy the following two properties:

$$
\begin{aligned}
\text { Property I } & : \int_{f_{0}-\Delta}^{f_{0}+\Delta}|\widehat{z}(f)|^{2} \mathrm{~d} f \geq(1-\epsilon) \int_{-\infty}^{+\infty}|\widehat{z}(f)|^{2} \mathrm{~d} f \\
\text { Property II } & : \int_{0}^{T}|z(t)|^{2} \mathrm{~d} t \geq(1-\epsilon) \int_{-\infty}^{+\infty}|z(t)|^{2} \mathrm{~d} t
\end{aligned}
$$

The main result of one-cluster recovery is to prove that the two properties in Definition 12.7.1 with a sufficiently small constant $\epsilon$ are sufficient to return $\widetilde{f}_{0}$ close to $f_{0}$ with high probability, which provides a black-box for $k$-cluster recovery algorithm.

We first prove that the pair of conditions, Property I and Property II in Definition 12.7.1, are sufficient to obtain an estimation of $e^{2 \pi \mathrm{i} f_{0}}$ in Section 12.7.2. We also provide the proof of the correctness of Procedures GetLegal1Sample and GetEmpirical1Engergy in Section 12.7.2.

Lemma 12.7.1. For a sufficiently small constant $\epsilon>0$, any $f_{0} \in[-F, F]$, and $\Delta>0$, given $\widehat{\beta} \approx \frac{1}{\Delta \sqrt{\Delta T}}$ and an $(\epsilon, \Delta)$-one-cluster signal $z(t)$ around $f_{0}$, Procedure GetLegal1Sample
in Algorithm 12.3 with any $\beta \leq 2 \widehat{\beta}$ takes $O\left((T \Delta)^{3}\right)$ samples to output $\alpha \in \mathbb{R}$ satisfying

$$
\left|z(\alpha+\beta)-z(\alpha) e^{2 \pi \mathrm{i} f_{0} \beta}\right| \leq 0.08(|z(\alpha)|+|z(\alpha+\beta)|),
$$

with probability at least 0.6.

The following lemma shows that for any $(\epsilon, \Delta)$-one-cluster signal $z(t)$ around $f_{0}$, we could use the above procedure to find a frequency $\widetilde{f}_{0}$ approximating $f_{0}$ with high probability.

Lemma 12.7.2. For a sufficiently small constant $\epsilon>0$, any $f_{0} \in[-F, F]$, and $\Delta>0$, given an $(\epsilon, \Delta)$-one-cluster signal $z(t)$ around $f_{0}$, Procedure FREQUENCyRECovery1Cluster in Algorithm 12.4 returns $\widetilde{f}_{0}$ with $\left|\widetilde{f}_{0}-f_{0}\right| \lesssim \Delta \cdot \sqrt{\Delta T}$ with probability at least $1-2^{-\Omega(k)}$.

We provide a proof of Lemma 12.7.2 in Section 12.7.4. We show $z(t)=\left(x^{*}(t)+g(t)\right)$. $H(t)$ satisfy Properties I and II (Definition 12.7.1) when all frequencies in $\widehat{x}^{*}$ are in a small range in Section 12.7.3.

Lemma 12.7.3. For any $f_{0} \in[-F, F], \Delta^{\prime}>0$, and $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathbf{i} f_{t}}$ with $\left|f_{j}-f_{0}\right| \leq \Delta^{\prime}$ for all $j \in[k]$, let $x(t)=x^{*}(t)+g(t)$ be our observable signal whose noise $\|g\|_{T}^{2} \leq c\left\|x^{*}\right\|_{T}^{2}$ for a sufficiently small constant $c$ and $H(t)$ be the filter function defined in Section 12.6 with $|\operatorname{supp}(\widehat{H})|=\Delta_{h}$. Then $z=H \cdot x$ is an $\left(O(\sqrt{c}), \Delta_{h}+\Delta^{\prime}\right)$-one-cluster signal around $f_{0}$.

From all discussion above, we summarize the result of frequency recovery when $\widehat{x}^{*}$ is in one cluster.

Theorem 12.7.4. For any $f_{0} \in[-F, F], \Delta^{\prime}>0$, and $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$ with $\left|f_{j}-f_{0}\right| \leq \Delta^{\prime}$ for all $j \in[k]$, let $x(t)=x^{*}(t)+g(t)$ be our observable signal whose noise $\|g\|_{T}^{2} \leq c\left\|x^{*}\right\|_{T}^{2}$ for a sufficiently small constant $c$ and $H(t)$ be the filter function defined in Section 12.6
with $|\operatorname{supp}(\widehat{H})|=\Delta_{h}$. Then Procedure FrequencyRecovery1Cluster in Algorithm 12.4 with $\Delta=\Delta^{\prime}+\Delta_{h}$ takes poly $(k, \log (1 / \delta)) \cdot \log (F T)$ samples, runs in poly $(k, \log (1 / \delta)) \cdot$ $\log ^{2}(F T)$ time, returns a frequency $\widetilde{f}_{0}$ satisfying $\left|\widetilde{f}_{0}-f_{0}\right| \lesssim \Delta \sqrt{\Delta T}$ with probability at least $1-2^{-\Omega(k)}$.
$k$-cluster recovery. Given any $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$, we plan to convolve the filter function $G(t)$ on $x(t) \cdot H(t)$ and use Lemma 12.7.2 as a black box to find a list of frequencies that covers $\left\{f_{1}, \cdots, f_{k}\right\}$.

We fix $\Delta=\operatorname{poly}(k, \log (1 / \delta)) / T, B=\Theta(k)$ and sample $\sigma$ uniformly at random from $\left[\frac{1}{B \Delta}, \frac{2}{B \Delta}\right]$ for $k$-cluster recovery. We will cover all $f^{*} \in[-F, F]$ with the following property :

$$
\begin{equation*}
\int_{f^{*}-\Delta}^{f^{*}+\Delta}|\widehat{x \cdot H}(f)|^{2} \mathrm{~d} f \geq T \mathcal{N}^{2} / k \tag{12.17}
\end{equation*}
$$

We consider one frequency $f^{*} \in[-F, F]$ satisfying (12.17) and use $j=h_{\sigma, b}\left(f^{*}\right)$ to denote its index in $[B]$ after hashing $(\sigma, b)$. Recall that for $j \in[B]$, any $\sigma>0$ and any $b$,

$$
G_{\sigma, b}^{(j)}(t)=\frac{1}{\sigma} G(t / \sigma) e^{2 \pi i t(j / B-\sigma b) / \sigma} \text { such that } \widehat{G}_{\sigma, b}^{(j)}(f)=\sum_{i \in \mathbb{Z}} \widehat{G}\left(i+\frac{j}{B}-\sigma f-\sigma b\right) .
$$

We set $\widehat{z}=\widehat{x \cdot H} \cdot \widehat{G}_{\sigma, b}^{(j)}$ and $z=(x \cdot H) * G_{\sigma, b}^{(j)}$ for $f^{*}$ and $j=h_{\sigma, b}\left(f^{*}\right)$. In Section 12.7.5, we show that with high probability over the hashing $(\sigma, b),(z, \widehat{z})$ satisfies Property I with $\left[f^{*}-\Delta, f^{*}+\Delta\right]$ and Property II in Definition 12.7.1 such that we could use Lemma 12.7.2 on $z$ to recover $f^{*}$.

Lemma 12.7.5. Let $f^{*} \in[-F, F]$ satisfy (12.17). For a random hashing $(\sigma, b)$, let $j=$ $h_{\sigma, b}\left(f^{*}\right)$ be the bucket that $f^{*}$ maps to under the hash such that $z=(x \cdot H) * G_{\sigma, b}^{(j)}$ and $\widehat{z}=\widehat{x \cdot H} \cdot \widehat{G}_{\sigma, b}^{(j)}$. With probability at least $0.9, z(t)$ is an $(\epsilon, \Delta)$-one-cluster signal around $f^{*}$.

Combining Lemma 12.7 .5 and Lemma 12.7.2, we could recover any heavy frequency $f^{*}$ satisfying (12.17) with probability at least 0.8 . Then we repeat this procedure to guarantee that we cover all heavy frequencies and finish the proof of the main frequency recovery Theorem 12.2.6 in Section 12.7.6.

### 12.7.2 Analysis of GetLegal1Sample and GetEmpirical1Energy

Let $I=\left[f_{0}-\Delta, f_{0}+\Delta\right]$ and $\bar{I}=(-\infty,+\infty) \backslash I$ in this proof. We define $\left(z^{I}(t), \widehat{z}^{I}(f)\right)$ and $\left(z^{\bar{I}}(t), \widehat{z}^{\bar{I}}(f)\right)$ as follows:

$$
\widehat{z}^{I}(f)=\left\{\begin{array}{ll}
\widehat{z}(f) & \text { if } f \in I \\
0 & \text { if } f \in \bar{I}
\end{array}, \quad \widehat{z}^{\bar{I}}(f)= \begin{cases}0 & \text { if } f \in I \\
\widehat{z}(f) & \text { if } f \in \bar{I}\end{cases}\right.
$$

We consider $z^{I}(t)$ as the "signal" to recover $f_{0}$ and treat $z^{\bar{I}}(t)$ as the "noise". We first show some basic properties of $z^{I}(t)$.

Claim 12.7.6. For $z^{\bar{I}}(t)$, we have $\int_{0}^{T}\left|z^{\bar{I}}(t)\right|^{2} \mathrm{~d} t \leq \epsilon \int_{-\infty}^{+\infty}|z(t)|^{2} \mathrm{~d} t$. For $z^{I}(t)$, we have

$$
\int_{0}^{T}\left|z^{I}(t)\right|^{2} \mathrm{~d} t \geq(1-5 \sqrt{\epsilon}) \int_{-\infty}^{+\infty}|z(t)|^{2} \mathrm{~d} t \text { and } \int_{0}^{T}\left|z^{I}(t)\right|^{2} \mathrm{~d} t \geq(1-6 \sqrt{\epsilon}) \int_{-\infty}^{+\infty}\left|z^{I}(t)\right|^{2} \mathrm{~d} t
$$

Proof. From the definition and Property I in Definition 12.7.1, we know

$$
z(t)=z^{I}(t)+z^{\bar{I}}(t) \quad \text { and } \quad \int_{-\infty}^{+\infty}\left|\bar{z}^{\bar{I}}(f)\right|^{2} \mathrm{~d} f \leq \epsilon \int_{-\infty}^{+\infty}|\widehat{z}(f)|^{2} \mathrm{~d} f .
$$

Notice that Property I(in Definition 12.7.1) indicates that

$$
\int_{0}^{T}\left|z^{\bar{I}}(t)\right|^{2} \mathrm{~d} t \leq \int_{-\infty}^{+\infty}\left|z^{\bar{I}}(t)\right|^{2} \mathrm{~d} t=\int_{-\infty}^{+\infty}\left|\bar{z}^{\bar{I}}(f)\right|^{2} \mathrm{~d} f \leq \epsilon \int_{-\infty}^{+\infty}|\widehat{z}(f)|^{2} \mathrm{~d} f
$$

On the other hand, from Property II(in Definition 12.7.1), we know

$$
\begin{aligned}
(1-\epsilon) \int_{-\infty}^{+\infty}|z(t)|^{2} \mathrm{~d} t & \leq \int_{0}^{T}\left|z^{I}(t)+z^{\bar{I}}(t)\right|^{2} \mathrm{~d} t \\
& \leq \int_{0}^{T}\left|z^{I}(t)\right|^{2} \mathrm{~d} t+2 \int_{0}^{T}\left|z^{I}(t)\right| \cdot\left|z^{\bar{I}}(t)\right| \mathrm{d} t+\int_{0}^{T}\left|z^{\bar{I}}(t)\right|^{2} \mathrm{~d} t .
\end{aligned}
$$

We have $\int_{0}^{T}\left|z^{I}(t)\right|^{2} \mathrm{~d} t \leq 2 \int_{-\infty}^{+\infty}|z(t)|^{2} \mathrm{~d} t$ from the above inequality. From $\int_{0}^{T}\left|z^{\bar{I}}(t)\right|^{2} \mathrm{~d} t \leq$ $\epsilon \int_{-\infty}^{+\infty}|\widehat{z}(f)|^{2} \mathrm{~d} f$, we bound

$$
\int_{0}^{T}\left|z^{I}(t)\right| \cdot\left|z^{\bar{I}}(t)\right| \mathrm{d} t \leq \sqrt{2 \epsilon} \int_{-\infty}^{+\infty}|z(t)|^{2} \mathrm{~d} t
$$

by the Cauchy-Schwartz inequality and have

$$
\begin{equation*}
\int_{0}^{T}\left|z^{I}(t)\right|^{2} \mathrm{~d} t \geq(1-5 \sqrt{\epsilon}) \int_{-\infty}^{+\infty}|z(t)|^{2} \mathrm{~d} t \tag{12.18}
\end{equation*}
$$

Because $\int_{-\infty}^{+\infty}\left|z^{\bar{I}}(t)\right|^{2} \mathrm{~d} t \leq \epsilon \int_{-\infty}^{+\infty}|z(t)|^{2} \mathrm{~d} t$, inequality (12.18) also indicates that

$$
\int_{0}^{T}\left|z^{I}(t)\right|^{2} \mathrm{~d} t \geq(1-6 \sqrt{\epsilon}) \int_{-\infty}^{+\infty}\left|z^{I}(t)\right|^{2} \mathrm{~d} t
$$

One useful property of $z^{I}(t)$ is that its maximum can be bounded by its average on $[0, T]$.

Claim 12.7.7. $\forall t \in[0, T],\left|z^{I}(t)\right| \leq 2 \sqrt{\Delta T} \cdot\left\|z^{I}\right\|_{T}$.
Proof. From the definition $\left|z^{I}(t)\right|$, it is upper bounded by $\int_{f_{0}-\Delta}^{f_{0}+\Delta}\left|\widehat{z^{I}}(f)\right| \mathrm{d} f$ for any $t \in[0, T]$. On the other hand,

$$
\begin{aligned}
\int_{f_{0}-\Delta}^{f_{0}+\Delta}\left|\hat{z}^{I}(f)\right| \mathrm{d} f & \leq \sqrt{2 \Delta}\left(\int_{f_{0}-\Delta}^{f_{0}+\Delta}\left|\widehat{z}^{I}(f)\right|^{2} \mathrm{~d} f\right)^{1 / 2} \\
& =\sqrt{2 \Delta}\left(\int_{-\infty}^{+\infty}\left|z^{I}(t)\right|^{2} \mathrm{~d} t\right)^{1 / 2} \\
& \leq 2 \sqrt{\Delta}\left(\int_{0}^{T}\left|z^{I}(t)\right|^{2} \mathrm{~d} t\right)^{1 / 2} \\
& =2 \sqrt{\Delta T}\left\|z^{I}\right\|_{T}
\end{aligned}
$$

Claim 12.7.8. Given $\widehat{\beta}=\frac{C_{\beta}}{\Delta \cdot \sqrt{\Delta T}}$ with a sufficiently small constant $C_{\beta}$, for any two $\widehat{\beta}$-close samples in $z^{I}(t)$, we have that

$$
\forall \alpha \in[0, T], \forall \beta \in[\widehat{\beta}, 2 \widehat{\beta}], \quad\left|z^{I}(\alpha) e^{2 \pi \mathbf{i} f_{0} \beta}-z^{I}(\alpha+\beta)\right| \leq 0.01 \cdot\left\|z^{I}\right\|_{T}
$$

Proof. From the definition of the Fourier transform, we have

$$
\begin{array}{rlrl}
\left|z^{I}(a+\beta)-z^{I}(a) e^{2 \pi \mathbf{i} f_{0} \beta}\right| & =\left|\int_{f_{0}-\Delta}^{f_{0}+\Delta} \widehat{z^{I}}(f) e^{2 \pi \mathbf{i}\left(f a+f_{0} \beta\right)}\left(e^{2 \pi \mathbf{i}\left(f-f_{0}\right) \beta}-1\right) \mathrm{d} f\right| & \\
& \leq 2 \cdot(2 \pi \Delta \beta) \cdot \int_{f_{0}-\Delta}^{f_{0}+\Delta}\left|\widehat{z^{I}}(f)\right| \mathrm{d} f & & \text { by Taylor expansion } \\
& \leq 4 \pi \beta \Delta \cdot \sqrt{2 \Delta}\left(\int_{f_{0}-\Delta}^{f_{0}+\Delta}\left|\widehat{z^{I}}(f)\right|^{2} \mathrm{~d} f\right)^{\frac{1}{2}} & & \text { by Hölder inequality } \\
& \leq 10 \pi \widehat{\beta} \Delta \cdot \sqrt{2 \Delta}\left(\int_{0}^{T}\left|z^{I}(t)\right|^{2} \mathrm{~d} t\right)^{\frac{1}{2}} & & \text { by inequality (12.18) } \\
& \leq 10^{-2}\left\|z^{I}\right\|_{T} . & &
\end{array}
$$

We consider how to output an $\alpha$ such that $e^{2 \pi \mathrm{i} \mathrm{f}_{0} \beta} \approx z(\alpha+\beta) / z(\alpha)$ with high probability in the rest of this section.

If we can sample from $z^{I}(t)$, we already know $\left|z^{I}(\alpha) e^{2 \pi \mathrm{i} f_{0} \beta}-z^{I}(\alpha+\beta)\right| \leq 0.01\left\|z^{I}\right\|_{T}$ from Claim 12.7.8. Then it is enough to find any $\alpha$ such that $\left|z^{I}(\alpha)\right| \geq 0.5\left\|z^{I}\right\|_{T}$. From Claim 12.7.7, we can take $O(\sqrt{\Delta T})$ samples $\left(z^{I}(\alpha), z^{I}(\alpha+\beta)\right)$ where each $\alpha$ is uniformly sampled from $[0, T]$ such that with high probability, the sample $z^{I}(\alpha)$ with the largest norm $\left|z^{I}(\alpha)\right|$ satisfies $\left|z^{I}(\alpha)\right| \geq 0.5\left\|z^{I}\right\|_{T}$. Then we have $e^{2 \pi \mathrm{i} f_{0} \beta} \approx z^{I}(\alpha+\beta) / z^{I}(\alpha)$.

Next, we move to $z(t)=z^{I}(t)+z^{\bar{I}}(t)$ and plan to output $\alpha \in[0, T]$ with probability at least 0.5 such that $\left|z^{\bar{I}}(\alpha)\right| \leq 0.1\left|z^{I}(\alpha)\right|$ and $\left|z^{\bar{I}}(\alpha+\beta)\right| \leq 0.1\left|z^{I}(\alpha+\beta)\right|$. Because the "noise" $z^{\bar{I}}(t)$ has $\left\|z^{\bar{I}}(t)\right\|_{T}^{2} \geq \epsilon\left\|z^{I}(t)\right\|_{T}^{2}$ for a constant $\epsilon$ and the bound $\sqrt{\Delta T}$ in Claim 12.7.7 is a polynomial in $k$, the approach for $z^{I}(t)$ cannot guarantee that $z(\alpha+\beta) / z(\alpha) \approx e^{2 \pi \mathrm{i} f_{0} \beta}$ with probability more than $1 / 2$.

The key observation is as follows:

Observation 12.7.9. For a sufficiently small $\epsilon$ and $\left\|z^{\bar{I}}\right\|_{T}^{2} \leq \epsilon\|z\|_{T}^{2}$, let $D_{T}$ be the weighted distribution on $[0, T]$ according to $|z(t)|^{2}$, i.e., $D_{T}(t)=\frac{|z(t)|^{2}}{T\|z\|_{T}^{2}}$. If we sample $\alpha \in[0, T]$ from the distribution $D_{T}$ instead of the uniform distribution on $[0, T],\left|z^{I}(\alpha)\right| \leq 0.01\left|z^{I}(\alpha)\right|$ with probability 0.9.

It follows from the fact that

$$
\underset{\alpha \sim D_{T}}{\mathbb{E}} \frac{\left|z^{\bar{I}}(\alpha)\right|^{2}}{|z(\alpha)|^{2}}=\int_{0}^{T} \frac{\left|z^{\bar{I}}(\alpha)\right|^{2}}{|z(\alpha)|^{2}} \cdot \frac{|z(\alpha)|^{2}}{T\|z\|_{T}^{2}} \mathrm{~d} \alpha=\frac{\int_{0}^{T}\left|z^{\bar{I}}(\alpha)\right|^{2} \mathrm{~d} \alpha}{T\|z\|_{T}^{2}} \leq \epsilon
$$

In Procedure GetLegaliSample, we collect $(\Delta T)^{2}$ samples (in expectation) $(z(\alpha), z(\alpha+$ $\beta))$ in $S_{\text {heavy }}$ with $|z(\alpha)| \geq 0.49\|z\|_{T}$ and resample one $\alpha$ from these samples according to their norm $|z(\alpha)|^{2}+|z(\alpha+\beta)|^{2}$. We show its correctness as follows.

Because we do not know $0.5\|z\|_{T}$, we use $z_{\mathrm{emp}}$ to approximate it.
Claim 12.7.10. Procedure GetEmpirical1Energy in Algorithm 12.3 takes $O\left((T \Delta)^{2}\right)$ samples to output $z_{\mathrm{emp}}$ such that $z_{\mathrm{emp}} \in\left[0.8\|z\|_{T}, 1.2\|z\|_{T}\right]$ with prob. 0.9.

Proof. We know $z_{\text {emp }}^{2}=\mathbb{E}_{i \in\left[R_{\text {est }}\right]}\left[\left|z\left(\alpha_{i}\right)\right|^{2}\right]=\mathbb{E}_{i \in\left[R_{\text {est }}\right]}\left[\left|z^{I}\left(\alpha_{i}\right)+z^{\bar{I}}\left(\alpha_{i}\right)\right|^{2}\right]$.
Notice that $\mathbb{E}_{i \in\left[R_{\text {est }}\right]}\left[\left|z^{I}\left(\alpha_{i}\right)\right|^{2}\right]$ is in $\left[0.99\left\|z^{I}\right\|_{T}, 1.01\left\|z^{I}\right\|_{T}\right]$ with prob. 0.99 from the Chernoff bound and Claim 12.7.7.

At the same time, $\mathbb{E}_{\alpha_{i}}\left[\left|z^{\bar{I}}\left(\alpha_{i}\right)\right|^{2}\right]=\left\|z^{\bar{I}}\right\|_{T}^{2}$. With prob. 0.92, $\mathbb{E}_{i \in\left[R_{\text {est }}\right]}\left[\left|z^{\bar{I}}\left(\alpha_{i}\right)\right|^{2}\right] \leq$ $13\left\|z^{\bar{I}}\right\|_{T}^{2}$. For a sufficiently small $\epsilon$ and $\left\|z^{\bar{I}}\right\|_{T}^{2} \leq \epsilon\left\|z^{I}\right\|_{T}^{2}, \mathbb{E}_{i \in\left[R_{\text {est }}\right]}\left[\left|z^{\bar{I}}\left(\alpha_{i}\right)\right|^{2}\right] \leq 13 \epsilon\left\|z^{I}\right\|_{T}^{2}$.

At last, we bound the cross terms of $\left|z^{I}\left(\alpha_{i}\right)+z^{\bar{I}}\left(\alpha_{i}\right)\right|^{2}$ by the Cauchy-Schwartz in-
equality,

$$
\begin{aligned}
& \underset{i \in R_{\text {est }}}{\mathbb{E}}\left[\left|\bar{z}^{I}\left(\alpha_{i}\right) z^{\bar{I}}\left(\alpha_{i}\right)\right|+\left|z^{I}\left(\alpha_{i}\right) \bar{z}^{\bar{I}}\left(\alpha_{i}\right)\right|\right] \\
\leq & 2 \underset{i \in R_{\text {est }}}{\mathbb{E}}\left[\left|z^{I}\left(\alpha_{i}\right)\right| \cdot\left|z^{\bar{I}}\left(\alpha_{i}\right)\right|\right] \\
\leq & 2\left(\underset{i \in\left[R_{\text {est }}\right]}{\mathbb{E}}\left[\left|z^{I}\left(\alpha_{i}\right)\right|^{2}\right] \cdot \underset{i \in\left[R_{\text {est }}\right]}{\mathbb{E}}\left[\left|z^{\bar{I}}\left(\alpha_{i}\right)\right|^{2}\right]\right)^{1 / 2} \\
\leq & 10 \sqrt{\epsilon}\left\|z^{I}\right\|_{T}^{2}
\end{aligned}
$$

For a sufficiently small $\epsilon$, we have $\mathbb{E}_{i \in\left[R_{\text {est }}\right]}\left[\left|z\left(\alpha_{i}\right)\right|^{2}\right]^{1 / 2}$ is in $\left[0.9\left\|z^{I}\right\|_{T}, 1.1\left\|z^{I}\right\|_{T}\right]$, which is also in $\left[0.8\|z\|_{T}, 1.2\|z\|_{T}\right]$ because of Property II.

We assume $z_{\mathrm{emp}} \in\left[0.8\|z\|_{T}, 1.2\|z\|_{T}\right]$ and focus on $U=\left\{t \in[0, T]| | z(t) \mid \geq 0.5 z_{\mathrm{emp}}\right\}$. Notice that

$$
\int_{U}|z(t)|^{2} \mathrm{~d} t=\int_{0}^{T}|z(t)|^{2} \mathrm{~d} t-\int_{[0, T] \backslash U}|z(t)|^{2} \mathrm{~d} t \geq\left(1-0.6^{2}\right) \int_{0}^{T}|z(t)|^{2} \mathrm{~d} t .
$$

Let $R_{\text {heavy }}=\left|S_{\text {heavy }}\right|$. From Claim 12.7 .7 and $\epsilon, \mathbb{E}\left[R_{\text {heavy }}\right] \geq R_{\text {repeat }} /(T \Delta)$. So we assume $R_{\text {heavy }} \geq 0.01 R_{\text {repeat }} /(T \Delta)=0.01(T \Delta)^{2}$ in the rest of this section and think each $\alpha_{i} \in S_{\text {heavy }}$ is a uniform sample from $U$ over the randomness on $S_{\text {heavy }}$.

Claim 12.7.11. With probability 0.95, $\sum_{i \in S_{\text {heavy }}}\left(\left|z^{\bar{I}}\left(\alpha_{i}\right)\right|^{2}+\left|z^{\bar{I}}\left(\alpha_{i}+\beta\right)\right|^{2}\right) \leq 10^{-4} \sum_{i \in S_{\text {heavy }}}\left(\left|z\left(\alpha_{i}\right)\right|^{2}+\right.$ $\left.\left|z\left(\alpha_{i}+\beta\right)\right|^{2}\right)$ for a sufficiently small $\epsilon$ and $\left\|z^{\bar{I}}\right\|_{T}^{2} \leq \epsilon\|z\|_{T}^{2}$.

Proof. At first,

$$
\underset{S_{\text {heavy }}}{\mathbb{E}}\left[\sum_{i \in S_{\text {heavy }}}\left(\left|z\left(\alpha_{i}\right)\right|^{2}+\left|z\left(\alpha_{i}+\beta\right)\right|^{2}\right)\right] \geq R_{\text {heavy }} \cdot \underset{t \sim U}{\mathbb{E}}\left[|z(t)|^{2}\right]=R_{\text {heavy }} \cdot \frac{\int_{U}|z(t)|^{2} \mathrm{~d} t}{|U|}
$$

At the same time,
$\underset{S_{\text {heavy }}}{\mathbb{E}}\left[\sum_{i \in S_{\text {heavy }}}\left[\left|z^{\bar{I}}\left(\alpha_{i}\right)\right|^{2}+\left|z^{\bar{I}}\left(\alpha_{i}+\beta\right)\right|^{2}\right]\right]=R_{\text {heavy }} \cdot \underset{t \sim U}{\mathbb{E}}\left[\left|z^{\bar{I}}(t)\right|^{2}+\left|z^{\bar{I}}(t+\beta)\right|^{2}\right] \leq \frac{2 \int_{T}^{0}\left|z^{\bar{I}}(t)\right|^{2} \mathrm{~d} t}{|U|}$.
From $\int_{U}|z(t)|^{2} \mathrm{~d} t \geq 0.64 \int_{0}^{T}|z(t)|^{2} \mathrm{~d} t$ and $\int_{T}^{0}\left|z^{\bar{I}}(t)\right|^{2} \mathrm{~d} t \leq \epsilon \int_{0}^{T}|z(t)|^{2} \mathrm{~d} t$, we get the conclusion.

We assume all results in the above claims hold and prove that the sample from $S_{\text {heavy }}$ is a good sample such that $z^{\bar{I}}(\alpha)$ is small.

Claim 12.7.12. If we sample $i \in S_{\text {heavy }}$ according to the weight $\left|z\left(\alpha_{i}\right)\right|^{2}+\left|z\left(\alpha_{i}+\beta\right)\right|^{2}$, with prob. at least $0.9,\left|z^{\bar{I}}\left(\alpha_{i}\right)\right|+\left|z^{\bar{I}}\left(\alpha_{i}+\beta\right)\right| \leq 0.05\left(\left|z\left(\alpha_{i}\right)\right|+\left|z\left(\alpha_{i}+\beta\right)\right|\right)$.

Proof. Similar to the proof of the key observation, we compute the expectation of $\frac{\left|z^{I}\left(\alpha_{i}\right)\right|^{2}+\left|z^{\bar{I}}\left(\alpha_{i}+\beta\right)\right|^{2}}{\left|z\left(\alpha_{i}\right)\right|^{2}+\left|z\left(\alpha_{i}+\beta\right)\right|^{2}}$ over the sampling in $S_{\text {heavy }}$ :

$$
\begin{aligned}
& \sum_{i \in S_{\text {heavy }}} \frac{\left|z\left(\alpha_{i}\right)\right|^{2}+\left|z\left(\alpha_{i}+\beta\right)\right|^{2}}{\sum_{j \in S_{\text {heavy }}}\left|z\left(\alpha_{j}\right)\right|^{2}+\left|z\left(\alpha_{j}+\beta\right)\right|^{2}} \cdot \frac{\left|z^{\bar{I}}\left(\alpha_{i}\right)\right|^{2}+\left|z^{\bar{I}}\left(\alpha_{i}+\beta\right)\right|^{2}}{\left|z\left(\alpha_{i}\right)\right|^{2}+\left|z\left(\alpha_{i}+\beta\right)\right|^{2}} \\
= & \frac{\sum_{i \in S_{\text {heavy }}}\left|z^{\bar{I}}\left(\alpha_{i}\right)\right|^{2}+\left|z^{\bar{I}}\left(\alpha_{i}+\beta\right)\right|^{2}}{\sum_{i \in S_{\text {heavy }}}\left|z\left(\alpha_{i}\right)\right|^{2}+\left|z\left(\alpha_{i}+\beta\right)\right|^{2}} \\
\leq & 10^{-4} .
\end{aligned}
$$

By Markov's inequality, when we sample $i \in S_{\text {heavy }}$ according to the weight $\left|z\left(\alpha_{i}\right)\right|^{2}+\mid z\left(\alpha_{i}+\right.$ $\beta)\left.\right|^{2}, \frac{\left|\bar{z}^{I}\left(\alpha_{i}\right)\right|^{2}+\left|z^{I}\left(\alpha_{i}+\beta\right)\right|^{2}}{\left|z\left(\alpha_{i}\right)\right|^{2}+\left|z\left(\alpha_{i}+\beta\right)\right|^{2}} \leq 10^{-3}$ with probability 0.9 . We have that with prob. at least 0.9, $\left|z^{\bar{I}}\left(\alpha_{i}\right)\right|+\left|z^{\bar{I}}\left(\alpha_{i}+\beta\right)\right| \leq 0.05\left(\left|z\left(\alpha_{i}\right)\right|+\left|z\left(\alpha_{i}+\beta\right)\right|\right)$.

We assume all above claims hold and finish the proof by setting $\alpha=\alpha_{i}$. From Claim
12.7.8, we know that

$$
\left|z^{I}(\alpha) e^{2 \pi \mathrm{i} f_{0} \beta}-z^{I}(\alpha+\beta)\right| \leq 0.01 \cdot \underset{t \in[0, T]}{\mathbb{E}}\left[\left|z^{I}(\alpha)\right|^{2}\right]^{1 / 2} \leq 0.03\left|z^{I}(\alpha)\right|
$$

Now we add back the noise $z^{\bar{I}}(\alpha)$ and $z^{\bar{I}}(\alpha+\beta)$ to get

$$
\begin{aligned}
\left|z(\alpha) e^{2 \pi \mathrm{i} f_{0} \beta}-z(\alpha+\beta)\right| & \leq\left|z^{I}(\alpha) e^{2 \pi \mathrm{i} \mathrm{f}_{0} \beta}-z^{I}(\alpha+\beta)\right|+\left|z^{\bar{I}}(\alpha)\right|+\left|z^{\bar{I}}(\alpha+\beta)\right| \\
& \leq 0.08(|z(\alpha)|+|z(\alpha+\beta)|) .
\end{aligned}
$$

### 12.7.3 A cluster of frequencies, times $H$, is a one-cluster signal per Definition 12.7.1

The goal of this section is to prove Lemma 12.7.3. Without loss of generality, we assume $g(t)=0$ for any $t \notin[0, T]$ and notice that $\operatorname{supp}\left(\widehat{H} * \widehat{x}^{*}\right) \subseteq f_{0}+[-\Delta, \Delta]$ for $\Delta=\Delta^{\prime}+\Delta_{h}$ from the definition of $\widehat{H}$. From the Property VI (presented in Lemma 12.6.4) of $(H, \widehat{H})$,

$$
\int_{0}^{T}\left|x^{*}(t)\right|^{2} \mathrm{~d} t=(1 \pm c) \int_{-\infty}^{+\infty}\left|H(t) \cdot x^{*}(t)\right|^{2} \mathrm{~d} t
$$

From the first two properties of $(H, \widehat{H})$, we bound the energy of $g \cdot H$ :

$$
\int_{-\infty}^{+\infty}|H(t) \cdot g(t)|^{2} \mathrm{~d} t \leq(1+c) \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t
$$

Let $z(t)=\left(x^{*}(t)+g(t)\right) H(t)$. We use the triangle inequality on the above two inequalities:

$$
\begin{aligned}
& \int_{0}^{T}|z(t)|^{2} \mathrm{~d} t \\
\geq & \int_{0}^{T}\left|H(t) \cdot x^{*}(t)\right|^{2} \mathrm{~d} t-\int_{0}^{T}|H(t) \cdot g(t)|^{2} \mathrm{~d} t-2 \int_{0}^{T}\left|H(t) \cdot x^{*}(t)\right| \cdot|H(t) \cdot g(t)| \mathrm{d} t \\
\geq & (1-c) \int_{0}^{T}\left|x^{*}(t)\right|^{2} \mathrm{~d} t-(1+c) \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t-2 \int_{0}^{T}\left|H(t) \cdot x^{*}(t)\right| \cdot|H(t) \cdot g(t)| \mathrm{d} t \\
\geq & (1-c) \int_{0}^{T}\left|x^{*}(t)\right|^{2} \mathrm{~d} t-(1+c) \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t-2\left((1+c)^{2} \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t \int_{0}^{T}\left|x^{*}(t)\right|^{2} \mathrm{~d} t \cdot\right)^{1 / 2} \\
\geq & (1-c) \int_{0}^{T}\left|x^{*}(t)\right|^{2} \mathrm{~d} t-(1+c) c \int_{0}^{T}\left|x^{*}(t)\right|^{2} \mathrm{~d} t-2 \sqrt{c}(1+c) \int_{0}^{T}\left|x^{*}(t)\right|^{2} \mathrm{~d} t \\
\geq & (1-5 \sqrt{c}) \int_{0}^{T}\left|x^{*}(t)\right|^{2} \mathrm{~d} t
\end{aligned}
$$

where the third step follows from Cauchy-Schwarz inequality, the fourth step follows from $\int_{0}^{T}|g(t)|^{2} \mathrm{~d} t \leq c \int_{0}^{T}\left|x^{*}(t)\right|^{2} \mathrm{~d} t$, and the last step follows from choosing $c \in(0,1)$ sufficiently small.

Similarly,

$$
\begin{aligned}
& \int_{-\infty}^{+\infty}|z(t)|^{2} \mathrm{~d} t \\
\leq & (1+c) \int_{0}^{T}\left|x^{*}(t)\right|^{2} \mathrm{~d} t+(1+c) \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t+2\left((1+c)^{2} \int_{0}^{T}\left|x^{*}(t)\right|^{2} \mathrm{~d} t \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t\right)^{1 / 2} \\
\leq & (1+5 \sqrt{c}) \int_{0}^{T}\left|x^{*}(t)\right|^{2} \mathrm{~d} t
\end{aligned}
$$

Hence we obtain Property II(in Definition 12.7.1) when $c \in(0,1)$ is sufficiently small.
Then we observe that

$$
\begin{aligned}
& \int_{f_{0}-\Delta_{h}}^{f_{0}+\Delta_{h}}|\widehat{z}(f)|^{2} \mathrm{~d} f \\
\geq & \int_{f_{0}-\Delta_{h}}^{f_{0}+\Delta_{h}} \mid H \cdot \widehat{\left.\left(x^{*}+g\right)\right|^{2} \mathrm{~d} f} \\
\geq & \int_{f_{0}-\Delta_{h}}^{f_{0}+\Delta_{h}}\left|\widehat{H \cdot x^{*}}\right|^{2}-|\widehat{H \cdot g}|^{2}-2\left|\widehat{H \cdot x^{*}}\right| \cdot|\widehat{H \cdot g}| \mathrm{d} f \\
\geq & \int_{f_{0}-\Delta_{h}}^{f_{0}+\Delta_{h}}\left|\widehat{H \cdot x^{*}}\right|^{2} \mathrm{~d} f-\int_{-\infty}^{+\infty}|\widehat{H \cdot g}|^{2} \mathrm{~d} f-2\left(\int_{f_{0}-\Delta_{h}}^{f_{0}+\Delta_{h}}\left|\widehat{H \cdot x^{*}}\right|^{2} \mathrm{~d} f \int_{-\infty}^{+\infty}|\widehat{H \cdot g}|^{2} \mathrm{~d} f\right)^{1 / 2} \\
= & \int_{-\infty}^{+\infty}\left|H \cdot x^{*}\right|^{2} \mathrm{~d} t-\int_{-\infty}^{+\infty}|H \cdot g|^{2} \mathrm{~d} t-2\left(\int_{f_{0}-\Delta_{h}}^{f_{0}+\Delta_{h}}\left|H \cdot x^{*}\right|^{2} \mathrm{~d} t \int_{-\infty}^{+\infty}|H \cdot g|^{2} \mathrm{~d} t\right)^{1 / 2} \\
\geq & \frac{(1-c)-c(1+c)-3 \sqrt{c}}{1+5 \sqrt{c}} \int_{-\infty}^{+\infty}|z(t)|^{2} \mathrm{~d} t .
\end{aligned}
$$

Thus we have Property I(in Definition 12.7.1) for $z$.

### 12.7.4 Frequency recovery of one-cluster signals

The goal of this section is prove Theorem 12.7.4. We first show the correctness of Procedure Locate1Inner. Second, we analyze the Procedure Locate1Signal. At end, we rerun Procedure Locate1Signal and use median analysis to boost the constant success probability. ${ }^{3}$

Lemma 12.7.13. Let $f_{0} \in \operatorname{region}\left(q^{\prime}\right)$. Let $\beta$ is sampled from $\left[\frac{s t}{4 \Delta}, \frac{s t}{2 \Delta l}\right]$ and let $\gamma$ denote the output of Procedure GetLegaliSample in Algorithm 12.4. Then using the pair of samples $z(\gamma+\beta)$ and $z(\gamma)$, we have
I. for the $q^{\prime}$ with probability at least $1-s, v_{q^{\prime}}$ will increase by one.
II. for any $q$ such that $\left|q-q^{\prime}\right|>3$, with probability at least $1-15 s$, $v_{q}$ will not increase.

Proof. We replace $f_{0}$ by $\theta$ in the rest of the proof. By Lemma 12.7.1, we have that for any $\widehat{\beta} \leq \beta \leq 2 \widehat{\beta}$, Procedure GetLegal1Sample outputs a $\gamma \in[0, T]$ satisfying

$$
\left|z(\gamma+\beta)-z(\gamma) e^{2 \pi \mathbf{i} f_{0} \beta}\right| \leq 0.1(|z(\gamma)|+|z(\gamma+\beta)|)
$$

with probability at least 0.6.
Furthermore, there exists such some constant $g \in(0,1)$ such that with probability $1-g$,

$$
\|\phi(z(\gamma+\beta))-(\phi(z(\gamma))-2 \pi \beta \theta)\|_{\circ} \lesssim \sin ^{-1}\left(\frac{1}{g}\right)
$$

where $\|x-y\|_{\bigcirc}=\min _{z \in \mathbb{Z}}|x-y+2 \pi z|$ denote the "circular distance" between $x$ and $y$. We can set $s=\Theta\left(g^{-1}\right)$. There exists some constant $p=\Theta(s)$, with probability at least $1-p$,

$$
\|o-2 \pi \beta \theta\|_{\bigcirc}<s \pi / 2
$$

[^61]where $o:=\phi(z(\gamma+\beta) / z(\gamma))$. The above equation shows that $o$ is a good estimate for $2 \pi \beta \theta$ with good probability. We will now show that this means the true region $Q_{q^{\prime}}$ gets a vote with large probability.

For each $q^{\prime}$ with $\theta \in\left[l-\frac{\Delta l}{2}+\frac{q^{\prime}-1}{t} \Delta l, l-\frac{\Delta l}{2}+\frac{q^{\prime}}{t} \Delta l\right] \subset[-F, F]$, we have that $\theta_{q^{\prime}}=$ $l-\frac{\Delta l}{2}+\frac{q^{\prime}-0.5}{t} \Delta l$ satisfies that

$$
\theta-\theta_{q^{\prime}} \leq \frac{\Delta l}{2 t}
$$

Note that we sample $\beta$ uniformly at random from $[\widehat{\beta}, 2 \widehat{\beta}]$, then $2 \widehat{\beta}=\frac{s t}{2 \Delta l} \leq \frac{c T}{10 A^{\frac{3}{2}}}$ (Note that $A$ is some constant $>1$ ), which implies that $2 \pi \beta \frac{\Delta l}{2 t} \leq \frac{s \pi}{2}$. Thus, we can show the observation $o$ is close to the true region in the following sense,

$$
\begin{aligned}
& \left\|o-2 \pi \beta \theta_{q^{\prime}}\right\|_{\bigcirc} \\
\leq & \|o-2 \pi \beta \theta\|_{\bigcirc}+\left\|2 \pi \beta \theta-2 \pi \beta \theta_{q^{\prime}}\right\|_{\bigcirc} \text { by triangle inequality } \\
\leq & \frac{s \pi}{2}+2 \pi\left\|\beta \theta-\beta \theta_{q^{\prime}}\right\|_{\bigcirc} \\
\leq & s \pi
\end{aligned}
$$

Thus, $v_{q^{\prime}}$ will increase in each round with probability at least $1-s$.
On the other side, consider $q$ with $\left|q-q^{\prime}\right|>3$. Then $\left|\theta-\theta_{q}\right| \geq \frac{7 \Delta l}{2 t}$, and (assuming $\left.\beta \geq \frac{s t}{4 \Delta l}\right)$ we have

$$
2 \pi \beta\left|\theta-\theta_{q}\right| \geq 2 \pi \frac{s t}{4 \Delta l}\left|\theta-\theta_{q}\right|=\frac{s \pi t}{2 \Delta l}\left|\theta-\theta_{q}\right| \geq \frac{7 s \pi}{4}>\frac{3 s \pi}{2} .
$$

There are two cases: $\left|\theta-\theta_{q}\right| \leq \frac{\Delta l}{s t}$ and $\left|\theta-\theta_{q}\right|>\frac{\Delta l}{s t}$.
First, if $\left|\theta-\theta_{q}\right| \leq \frac{\Delta l}{s t}$. In this case, from the definition of $\beta$ it follows that

$$
2 \pi \beta\left|\theta-\theta_{q}\right| \leq \frac{s \pi t}{\Delta l}\left|\theta-\theta_{q}\right| \leq \pi
$$

Combining the above equations implies that

$$
\operatorname{Pr}\left[2 \pi \beta\left(\theta-\theta_{q}\right) \quad(\bmod 2 \pi) \in\left[-\frac{3 s}{4} 2 \pi, \frac{3 s}{4} 2 \pi\right]\right]=0
$$

Second, if $\left|\theta-\theta_{q}\right|>\frac{\Delta l}{s t}$. We show this claim is true : $\operatorname{Pr}\left[2 \pi \beta\left(\theta-\theta_{q}\right)(\bmod 2 \pi) \in\left[-\frac{3 s}{4} 2 \pi, \frac{3 s}{4} 2 \pi\right]\right] \lesssim$ $s$. To prove it, we apply Lemma 12.6 .3 by setting $\widetilde{T}=2 \pi, \widetilde{\sigma}=2 \pi \beta, \widetilde{\delta}=0, \epsilon=\frac{3 s}{4} 2 \pi$, $A=2 \pi \widehat{\beta}, \Delta f=\left|\theta-\theta_{q}\right|$. By upper bound of Lemma 12.6.3, the probability is at most

$$
\frac{2 \widetilde{\epsilon}}{\widetilde{T}}+\frac{4 \widetilde{\epsilon}}{A \Delta f}=\frac{3 s}{2}+\frac{3 s}{\widehat{\beta} \Delta f} \leq \frac{3 s}{2}+\frac{3 s}{\frac{s t}{4 \Delta l} \frac{\Delta l}{s t}}<15 s
$$

Then in either case, with probability at least $1-15 s$, we have

$$
\left\|2 \pi \beta \theta_{q}-2 \pi \beta \theta\right\|_{\bigcirc}>\frac{3 s}{4} 2 \pi
$$

which implies that $v_{q}$ will not increase.

Lemma 12.7.14. Procedure Locate1Inner in Algorithm 12.4 uses $R_{\mathrm{loc}}$ "legal" samples, and then after Procedure Locate1Signal in Algorithm 12.4 running Procedure Locate1Inner $D_{\text {max }}$ times, it outputs a frequency $\widetilde{f}_{0}$ such that

$$
\left|\widetilde{f}_{0}-f_{0}\right| \lesssim \Delta \cdot \sqrt{T \Delta}
$$

with arbitrarily large constant probability.

Proof. For each observation, $v_{q^{\prime}}$ incremented with probability at least $1-p$ and $v_{q}$ is incremented with probability at most $15 s+p$ for $\left|q-q^{\prime}\right|>3$. The probabilities corresponding to different observations are independent. Then after $R_{\text {loc }}$ observations, there exists some
constant $c<\frac{1}{2}$, for any $q$ such that $\left|q-q^{\prime}\right|>3$,

$$
\begin{aligned}
& \operatorname{Pr}[\text { False region gets more than half votes }] \\
= & \operatorname{Pr}\left[v_{j, q}>R_{\mathrm{loc}} / 2\right] \\
\leq & \binom{R_{\mathrm{loc}}}{R_{\mathrm{loc}} / 2}(15 s+p)^{R_{\mathrm{loc}} / 2} \\
\leq & c^{\Omega\left(R_{\mathrm{loc})}\right)}
\end{aligned}
$$

Similarly, on the other side,

$$
\begin{aligned}
& \operatorname{Pr}[\text { True region gets less than half votes }] \\
= & \operatorname{Pr}\left[v_{j, q^{\prime}}<R_{\mathrm{loc}} / 2\right] \\
\leq & \binom{R_{\mathrm{loc}}}{R_{\mathrm{loc}} / 2}(p)^{R_{\mathrm{loc}} / 2} \\
\leq & c^{\Omega\left(R_{\mathrm{loc}}\right)}
\end{aligned}
$$

Taking the union bound over all the $t$ regions, it gives with probability at least $1-t f^{\Omega\left(R_{\text {loc }}\right)}$ we can find some region $q$ such that $\left|q-q^{\prime}\right|<3$.

If we repeat the above procedure $D_{\max }$ rounds, each round we choose the "False" region with probability at most $1-t c^{\Omega\left(R_{\text {loc }}\right)}$. Thus, taking the union bound over all the $D_{\max }$ rounds, we will report a region has size $\sim \Delta \sqrt{\Delta T}$ and contains $f_{0}$ with probability at least $1-D_{\max } t c^{\Omega\left(R_{\mathrm{loc}}\right)}$.

The reason for not ending up with region that has size $\sim \Delta$ is, the upper bound of the sample range of $\beta$ force us to choose $\beta$ is at most $\lesssim \frac{T}{(\Delta T)^{\frac{3}{2}}}$ by Claim 12.7.8

It remains to explain how to set $D_{\max }, t$, and $R_{\text {loc }}$. At the beginning of the first round, we start with frequency interval of length $2 F$, at the beginning of the last round, we start
with frequency interval of length $t \cdot \Delta \sqrt{T \Delta}$. Each round we do a $t$-ary search, thus

$$
D_{\max }=\log _{t}\left(\frac{2 F}{t \Delta \sqrt{T \Delta}}\right) \leq \log _{t}(F / \Delta)
$$

We can set $R_{\text {loc }} \approx \log _{1 / c}(t / c)$ and $t>D_{\max }$, e.g. $t=\log (F / \Delta)$. Thus, the probability becomes,

$$
1-D_{\max } t c^{\Omega\left(R_{\mathrm{loc}}\right)} \geq 1-t^{2} c^{\Omega\left(R_{\mathrm{loc}}\right)} \geq 1-\operatorname{poly}(1 / t, c)
$$

which is larger than any constant probability.

Using the same parameters setting in the proof of Lemma 12.7.14, we show the running time and sample complexity of Procedure Locate1Signal,

Lemma 12.7.15. Procedure Locate1Signal in Algorithm 12.4 uses
$O(\operatorname{poly}(k, \log (1 / \delta))) \cdot \log (F T)$ samples and runs in $O(\operatorname{poly}(k, \log (1 / \delta))) \cdot \log ^{2}(F T)$ time.

Proof. The number of "legal" observations is

$$
D_{\max } R_{\mathrm{loc}}=O\left(\log _{t}(F / \Delta) \log _{1 / c}(t / c)\right)=O(\log (F / \Delta))
$$

The total number of samples is

$$
R_{\text {est }}+R_{\text {repeat }} D_{\text {max }} R_{\text {loc }}=O\left(T \Delta_{h}\right)^{2}+\left(T \Delta_{h}\right)^{3} \cdot \log (F T)=\operatorname{poly}(k, \log (1 / \delta)) \cdot \log (F T)
$$

where the first step follows by Claim 12.7.10 and Lemma 12.7.1 and the last step follows by the setting of $\Delta_{h}$ in Section 12.12.3.

The running time includes two parts, one is approximately computing $H(t)$ for all the samples, each sample takes poly $(k, \log (1 / \delta))$ time according to Lemma 12.12.5; the other is for each legal sample we need to assign vote to some regions.

$$
\operatorname{poly}(k, \log (1 / \delta)) \cdot\left(R_{\text {est }}+R_{\text {repeat }} D_{\max } R_{\mathrm{loc}}\right)+D_{\max } R_{\mathrm{loc}} t=\operatorname{poly}(k, \log (1 / \delta)) \log ^{2}(F T)
$$

Lemma 12.7.16 only achieves constant success probability, using median analysis we can boost the success probability,

Lemma 12.7.16. Let $\widetilde{f}_{0}$ denote the frequency output by Procedure FrequencyRecovERY1CLUSTER in Algorithm 12.5, then with probability at least $1-2^{-\Omega(k)}$,

$$
\left|\widetilde{f}_{0}-f_{0}\right| \lesssim \Delta \sqrt{T \Delta}
$$

Proof. Because of Procedure FrequencyRecovery1Cluster taking the median of $O(k)$ independent results by repeating algorithm Locate1Signal $O(k)$ times. Each sample $L_{r}$ is close to $\widetilde{f}_{0}$ with sufficiently large probability. Thus, using the Chernoff bound will output $\widetilde{f}_{0}$ with probability $1-2^{-\Omega(k)}$ such that

$$
\left|\widetilde{f}_{0}-f_{0}\right| \lesssim \Delta \sqrt{T \Delta}
$$

Combining Lemma 12.7.16 with the sample complexity and running time in Lemma 12.7.14, we are able to finish the proof of Theorem 12.7.4.
12.7.5 The full signal, after multiplying by $H$ and convolving with $G$, is oneclustered.

The goal of this section is to prove Lemma 12.7.5. We fix $f^{*} \in[-F, F]$ satisfying (12.17) in this section. We first define a good hashing $(\sigma, b)$ of $f^{*}$ as follows.

Definition 12.7.2. We say that a frequency $f^{*}$ is well-isolated under the hashing $(\sigma, b)$ if, for $j=h_{\sigma, b}\left(f^{*}\right)$, we have that the signal

$$
\widehat{z}^{(j)}=\widehat{x \cdot H} \cdot \widehat{G}_{\sigma, b}^{(j)}
$$

satisfies, over the interval $\overline{I_{f^{*}}}=(-\infty, \infty) \backslash\left(f^{*}-\Delta, f^{*}+\Delta\right)$,

$$
\int_{\overline{I_{f^{*}}}}\left|\widehat{z}^{(j)}(f)\right|^{2} d f \lesssim \epsilon \cdot T \mathcal{N}^{2} / k
$$

For convenience, we simplify $z^{(j)}$ by using $z$ in the rest of this section.

Lemma 12.7.17. Let $f^{*}$ be any frequency. Then $f^{*}$ is well-isolated by a hashing $(\sigma, b)$ with probability $\geq 0.9$ given $B=\Theta(k)$ and $\sigma \in\left[\frac{1}{B \Delta}, \frac{2}{B \Delta}\right]$ chosen uniformly at random.

Proof. For any other frequency $f^{\prime}$ in $x^{*}$, its contribution in $\widehat{z}$ depends on how far it is from $f^{*}$. Either it is:

- Within $\Delta$ of $f^{*}, f^{\prime}$ and $f^{*}$ will be mapped into the same bucket with probability at least 0.99.
- Between $\Delta$ and $1 / \sigma$ far, from Claim 12.6.2, $f^{\prime}$ and $f^{*}$ will always mapped into different buckets. Hence $f^{\prime}$ always contributes in the $\frac{\epsilon \delta}{k}$ region of Property III in Lemma 12.6.5
about filter function $(G(t), \widehat{G}(f))$, i.e., it contributes at most $\frac{\epsilon \delta}{k} \cdot \int_{f^{\prime}-\Delta}^{f^{\prime}+\Delta}|\widehat{x \cdot H}|^{2} \mathrm{~d} f$. Overall it will contribute

$$
\frac{\epsilon \delta}{k} \cdot \int|\widehat{x \cdot H}|^{2} \mathrm{~d} f=\frac{\epsilon \delta}{k} \int|x \cdot H|^{2} \mathrm{~d} t .
$$

- More than $1 / \sigma$ far, in which case they contribute in the same region with probability at most $3 / B$. By a union bound, it is at most $3 k / B \leq 0.01$

Without loss of generality, we assume $\operatorname{supp}(\widehat{g \cdot H}) \cap \operatorname{supp}\left(\widehat{x^{*} \cdot H}\right)=\emptyset$, otherwise we treat it as a part of $x^{*} \cdot H$. We first consider frequency $f^{*} \in \widehat{x^{*} \cdot H}$ under $G_{\sigma, b}^{(j)}$.

Lemma 12.7.18. Let $f^{*}$ satisfying $\int_{f^{*}-\Delta}^{f^{*}+\Delta}\left|\widehat{x^{*} \cdot H}(f)\right|^{2} \mathrm{~d} f \geq T \mathcal{N}^{2} / k$ and $\widehat{z}=\widehat{x^{*} \cdot H} \cdot \widehat{G}_{\sigma, b}^{(j)}$ where $j=h_{\sigma, b}\left(f^{*}\right)$. If $f^{*}$ is well-isolated, then $z$ and $\widehat{z}$ satisfying Property I(in Definition 12.7.1), i.e.,

$$
\int_{0}^{T}|z(t)|^{2} \mathrm{~d} t \geq(1-\epsilon) \int_{-\infty}^{+\infty}|z(t)|^{2} \mathrm{~d} t
$$

Proof. We first notice that $z(t)=x^{*}(t) \cdot H(t) * G_{\sigma, b}^{(j)}(t)$ and lower bound $\int_{-\infty}^{+\infty}|z(t)|^{2} \mathrm{~d} t$ as
follows:

$$
\begin{align*}
& \int_{-\infty}^{+\infty}\left|x^{*}(t) \cdot H(t) * G_{\sigma, b}^{(j)}(t)\right|^{2} \mathrm{~d} t \\
= & \int_{-\infty}^{+\infty}\left|\widehat{x^{*} \cdot H}(f) \cdot \widehat{G}_{\sigma, b}^{(j)}(f)\right|^{2} \mathrm{~d} f \quad \text { by FT } \\
\geq & \left.\int_{f_{0}-\Delta}^{f_{0}+\Delta} \widehat{\mid x^{*} \cdot H}(f) \cdot \widehat{G}_{\sigma, b}^{(j)}(f)\right|^{2} \mathrm{~d} f \\
\geq & (1-\delta)^{2} \int_{f_{0}-\Delta}^{f_{0}+\Delta}\left|\widehat{x^{*} \cdot H}(f)\right|^{2} \mathrm{~d} f \\
\geq & (1-\delta)^{2} T \mathcal{N}^{2} / k \\
\geq & 0.9 \frac{\delta}{k} \int_{0}^{T}\left|x^{*}(t)\right|^{2} \mathrm{~d} t \tag{12.19}
\end{align*}
$$

We give an upper bound $\int_{-\infty}^{0}|z(t)|^{2} \mathrm{~d} t+\int_{T}^{+\infty}|z(t)|^{2} \mathrm{~d} t \lesssim \epsilon \frac{\delta}{k} \int_{0}^{T}\left|x^{*}(t) H(t)\right|^{2} \mathrm{~d} t$ in the rest of this proof.

Consider the case $t<0$, by definition of Convolution,

$$
z^{(j)}(t)=x^{*}(t) \cdot H(t) * G_{\sigma, b}^{(j)}(t)=\int_{-\infty}^{+\infty} G_{\sigma, b}^{(j)}(t-\tau) \cdot\left(x^{*} \cdot H\right)(\tau) \mathrm{d} \tau
$$

Without loss of generality, we can shift the original signal and $H(t)$ from $[0, T]$ to $[-T / 2, T / 2]$, by Property of $H(t)$, we know that if $s_{3} T / 2 \leq|t| \leq T / 2$, then $H(t) \leq 2^{-O \Theta(\ell)}$. Note that $G(t)$ is compact and has support $D B$, we also assume its compact region is $[-D B / 2, D B / 2]$ (Recall that $D=\frac{l}{\alpha \pi}$ ).

Thus, by definition of convolution,

$$
\begin{aligned}
& z(t) \\
= & \int_{-D B \sigma / 2}^{D B \sigma / 2} G_{\sigma, b}^{(j)}(s) \cdot(x \cdot H)(t-\tau) \mathrm{d} \tau \\
= & \frac{1}{\sigma} \int_{-D B \sigma / 2}^{D B \sigma / 2} G(s / \sigma) e^{2 \pi \mathrm{i} s(j / B-\sigma b) / \sigma} \cdot(x \cdot H)(t-\tau) \mathrm{d} \tau \\
\leq & \frac{1}{\sigma} \int_{-D B \sigma / 2}^{D B \sigma / 2}|G(\tau / \sigma)| \cdot|(x \cdot H)(t-\tau)| \mathrm{d} \tau \\
\leq & \left(\frac{1}{\sigma} \int_{-D B \sigma / 2}^{D B \sigma / 2}|G(\tau / \sigma)| \mathrm{d} \tau\right) \cdot\left(\max _{|\tau| \leq D B \sigma / 2}|(x \cdot H)(t-\tau)|\right)
\end{aligned}
$$

So, if $t \notin[-T / 2, T / 2]$, then $t-s \notin[-T / 2+D B \sigma / 2, T / 2-D B \sigma / 2]$. By Property V of $G(t),|G(t)| \leq \operatorname{poly}(k, \log (1 / \delta))$. Because of the parameter setting ${ }^{4}$, we have the fact $\left[-T s_{3} / 2, T s_{3} / 2\right] \subseteq[-T / 2+D B \sigma / 2, T / 2-D B \sigma / 2] \subseteq[-T / 2, T / 2]$. Thus, we know $T(1-$ $\left.s_{3}\right) / 2>D B \sigma / 2$, then for any $t-\tau \in[-T / 2,-T / 2+D B \sigma / 2] \cup[T / 2-D B \sigma / 2, T / 2]=S$, then
$|z(t)|^{2} \lesssim\left(D B \sigma \cdot \frac{1}{\sigma} \cdot \operatorname{poly}(k, \log (1 / \delta))\right)^{2} \cdot 2^{-\Theta(\ell)} \cdot k^{4} \cdot\left\|x^{*}(t)\right\|_{T}^{2} \lesssim \operatorname{poly}(k, \log (1 / \delta)) \cdot 2^{-\Theta(\ell)} \cdot\left\|x^{*}(t)\right\|_{T}^{2}$.

Thus, taking the integral over $S$,

$$
\int_{S}|z(t)|^{2} \mathrm{~d} t \lesssim|S| \cdot 2^{-\Theta(\ell)} \operatorname{poly}(k, \log (1 / \delta)) \cdot\left\|x^{*}(t)\right\|^{2} \lesssim 2^{-\Theta(\ell)} T\left\|x^{*}(t) \cdot H(t)\right\|_{T}^{2}
$$

By property of filter function $H(t), \widehat{H}(f)$, we have

$$
|(x \cdot H)(t)|^{2} \leq\left(\frac{t}{T}\right)^{-\ell}\left\|x^{*}(t) \cdot H(t)\right\|_{T}^{2} \text { if } t \geq 3 T
$$

[^62]Thus for any constant $\epsilon$,

$$
\begin{equation*}
\int_{-\infty}^{-T / 2}|z(t)|^{2} \mathrm{~d} t+\int_{T / 2}^{+\infty}|z(t)|^{2} \mathrm{~d} t \lesssim 2^{-\ell} T\left\|x^{*}(t) \cdot H(t)\right\|_{T}^{2} \leq 0.9 \epsilon \cdot \frac{\delta}{k} \int_{-T / 2}^{T / 2}\left|x^{*}(t)\right|^{2} \mathrm{~d} t \tag{12.20}
\end{equation*}
$$

where the last inequality follows by $\ell \gtrsim k \log (k / \delta)$. Shifting the interval from $[-T / 2, T / 2]$ to $[0, T]$, the same result is still holding. Combining Equation (12.19) and (12.20) completes the proof of Property II.

We consider frequency $f^{*} \in \widehat{g \cdot H}$ under $G_{\sigma, b}^{(j)}$ and show the energy of noise $g(t)$ is evenly distributed over $B$ bins on expectation.

Lemma 12.7.19. Given any noise $g(t):[0, T] \rightarrow \mathbb{C}$ and $g(t)=0, \forall t \notin[0, T]$. We have, $\forall j \in[B]$,

$$
\underset{\sigma, b}{\mathbb{E}}\left[\int_{-\infty}^{+\infty}\left|g(t) H(t) * G_{\sigma, b}^{(j)}(t)\right|^{2} \mathrm{~d} t\right] \lesssim \frac{1}{B} \int_{-\infty}^{+\infty}|g(t) H(t)|^{2} \mathrm{~d} t
$$

Proof. Because of Fourier Transform preserves $\ell_{2}$ norm, it suffices to prove

$$
\underset{\sigma, b}{\mathbb{E}}\left[\int_{-\infty}^{+\infty}\left|\widehat{g \cdot H}(f) \cdot \widehat{G}_{\sigma, b}^{(j)}(f)\right|^{2} \mathrm{~d} f\right] \lesssim \frac{1}{B} \int_{-\infty}^{+\infty}|\widehat{g \cdot H}(f)|^{2} \mathrm{~d} f
$$

Since $\widehat{G}_{\sigma, b}^{(j)}(f)$ is a periodic function and outputs at most 1 on $O(1 / B)$ fraction of the period, and outputs $\leq \delta$ on other part. Thus, for any frequency $f$, we have

$$
\underset{\sigma, b}{\mathbb{E}}\left[\left|\widehat{G}_{\sigma, b}^{(j)}(f)\right|^{2}\right] \lesssim \frac{1}{B}
$$

Thus, we have

$$
\begin{aligned}
& \underset{\sigma, b}{\mathbb{E}}\left[\int_{-\infty}^{+\infty}\left|\widehat{g \cdot H}(f) \cdot \widehat{G}_{\sigma, b}^{(j)}(f)\right|^{2} \mathrm{~d} f\right] \\
\leq & \underset{\sigma, b}{\mathbb{E}}\left[\int_{-\infty}^{+\infty}|\widehat{g \cdot H}(f)|^{2} \cdot\left|\widehat{G}_{\sigma, b}^{(j)}(f)\right|^{2} \mathrm{~d} f\right] \\
= & \int_{-\infty}^{+\infty}|\widehat{g \cdot H}(f)|^{2} \cdot \underset{\sigma, b}{\mathbb{E}}\left[\left|\widehat{G}_{\sigma, b}^{(j)}(f)\right|^{2}\right] \mathrm{d} f \\
\leq & \int_{-\infty}^{+\infty}|\widehat{g \cdot H}(f)|^{2} \mathrm{~d} f \cdot \max _{f}\left[\underset{\sigma, b}{\mathbb{E}}\left|\widehat{G}_{\sigma, b}^{(j)}(f)\right|^{2}\right] \\
\lesssim & \frac{1}{B} \int_{-\infty}^{+\infty}|\widehat{g \cdot H}(f)|^{2} \mathrm{~d} f,
\end{aligned}
$$

which completes the proof.

Proof of Lemma 12.7.5. Let $j=h_{\sigma, b}\left(f^{*}\right)$, signal

$$
\begin{equation*}
\widehat{z}=\widehat{x \cdot H} \cdot \widehat{G}_{\sigma, b}^{(j)} \tag{12.21}
\end{equation*}
$$

and region $I_{f^{*}}=\left(f^{*}-\Delta, f^{*}+\Delta\right)$ with complement $\overline{I_{f^{*}}}=(-\infty, \infty) \backslash I_{f^{*}}$. From Property I of $G$ in Lemma 12.6.5, we have that

$$
\widehat{G}_{\sigma, b}^{(l)}(f) \gtrsim 1
$$

for all $f \in I_{f^{*}}$, so by (12.17)

$$
\int_{I_{f^{*}}}|\widehat{z}(f)|^{2} \mathrm{~d} f \geq T \mathcal{N}^{2} / k
$$

On the other hand, $f^{*}$ is will-isolated with probability 0.9 :

$$
\int_{\overline{I_{f^{*}}}}|\widehat{z}(f)|^{2} \mathrm{~d} f \lesssim \epsilon T \mathcal{N}^{2} / k
$$

Hence, $\widehat{z}$ satisfies the Property I(in Definition 12.7.1) of one-mountain recovery. Combining Lemma 12.7.18 and Lemma 12.7.19, we know that $\left(x^{*} \cdot H\right) * G_{\sigma, b}^{(j)}$ always satisfies Property $\mathrm{II}\left(\right.$ in Definition 12.7.1) and $\int_{-\infty}^{+\infty}\left|g(t) H(t) * G_{\sigma, b}^{(j)}(t)\right|^{2} \mathrm{~d} t$ is less than $20 T \mathcal{N}^{2} / B \leq \epsilon T \mathcal{N}^{2} / k$ with probability at least 0.95 , which indicates that $z=\left(x^{*}+g\right) \cdot H * G_{\sigma, b}^{(j)}$ satisfies Property II(in Definition 12.7.1) with probability 0.95 .

### 12.7.6 Frequency recovery of $k$-clustered signals

The goal of this section is to prove that the frequencies found by Procedure FreQUENCYRECOVERYKCLUSTER in Algorithm 12.8 have some reasonable guarantee.

We first notice that Lemma 12.7.5 and Lemma 12.7.2 imply the following lemma by a union bound.

Lemma 12.7.20. Let $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$. We observe $x(t)=x^{*}(t)+g(t)$, where $\|g(t)\|_{T}^{2} \leq$ $c\left\|x^{*}(t)\right\|_{T}^{2}$ for a sufficiently small constant $c$ and define $\mathcal{N}^{2}:=\|g(t)\|_{T}^{2}+\delta\left\|x^{*}(t)\right\|_{T}^{2}$. Then Procedure OneStage returns a set $L$ of $O(k)$ frequencies that covers the heavy frequencies of $x^{*}$. In particular, for any $f^{*}$ with

$$
\begin{equation*}
\int_{f^{*}-\Delta}^{f^{*}+\Delta}|\widehat{x \cdot H}(f)|^{2} \mathrm{~d} f \geq T \mathcal{N}^{2} / k \tag{12.22}
\end{equation*}
$$

there will exist an $\tilde{f} \in L$ satisfying $\left|f^{*}-\widetilde{f}\right| \lesssim \sqrt{T \Delta} \cdot \Delta T$ with probability 0.99 .
Lemma 12.7.21. Let $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$ and $R=O(k)$. We observe $x(t)=x^{*}(t)+g(t)$, where $\|g(t)\|_{T}^{2} \leq c\left\|x^{*}(t)\right\|_{T}^{2}$ for a sufficiently small constant $c$ and choose $\mathcal{N}^{2}:=\|g(t)\|_{T}^{2}+$ $\delta\left\|x^{*}(t)\right\|_{T}^{2}$. Then Algorithm MultipleStages returns a set $L$ of $O(k)$ frequencies that approximates the heavy frequencies of $x^{*}$. In particular, with probability $1-2^{-\Omega(k)}$, for any $f^{*}$ such that

$$
\begin{equation*}
\int_{f^{*}-\Delta}^{f^{*}+\Delta}|\widehat{x \cdot H}(f)|^{2} \mathrm{~d} f \geq T \mathcal{N}^{2} / k \tag{12.23}
\end{equation*}
$$

there will exist an $\tilde{f} \in L$ satisfying $\left|f^{*}-\widetilde{f}\right| \lesssim \sqrt{T \Delta} \Delta$.

Proof. Let $A \subset[-F, F]$ denote the set of frequencies $f^{*}$ satisfying Equation (12.22). Let $A^{\prime} \subset[-F, F]$ denote a net of $A$ of distance 2 $\Delta$, so the intervals used in Equation (12.22) for
each $f^{*} \in A^{\prime}$ are disjoint. Then

$$
\left|A^{\prime}\right| \leq 2 k+k=3 k
$$

because each frequency in $x^{*}$ contributes to at most two of the intervals, and the total mass of $\widehat{g}$ is at most $k$ times the threshold $T \mathcal{N}^{2}$.

Let $L_{1}, \ldots, L_{R}$ be the results of $R$ rounds of Algorithm OneStage. We say that a frequency $f \in A^{\prime}$ is successfully recovered in round $r$ if there exists an $\tilde{f} \in L_{r}$ such that $|f-\widetilde{f}| \leq \Delta_{a}$, where

$$
\Delta_{a}=\Delta \sqrt{T \Delta} \lesssim \sqrt{T \Delta} \Delta .
$$

By Lemma 12.7.20, each frequency is successfully recovered with 0.8 probability in each round. Then by the Chernoff bound, with $1-2^{-\Omega(k)}$ probability, every $f \in A^{\prime}$ will be successfully recovered in at least $0.6 R$ rounds.

Then, by Lemma 12.7.22, we output a set $L$ of $O(B)$ frequencies such that every $f \in A^{\prime}$ is within $\Delta_{a}$ of some $\tilde{f} \in L$. Hence every $f \in A$ is within $2 \Delta_{a}$ of some $\tilde{f} \in L$.

Lemma 12.7.22. Let $L_{1}, \ldots, L_{R}$ by sets of frequencies and $f^{*}$ be any frequency. Then $L=\operatorname{MERGEDStagES}\left(L_{1}, \ldots, L_{R}\right)$ is a set of $2 \frac{\sum\left|L_{r}\right|}{R}$ frequencies satisfying

$$
\min _{\widetilde{f} \in L}\left|f^{*}-\widetilde{f}\right| \leq \operatorname{median}_{r \in[R]} \min _{f \in L_{r}}\left|f^{*}-f\right|
$$

Proof. The algorithm is to take the union, sort, and take every $\frac{R}{2}$ th entry of the sorted list.
Let $\Delta=\operatorname{median}_{r \in[R]} \min _{f \in L_{r}}\left|f^{*}-f\right|$. We have that at least $R / 2$ different $f \in \bigcup_{r} L_{r}$ lie within $\Delta$ of $f^{*}$. This set forms a sequential subsequence of the sorted list of frequencies, so our output will include one.

### 12.7.7 Time and sample complexity of frequency recovery of $k$-clustered signals

 The goal of this section is to show that Procedure FrequencyRecoveryKCluster takes$\operatorname{poly}(k, \log (1 / \delta)) \log (F T)$ samples, and runs in $\operatorname{poly}(k, \log (1 / \delta)) \log ^{2}(F T)$ time.
In order to analyze the running time and sample complexity. We need to extend the one-cluster version Procedure GetLegal1Sample and GetEmpirical1Energy (in Algorithm 12.3) to $k$-cluster version GetLegalKSample and GetEmpiricalKEnergy (in Algorithm 12.7) ${ }^{5}$,

Lemma 12.7.23. Procedure GetLegalKSample in Algorithm 12.7 runs Procedure HashToBins $R_{\text {repeat }}=O\left((T \Delta)^{3}\right)$ times to output two vectors $\widehat{v}, \widehat{v}^{\prime} \in \mathbb{C}^{B}$ such that, for each $j \in[B]$,

$$
\left|\widehat{v}_{j}-\widehat{v}_{j}^{\prime} e^{2 \pi \mathrm{i} f_{j} \beta}\right| \leq 0.08\left(\left|\widehat{v}_{j}\right|+\left|\widehat{v}_{j}^{\prime}\right|\right)
$$

holds with probability at least 0.6.

Using the definition of $z$ in Definition 12.7.2.

Claim 12.7.24. Procedure GetEmpiricalKEnergy in Algorithm 12.7 runs Procedure HashTobins $R_{\mathrm{est}} O\left((T \Delta)^{2}\right)$ times to output a vector $z_{\mathrm{emp}} \in \mathbb{R}^{B}$ such that, for each $j \in[B]$,

$$
z_{\mathrm{emp}}^{j} \in\left[0.8\left\|z^{(j)}\right\|_{T}, 1.2\left\|z^{(j)}\right\|_{T}\right]
$$

holds with probability at least 0.9.

[^63]Claim 12.7.25. Algorithm LocateKSignal in Algorithm 12.6 uses $O(\operatorname{poly}(k, \log (1 / \delta))$ • $\log (F T))$, and runs in $O\left(\operatorname{poly}(k, \log (1 / \delta)) \cdot \log ^{2}(F T)\right)$.

Proof. We first calculate the number of samples. All the samples is basically all the Fourier samples, each time needs $B \log (k / \delta)$. In total it calls HashToBins $O\left(R_{\text {est }}+R_{\text {repeat }} D_{\text {max }} R_{\text {loc }}\right)$ times where $D_{\max } R_{\text {loc }}=\Theta(\log (F T))$ by similar analysis as one-cluster frequency recovery. Thus, the total number of samples is

$$
\left(R_{\text {est }}+R_{\text {repeat }} D_{\max } R_{\text {loc }}\right) B \log (k / \delta)=\operatorname{poly}(k, \log (1 / \delta)) \cdot \log (F T)
$$

Then, we analyze the running time.
The expected running time includes the following parts: the first part is running Procedure HashToBins $O\left(R_{\text {est }}+R_{\text {repeat }} D_{\text {max }} R_{\text {loc }}\right)$ times, each run takes $O(B \log (k / \delta)+$ $B \log B)$ samples. For each such sample we need $\operatorname{poly}(k, \log (1 / \delta))$ time to compute $H(t)$ according to Lemma 12.12 .5 and there are $\operatorname{poly}(k, \log (1 / \delta)) \log (F T))$ many samples; the second part is updating the counter $v$, which takes $O\left(D_{\max } R_{l o c} B t\right)$ time. Thus, in total

$$
\begin{aligned}
& \operatorname{poly}(k, \log (1 / \delta)) \cdot O\left(R_{\text {est }}+R_{\text {repeat }} D_{\max } R_{\mathrm{loc}}\right) \cdot O(B \log (k / \delta)+B \log B)+O\left(D_{\max } R_{\text {loc }} B t\right) \\
= & \operatorname{poly}(k, \log (1 / \delta)) \cdot \log ^{2}(F T),
\end{aligned}
$$

where by similar analysis as one-cluster recovery, $t=\Theta(\log (F T))$ and $D_{\max } R_{\text {loc }}=\Theta(\log (F T))$.

To boost the success probability, Procedure MultipleStages reruns Procedure LocateKSignal $O(k)$ times. At the end, Procedure FrequencyRecoveryKCluster combining Procedure MultipleStages and MergedStages directly, and the running time
and sample complexity of MultipleStages are dominating MergedStages. Thus we have

Lemma 12.7.26. Procedure FrequencyRecoveryKCluster in Algorithm 12.8 uses

$$
\operatorname{poly}(k, \log (1 / \delta)) \cdot \log (F T)
$$

samples and runs in $\operatorname{poly}(k, \log (1 / \delta)) \cdot \log ^{2}(F T)$ time.

### 12.8 One-cluster Signal Recovery

### 12.8.1 Overview

In this section, we consider $x^{*}$ whose frequencies in $\widehat{x}^{*}$ are in the range $\left[f_{0}-\Delta^{\prime}, f_{0}+\right.$ $\left.\Delta^{\prime}\right]$ for some frequency $f_{0}$ and $\Delta^{\prime}>0$ and provide an algorithm to approximate it by a polynomial.

We fix $T$ in this section and recall that $\langle f(t), g(t)\rangle_{T}:=\frac{1}{T} \int_{0}^{T} f(t) \bar{g}(t) \mathrm{d} t$ such that $\left\|e^{2 \pi \mathbf{i} f_{i} t}\right\|_{T}=\sqrt{\left\langle e^{2 \pi \mathbf{i} f_{i} t}, e^{2 \pi \mathbf{i} f_{i} t}\right\rangle_{T}}=1$. For convenience, given $\sum_{j=1}^{k} v_{j} e^{2 \pi \mathbf{i} f_{j} t}$, we say the frequency gap of this signal is $\min _{i \neq j}\left|f_{i}-f_{j}\right|$.

For simplicity, we first consider frequencies clustered around 0 . The main technical lemma in this section is that any signal $x^{*}$ with bounded frequencies in $\widehat{x}^{*}$ can be approximated by a low-degree polynomial on $[0, T]$.

Lemma 12.2.3. For any $\Delta>0$ and any $\delta>0$, let $x^{*}(t)=\sum_{j \in[k]} v_{j} e^{2 \pi \mathbf{i} f_{j} t}$ where $\left|f_{j}\right| \leq \Delta$ for each $j \in[k]$. There exists a polynomial $P(t)$ of degree at most

$$
d=O\left(T \Delta+k^{3} \log k+k \log 1 / \delta\right)
$$

such that

$$
\left\|P(t)-x^{*}(t)\right\|_{T}^{2} \leq \delta\left\|x^{*}\right\|_{T}^{2}
$$

One direct corollary is that when $\widehat{x}^{*}$ are in the range $\left[f_{0}+\Delta^{\prime}, f_{0}+\Delta^{\prime}\right]$, we can approximate $x^{*}$ by $P(t) \cdot e^{2 \pi \mathbf{i} f_{0} t}$ for some low degree polynomial $P$.

We give an overview of this section first. We first show some technical tools in Section 12.8.2, 12.8.3. In Section 12.8.4, using those tools, we can show for any $k$-Fourier-sparse signal, there exists another $k$-Fourier-sparse signal with bounded frequency gap close to the
original signal. In Section 12.8.5, we show that for any $k$-Fourier-sparse signal with bounded frequency gap, then there exists a low degree polynomial close to it. In Section 12.8.6, we show how to transfer low degree polynomial back to a Fourier-sparse signal. Combining all the above steps finishes the proof of Lemma 12.2.3.

We apply Theorem 12.7 .4 of frequency estimation on $x^{*}$ to obtain an estimation $\tilde{f}_{0}$ of $f_{0}$ and use Theorem 12.4.5 on the approximation $Q(t) e^{2 \pi \mathrm{i} \tilde{\mathrm{f}}_{0} t}$ of $x^{*}$ to recover the signal. We summarize this result as follows.

Theorem 12.8.1 (One-cluster Signal Recovery). Let $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$ where $\forall j \in[k], \mid f_{j}-$ $f_{0} \mid \leq \Delta$ and $x(t)=x^{*}(t)+g(t)$ be our observable signal. For any $\delta>0$ and any $T>0$, let $\mathcal{N}^{2}:=\|g\|_{T}^{2}+\delta\left\|x^{*}\right\|_{T}^{2}$. Procedure CFT1Culster in Algorithm 12.5 finds a polynomial $P(t)$ of degree at most $d=O\left(\left(T \Delta_{h}+T \Delta\right)^{1.5}+k^{3} \log k+k \log 1 / \delta\right)$ and a frequency $\widetilde{f}_{0}$ such that

$$
\begin{equation*}
\left\|P(t) \cdot e^{2 \pi \mathrm{i} \tilde{\mathrm{f}}_{0} t}-x^{*}(t)\right\|_{T}^{2} \lesssim \mathcal{N}^{2} \tag{12.24}
\end{equation*}
$$

The algorithm uses $O(k d)+\operatorname{poly}(k, \log (1 / \delta)) \log (F T)$ samples, run in $O\left(k d^{\omega}\right)+\operatorname{poly}(k, \log (1 / \delta)) \log ^{2}(F T)$ time, and succeeds with probability at least $1-2^{-\Omega(k)}$.

Proof. We apply the algorithm in Theorem 12.7.4 to obtain an estimation $\widetilde{f}_{0}$ with poly $(k) \log (F T)$ samples and poly $(k) \log ^{2}(F T)$ running time such that $\left|\widetilde{f}_{0}-f_{0}\right| \lesssim\left(\Delta_{h}+\Delta\right) \sqrt{T\left(\Delta_{h}+\Delta\right)}$ holds with probability at least $1-2^{-\Omega(k)}$. Notice that $\left|f_{j}-\widetilde{f}_{0}\right| \leq\left|f_{j}-f_{0}\right|+\left|\widetilde{f}_{0}-f_{0}\right| \lesssim\left(T\left(\Delta_{h}+\Delta\right)\right)^{1.5}$.

We consider $x^{\prime}(t)=e^{-2 \pi \mathbf{i} \tilde{f}_{0} t} x(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathbf{i}\left(f_{j}-\widetilde{f_{0}}\right) t}$. By Lemma 12.2.3, there exists a polynomial $P(t)$ of degree at most

$$
d=O\left(\left(T \Delta_{h}+T \Delta\right)^{1.5}+k^{3} \log k+k \log 1 / \delta\right)
$$

such that it approximates $x^{\prime}$ by

$$
\left\|P(t)-x^{\prime}(t)\right\|_{T} \leq \frac{\delta}{4}\left\|x^{\prime}(t)\right\|_{T}=\frac{\delta}{4}\left\|x^{*}(t)\right\|_{T}
$$

which indicates $\left\|Q(t)-e^{-2 \pi \mathbf{i} \tilde{f}_{0} t} \cdot x^{*}(t)\right\|_{T} \leq \frac{\delta}{4}\left\|x^{*}(t)\right\|_{T}$.
Because we can sample $x(t)$, we can also sample $e^{-2 \pi \mathrm{i} \widetilde{\mathrm{f}}_{0} t} \cdot x(t)=Q(t)+g^{\prime}(t)$ for $g^{\prime}(t)=e^{-2 \pi \mathbf{i} \tilde{\mathrm{f}}_{0} t} \cdot g(t)+\left(e^{-2 \pi \mathbf{i} \widetilde{f}_{0} t} \cdot x^{*}(t)-Q(t)\right)$. Hence we apply the algorithm in Theorem 12.4.5 and choose $R=O(k)$ in that proof. Then Procedure RobustPolynomialLearning ${ }^{+}$ takes $O(k d)$ samples and $O\left(k d^{\omega}\right)$ time to find a degree $d$ polynomial $P(t)$ approximating $Q(t)$ such that

$$
\|P(t)-Q(t)\|_{T} \lesssim\left\|g^{\prime}(t)\right\|_{T},
$$

holds with probability at least $1-2^{-\Omega(k)}$. It indicates

$$
\left\|P(t)-e^{-2 \pi \mathrm{i} \tilde{\mathrm{f}}_{0} t} \cdot x^{*}(t)\right\|_{T} \lesssim\|P(t)-Q(t)\|_{T}+\left\|Q(t)-x^{*}(t)\right\| \lesssim \delta\left\|x^{*}(t)\right\|_{T}+\|g(t)\|_{T} \bar{\sim} \mathcal{N}
$$

Therefore we know $\left\|e^{2 \pi \mathrm{i} \widetilde{\mathrm{f}}_{0} t} \cdot P(t)-x^{*}(t)\right\|_{T}^{2} \lesssim \mathcal{N}^{2}$.

### 12.8.2 Bounding the Gram matrix determinant

We define Gram matrix for $e^{2 \pi \mathbf{i} f_{1} t}, e^{2 \pi \mathbf{i} f_{2} t}, \cdots, e^{2 \pi \mathbf{i} f_{k} t}$ and provide lower/upper bounds for its determinant.

Definition 12.8.1 (Gram matrix). We define $\operatorname{Gram}_{f_{1}, \cdots, f_{k}}$ to be

$$
\left[\begin{array}{cccc}
\left\langle e^{2 \pi \mathbf{i} f_{1} t}, e^{2 \pi \mathbf{i} f_{1} t}\right\rangle_{T} & \left\langle e^{2 \pi \mathbf{i} f_{1} t}, e^{2 \pi \mathbf{i} f_{2} t}\right\rangle_{T} & \cdots & \left\langle e^{2 \pi \mathbf{i} f_{1} t}, e^{2 \pi \mathbf{i} f_{k} t}\right\rangle_{T} \\
\left\langle e^{2 \pi \mathbf{i} \mathbf{f}_{2} t}, e^{2 \pi \mathbf{i} f_{1} t}\right\rangle_{T} & \left\langle e^{2 \pi \mathbf{i} f_{2} t}, e^{2 \pi \mathbf{i} f_{2} t}\right\rangle_{T} & \cdots & \left\langle e^{2 \pi \mathbf{i} \mathrm{f}_{2} t}, e^{2 \pi \mathbf{i} \mathrm{f}_{k} t}\right\rangle_{T} \\
\cdots & \cdots & \cdots & \cdots \\
\left\langle e^{2 \pi \mathbf{i} f_{k} t}, e^{2 \pi \mathbf{i} f_{1} t}\right\rangle_{T} & \left\langle e^{2 \pi \mathbf{i} \mathrm{f}_{k} t}, e^{2 \pi \mathbf{i} \mathrm{i}_{2} t}\right\rangle_{T} & \cdots & \left\langle e^{2 \pi \mathbf{i} f_{k} t}, e^{2 \pi \mathbf{i} f_{k} t}\right\rangle_{T}
\end{array}\right]
$$

Note that the above matrix is a Hermitian matrix with complex entries, thus both its determinant and all eigenvalues are in $\mathbb{R}$.

We defer the proof of the following Theorem to Section 12.10.1.
Theorem 12.8.2. For real numbers $\xi_{1}, \ldots, \xi_{k}$, let $G_{\xi_{1}, \ldots, \xi_{k}}$ be the matrix whose $(i, j)$-entry is

$$
\int_{-1}^{1} e^{2 \pi \mathbf{i}\left(\xi_{i}-\xi_{j}\right) t} d t
$$

Then

$$
\operatorname{det}\left(G_{\xi_{1}, \ldots, \xi_{k}}\right)=2^{\widetilde{O}\left(k^{2}\right)} \prod_{i<j} \min \left(\left|\xi_{i}-\xi_{j}\right|^{2}, 1\right)
$$

We use the following corollary in this section.
Corollary 12.8.3. There exists a universal constant $\alpha>0$ such that, for any $T>0$ and real numbers $f_{1}, \cdots, f_{k}$, the $k \times k$ Gram matrix of $e^{2 \pi \mathbf{i} \mathrm{f}_{1} t}, e^{2 \pi \mathbf{i} f_{2} t}, \cdots, e^{2 \pi \mathbf{i} f_{k} t}$ whose $(i, j)$-entry is

$$
\operatorname{Gram}_{f_{1}, \cdots, f_{k}}(i, j)=\left\langle e^{2 \pi \mathbf{i} f_{i} t}, e^{2 \pi \mathbf{i} \mathrm{f}_{j} t}\right\rangle_{T}=\frac{1}{T} \int_{0}^{T} e^{2 \pi \mathbf{i}\left(f_{i}-f_{j}\right) t} \mathrm{~d} t
$$

satisfies

$$
k^{-\alpha k^{2}} \prod_{i<j} \min \left(\left(\left|f_{i}-f_{j}\right| T\right)^{2}, 1\right) \leq \operatorname{det}\left(\operatorname{Gram}_{f_{1}, \cdots, f_{k}}\right) \leq k^{\alpha k^{2}} \prod_{i<j} \min \left(\left(\left|f_{i}-f_{j}\right| T\right)^{2}, 1\right) .
$$

Based on Corollary 12.8.3, we show the coefficients of a $k$-Fourier-sparse signal can be upper bounded by the energy $\|x\|_{T}^{2}$.

Lemma 12.2.2. There exists a universal constant $c>0$ such that for any $x(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$ with frequency gap $\eta=\min _{i \neq j}\left|f_{i}-f_{j}\right|$,

$$
\|x(t)\|_{T}^{2} \geq k^{-c k^{2}} \min \left((\eta T)^{2 k}, 1\right) \sum_{j=1}^{k}\left|v_{j}\right|^{2}
$$

Proof. Let $\overrightarrow{v_{i}}$ denote the vector $e^{2 \pi \mathrm{i} f_{i} t}$ and $V=\left\{\overrightarrow{v_{1}}, \cdots, \overrightarrow{v_{k}}\right\}$. Notice that $\left\|\overrightarrow{v_{i}}\right\|_{T}^{2}=\left\langle\overrightarrow{v_{i}}, \overrightarrow{v_{i}}\right\rangle=1$. For each $\overrightarrow{v_{i}}$, we define $\overrightarrow{v_{i}}$ to be the projection of $\overrightarrow{v_{i}}$ into the linear subspace $\operatorname{span}\left\{V \backslash \overrightarrow{v_{i}}\right\}=$ $\operatorname{span}\left\{\vec{v}_{1}, \cdots, \vec{v}_{i-1}, \vec{v}_{i+1}, \cdots, \vec{v}_{k}\right\}$ and $\vec{v}_{i}^{\perp}=\vec{v}_{i}-\vec{v}_{i}^{\|}$which is orthogonal to $\operatorname{span}\left\{V \backslash \vec{v}_{i}\right\}$ by the definition.

Therefore from the orthogonality,

$$
\|x(t)\|_{T}^{2} \geq \max _{j \in[k]}\left\{\left|v_{j}\right|^{2} \cdot\left\|\vec{v}_{j}^{\perp}\right\|_{T}^{2}\right\} \geq \frac{1}{k} \sum_{j=1}^{k}\left|v_{j}\right|^{2} \cdot\left\|\vec{v}_{j}^{\perp}\right\|_{T}^{2}
$$

It is enough to estimate $\left\|\vec{v}_{j}^{\perp}\right\|_{T}^{2}$ from Claim 12.3.11:

$$
\left\|\vec{v}_{j}^{\perp}\right\|_{T}^{2}=\frac{\operatorname{det}(\operatorname{Gram}(V))}{\operatorname{det}\left(\operatorname{Gram}\left(V \backslash \vec{v}_{i}\right)\right)} \geq k^{-2 \alpha k^{2}} \prod_{j \neq i} \min \left(\left(f_{j}-f_{i}\right) T, 1\right)^{2} \geq k^{-2 \alpha k^{2}}(\eta T)^{2 k-2}
$$

where we use Corollary 12.8.3 to lower bound it in the last step.

### 12.8.3 Perturbing the frequencies does not change the subspace much

We show that for a $k$-Fourier-sparse signal with unboundedly close frequency gap, there always exists another $k$-Fourier-sparse signal with slightly separated gap.

Lemma 12.8.4 (Slightly Shifting one Frequency). There is a universal constant $C_{0}>0$ such that for any $x(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathbf{i} f_{j} t}$ and any frequency $f_{k+1}$, there always exists

$$
x^{\prime}(t)=\sum_{j=1}^{k-1} v_{j}^{\prime} e^{2 \pi \mathrm{i} f_{j} t}+v_{k+1}^{\prime} e^{2 \pi \mathbf{i} f_{k+1} t}
$$

with $k$ coefficients $v_{1}^{\prime}, v_{2}^{\prime}, \cdots, v_{k-1}^{\prime}, v_{k+1}^{\prime}$ satisfying

$$
\left\|x^{\prime}(t)-x(t)\right\|_{T} \leq k^{C_{0} k^{2}} \cdot\left(\left|f_{k}-f_{k+1}\right| T\right) \cdot\|x(t)\|_{T}
$$

Proof. We abuse the notation $e^{2 \pi \mathrm{i} f_{j} t}$ to denote a vector in the linear subspace. We plan to shift $f_{k}$ to $f_{k+1}$ and define

$$
\begin{aligned}
V & =\left\{e^{2 \pi \mathrm{i} \mathrm{f}_{1} t}, \cdots, e^{2 \pi \mathbf{i} f_{k-1} t}, e^{2 \pi \mathbf{i} f_{k} t}\right\} \\
V^{\prime} & =\left\{e^{2 \pi \mathbf{i} f_{1} t}, \cdots, e^{2 \pi \mathrm{i} f_{k-1} t}, e^{2 \pi \mathbf{i} f_{k+1} t}\right\} \\
U & =\left\{e^{2 \pi \mathbf{i} f_{1} t}, \cdots, e^{2 \pi \mathbf{i} f_{k-1} t}\right\} \\
W & =\left\{e^{2 \pi \mathbf{i} f_{1} t}, \cdots, e^{2 \pi \mathbf{i} f_{k-1} t}, e^{2 \pi \mathbf{i} f_{k} t}, e^{2 \pi \mathbf{i} f_{k+1} t}\right\}
\end{aligned}
$$

where $f_{1}, f_{2}, \cdots, f_{k}$ are original frequencies in $x$. The idea is to show that any vector in the linear subspace $\operatorname{span}\{V\}$ is close to some vector in the linear subspace span $\left\{V^{\prime}\right\}$.

For convenience, we use $\vec{u}^{\|}$to denote the projection of vector $e^{2 \pi \mathrm{i} f_{k} t}$ to the linear subspace $\operatorname{span}\{U\}=\operatorname{span}\left\{e^{2 \pi \mathbf{i} f_{1} t}, \cdots, e^{2 \pi \mathbf{i} f_{k-1} t}\right\}$ and $\vec{w}^{\|}$denote the projection of vector $e^{2 \pi \mathbf{i} f_{k+1} t}$
to this linear subspace $\operatorname{span}\{U\}$. Let $\vec{u}^{\perp}=e^{2 \pi \mathbf{i} f_{k} t}-\vec{u}^{\|}$and $\vec{w}^{\perp}=e^{2 \pi \mathbf{i} f_{k+1} t}-\vec{w}^{\|}$be their orthogonal part to span $\{U\}$.

From the definition $e^{2 \pi \mathbf{i} f_{k} t}=\vec{u}^{\|}+\vec{u}^{\perp}$ and $\vec{u}^{\|} \in \operatorname{span}\{U\}=\operatorname{span}\left\{e^{2 \pi \mathbf{i} f_{1} t}, \cdots, e^{2 \pi \mathbf{i} f_{k-1} t}\right\}$, we rewrite the linear combination

$$
x(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathbf{i} f_{j} t}=\sum_{j=1}^{k-1} \alpha_{j} e^{2 \pi \mathbf{i} f_{j} t}+v_{k} \cdot \vec{u}^{\perp}
$$

for some scalars $\alpha_{1}, \cdots, \alpha_{k-1}$.
We will substitute $\vec{u}^{\perp}$ by $\vec{w}^{\perp}$ in the above linear combination and find a set of new coefficients. Let $\vec{w}^{\perp}=\vec{w}_{1}+\vec{w}_{2}$ where $\vec{w}_{1}=\frac{\left\langle\vec{u}^{\perp}, \vec{w}^{\perp}\right\rangle}{\left\|\vec{u}^{\perp}\right\|_{T}^{2}} \vec{u}^{\perp}$ is the projection of $\vec{w}^{\perp}$ to $\vec{u}^{\perp}$. Therefore $\vec{w}_{2}$ is the orthogonal part of the vector $e^{2 \pi \mathbf{i} f_{k+1} t} \operatorname{to~} \operatorname{span}\{V\}=\operatorname{span}\left\{e^{2 \pi \mathbf{i} f_{1} t}, \cdots, e^{2 \pi \mathbf{i} f_{k-1} t}, e^{2 \pi \mathbf{i} f_{k} t}\right\}$. We use $\delta=\frac{\left\|\vec{w}_{2}\right\|_{T}}{\left\|\vec{w}^{\perp}\right\|_{T}}$ for convenience.

Notice that the $\min _{\beta \in \mathbb{C}} \frac{\left\|\vec{u}^{\perp}-\beta \cdot \vec{w}^{\perp}\right\|_{T}}{\left\|\vec{u}^{\perp}\right\|_{T}}=\delta$ and $\beta^{*}=\frac{\left\langle\vec{u}^{\perp}, \vec{w}^{\perp}\right\rangle}{\left\|\vec{w}^{\perp}\right\|_{T}^{2}}$ is the optimal choice. Therefore we set

$$
x^{\prime}(t)=\sum_{j=1}^{k-1} \beta_{j} e^{2 \pi \mathbf{i} f_{j} t}+v_{k} \cdot \beta^{*} \cdot \vec{w}^{\perp} \in \operatorname{span}\left\{e^{2 \pi \mathbf{i} f_{1} t}, \cdots, e^{2 \pi \mathbf{i} f_{k-1} t}, e^{2 \pi \mathbf{i} f_{k+1} t}\right\}
$$

where the coefficients $\beta_{1}, \cdots, \beta_{k-1}$ guarantee that the projection of $x^{\prime}$ onto $\operatorname{span}\{U\}$ is as same as the projection of $x$ onto $\operatorname{span}\{U\}$. From the choice of $\beta^{*}$ and the definition of $x^{\prime}$,

$$
\left\|x(t)-x^{\prime}(t)\right\|_{T}^{2}=\delta^{2} \cdot\left|v_{k}\right|^{2} \cdot\left\|\vec{u}^{\perp}\right\|_{T}^{2} \leq \delta^{2} \cdot\|x(t)\|_{T}^{2} .
$$

Eventually, we show an upper bound for $\delta^{2}$ from Claim 12.3.11.

$$
\begin{aligned}
\delta^{2} & =\frac{\left\|\vec{w}_{2}\right\|_{T}^{2}}{\left\|\vec{w}^{\perp}\right\|_{T}^{2}} \\
& =\frac{\operatorname{det}\left(\operatorname{Gram}_{W}\right)}{\operatorname{det}\left(\operatorname{Gram}_{V}\right)} / \frac{\operatorname{det}\left(\operatorname{Gram}_{V^{\prime}}\right)}{\operatorname{det}\left(\operatorname{Gram}_{U}\right)} \text { by Claim 12.3.11 } \\
& =\frac{\operatorname{det}\left(\operatorname{Gram}_{W}\right)}{\operatorname{det}\left(\operatorname{Gram}_{V}\right)} \cdot \frac{\operatorname{det}\left(\operatorname{Gram}_{U}\right)}{\operatorname{det}\left(\operatorname{Gram}_{V^{\prime}}\right)} \text { by Corollary 12.8.3 } \\
& \leq k^{4 \alpha k^{2}} \cdot \frac{\prod_{i=1}^{k+1 k+1} \prod_{\substack{ \\
j \neq i}}^{k} \min \left(\left|f_{i}-f_{j}\right| T, 1\right)}{\prod_{i=1}^{k} \prod_{\substack{j=1 \\
j \neq i}}^{k} \min \left(\left|f_{i}-f_{j}\right| T, 1\right)} \cdot \frac{\prod_{i=1}^{k-1} \prod_{\substack{j=1 \\
j \neq i}}^{k-1} \min \left(\left|f_{i}-f_{j}\right| T, 1\right)}{\prod_{i=1}^{k-1} \min \left(\left|f_{i}-f_{j}\right| T, 1\right)} \cdot \prod_{i=1}^{k-1} \min \left(\left|f_{i}-f_{k+1}\right|^{2} T^{2}, 1\right) \\
& =k^{4 \alpha k^{2}\left|f_{k}-f_{k+1}\right|^{2} T^{2}}
\end{aligned}
$$

Lemma 12.8.5. For any $k$ frequencies $f_{1}<f_{2}<\cdots<f_{k}$, there exists $k$ frequencies $f_{1}^{\prime}, \cdots, f_{k}^{\prime}$ such that $\min _{i \in[k-1]} f_{i+1}^{\prime}-f_{i}^{\prime} \geq \eta$ and for all $i \in[k],\left|f_{i}^{\prime}-f_{i}\right| \leq k \eta$.

Proof. We define the new frequencies $f_{i}^{\prime}$ as follows: $f_{1}^{\prime}=f_{1}$ and $f_{i}^{\prime}=\max \left\{f_{i-1}^{\prime}+\eta, f_{i}\right\}$ for $i \in\{2,3, \cdots, k\}$.

### 12.8.4 Existence of nearby $k$-Fourier-sparse signal with frequency gap bounded away from zero

We combine the results in the above section to finish the proof of Lemma 12.2.3. We first prove that for any $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathbf{i} f_{j} t}$, there always exists another $k$-Fourier-sparse signal $x^{\prime}$ close to $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathbf{i} f_{j} t}$ such that the frequency gap in $x^{\prime}$ is at least $\eta \geq 2^{-\operatorname{poly}(k)}$. Then we show how to find a low degree polynomial $P(t)$ approximating $x^{\prime}(t)$. Lemma 12.2.1. There is a universal constant $C_{1}>0$ such that, for any $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$ and any $\delta>0$, there always exist $\eta \geq \frac{\delta}{T} \cdot k^{-C_{1} k^{2}}$ and $x^{\prime}(t)=\sum_{j=1}^{k} v_{j}^{\prime} e^{2 \pi \mathrm{i} f_{j}^{\prime} t}$ satisfying

$$
\left\|x^{\prime}(t)-x^{*}(t)\right\|_{T} \leq \delta\left\|x^{*}(t)\right\|_{T}
$$

with $\min _{i \neq j}\left|f_{i}^{\prime}-f_{j}^{\prime}\right| \geq \eta$ and $\max _{j \in[k]}\left\{\left|f_{j}^{\prime}-f_{j}\right|\right\} \leq k \eta$.

Proof. Using Lemma 12.8.5 on frequencies $f_{1}, \cdots, f_{k}$, we obtain $k$ new frequencies $f_{1}^{\prime}, \cdots, f_{k}^{\prime}$ such that their gap is at least $\eta$ and $\max _{i}\left|f_{i}-f_{i}^{\prime}\right| \leq k \eta$. Next we use the hybrid argument to find $x^{\prime}$.

Let $x^{(0)}(t)=x^{*}(t)$. For $i=1, \cdots, t$, we apply Lemma 12.8 .4 to shift $f_{i}$ to $f_{i}^{\prime}$ and obtain

$$
x^{(i)}(t)=\sum_{j=i+1}^{k} v_{j}^{(i)} e^{2 \pi \mathrm{i} f_{j} t}+\sum_{j=1}^{i} v_{j}^{(i)} e^{2 \pi \mathrm{i} f_{j}^{\prime} t}
$$

From Lemma 12.8.4, we know $\left\|x^{(i)}(t)-x^{(i-1)}(t)\right\|_{T} \leq k^{C_{0} k^{2}}\left(\left|f_{i}-f_{i}^{\prime}\right| T\right)\left\|x^{(i-1)}\right\|_{T}$. Thus we obtain

$$
\left(1-k^{C_{0} k^{2}}(k \eta T)\right)^{i}\left\|x^{(0)}(t)\right\|_{T} \leq\left\|x^{(i)}(t)\right\|_{T} \leq\left(1+k^{C_{0} k^{2}}(k \eta T)\right)^{i}\left\|x^{(0)}(t)\right\|_{T}
$$

which is between $\left[\left(1-i \cdot k^{C_{0} k^{2}}(k \eta T)\right)\left\|x^{(0)}(t)\right\|_{T},\left(1+2 i \cdot k^{C_{0} k^{2}}(k \eta T)\right)\left\|x^{(0)}(t)\right\|_{T}\right]$ for $\eta \leq$ $\frac{1}{5 T} \cdot k^{-C_{1} k^{2}}$ with some $C_{1}>C_{0}$.

At last, we set $x^{\prime}(t)=x^{(k)}(t)$ and bound the distance between $x^{\prime}(t)$ and $x^{*}(t)$ by

$$
\begin{array}{rlr}
\left\|x^{(k)}(t)-x^{(0)}(t)\right\|_{T} & \leq \sum_{i=1}^{k}\left\|x^{(i)}(t)-x^{(i-1)}(t)\right\|_{T} & \text { by triangle inequality } \\
& \leq \sum_{i=1}^{k} k^{C_{0} k^{2}}\left(\left|f_{i}-f_{i}^{\prime}\right| T\right)\left\|x^{(i-1)}(t)\right\|_{T} & \\
& \leq \sum_{i=1}^{k} 2 k^{C_{0} k^{2}}(k \eta T)\left\|x^{(i-1)}(t)\right\|_{T} & \text { by Lemma } 12.8 .4 \\
& \leq k \cdot 2 k^{C_{0} k^{2}}(k \eta T)\left\|x^{*}(t)\right\|_{T} & \\
& \leq \delta\left\|x^{*}(t)\right\|_{T} &
\end{array}
$$

where the last inequality follows by the sufficiently small $\eta$.

### 12.8.5 Approximating $k$-Fourier-sparse signals by polynomials

For any $k$-Fourier-sparse signal with frequency gap bounded away from zero, we show that there exists a low degree polynomial which is close to the original $k$-Fourier-sparse signal in $\|\cdot\|_{T}$ distance.

Lemma 12.8.6 (Existence of low degree polynomial). Let $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i}_{j} t}$, where $\forall j \in$ $[k],\left|f_{j}\right| \leq \Delta$ and $\min _{i \neq j}\left|f_{i}-f_{j}\right| \geq \eta$. There exists a polynomial $Q(t)$ of degree

$$
d=O\left(T \Delta+k \log 1 /(\eta T)+k^{2} \log k+k \log (1 / \delta)\right)
$$

such that,

$$
\begin{equation*}
\left\|Q(t)-x^{*}(t)\right\|_{T}^{2} \leq \delta\left\|x^{*}(t)\right\|_{T}^{2} \tag{12.25}
\end{equation*}
$$

Proof. For each frequency $f_{j}$, let $Q_{j}(t)=\sum_{k=0}^{d-1} \frac{\left(2 \pi \mathrm{i} f_{j} t\right)^{k}}{k!}$ be the first $d$ terms in the Taylor Expansion of $e^{2 \pi \mathrm{i} f_{j} t}$. For any $t \in[0, T]$, we know the difference between $Q_{j}(t)$ and $e^{2 \pi \mathbf{i} f_{j} t}$ is at most

$$
\left|Q_{j}(t)-e^{2 \pi \mathbf{i} f_{j} t}\right| \leq\left|\frac{\left(2 \pi \mathbf{i} f_{j} T\right)^{d}}{d!}\right| \leq\left(\frac{2 \pi T \Delta \cdot e}{d}\right)^{d}
$$

We define

$$
Q(t)=\sum_{j=1}^{k} v_{j} Q_{j}(t)
$$

and bound the distance between $Q$ and $x^{*}$ from the above estimation:

$$
\begin{array}{rlr}
\left\|Q(t)-x^{*}(t)\right\|_{T}^{2} & =\frac{1}{T} \int_{0}^{T}\left|Q(t)-x^{*}(t)\right|^{2} \mathrm{~d} t & \\
& =\frac{1}{T} \int_{0}^{T}\left|\sum_{j=1}^{k} v_{j}\left(Q_{j}(t)-e^{2 \pi \mathrm{i}_{j} t}\right)\right|^{2} \mathrm{~d} t & \\
& \leq 2 k \sum_{j=1}^{k} \frac{1}{T} \int_{0}^{T}\left|v_{j}\right|^{2} \cdot\left|Q_{j}(t)-e^{2 \pi \mathrm{i} \mathrm{f}_{j} t}\right|^{2} \mathrm{~d} t & \\
\text { by triangle inequality } \\
& \leq k \sum_{j=1}^{k}\left|v_{j}\right|^{2} \cdot\left(\frac{2 \pi T \Delta \cdot e}{d}\right)^{2 d} &
\end{array}
$$

On the other hand, from Lemma 12.2.2, we know

$$
\left\|x^{*}(t)\right\|_{T}^{2} \geq(\eta T)^{2 k} \cdot k^{-c k^{2}} \sum_{j}\left|v_{j}\right|^{2}
$$

Because $d=10 \cdot \pi e\left(T \Delta+k \log 1 /(\eta T)+k^{2} \log k+k \log (1 / \delta)\right)$ is large enough, we have $k\left(\frac{2 \pi T \Delta \cdot e}{d}\right)^{2 d} \leq \delta(\eta T)^{2 k} \cdot k^{-c k^{2}}$, which indicates that $\left\|Q(t)-x^{*}(t)\right\|_{T}^{2} \leq \delta\left\|x^{*}\right\|_{T}^{2}$ from all discussion above.

### 12.8.6 Transferring degree- $d$ polynomial to $(d+1)$-Fourier-sparse signal

In this section, we show how to transfer a degree- $d$ polynomial to ( $d+1$ )-Fourier-sparse signal.

Lemma 12.8.7. For any degree-d polynomial $Q(t)=\sum_{j=0}^{d} c_{j} t^{j}$, any $T>0$ and any $\epsilon>0$, there always exist $\gamma>0$ and

$$
x^{*}(t)=\sum_{i=1}^{d+1} \alpha_{i} e^{2 \pi \mathbf{i}(\gamma i) t}
$$

with some coefficients $\alpha_{0}, \cdots, \alpha_{d}$ such that

$$
\forall t \in[0, T],\left|x^{*}(t)-Q(t)\right| \leq \epsilon
$$

Proof. We can rewrite $x^{*}(t)$,

$$
\begin{aligned}
x^{*}(t) & =\sum_{i=1}^{d+1} \alpha_{i} e^{2 \pi \mathbf{i} \gamma i t} \\
& =\sum_{i=1}^{d+1} \alpha_{i} \sum_{j=0}^{\infty} \frac{(2 \pi \mathbf{i} \gamma i t)^{j}}{j!} \\
& =\sum_{j=0}^{\infty} \frac{(2 \pi \mathbf{i} \gamma t)^{j}}{j!} \sum_{i=1}^{d+1} \alpha_{i} \cdot i^{j} \\
& =\sum_{j=0}^{d} \frac{(2 \pi \mathbf{i} \gamma t)^{j}}{j!} \sum_{i=1}^{d+1} \alpha_{i} \cdot i^{j}+\sum_{j=d+1}^{\infty} \frac{(2 \pi \mathbf{i} \gamma t)^{j}}{j!} \sum_{i=1}^{d+1} \alpha_{i} \cdot i^{j} \\
& =Q(t)+\underbrace{\left(\sum_{j=0}^{d} \frac{(2 \pi \mathbf{i} \gamma t)^{j}}{j!} \sum_{i=1}^{d+1} \alpha_{i} \cdot i^{j}-Q(t)\right)}_{C_{1}}+\underbrace{\left(\sum_{j=d+1}^{\infty} \frac{(2 \pi \mathbf{i} \gamma t)^{j}}{j!} \sum_{i=1}^{d+1} \alpha_{i} \cdot i^{j}\right)}_{C_{2}} .
\end{aligned}
$$

Our goal is to show there exists some parameter $\gamma$ and coefficients $\left\{\alpha_{0}, \alpha_{1}, \cdots, \alpha_{d}\right\}$ such that
the term $C_{1}=0$ and $\left|C_{2}\right| \leq \epsilon$. Let's consider $C_{1}$,

$$
C_{1}=\sum_{j=0}^{d}\left(\frac{t}{T}\right)^{j}\left(\frac{(2 \pi \mathbf{i} \gamma T)^{j}}{j!} \sum_{i=1}^{d+1} \alpha_{i} i^{j}-c_{j}\right)
$$

To guarantee $C_{1}=0$, we need to solve a linear system with $d+1$ unknown variables and $d+1$ constraints,

$$
\begin{aligned}
\text { Find } & \alpha_{1}, \alpha_{2}, \cdots \alpha_{d+1} \\
\text { s.t. } & \frac{(2 \pi \mathbf{i} \gamma T)^{j}}{j!} \sum_{i=1}^{d+1} \alpha_{i} i^{j}-c_{j}=0, \forall j \in\{0,1, \cdots, d\}
\end{aligned}
$$

Define $c_{j}^{\prime}=c_{j} j!/(2 \pi \mathbf{i} \gamma)^{j}$, let $\alpha$ and $c^{\prime}$ be the length- $(d+1)$ column vectors with $\alpha_{i}$ and $c_{j}^{\prime}$. Let $A \in \mathbb{R}^{d+1 \times d+1}$ denote the Vandermonde matrix where $A_{i, j}=i^{j}, \forall i, j \in[d+1] \times\{0,1, \cdots, d\}$. Then we need to guarantee $A \alpha=c^{\prime}$. Using the definition of determinant, $\operatorname{det}(A)=\prod_{i<j}|i-j| \leq$ $2^{O\left(d^{2} \log d\right)}$. Thus $\sigma_{\max }(A) \leq 2^{O\left(d^{2} \log d\right)}$ and then

$$
\sigma_{\min }(A)=\frac{\operatorname{det}(A)}{\prod_{i=1}^{d-1} \sigma_{i}} \geq 2^{-O\left(d^{3} \log d\right)} .
$$

We show how to upper bound $\left|\alpha_{i}\right|$,

$$
\max _{i \in[d+1]}\left|\alpha_{i}\right| \leq\|\alpha\|_{2}=\left\|A^{\dagger} c^{\prime}\right\|_{2} \leq\left\|A^{\dagger}\right\|_{2} \cdot\left\|c^{\prime}\right\|_{2} \leq \frac{1}{\sigma_{\min }(A)} \sqrt{d+1} \max _{0 \leq j \leq d} \frac{\left|c_{j}\right| j!}{(2 \pi \gamma T)^{j}}
$$

Plugging the above equation into $C_{2}$, we have

$$
\begin{aligned}
\left|C_{2}\right| & =\left|\sum_{j=d+1}^{\infty} \frac{(2 \pi \mathbf{i} \gamma t)^{j}}{j!} \sum_{i=1}^{d+1} \alpha_{i} \cdot i^{j}\right| \\
& \leq \sum_{j=d+1}^{\infty} \frac{(2 \pi \gamma t)^{j}}{j!} \sum_{i=1}^{d+1}\left|\alpha_{i}\right| \cdot i^{j} \\
& \leq \sum_{j=d+1}^{\infty} \frac{(2 \pi \gamma t)^{j}}{j!}(d+1)^{d+1} \max _{i \in[d+1]}\left|\alpha_{i}\right| \\
& \leq \sum_{j=d+1}^{\infty} \frac{(2 \pi \gamma t)^{j}}{j!}(d+1)^{d+2} \frac{1}{\sigma_{\min }(A)} \frac{d!}{(2 \pi \gamma T)^{d}} \max _{0 \leq j \leq d}\left|c_{j}\right| \\
& \leq \epsilon
\end{aligned}
$$

where the last step follows by choosing sufficiently small

$$
\gamma \lesssim \epsilon /\left(T 2^{\Theta\left(d^{3} \log d\right)} \max _{0 \leq j \leq d}\left|c_{j}\right|\right)
$$

## $12.9 \quad k$-cluster Signal Recovery

### 12.9.1 Overview

In this section, we prove Lemma 12.9.1 as the main technical lemma to finish the proof of main Theorem 12.1.1, which shows how to learn $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$ with noise. Lemma 12.9.1. Let $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathbf{i} f_{j} t}$ and $x(t)=x^{*}(t)+g(t)$ be our observation. For any $\delta>0$ and $T>0$, let $\mathcal{N}^{2}:=\frac{1}{T} \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t+\delta \cdot \frac{1}{T} \int_{0}^{T}\left|x^{*}(t)\right|^{2} \mathrm{~d} t$. For $\Delta=$ $\operatorname{poly}(k, \log (1 / \delta)) / T$, Procedure SignalRecoveryKCluster ${ }^{+}$in Algorithm 12.8 takes $l=$ $O(k)$ frequencies $\widetilde{f}_{1}, \cdots, \widetilde{f}_{l}$ as input and finds l polynomials $Q_{1}, \cdots, Q_{l}$ of degree $d=O\left((T \Delta)^{1.5}+\right.$ $\left.k^{3} \log k+k \log 1 / \delta\right)$ such that

$$
\begin{equation*}
\widetilde{x}(t)=\sum_{j \in[l]} Q_{j}(t) e^{2 \pi \mathrm{i} \tilde{\mathrm{f}}_{j} t} \text { satisfies }\left\|\widetilde{x}(t)-x^{*}(t)\right\|_{T}^{2} \lesssim \mathcal{N}^{2} \tag{12.26}
\end{equation*}
$$

The procedure succeeds with probability at least $1-2^{-\Omega(k)}$, uses poly $(k, \log (1 / \delta)) \cdot \log (F T)$ samples, and runs in $\operatorname{poly}(k, \log (1 / \delta)) \cdot \log ^{2}(F T)$ time.

For any set $W=\left\{t_{1}, \cdots, t_{m}\right\}$ where each $t_{i} \in[0, T]$, we use

$$
\|\vec{v}\|_{W}=\sqrt{\frac{\sum_{i \in W}\left|\vec{v}\left(t_{i}\right)\right|^{2}}{|W|}} \text { for any } \vec{v}:[0, T] \rightarrow \mathbb{C}
$$

in this section. We first show that Procedure SignalRecoveryKCluster succeeds with constant probability, then prove that Procedure SignalRecoveryKCluster ${ }^{+}$succeeds with probability at least $1-2^{-\Omega(k)}$.

### 12.9.2 Heavy clusters separation

Recall the definition of "heavy" clusters.
Definition 12.2.4. Given $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i}_{j} t}$, any $\mathcal{N}>0$, and a filter function $(H, \widehat{H})$ with bounded support in frequency domain. Let $L_{j}$ denote the interval of $\operatorname{supp}\left(\widehat{e^{2 \pi \mathrm{i} f_{j} t} \cdot} H\right)$ for each $j \in[k]$.

Define an equivalence relation $\sim$ on the frequencies $f_{i}$ by the transitive closure of the relation $f_{i} \sim f_{j}$ if $L_{i} \cap L_{j} \neq \emptyset$. Let $S_{1}, \ldots, S_{n}$ be the equivalence classes under this relation.

Define $C_{i}=\underset{f \in S_{i}}{ } L_{i}$ for each $i \in[n]$. We say $C_{i}$ is a "heavy" cluster iff $\int_{C_{i}}\left|\widehat{H \cdot x^{*}}(f)\right|^{2} \mathrm{~d} f \geq$ $T \cdot \mathcal{N}^{2} / k$.

By reordering $C_{i}$, we can assume $\left\{C_{1}, C_{2}, \cdots, C_{l}\right\}$ are heavy clusters, where $l \leq n \leq k$. Claim 12.2.5. Given $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$ and any $\mathcal{N}>0$, let $H$ be the filter function defined in Section 12.12 .1 and $C_{1}, \cdots, C_{l}$ be the heavy clusters from Definition 12.2.4. For

$$
S=\left\{j \in[k] \mid f_{j} \in C_{1} \cup \cdots C_{l}\right\},
$$

we have $x^{(S)}(t)=\sum_{j \in S} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$ approximating $x^{*}$ within distance $\left\|x^{(S)}(t)-x^{*}(t)\right\|_{T}^{2} \lesssim \mathcal{N}^{2}$.
Proof. Let $x^{(\bar{S})}(t)=\sum_{j \in[k] \backslash S} v_{j} e^{2 \pi \mathbf{i} j_{j} t}$. Notice that $\left\|x^{*}-x^{(S)}\right\|_{T}^{2}=\left\|x^{(\bar{S})}\right\|_{T}^{2}$.
From the property VI of filter function $(H, \widehat{H})$ in Section 12.12.1, we have

$$
\int_{-\infty}^{+\infty}\left|x^{(\bar{S})}(t) \cdot H(t)\right|^{2} \mathrm{~d} t \geq 0.9 \int_{0}^{T}\left|x^{(\bar{S})}(t)\right|^{2} \mathrm{~d} t=0.9 \cdot T\left\|x^{\bar{S}}\right\|_{T}^{2}
$$

From Definition 12.2.4, we have

$$
\begin{aligned}
\int_{-\infty}^{+\infty}\left|x^{(\bar{S})}(t) \cdot H(t)\right|^{2} \mathrm{~d} t & =\int_{-\infty}^{+\infty}\left|\widehat{x^{(\bar{S})} \cdot H}(f)\right|^{2} \mathrm{~d} f \\
& =\int_{[-\infty,+\infty] \backslash C_{1} \cup \ldots \cup C_{l}}\left|\widehat{x^{*} \cdot H}(f)\right|^{2} \mathrm{~d} f \\
& \leq k \cdot T \mathcal{N}^{2} / k
\end{aligned}
$$

Overall, we have $\left\|x^{(\bar{S})}\right\|_{T}^{2} \lesssim \mathcal{N}^{2}$.

From the guarantee of Theorem 12.2.6, for any $j \in S, \min _{i \in[l]}\left|f_{j}-\widetilde{f}_{i}\right| \leq \Delta \sqrt{\Delta T}$. From now on, we focus on the recovery of $x^{(S)}$, which is enough to approximate $x^{*}$ from the above claim. Because we are looking for $\widetilde{x}$ approximating $x^{(S)}$ within distance $O\left(\mathcal{N}^{2}\right)$, from Lemma 12.2.1, we can assume there is a frequency gap $\eta \geq \frac{\delta}{10 T} k^{-O\left(k^{2}\right)}$ among $x^{(S)}$.

### 12.9.3 Approximating clusters by polynomials

In this section, we show how to approximate $x^{(S)}$ by $x^{\prime}(t)=\sum_{i \in[l]} e^{2 \pi \mathbf{i} \tilde{f_{i}} t} P_{i}(t)$ where $P_{1}, \cdots, P_{l}$ are low degree polynomials.

Claim 12.9.2. For any $x^{(S)}(t)=\sum_{j \in S} v_{j} e^{2 \pi \mathbf{i} f_{j} t}$ with a frequency gap $\eta=\min _{i \neq j}\left|f_{i}-f_{j}\right|$ and $l$ frequencies $\tilde{f}_{1}, \cdots, \widetilde{f_{l}}$ with the property $\forall j \in S, \min _{i \in[l]}\left|f_{j}-\widetilde{f}_{i}\right| \leq \Delta \sqrt{\Delta T}$, let

$$
d=5 \pi\left((T \Delta)^{1.5}+k^{3} \log k+\log 1 / \delta\right) \text { and } V=\left\{t^{j} e^{2 \pi i \tilde{f}_{i} t} \mid i \in[l], j \in\{0, \cdots, d\}\right\} .
$$

There exists $x^{\prime}(t) \in \operatorname{span}\{V\}$ that approximates $x^{(S)}(t)$ as follows:

$$
\forall t \in[0, T],\left|x^{\prime}(t)-x^{(S)}(t)\right| \leq \delta\left\|x^{(S)}\right\|_{T}
$$

Proof. From Lemma 12.2.2, we know

$$
\left\|x^{(S)}\right\|_{T}^{2} \geq(\eta T)^{2 k} \cdot k^{-c k^{2}} \sum_{j \in S}\left|v_{j}\right|^{2}
$$

For each frequency $f_{j}$, we use $p_{j}$ to denote the index in $[l]$ such that $\left|f_{j}-\widetilde{f}_{p_{j}}\right| \leq \Delta \sqrt{\Delta T}$. We rewrite

$$
x^{(S)}(t)=\sum_{i=1}^{l} e^{2 \pi \mathbf{i} \tilde{f}_{i}}\left(\sum_{j \in S: p_{j}=i} v_{j} e^{2 \pi \mathbf{i}\left(f_{j}-\tilde{f}_{i}\right) t}\right)
$$

For $d=5 \pi\left((T \Delta)^{1.5}+k^{3} \log k+\log 1 / \delta\right)$ and each $e^{2 \pi \mathbf{i}\left(f_{j}-\tilde{f}_{p_{j}}\right) t}$, let $Q_{j}(t)=\sum_{i=0}^{d-1} \frac{\left(2 \pi \mathbf{i}\left(f_{j}-\tilde{f}_{p_{j}}\right) t\right)^{i}}{i!}$ be the first $d$ terms in the Taylor Expansion of $e^{2 \pi \mathbf{i}\left(f_{j}-\widetilde{f}_{p_{j}}\right) t}$. For any $t \in[0, T]$, we know the difference between $Q_{j}(t)$ and $e^{2 \pi \mathbf{i}\left(f_{j}-\widetilde{f}_{p_{j}}\right) t}$ is at most

$$
\forall t \in[0, T],\left|Q_{j}(t)-e^{2 \pi \mathbf{i}\left(f_{j}-\tilde{f}_{p_{j}}\right) t}\right| \leq\left|\frac{\left(2 \pi \mathbf{i}\left(f_{j}-\widetilde{f}_{p_{j}}\right) T\right)^{d}}{d!}\right| \leq\left(\frac{8 \pi(\Delta T)^{1.5}}{d}\right)^{d} .
$$

Let $x^{\prime}=\sum_{i=1}^{l} e^{2 \pi \mathrm{i} \tilde{f}_{i} t}\left(\sum_{j \in S: p_{j}=i} v_{j} Q_{j}(t)\right)$. From all discussion above, we know for any $t \in[0, T]$,

$$
\begin{aligned}
\left|x^{\prime}(t)-x^{(S)}(t)\right|^{2} & \leq\left(\sum_{j \in S}\left|v_{j}\right|\left(\frac{8 \pi(T \Delta)^{1.5}}{d}\right)^{d}\right)^{2} \\
& \leq k\left(\frac{8 \pi(T \Delta)^{1.5}}{d}\right)^{2 d} \sum_{j}\left|v_{j}\right|^{2} \\
& \leq \frac{k\left(\frac{8 \pi(T \Delta)^{1.5}}{d}\right)^{2 d}}{(\eta T)^{2 k} \cdot k^{-c k^{2}}}\left\|x^{(S)}\right\|_{T}^{2} \\
& \leq \delta^{2}\left\|x^{(S)}\right\|_{T}^{2}
\end{aligned}
$$

We provide a property of functions in $\operatorname{span}\{V\}$ such that we can use the Chernoff bound and the $\epsilon$-net argument on vectors in $\operatorname{span}\{V\}$.

Claim 12.2.7. For any $\vec{u} \in \operatorname{span}\left\{e^{2 \pi \mathrm{i} \tilde{f}_{i} t} \cdot t^{j} \mid j \in\{0, \cdots, d\}, i \in[l]\right\}$, there exists some universal constants $C_{1} \leq 4$ and $C_{2} \leq 3$ such that

$$
\max _{t \in[0, T]}\left\{|\vec{u}(t)|^{2}\right\} \lesssim(l d)^{C_{1}} \log ^{C_{2}}(l d) \cdot\|\vec{u}\|_{T}^{2}
$$

Proof. From Lemma 12.8.7, we can approximate each polynomial in $\vec{u}$ by a linear combination of $\left\{1, e^{2 \pi \mathbf{i} \cdot \gamma t}, \cdots, e^{2 \pi \mathbf{i} \cdot(\gamma d) t}\right\}$ such that we obtain $u^{*} \in \operatorname{span}\left\{e^{2 \pi \mathbf{i} \cdot(\gamma j) t} \cdot e^{2 \pi \mathbf{i} \tilde{\mathrm{f}}_{i} t} \mid i \in[l], j \in\{0, \cdots, d+1\}\right\}$ for some small $\gamma$ such that $\forall t \in[0, T],\left|\vec{u}(t)-u^{*}(t)\right| \leq 0.01\|\vec{u}\|_{T}$.

From Lemma 12.5.1, we know

$$
\max _{t \in[0, T]}\left|u^{*}(t)\right|^{2} \leq C \cdot\left((l d+1)^{4} \cdot \log ^{3}(l d+1)\right)\left\|u^{*}\right\|_{T}^{2}
$$

For some constant $C^{\prime}$, we have

$$
\max _{t \in[0, T]}|\vec{u}(t)|^{2} \leq C^{\prime}\left((k d)^{C_{1}} \log ^{C_{2}} d\right)\|\vec{u}\|_{T}^{2}
$$

### 12.9.4 Main result, with constant success probability

In this section, we show that the output $\widetilde{x}$ is close to $x^{\prime}$ with high probability using the $\epsilon$-net argument, which is enough to prove $\|\widetilde{x}-x\|_{T} \lesssim \mathcal{N}^{2}$ from all discussion above. Because we can prove Lemma 12.9.6(which is the main goal of this section), then combining $\left\|x^{\prime}-x^{*}\right\|_{T} \leq\left\|x^{\prime}-x^{(S)}\right\|_{T}+\left\|x^{(S)}-x^{*}\right\|_{T} \lesssim \delta\left\|x^{*}\right\|_{T}$ and Lemma 12.9.6, we have $\| x^{*}-$ $\widetilde{x}\left\|_{T} \lesssim\right\| g\left\|_{T}+\delta\right\| x^{*} \|_{T}$, which finishes the proof of Procedure SignalRecoveryKCluster in Algorithm 12.8 achieving the Equation (12.26) with constant success probability but not $1-2^{-\Omega(k)}$. We will boost the success probability in Section 12.9.5.

We first provide an $\epsilon$-net $\mathcal{P}$ for the unit vectors $\mathcal{Q}=\left\{\vec{u} \in \operatorname{span}\{V\} \mid\|\vec{u}\|_{T}^{2}=1\right\}$ in the linear subspace span $\{V\}$ where $V=\left\{t^{j} \cdot e^{2 \pi i \tilde{f}_{i} t} \mid j \in\{0,1, \cdots, d\}, i \in[l]\right\}$ from the above discussion. Notice that the dimension of $\operatorname{span}\{V\}$ is at most $l(d+1)$.

Claim 12.9.3. There exists an $\epsilon$-net $\mathcal{P} \subset \operatorname{span}\{V\}$ such that

1. $\forall \vec{u} \in Q, \exists \vec{w} \in \mathcal{P},\|\vec{u}-\vec{w}\|_{T} \leq \epsilon$.
2. $|\mathcal{P}| \leq\left(5 \frac{l(d+1)}{\epsilon}\right)^{2 l(d+1)}$.

Proof. Let $\mathcal{P}^{\prime}$ be an $\frac{\epsilon}{l(d+1)}$-net in the unit circle of $\mathbb{C}$ with size at most $\left(4 \frac{l(d+1)}{\epsilon}+1\right)^{2}$, i.e.,

$$
\mathcal{P}^{\prime}=\left\{\frac{\epsilon}{2 l(d+1)} j_{1}+\mathbf{i} \frac{\epsilon}{2 l(d+1)} j_{2}\left|j_{1}, j_{2} \in \mathbb{Z},\left|j_{1}\right| \leq \frac{2 l(d+1)}{\epsilon},\left|j_{2}\right| \leq \frac{2 l(d+1)}{\epsilon}\right\}\right.
$$

Observe that the dimension of $\operatorname{span}\{V\}$ is at most $l(d+1)$. Then we take an orthogonal basis $\vec{w}_{1}, \cdots, \vec{w}_{l(d+1)}$ in $\operatorname{span}\{V\}$ and set

$$
P=\left\{\sum_{i=1}^{l(d+1)} \alpha_{i} \vec{w}_{i} \mid \forall i \in[l(d+1)], \alpha_{i} \in \mathcal{P}^{\prime}\right\} .
$$

Therefore $\mathcal{P}$ is an $\epsilon$-net for $Q$ and $|\mathcal{P}| \leq\left(5 \frac{l(d+1)}{\epsilon}\right)^{2 l(d+1)}$.

We first prove that $W$ is a good estimation for all functions in the $\epsilon$-net $\mathcal{P}$.

Claim 12.9.4. For any $\epsilon>0$, there exists a universal constant $C_{3} \leq 5$ such that for a set $S$ of i.i.d. samples chosen uniformly at random over $[0, T]$ of size $|S| \geq \frac{3(k d)^{C_{3}} \log ^{C_{3}} d / \epsilon}{\epsilon^{2}}$, then with probability at least $1-k^{-k}$, for all $\vec{w} \in \mathcal{P}$, we have

$$
\|\vec{w}\|_{W} \in\left[(1-\epsilon)\|\vec{w}\|_{T},(1+\epsilon)\|\vec{w}\|_{T}\right]
$$

Proof. From Claim 12.2.7 and Lemma 12.3.5, for each $\vec{w} \in \mathcal{P}$,

$$
\operatorname{Pr}\left[\|\vec{w}(t)\|_{W} \notin\left[(1-\epsilon)\|\vec{w}\|_{T},(1+\epsilon)\|\vec{w}\|_{T}\right]\right] \leq 2^{-\frac{|W| \epsilon^{2}}{3(k d)^{C_{11} \log _{2}+0.5}{ }_{d}}} \leq 2^{-k d \log ^{1.5} \frac{d}{\epsilon}}
$$

From the union bound, $\|\vec{w}\|_{W} \in\left[(1-\epsilon)\|\vec{w}\|_{T},(1+\epsilon)\|\vec{w}\|_{T}\right]$ for any $\vec{w} \in \mathcal{P}$ with probability at least $1-\left(\frac{d}{\epsilon}\right)^{-k d \log ^{0.5} d} \cdot|\mathcal{P}| \geq 1-d^{-d}$.

Then We prove that $W$ is a good estimation for all functions in $\operatorname{span}\{V\}$ using the property of $\epsilon$-nets.

Claim 12.9.5. For any $\epsilon>0$, there exists a universal constant $C_{3} \leq 5$ such that for a set $W$ of i.i.d. samples chosen uniformly at random over $[0, T]$ of size $|W| \geq \frac{3(k d)^{C_{3}} \log ^{C_{3}} d / \epsilon}{\epsilon^{2}}$, then with probability at least $1-d^{-d}$, for all $u \in \operatorname{span}\{V\}$, we have

$$
\|\vec{u}\|_{W} \in\left[(1-3 \epsilon)\|\vec{u}\|_{T},(1+3 \epsilon)\|\vec{u}\|_{T}\right]
$$

Proof. We assume that the above claim is true for any $\vec{w} \in \mathcal{P}$. Without loss of generality, we consider $\vec{u} \in Q$ such that $\|\vec{u}\|_{T}=1$.

Let $\vec{w}_{0}$ be the vector in $\mathcal{P}$ that minimizes $\|\vec{w}-\vec{u}\|_{T}$ for all $\vec{w} \in \mathcal{P}$, i.e., $\vec{w}_{0}=\underset{\vec{w} \in \mathcal{P}}{\arg \min } \| \vec{w}-$ $\vec{u} \|_{T}$. Define $\vec{u}_{1}=\vec{u}-\vec{w}_{0}$ and notice that $\left\|\vec{u}_{1}\right\|_{T} \leq \epsilon$ because $\mathcal{P}$ is a $\epsilon$-net. If $\left\|\vec{u}_{1}\right\|_{T}=0$, then we skip the rest of this procedure. Otherwise, we define $\alpha_{1}=\left\|\vec{u}_{1}\right\|_{T}$ and normalize $\widetilde{u}_{1}=\vec{u}_{1} / \alpha_{1}$.

Then we choose $\vec{w}_{1}$ to be the vector in $\mathcal{P}$ that minimizes $\left\|\vec{w}-\widetilde{u}_{1}\right\|_{T}$ for all $\vec{w} \in \mathcal{P}$. Similarly, we set $\vec{u}_{2}=\widetilde{u}_{1}-\vec{w}_{1}$ and $\alpha_{2}=\left\|\vec{u}_{2}\right\|_{T}$. Next we repeat this process for $\widetilde{u}_{2}=\vec{u}_{2} / \alpha_{2}$ and so on. The recursive definition can be summarized in the following sense,

$$
\text { initial : } \quad \widetilde{u}_{0}=\vec{u} \text { and } m=10 \log _{1 / \epsilon}(l d)+1,
$$

For $i \in\{0,1,2, \cdots, m\}: \quad \vec{w}_{i}=\underset{\vec{w} \in \mathcal{P}}{\arg \min }\left\|\vec{w}-\widetilde{u}_{i}\right\|_{T}$,

$$
\vec{u}_{i+1}=\widetilde{u}_{i}-\vec{w}_{i} \text { and } \alpha_{i+1}=\left\|\vec{u}_{i+1}\right\|_{T},
$$

if $\alpha_{i+1}=0$, stop.
if $\alpha_{i+1} \neq 0, \widetilde{u}_{i+1}=\vec{u}_{i+1} / \alpha_{i+1}$ and continue,

Eventually, we have $\vec{u}=\vec{w}_{0}+\alpha_{1} \vec{w}_{1}+\alpha_{1} \alpha_{2} \vec{w}_{2}+\cdots+\prod_{j=1}^{m} \alpha_{j}\left(\vec{w}_{m}+\vec{u}_{m+1}\right)$ where each $\left|\alpha_{i}\right| \leq \epsilon$ and each $\vec{w}_{i}$ is in the $\epsilon$-net $\mathcal{P}$. Notice that $\left\|\vec{u}_{m+1}\right\|_{T} \leq 1$ and $\left\|\vec{u}_{m+1}\right\|_{W} \leq$ $(l d+1)^{3} \cdot\left\|\vec{u}_{m+1}\right\|_{T}$ from Claim 12.2.7. We prove a lower bound for $\|\vec{u}\|_{W}$,

$$
\begin{aligned}
\|\vec{u}\|_{W} & =\left\|\vec{w}_{0}+\alpha_{1} \vec{w}_{1}+\alpha_{1} \alpha_{2} \vec{w}_{2}+\cdots+\prod_{j=1}^{m} \alpha_{j}\left(\vec{w}_{m}+\vec{u}_{m+1}\right)\right\|_{W} \\
& \geq\left\|\vec{w}_{0}\right\|_{W}-\left\|\alpha_{1} \vec{w}_{1}\right\|_{W}-\left\|\alpha_{1} \alpha_{2} \vec{w}_{2}\right\|_{W}-\cdots-\left\|\prod_{j=1}^{m} \alpha_{j} \vec{w}_{m}\right\|_{W}-\left\|\prod_{j=1}^{m} \alpha_{j} \vec{u}_{m+1}\right\|_{W} \\
& \geq(1-\epsilon)-\epsilon(1+\epsilon)-\epsilon^{2}(1+\epsilon)-\cdots-\epsilon^{m}(1+\epsilon)-\epsilon^{m}\left\|\vec{u}_{m+1}\right\|_{W} \\
& \geq 1-\epsilon-\frac{(1+\epsilon) \epsilon}{1-\epsilon}-\epsilon^{m} \cdot(l d+1)^{3} \geq 1-3 \epsilon
\end{aligned}
$$

Similarly, we have $\|\vec{u}\|_{W} \leq 1+3 \epsilon$.

Lemma 12.9.6. With probability at least 0.99 over the $m$ i.i.d samples in $W$,

$$
\left\|x^{\prime}(t)-\widetilde{x}(t)\right\|_{T} \leq 2200\left(\|g(t)\|_{T}+\left\|x^{(S)}(t)-x^{\prime}(t)\right\|_{T}\right) .
$$

Proof. Let $g^{\prime}(t)=g(t)+x^{*}(t)-x^{\prime}(t)$ such that $x(t)=x^{\prime}(t)+g^{\prime}(t)$. Then we choose $\epsilon=0.03$ and bound:

$$
\begin{aligned}
& \left\|x^{\prime}(t)-\widetilde{x}(t)\right\|_{T} \\
\leq & (1+3 \epsilon)\left\|x^{\prime}(t)-\widetilde{x}(t)\right\|_{W} \\
=1.09\left\|x^{\prime}(t)-\widetilde{x}(t)\right\|_{W} & \text { with prob. } 1-2^{-\Omega(d \log d)} \text { by Claim } 12.9 .5 \\
=1.09\left\|x(t)-g^{\prime}(t)-\widetilde{x}(t)\right\|_{W} & \text { by } \epsilon=0.03 \\
\leq 1.09\|x(t)-\widetilde{x}(t)\|_{W}+1.09\left\|g^{\prime}(t)\right\|_{W} & \text { by } x^{\prime}(t)=x(t)-g^{\prime}(t) \\
\leq 1.09\left\|x(t)-x^{\prime}(t)\right\|_{W}+1.09\left\|g^{\prime}(t)\right\|_{W} & \text { by triangle inequality } \\
=2.18\left\|g^{\prime}(t)\right\|_{W} . & \text { by } \widetilde{x}=\underset{y \in \operatorname{span}\{V\}}{\arg \min }\|x-y\|_{W} \\
= & \text { by } x(t)-x^{\prime}(t)=g(t)
\end{aligned}
$$

From the fact that $\mathbb{E}_{W}\left[\left\|g^{\prime}\right\|_{W}\right]=\left\|g^{\prime}\right\|_{T},\left\|g^{\prime}\right\|_{W} \leq 1000\left\|g^{\prime}\right\|_{T}$ with probability at least .999. It indicates $\left\|x^{\prime}(t)-\widetilde{x}(t)\right\|_{T} \leq 2200\left\|g^{\prime}\right\|_{T}$ with probability at least 0.99 from all discussion above.

### 12.9.5 Boosting the success probability

In order to achieve $1-2^{-\Omega(k)}$ for the main theorem, we cannot combine Procedure SignalRecoveryKCluster with FrequencyRecoveryKCluster directly. However, using the similar proof technique in Theorem 12.4.5, we are able to boost the success probability by using Procedure SignalRecoveryKCluster ${ }^{+}$in Algorithm 12.8. It runs Procedure SignalRecoveryKCluster $R=O(k)$ times in parallel for independent fresh samples and report $R$ different $d$-Fourier-sparse signals $\widetilde{x}_{i}(t)$. Then, taking $m=\operatorname{poly}(k)$ new locations $\left\{t_{1}, t_{2}, \cdots, t_{m}\right\}$, and computing $\widetilde{A}$ as before and $\widetilde{b}_{j}$ by taking the median of $\left\{\widetilde{x}_{1}\left(t_{j}\right), \cdots, \widetilde{x}_{R}\left(t_{j}\right)\right\}$. At the end, solving the linear regression for matrix $\widetilde{A}$ and vector $\widetilde{b}$. Thus, we complete the proof of Lemma 12.9.1.

Because we can transfer a degree- $d$ polynomial to a $d$-Fourier-sparse signal by Lemma 12.8.7, the output of Procedure CFTKCluster in Algorithm 12.8 matches the main theorem,

Theorem 12.1.1. Let $x(t)=x^{*}(t)+g(t)$, where $x^{*}$ is $k$-Fourier-sparse signal with frequencies in $[-F, F]$. Given samples of $x$ over $[0, T]$ we can output $\widetilde{x}(t)$ such that with probability at least $1-2^{-\Omega(k)}$,

$$
\left\|\widetilde{x}-x^{*}\right\|_{T} \lesssim\|g\|_{T}+\delta\left\|x^{*}\right\|_{T} .
$$

Our algorithm uses poly $(k, \log (1 / \delta)) \cdot \log (F T)$ samples and poly $(k, \log (1 / \delta)) \cdot \log ^{2}(F T)$ time. The output $\widetilde{x}$ is $\operatorname{poly}(k, \log (1 / \delta))$-Fourier-sparse signal.

### 12.10 Technical Proofs

### 12.10.1 Proof of Theorem 12.8.2

We prove the following Theorem
Theorem 12.8.2. For real numbers $\xi_{1}, \ldots, \xi_{k}$, let $G_{\xi_{1}, \ldots, \xi_{k}}$ be the matrix whose $(i, j)$-entry is

$$
\int_{-1}^{1} e^{2 \pi \mathbf{i}\left(\xi_{i}-\xi_{j}\right) t} d t
$$

Then

$$
\operatorname{det}\left(G_{\xi_{1}, \ldots, \xi_{k}}\right)=2^{\widetilde{O}\left(k^{2}\right)} \prod_{i<j} \min \left(\left|\xi_{i}-\xi_{j}\right|^{2}, 1\right)
$$

First, we note by the Cauchy-Binet formula that the determinant in question is equal to

$$
\begin{equation*}
\int_{-1}^{1} \int_{-1}^{1} \ldots \int_{-1}^{1}\left|\operatorname{det}\left(\left[e^{2 \pi i \xi_{i} t_{j}}\right]_{i, j}\right)\right|^{2} d t_{1} d t_{2} \ldots d t_{k} \tag{12.27}
\end{equation*}
$$

We next need to consider the integrand in the special case when $\sum\left|\xi_{i}\right| \leq 1 / 8$.

Lemma 12.10.1. If $\xi_{i} \in R$ and $t_{j} \in \mathbb{R}, \sum_{i}\left|\xi_{i}\right|\left(\max _{i}\left|t_{i}\right|\right) \leq 1 / 8$ then

$$
\left|\operatorname{det}\left(\left[e^{2 \pi i \xi_{i} t_{j}}\right]_{i, j}\right)\right|=\Theta\left(\frac{(2 \pi)^{\binom{k}{2}} \prod_{i<j}\left|t_{i}-t_{j}\right|\left|\xi_{i}-\xi_{j}\right|}{1!2!\cdots k!}\right)
$$

Proof. Firstly, by adding a constant to all the $t_{j}$ we can make them non-negative. This multiplies the determinant by a root of unity, and at most doubles $\sum_{i}\left|\xi_{i}\right|\left(\max _{i}\left|t_{i}\right|\right)$.

By continuity, it suffices to consider the $t_{i}$ to all be multiples of $1 / N$ for some large integer $N$. By multiplying all the $t_{j}$ by $N$ and all $\xi_{i}$ by $1 / N$, we may assume that all of the $t_{j}$ are non-negative integers with $t_{1} \leq t_{2} \leq \ldots \leq t_{k}$.

Let $z_{i}=\exp \left(2 \pi i \xi_{i}\right)$. Then our determinant is

$$
\operatorname{det}\left(\left[z_{i}^{t_{j}}\right]_{i, j}\right)
$$

which is equal to the Vandermonde determinant times the Schur polynomial $s_{\lambda}\left(z_{i}\right)$ where $\lambda$ is the partition $\lambda_{j}=t_{j}-(j-1)$.

Therefore, this determinant equals

$$
\prod_{i<j}\left(z_{i}-z_{j}\right) s_{\lambda}\left(z_{1}, z_{2}, \ldots, z_{k}\right)
$$

The absolute value of

$$
\prod_{i<j}\left(z_{i}-z_{j}\right)
$$

is approximately $\prod_{i<j}(2 \pi i)\left(\xi_{i}-\xi_{j}\right)$, which has absolute value $(2 \pi)^{\binom{k}{2}} \prod_{i<j}\left|\xi_{i}-\xi_{j}\right|$. We have left to evaluate the size of the Schur polynomial.

By standard results, $s_{\lambda}$ is a polynomial in the $z_{i}$ with non-negative coefficients, and all exponents at most $\max _{j}\left|t_{j}\right|$ in each variable. Therefore, the monomials with non-zero coefficients will all have real part at least $1 / 2$ and absolute value 1 when evaluated at the $z_{i}$. Therefore,

$$
\left|s_{\lambda}\left(z_{1}, \ldots, z_{k}\right)\right|=\Theta\left(\left|s_{\lambda}(1,1, \ldots, 1)\right|\right) .
$$

On the other hand, by the Weyl character formula

$$
s_{\lambda}(1,1, \ldots, 1)=\prod_{i<j} \frac{t_{j}-t_{i}}{j-i}=\frac{\prod_{i<j}\left|t_{i}-t_{j}\right|}{1!2!\ldots k!} .
$$

This completes the proof.

Next we prove our Theorem when the $\xi$ have small total variation.

Lemma 12.10.2. If there exists a $\xi_{0}$ so that $\sum\left|\xi_{i}-\xi_{0}\right|<1 / 8$, then

$$
\operatorname{det}\left(G_{\xi_{1}, \ldots, \xi_{k}}\right)=\Theta\left(\frac{2^{3 k(k-1) / 2} \pi^{k(k-1)} \prod_{i<j}\left|\xi_{i}-\xi_{j}\right|^{2}}{(k!)^{3} \prod_{n=0}^{k-1}(2 n)!}\right)
$$

Proof. By translating the $\xi_{i}$ we can assume that $\xi_{0}=0$.
By the above we have

$$
\Theta\left(\frac{(2 \pi)^{k(k-1)} \prod_{i<j}\left|\xi_{i}-\xi_{j}\right|^{2}}{(1!2!\cdots k!)^{2}}\right) \int_{-1}^{1} \ldots \int_{-1}^{1} \prod_{i<j}\left|t_{i}-t_{j}\right|^{2} d t_{1} \ldots d t_{k} .
$$

We note that by the Cauchy-Binet formula the latter term is the determinant of the matrix $M$ with $M_{i, j}=\int_{-1}^{1} t^{i+j} d t$. This is the Graham matrix associated to the polynomials $t^{i}$ for $0 \leq i \leq k-1$. Applying Graham-Schmidt (without the renormalization step) to this set yields the basis $P_{n} \alpha_{n}$ where $\alpha_{n}=\frac{2^{n}(n!)^{2}}{(2 n)!}$ is the inverse of the leading term of $P_{n}$. This polynomial has norm $\alpha_{n}^{2} 2 /(2 n+1)$. Therefore, the integral over the $t_{i}$ yields

$$
\prod_{n=0}^{k-1} \frac{2^{n+1}(n!)^{2}}{(n+1)(2 n)!}
$$

This completes the proof.

Next we extend this result to the case that all the $\xi$ are within $\operatorname{poly}(k)$ of each other.

Proposition 12.10.3. If there exists a $\xi_{0}$ so that $\left|\xi_{i}-\xi_{0}\right|=\operatorname{poly}(k)$ for all $i$, then

$$
\operatorname{det}\left(G_{\xi_{1}, \ldots, \xi_{k}}\right)=2^{\widetilde{O}\left(k^{2}\right)} \prod_{i<j} \min \left(\left|\xi_{i}-\xi_{j}\right|^{2}, 1\right)
$$

Proof. We begin by proving the lower bound. We note that for $0<x<1$,

$$
\operatorname{det}\left(G_{\xi_{1}, \ldots, \xi_{k}}\right) \geq \int_{-x}^{x} \int_{-x}^{x} \ldots \int_{-1}^{1}\left|\operatorname{det}\left(\left[e^{2 \pi i \xi_{i} t_{j}}\right]_{i, j}\right)\right|^{2} d t_{1} d t_{2} \ldots d t_{k}=x^{k} \operatorname{det}\left(G_{\xi_{1} / x, \xi_{2} / x, \ldots, \xi_{k} / k}\right)
$$

Taking $x=1 / \operatorname{poly}(k)$, we may apply the above Lemma to compute the determinant on the right hand side, yielding an appropriate lower bound.

To prove the lower bound, we note that we can divide our $\xi_{i}$ into clusters, $\mathcal{C}_{i}$, where for any $i, j$ in the same cluster $\left|\xi_{i}-\xi_{j}\right|<1 / k$ and for $i$ and $j$ in different clusters $\left|\xi_{i}-\xi_{j}\right| \geq 1 / k^{2}$. We then note as a property of Graham matrices that

$$
\operatorname{det}\left(G_{\xi_{1}, \ldots, \xi_{k}}\right) \leq \prod_{\mathcal{C}_{i}} \operatorname{det}\left(G_{\left\{\xi_{j} \in \mathcal{C}_{i}\right\}}\right)=2^{\widetilde{O}\left(k^{2}\right)} \prod_{i<j, \text { in same cluster }}\left|\xi_{i}-\xi_{j}\right|^{2}=2^{\widetilde{O}\left(k^{2}\right)} \prod_{i<j}\left|\xi_{i}-\xi_{j}\right|^{2}
$$

This completes the proof.

Finally, we are ready to prove our Theorem.

Proof. Let $I(t)$ be the indicator function of the interval $[-1,1]$.
Recall that there is a function $h(t)$ so that for any function $f$ that is a linear combination of at most $k$ complex exponentials that $|h(t) f(t)|_{2}=\Theta\left(|I(t) f(t)|_{2}\right)$ and so that $\widehat{h}$ is supported on an interval of length $\operatorname{poly}(k)<k^{C}$ about the origin.

Note that we can divide our $\xi_{i}$ into clusters, $\mathcal{C}$, so that for $i$ and $j$ in a cluster $\left|\xi_{i}-\xi_{j}\right|<k^{C+1}$ and for $i$ and $j$ in different clusters $\left|\xi_{i}-\xi_{j}\right|>k^{C}$.

Let $\widetilde{G}_{\xi_{1}, \xi_{2}, \ldots, \xi_{k}^{\prime}}$ be the matrix with $(i, j)$-entry $\int_{\mathbb{R}}|h(t)|^{2} e^{(2 \pi i)\left(\xi_{i}-\xi_{j}\right) t} d t$.
We claim that for any $k^{\prime} \leq k$ that

$$
\operatorname{det}\left(\widetilde{G}_{\xi_{1}, \xi_{2}, \ldots, \xi_{k}^{\prime}}\right)=2^{O\left(k^{\prime}\right)} \operatorname{det}\left(G_{\xi_{1}, \xi_{2}, \ldots, \xi_{k}^{\prime}}\right)
$$

This is because both are Graham determinants, one for the set of functions $I(t) \exp \left((2 \pi i) \xi_{j} t\right)$ and the other for $h(t) \exp \left((2 \pi i) \xi_{j} t\right)$. However since any linear combination of the former has
$L^{2}$ norm a constant multiple of that the same linear combination of the latter, we have that

$$
\widetilde{G}_{\xi_{1}, \xi_{2}, \ldots, \xi_{k}^{\prime}}=\Theta\left(G_{\xi_{1}, \xi_{2}, \ldots, \xi_{k}^{\prime}}\right)
$$

as self-adjoint matrices. This implies the appropriate bound.
Therefore, we have that

$$
\operatorname{det}\left(G_{\xi_{1}, \ldots, \xi_{k}}\right)=2^{O(k)} \operatorname{det}\left(\widetilde{G}_{\xi_{1}, \ldots, \xi_{k}}\right)
$$

However, note that by the Fourier support of $h$ that

$$
\int_{\mathbb{R}}|h(t)|^{2} e^{(2 \pi i)\left(\xi_{i}-\xi_{j}\right) t} d t=0
$$

if $\left|\xi_{i}-\xi_{j}\right|>k^{C}$, which happens if $i$ and $j$ are in different clusters. Therefore $\widetilde{G}$ is block diagonal and hence its determinant equals

$$
\operatorname{det}\left(\widetilde{G}_{\xi_{1}, \ldots, \xi_{k}}\right)=\prod_{\mathcal{e}} \operatorname{det}\left(\widetilde{G}_{\left\{\xi_{j} \in \mathcal{C}_{i}\right\}}\right)=2^{O(k)} \prod_{\mathcal{e}} \operatorname{det}\left(G_{\left\{\xi_{j} \in \mathcal{C}_{i}\right\}}\right)
$$

However the Proposition above shows that

$$
\prod_{\mathcal{C}} \operatorname{det}\left(G_{\left\{\xi_{j} \in \mathcal{C}_{i}\right\}}\right)=2^{\widetilde{O}\left(k^{2}\right)} \prod_{i<j} \min \left(1,\left|\xi_{i}-\xi_{j}\right|^{2}\right)
$$

This completes the proof.

### 12.10.2 Proofs of Lemma 12.5.3 and Lemma 12.5.4

We fix $z_{1}, \cdots, z_{k}$ to be complex numbers on the unit circle and use $Q(z)$ to denote the degree- $k$ polynomial $\prod_{i=1}^{k}\left(z-z_{i}\right)$.
Lemma 12.5.3. Let $Q(z)$ be a degree $k$ polynomial, all of whose roots are complex numbers with absolute value 1. For any integer $n$, let $r_{n, k}(z)=\sum_{l=0}^{k-1} r_{n, k}^{(l)} \cdot z^{l}$ denote the residual polynomial of

$$
r_{n, k}(z) \equiv z^{n} \quad(\bmod Q(z))
$$

Then, each coefficient of $r_{n, k}$ is bounded: $\left|r_{n, k}^{(l)}\right| \leq 2^{k} n^{k-1}$ for any $l$.

Proof. By definition, $r_{n, k}\left(z_{i}\right)=z_{i}^{n}$. From the polynomial interpolation, we have

$$
r_{n, k}(z)=\sum_{i=1}^{k} \frac{\prod_{j \in[k] \backslash i}\left(z-z_{j}\right) z_{i}^{n}}{\prod_{j \in[k] \backslash i}\left(z_{i}-z_{j}\right)} .
$$

Let $\operatorname{Sym}_{S, i}$ be the symmetry polynomial of $z_{1}, \cdots, z_{k}$ with degree $i$ among subset $S \subseteq[k]$, i.e., $\operatorname{Sym}_{S, i}=\sum_{S^{\prime} \subseteq\binom{S}{i}} \prod_{j \in S^{\prime}} z_{j}$. Then

$$
r_{n, k}^{(l)}=(-1)^{k-1-l} \sum_{i=1}^{k} \frac{\operatorname{Sym}_{[k] \backslash i, k-1-l} \cdot z_{i}^{n}}{\prod_{j \in[k] \backslash i}\left(z_{i}-z_{j}\right)} .
$$

We omit $(-1)^{k-1-l}$ in the rest of proof and use induction on $n, k$, and $l$ to prove $\left|r_{n, k}^{(l)}\right| \leq$ $\binom{k-1}{l}\binom{n}{k-1}$.

Base Case of $n$ : For any $n<k$, from the definition, $r(z)=z^{n}$ and $\left|r_{n, k}^{(l)}\right| \leq 1$.
Suppose it is true for any $n<n_{0}$. We consider $r_{n_{0}, k}^{l}$ from now on. When $k=1$, $r_{n, 0}=z_{1}^{n}$ is bounded by 1 because $z_{1}$ is on the unit circle of $\mathbb{C}$.

Given $n_{0}$, suppose the induction hypothesis is true for any $k<k_{0}$ and any $l<k$. For $k=k_{0}$, we first prove that $\left|r_{n_{0}, k_{0}}^{\left(k_{0}-1\right)}\right| \leq\binom{ n_{0}}{k_{0}-1}$ then prove that $\left|r_{n_{0}, k_{0}}^{(l)}\right| \leq\binom{ k_{0}-1}{l}\binom{n_{0}}{k_{0}-1}$ for $l=k_{0}-2, \cdots, 0$.

$$
\begin{aligned}
r_{n_{0}, k_{0}}^{\left(k_{0}-1\right)} & =\sum_{i=1}^{k_{0}} \frac{z_{i}^{n_{0}}}{\prod_{j \in\left[k_{0}\right] \backslash i}\left(z_{i}-z_{j}\right)} \\
& =\sum_{i=1}^{k_{0}-1} \frac{z_{i}^{n_{0}}}{\prod_{j \in\left[k_{0} \backslash \backslash i\right.}\left(z_{i}-z_{j}\right)}+\frac{z_{k_{0}}^{n_{0}}}{\prod_{j \in\left[k_{0}\right] \backslash k_{0}}\left(z_{k_{0}}-z_{j}\right)} \\
& =\sum_{i=1}^{k_{0}-1} \frac{z_{i}^{n_{0}}-z_{i}^{n_{0}-1} z_{k_{0}}+z_{i}^{n_{0}-1} z_{k_{0}}}{\prod_{j \in\left[k_{0}\right] \backslash i}\left(z_{i}-z_{j}\right)}+\frac{z_{k_{0}}^{n_{0}}}{\prod_{j \in k_{0} \backslash k_{0}}\left(z_{k_{0}}-z_{j}\right)} \\
& =\sum_{i=1}^{k_{0}-1}\left(\frac{z_{i}^{n_{0}-1}}{\prod_{j \in\left[k_{0}-1\right] \backslash i}\left(z_{i}-z_{j}\right)}+\frac{z_{i}^{n_{0}-1} z_{k_{0}}}{\prod_{j \in k_{0} \backslash i}\left(z_{i}-z_{j}\right)}\right)+\frac{z_{k_{0}}^{n_{0}}}{\prod_{j \in k_{0} \backslash k_{0}}\left(z_{k_{0}}-z_{j}\right)} \\
& =\left(\sum_{i=1}^{k_{0}-1} \frac{z_{i}^{n_{0}-1}}{\prod_{j \in\left[k_{0}-1\right] \backslash i}\left(z_{i}-z_{j}\right)}\right)+\left(z_{k_{0}} \sum_{i=1}^{k_{0}} \frac{z_{i}^{n_{0}-1}}{\prod_{j \in k_{0} \backslash i}\left(z_{i}-z_{j}\right)}\right) \\
& \left.=r_{n_{0}-1, k_{0}-1}^{\left(k_{0}-2\right)}\right)
\end{aligned}
$$

Hence $\left|r_{n_{0}, k_{0}}^{\left(k_{0}-1\right)}\right| \leq\left|r_{n_{0}-1,\left[k_{0}-1\right]}^{\left(k_{0}-2\right)}\right|+\left|r_{n_{0}-1, k_{0}}^{\left(k_{0}-1\right)}\right| \leq\binom{ n_{0}-2}{k_{0}-2}+\binom{n_{0}-2}{k_{0}-1}=\binom{n_{0}-1}{k_{0}-1}$. For $l<k_{0}-1$, we have

$$
\begin{aligned}
r_{n_{0}, k_{0}}^{(l)}= & \sum_{i=1}^{k_{0}} \frac{\operatorname{Sym}_{\left[k_{0}\right] \backslash i, k_{0}-1-l} \cdot z_{i}^{n_{0}}}{\prod_{j \in\left[k_{0}\right] \backslash i}\left(z_{i}-z_{j}\right)} \text { let } l^{\prime}=k_{0}-1-l \\
= & \sum_{i=1}^{k_{0}-1} \frac{\left(\operatorname{Sym}_{\left[k_{0}-1\right] \backslash i, l^{\prime}}+\operatorname{Sym}_{\left[k_{0}-1\right] \backslash i, l^{\prime}-1} \cdot z_{k_{0}}\right) z_{i}^{n_{0}}}{\prod_{j \in\left[k_{0}\right] \backslash i}\left(z_{i}-z_{j}\right)}+\frac{\operatorname{Sym}_{\left[k_{0}-1\right], l^{\prime}} \cdot z_{k_{0}}^{n_{0}}}{\prod_{j<k_{0}}\left(z_{k_{0}}-z_{j}\right)} \\
= & \sum_{i=1}^{k_{0}-1} \frac{\operatorname{Sym}_{\left[k_{0}-1\right] \backslash i, l^{\prime}} \cdot\left(z_{i}-z_{k_{0}}\right) z_{i}^{n_{0}-1}+\operatorname{Sym}_{\left[k_{0}-1\right] \backslash i, l^{\prime}} \cdot z_{k_{0}} z_{i}^{n_{0}-1}+\operatorname{Sym}_{\left[k_{0}-1\right] \backslash i, l^{\prime}-1} \cdot z_{k_{0}} z_{i}^{n_{0}}}{\prod_{j \in\left[k_{0}\right] \backslash i}\left(z_{i}-z_{j}\right)} \\
& +\frac{\operatorname{Sym}_{\left[k_{0}-1\right], l^{\prime}} \cdot z_{k_{0}}^{n_{0}}}{\prod_{j<k_{0}}\left(z_{k_{0}}-z_{j}\right)} \\
= & \sum_{i=1}^{k_{0}-1} \frac{\operatorname{Sym}_{\left[k_{0}-1\right] \backslash i, l^{\prime}} \cdot\left(z_{i}-z_{k_{0}}\right) z_{i}^{n_{0}-1}+\operatorname{Sym}_{\left[k_{0}-1\right], l^{\prime}} \cdot z_{k_{0}} z_{i}^{n_{0}-1}}{\prod_{j \in\left[k_{0}\right] \backslash i}\left(z_{i}-z_{j}\right)}+\frac{\operatorname{Sym}_{\left[k_{0}-1\right], l^{\prime}} \cdot z_{k_{0}}^{n_{0}}}{\prod_{j<k_{0}}\left(z_{k_{0}}-z_{j}\right)} \\
= & \sum_{i=1}^{k_{0}-1} \frac{\operatorname{Sym}_{\left[k_{0}-1\right] \backslash i, l^{\prime}} z_{i}^{n_{0}-1}}{\prod_{j \in\left[k_{0}-1\right] \backslash i}\left(z_{i}-z_{j}\right)}+\sum_{i=1}^{k_{0}-1} \frac{\operatorname{Sym}_{\left[k_{0}-1\right], l^{\prime}} \cdot z_{k_{0}} z_{i}^{n_{0}-1}}{\prod_{j \in\left[k_{0}\right] \backslash i}\left(z_{i}-z_{j}\right)}+\frac{\operatorname{Sym}_{\left[k_{0}-1\right], l^{\prime}} \cdot z_{k_{0}}^{n_{0}}}{\prod_{j<k_{0}}\left(z_{k_{0}}-z_{j}\right)} \\
= & r_{n_{0}-1, k_{0}-1}^{(l-1)}+\operatorname{Sym}_{\left[k_{0}-1\right], k_{0}-1-l} \cdot z_{k_{0}} \cdot r_{n_{0}-1, k_{0}}^{\left(k_{0}-1\right)}
\end{aligned}
$$

By induction hypothesis, $\left|r_{n_{0}, k_{0}}^{(l)}\right| \leq\binom{ k_{0}-2}{l-1}\binom{n_{0}-1}{k_{0}-2}+\binom{k_{0}-1}{l}\binom{n_{0}-1}{k_{0}-1} \leq\binom{ k_{0}-1}{l}\binom{n_{0}}{k_{0}-1}$.

Now we finish the proof of Lemma 12.5.4.
Lemma 12.5.4. For any $k \in \mathbb{Z}$ and any $z_{1}, \cdots, z_{k}$ on the unit circle of $\mathbb{C}$, there always exists a degree $m=O\left(k^{2} \log k\right)$ polynomial $P(z)=\sum_{j=0}^{m} c_{j} z^{j}$ with the following properties:

Property I $\quad P\left(z_{i}\right)=0, \forall i \in\{1, \cdots, k\}$,
Property II $\quad c_{0}=1$,
Property III $\quad\left|c_{j}\right| \leq 11, \forall j \in\{1, \cdots, m\}$.

Let $m=10 k^{2} \log k$ and $\mathcal{P}$ denote a set of polynomials that has degree at most $m$, and all the coefficients are integers chosen from $\{-5, \cdots,-1,0,1, \cdots, 5\}$, i.e.,

$$
\mathcal{P}:=\left\{P(z)=\sum_{i=0}^{m} \alpha_{i} z^{i}\left|\forall i \in\{0,1, \cdots, m\},\left|\alpha_{i}\right| \leq 2\right\} .\right.
$$

Claim 12.10.4. There exists $P^{*}(z)=\sum_{i=0}^{m} \alpha_{i} z^{i}$ with coefficient $\left|\alpha_{i}\right| \leq 10$ for any $i \in\{0,1, \cdots, m\}$, such that every coefficient of $P^{*}(z) \bmod Q(z)$ is bounded by $2^{-m}$.

Proof. For $P(z)=\sum_{i=0}^{m} \alpha_{i} z^{i} \in \mathcal{P}, P(z) \bmod Q(z) \equiv \sum_{i=0}^{m} \alpha_{i} r_{n, k}(z)$ from the definition $r_{n, k}(z)$. Hence

$$
P(z) \quad \bmod Q(z)=\sum_{i=0}^{m} \alpha_{i} \sum_{l=0}^{k-1} r_{i, k}^{(l)} z^{l}=\sum_{l=0}^{k-1} z^{l} \sum_{i=0}^{m} \alpha_{i} r_{i, k}^{(l)} .
$$

Each coefficient in $P(z) \bmod Q(z)$ is bounded by $\left|\sum_{i=0}^{m} \alpha_{i} r_{i, k}^{(l)}\right| \leq 5 \sum_{i=0}^{m} 2^{k} i^{k-1} \leq 2^{k} m^{k}$.
At the same time, $|\mathcal{P}|=11^{m}$. From the pigeonhole principle and $\left(\frac{2^{k} m^{k}}{\left(2^{-m}\right)^{2}}\right)^{k}<11^{m}$, there exists $P_{1}, P_{2} \in \mathcal{P}$ such that for $P^{*}(z)=P_{1}(z)-P_{2}(z), P^{*}(z) \bmod Q(z)=\sum_{i=0}^{k-1} \gamma_{i} z^{i}$ where each coefficient $\left|\gamma_{i}\right| \leq 2^{-m}$.

Let $r(z)=\sum_{i=0}^{k-1} \gamma_{i} z^{i}=P^{*}(z) \bmod Q(z)$ for convenience. If $P^{*}(0)$ (the constant term of $\left.P^{*}\right)$ is nonzero, then $\left|P^{*}(0)-r(0)\right| \geq 0.99$ from the above lemma. Therefore the polynomial $\frac{P^{*}(z)-r(z)}{P^{*}(0)-r(0)}$ satisfies the three properties in Lemma 12.5.4.

Otherwise, we assume $z^{l}$ is the first term in $P^{*}(z)$ with a non-zero coefficient. Let

$$
r_{-l, k}(z)=z^{-l} \quad \bmod Q(z)=\sum_{i=1}^{k} \frac{\prod_{j \in[k] \backslash i}\left(z-z_{j}\right) z_{i}^{-l}}{\prod_{j \in[k] \backslash i}\left(z_{i}-z_{j}\right)} .
$$

For convenience, we use $z_{S}=\prod_{i \in S} z_{i}$ for any subset $S \subseteq[k]$. Notice that $z_{i}^{-l}=\frac{z_{|k| \backslash i}^{l}}{z_{[k]}^{l}}$. Hence $r_{-l, k}(z)=r_{l, k}^{\prime}(z) / z_{[k]}^{l}$ where $r^{\prime}$ is the polynomial for $k$ units roots $z_{[k] \backslash 1}, \cdots, z_{[k] \backslash k}$. So each coefficients of $r$ is still bounded by $2^{k} l^{k}$, which is less than $2^{-m / 2}$.

Eventually we choose $P^{*}(z) / z^{l}-r(z) \cdot r_{-l, k}(z)$ and renormalize it to satisfy the three properties in Lemma 12.5.4.

### 12.10.3 Proof of Lemma 12.4.3

Lemma 12.4.3. For any degree $d$ polynomial $P(t): \mathbb{R} \rightarrow \mathbb{C}$ with derivative $P^{\prime}(t)$, we have,

$$
\begin{equation*}
\int_{-1}^{1}\left(1-t^{2}\right)\left|P^{\prime}(t)\right|^{2} \mathrm{~d} t \leq 2 d^{2} \int_{-1}^{1}|P(t)|^{2} \mathrm{~d} t . \tag{12.28}
\end{equation*}
$$

Given a degree $d$ polynomial $P(x)$, we rewrite $P(x)$ as a linear combination of the Legendre polynomials:

$$
P(x)=\sum_{i=0}^{d} \alpha_{i} L_{i}(x) .
$$

We use $F_{i}(x)=\left(1-x^{2}\right) L_{i}^{\prime}(x)$ for convenience. From the definition of the Legendre polynomials in the Equation (12.3), $F_{i}^{\prime}(x)=-i(i+1) \cdot L_{i}(x)$ and $F_{i}^{\prime \prime}(x)=-i(i+1) \cdot L_{i}^{\prime}(x)$.

Hence we have

$$
\begin{aligned}
\int_{1}^{-1}\left(1-x^{2}\right)\left|P^{\prime}(x)\right|^{2} \mathrm{~d} x & =\int_{1}^{-1}\left(1-x^{2}\right) P^{\prime}(x) \cdot \overline{P^{\prime}}(x) \mathrm{d} x \\
& =\int_{1}^{-1}\left(\sum_{i \in[d]} \alpha_{i} F_{i}(x)\right) \cdot\left(\sum_{i \in[d]} \overline{\alpha_{i}} \frac{-F_{i}^{\prime \prime}(x)}{i(i+1)}\right) \mathrm{d} x \\
& =\left.\left(\sum_{i \in[d]} \alpha_{i} F_{i}(x)\right) \cdot\left(\sum_{i \in[d]} \overline{\alpha_{i}} \frac{-F_{i}^{\prime}(x)}{i(i+1)}\right)\right|_{-1} ^{1} \\
& +\int_{1}^{-1}\left(\sum_{i \in[d]} \alpha_{i} F_{i}^{\prime}(x)\right) \cdot\left(\sum_{i \in[d]} \overline{\alpha_{i}} \frac{F_{i}^{\prime}(x)}{i(i+1)}\right) \mathrm{d} x \\
& =\int_{1}^{-1}\left(\sum_{i \in[d]} \alpha_{i} \cdot i(i+1) \cdot L_{i}(x)\right) \cdot\left(\sum_{i \in[d]} \overline{\alpha_{i}} \frac{i(i+1) \cdot L_{i}(x)}{i(i+1)}\right) \mathrm{d} x \\
& =\sum_{i \in[d]}\left|\alpha_{i}\right|^{2} i(i+1)\left\|L_{i}\right\|_{T}^{2} \\
& \leq d(d+1)\|P\|_{T}^{2}
\end{aligned}
$$

### 12.10.4 Proof of Lemma 12.6.1

Lemma 12.6.1. $\widehat{P_{\sigma, a, b} x}(\sigma(f-b))=\frac{1}{\sigma} e^{-2 \pi \mathbf{i} \sigma a f} \widehat{x}(f)$ and $\widehat{P_{\sigma, a, b} x}(f)=\frac{1}{\sigma} e^{-2 \pi \mathbf{i} \sigma a(f / \sigma+b)} \widehat{x}(f / \sigma+b)$

Proof. Let's compute the Fourier Transform of $\left(P_{\sigma, a, b} x\right)(t)$,

$$
\begin{array}{rlr} 
& \widehat{P_{\sigma, a, b} x}(f) & \\
= & \int_{-\infty}^{+\infty}\left(P_{\sigma, a, b}(x)\right)(t) e^{-2 \pi \mathbf{i} f t} \mathrm{~d} t & \\
= & \int_{-\infty}^{+\infty} x(\sigma(t-a)) e^{-2 \pi \mathbf{i} \sigma b t} e^{-2 \pi \mathbf{i} f t} \mathrm{~d} t & \\
= & e^{-2 \pi \mathbf{i}(\sigma a b+f a)} \int_{-\infty}^{+\infty} x(\sigma(t-a)) e^{-2 \pi \mathbf{i} \sigma b(t-a)} e^{-2 \pi \mathbf{i} f(t-a)} \mathrm{d} t & \text { by shifting } t \text { by } a \\
= & e^{-2 \pi \mathbf{i}(\sigma a b+f a)} \int_{-\infty}^{+\infty} x(\sigma t) e^{-2 \pi \mathbf{i} \sigma b t} e^{-2 \pi \mathbf{i} f t} \mathrm{~d} t & \text { by replacing } t-a \text { by } t \\
= & \frac{1}{\sigma} e^{-2 \pi \mathbf{i}(\sigma a b+f a)} \int_{-\infty}^{+\infty} x(\sigma t) e^{-2 \pi \mathbf{i} b \sigma t} e^{-2 \pi \mathbf{i} f \sigma t / \sigma} \mathrm{d} \sigma t & \\
= & \frac{1}{\sigma} e^{-2 \pi \mathbf{i}(\sigma a b+f a)} \int_{-\infty}^{+\infty} x(t) e^{-2 \pi \mathbf{i}(b+f / \sigma) t} \mathrm{~d} t & \text { by replacing } t \sigma \text { by } t \\
= & \frac{1}{\sigma} e^{-2 \pi \mathbf{i} a \sigma(f / \sigma+b)} \widehat{x}(f / \sigma+b) & \text { by definition of } \mathrm{FT}
\end{array}
$$

The first result follows immediately by replacing $f / \sigma+b$ by $f^{\prime}$, which gives

$$
\widehat{P_{\sigma, a, b} x}\left(\sigma\left(f^{\prime}-b\right)\right)=\frac{1}{\sigma} e^{-2 \pi \mathbf{i} a \sigma f^{\prime}} \widehat{x}\left(f^{\prime}\right)
$$

Thus, we complete the proof of this Lemma.

### 12.10.5 Proof of Lemma 12.3.5

Lemma 12.3.5. Given any function $x(t): \mathbb{R} \rightarrow \mathbb{C}$ with $\max _{t \in[0, T]}|x(t)|^{2} \leq d\|x(t)\|_{T}^{2}$. Let $S$ denote a set of points from 0 to $T$. If each point of $S$ is chosen uniformly at random from $[0, T]$, we have

$$
\left.\left.\operatorname{Pr}\left[\left.\left|\frac{1}{|S|} \sum_{i \in S}\right| x\left(t_{i}\right)\right|^{2}-\|x(t)\|_{T}^{2}\right] \right\rvert\, \geq \epsilon\|x(t)\|_{T}^{2}\right] \leq e^{-\Omega\left(\epsilon^{2}|S| / d\right)}
$$

Proof. Let $M$ denote $\max _{t \in[0, T]}|x(t)|^{2}$. Replacing $X_{i}$ by $\frac{\left|x\left(t_{i}\right)\right|^{2}}{M}$ and $n$ by $|S|$ in Lemma A.1.1, we obtain that

$$
\begin{aligned}
& \operatorname{Pr}[|X-\mu|>\epsilon \mu] \leq 2 \exp \left(-\frac{\epsilon^{2}}{3} \mu\right) \\
\Longrightarrow & \operatorname{Pr}\left[\left|\sum_{i \in S} \frac{\left|x\left(t_{i}\right)\right|^{2}}{M}-|S| \frac{\|x(t)\|_{T}^{2}}{M}\right|>\epsilon|S| \frac{\|x(t)\|_{T}^{2}}{M}\right] \leq 2 \exp \left(-\frac{\epsilon^{2}}{3} \mu\right) \\
\Longrightarrow & \operatorname{Pr}\left[\left.\left.\left|\frac{1}{|S|} \sum_{i \in S}\right| x\left(t_{i}\right)\right|^{2}-\|x(t)\|_{T}^{2} \right\rvert\, \geq \epsilon\|x(t)\|_{T}^{2}\right] \leq 2 \exp \left(-\frac{\epsilon^{2}}{3} \mu\right) \\
\Longrightarrow & \operatorname{Pr}\left[\left.\left.\left|\frac{1}{|S|} \sum_{i \in S}\right| x\left(t_{i}\right)\right|^{2}-\|x(t)\|_{T}^{2} \right\rvert\, \geq \epsilon\|x(t)\|_{T}^{2}\right] \leq 2 \exp \left(-\frac{\epsilon^{2}}{3}|S| \frac{\|x(t)\|_{T}^{2}}{M}\right)
\end{aligned}
$$

which is less than $2 \exp \left(-\frac{\epsilon^{2}}{3}|S| / d\right)$, thus completes the proof.

### 12.10.6 Proof of Lemma 12.3.9

Lemma 12.3.9. For any polynomial $P(t)$ of degree at most $d$ from $R$ to $\mathbb{C}$, for any interval $[S, T]$,

$$
\max _{t \in[S, T]}|P(t)|^{2} \leq(d+1)^{2} \cdot \frac{1}{T-S} \int_{S}^{T}|P(t)|^{2} \mathrm{~d} x
$$

Proof. Let $t^{*}=\underset{t \in[S, T]}{\arg \max }|P(t)|^{2}$. If $t^{*} \in(S, T)$, then it is enough to prove that

$$
\left|P\left(t^{*}\right)\right|^{2} \leq(d+1)^{2} \frac{1}{t^{*}-S} \int_{S}^{t^{*}}|P(x)|^{2} \mathrm{~d} x \text { and }\left|P\left(t^{*}\right)\right|^{2} \leq(d+1)^{2} \frac{1}{T-t^{*}} \int_{t^{*}}^{T}|P(x)|^{2} \mathrm{~d} x
$$

on the two intervals $\left[S, t^{*}\right]$ and $\left[t^{*}, T\right]$ separately.
Without loss of generality, we will prove the inequality for $S=-1$ and $t^{*}=T=1$. We find the minimum $\|P(x)\|_{T}^{2}$ assuming $|P(1)|^{2}=1$. Because the first $(d+1)$ Legendre polynomials provide a basis of polynomials of degree at most $d$ and their evaluation $L_{n}(1)=1$ for any $n$, we consider:

$$
\begin{aligned}
\min _{\alpha_{0}, \alpha_{1}, \cdots, \alpha_{d} \in \mathbb{C}} & \int_{-1}^{1}|P(x)|^{2} \mathrm{~d} x \\
\text { s.t. } & P(x)=\sum_{i=0}^{d} \alpha_{i} L_{i}(x) \\
& |P(1)|=\left|\sum_{i=0}^{d} \alpha_{i}\right|=1 .
\end{aligned}
$$

We simplify the integration of $P(x)^{2}$ over $[-1,1]$ by the orthogonality of Legendre polyno-
mials:

$$
\begin{aligned}
\int_{-1}^{1}|P(x)|^{2} \mathrm{~d} x & =\int_{-1}^{1}\left(\sum_{i=0}^{d} \alpha_{i} L_{i}(x)\right) \cdot\left(\sum_{j=0}^{d} \overline{\alpha_{j}} \overline{L_{j}}(x)\right) \mathrm{d} x \\
& =\int_{-1}^{1} \sum_{i=0}^{d}\left|\alpha_{i}\right|^{2} L_{i}(x)^{2}+\sum_{i \neq j} \alpha_{i} \overline{\alpha_{j}} L_{i}(x) \overline{L_{j}}(x) \mathrm{d} x \\
& =\sum_{i=0}^{d}\left|\alpha_{i}\right|^{2} \frac{2}{2 i+1} \text { by Fact } 12.3 .8
\end{aligned}
$$

Using $\int_{-1}^{1}|P(x)|^{2} \mathrm{~d} x=\sum_{i=0}^{d}\left|\alpha_{i}\right|^{2} \frac{2}{2 i+1}$, we simplify the optimization problem to

$$
\begin{aligned}
\min _{\alpha_{0}, \alpha_{1}, \cdots, \alpha_{d} \in \mathbb{C}} & \sum_{i=0}^{d}\left|\alpha_{i}\right|^{2} \frac{2}{2 i+1} \\
\text { s.t. } & \left|\sum_{i=0}^{d} \alpha_{i}\right|=1
\end{aligned}
$$

From the Cauchy-Schwarz inequality, we have

$$
\left|\sum_{i=0}^{d} \alpha_{i}\right|^{2} \leq\left(\sum_{i=0}^{d}\left|\alpha_{i}\right|^{2} \frac{2}{2 i+1}\right)\left(\sum_{i=0}^{d} \frac{2 i+1}{2}\right)
$$

Therefore $\sum_{i=0}^{d}\left|\alpha_{i}\right|^{2} \frac{2}{2 i+1} \geq \frac{2}{(d+1)^{2}}$ and $|P(1)|^{2} \leq(d+1)^{2} \cdot \frac{1}{2} \int_{-1}^{1}|P(x)|^{2} \mathrm{~d} x$.

```
Algorithm 12.1 Linear Regression Algorithms
    procedure LinearRegression \((A, b\),\() - Fact 12.11.1\)
        \(x^{\prime} \leftarrow \arg \min \|A x-b\|_{2}\).
        return \(\stackrel{x}{x^{\prime}}\)
    end procedure
    procedure LinearRegressionW \((A, b, w)\) - Fact 12.11.1
        \(x^{\prime} \leftarrow \underset{x}{\arg \min } \sum_{i=1}^{d} w_{i}\left|(A x)_{i}-b_{i}\right|^{2}\).
        return \(x^{\prime}\)
    end procedure
```


### 12.11 Known Facts

This section provides a list of well-known facts existing in literature.

### 12.11.1 Linear regression

Given a linear subspace $\operatorname{span}\left\{\vec{v}_{1}, \cdots, \vec{v}_{d}\right\}$ and $n$ points, we always use $\ell_{2}$-regression to find a vector as the linear combination of $\vec{v}_{1}, \cdots, \vec{v}_{d}$ that minimizes the distance of this vector to those $n$ points.

Fact 12.11.1. Given an $n \times d$ matrix $A$ and an $n \times 1$ column vector $b$, it takes $O\left(n d^{\omega-1}\right)$ time to output an $x^{\prime}$ such that

$$
x^{\prime}=\underset{x}{\arg \min }\|A x-b\|_{2} .
$$

where $\omega$ is the exponent of matrix multiplication[Wil12].

Notice that weighted linear regression can be solved by linear regression solver as a black-box.

```
Algorithm 12.2 Multipoint Evaluation of a Polynomial
    procedure MultipointEvaluation \(\left(P,\left\{t_{1}, t_{2}, \cdots, t_{d}\right\}\right)\) - Fact 12.11.2
        return \(P\left(t_{1}\right), P\left(t_{2}\right), \cdots, P\left(t_{d}\right)\)
    end procedure
```


### 12.11.2 Multipoint evaluation of a polynomial

Given a degree- $d$ polynomial, and $n$ locations. The naive algorithm of computing the evaluations at those $n$ locations takes $O(n d)$. However, the running time can be improved to $O(n$ poly $(\log d))$ by using this well-known result,

Fact 12.11.2 ([BS12]). Let $c>1$ denote some fixed constant. Given a degree-d polynomial $P(t)$, and a set of $d$ locations $\left\{t_{1}, t_{2}, \cdots, t_{d}\right\}$. There exists an algorithm that takes $O\left(d \log ^{c} d\right)$ time to output the evaluations

$$
\left\{P\left(t_{1}\right), P\left(t_{2}\right), \cdots, P\left(t_{d}\right)\right\}
$$

### 12.12 Analysis of Hash Functions and Filter Functions

### 12.12.1 Analysis of filter function $(H(t), \widehat{H}(f))$

We construct the Filter function $(H(t), \widehat{H}(f))$ in this section.
We fix the interval to be $\operatorname{supp}\left(\right.$ rect $\left._{1}\right)=[-1 / 2,1 / 2]$ instead of $[0, T]$ for convenience. We first define the filter function $H_{1}(t)$ which preserves the energy of a $k$-Fourier-sparse signal $x^{*}$ on $[-1 / 2,1 / 2]$ to the signal $H_{1} \cdot x^{*}$ on $[-\infty,+\infty]$.

Definition 12.12.1. Let $s_{1}=\Theta\left(k^{4} \log ^{4} k\right), \ell=\Omega(k \log k / \delta)$ be a even number, and $s_{0}=$ $C_{0} s_{1} \sqrt{\ell}$ for some constant $C_{0}$ that will normalize $H_{1}(0)=1$. Recall that $\operatorname{rect}_{s}(t)=1 \mathrm{iff}$ $|t| \leq s / 2$ and $\widehat{\operatorname{rect}_{s}(f)}=\operatorname{sinc}(f s)=\frac{\sin (\pi f s)}{\pi f s}$.

We define the filter function $H_{1}(t)$ and its Fourier transform $\widehat{H}_{1}(f)$ as follows:

$$
\widehat{H}_{1}(f)=s_{0} \cdot\left(\operatorname{rect}_{s_{1}}(f)\right)^{* \ell} \cdot \operatorname{sinc}\left(f s_{2}\right)
$$

$$
=s_{0} \cdot(\overbrace{-s_{1} / 2}^{s_{1} / 2})^{* \ell} \cdot \underbrace{1 / s_{2}}_{-1 / s_{2}}
$$

$$
H_{1}(t)=s_{0} \cdot\left(\operatorname{sinc}\left(s_{1} t\right)\right)^{\cdot \ell} * \operatorname{rect}_{s_{2}}(t)
$$


where $s_{0}$ is a fixed parameter s.t. $H_{1}(0)=1$.

We provide some basic properties about our filter function. Notice that $\operatorname{sinc}(t)=$ $\frac{\sin (\pi t)}{\pi t}(\operatorname{sinc}(0)$ is defined to be 1$)$ has the following properties (shown in Figure 12.4):

1. $\forall t \in \mathbb{R}, 1-\frac{(\pi t)^{2}}{3!} \leq \operatorname{sinc}(t) \leq 1$.


Figure 12.4: The Property of $\operatorname{sinc}(t)$.
2. $\forall|t| \leq 1.2 / \pi, \operatorname{sinc}(t) \leq 1-\frac{t^{2}}{8}$.
3. $\forall|t|>1.2 / \pi,|\operatorname{sinc}(t)| \leq \frac{1}{\pi|t|}$.

Claim 12.12.1. $\int_{-\frac{1.2}{\pi s_{1}}}^{\frac{1.2}{\pi s_{1}}}\left(\operatorname{sinc}\left(s_{1} t\right)\right)^{\ell} \mathrm{d} t \bar{\sim} \frac{1}{s_{1} \sqrt{\ell}}$.

Proof. We use the above properties for the sinc function to prove the upper bound:

$$
\begin{aligned}
\int_{-\frac{1.2}{\pi s_{1}}}^{+\frac{1.2}{\pi s_{1}}}\left(\operatorname{sinc}\left(s_{1} t\right)\right)^{\ell} \mathrm{d} t & =\frac{1}{s_{1}} \int_{-1.2 / \pi}^{+1.2 / \pi}(\operatorname{sinc}(t))^{\ell} \mathrm{d} t \\
& =\frac{2}{s_{1}}\left(\int_{0}^{\sqrt{8 / \ell}}(\operatorname{sinc} t)^{\ell} \mathrm{d} t+\int_{\sqrt{8 / \ell}}^{1.2 / \pi}(\operatorname{sinc}(t))^{\ell} \mathrm{d} t\right) \\
& \leq \frac{2}{s_{1}}\left(\sqrt{8 / \ell}+\sum_{i=1}^{\frac{1.2 / \pi}{\sqrt{8 / \ell}}-1} \int_{i \sqrt{8 / \ell}}^{(i+1) \sqrt{8 / \ell}}\left(1-x^{2} / 8\right)^{\ell} \mathrm{d} x\right) \\
& \leq \frac{2}{s_{1}}\left(\sqrt{8 / \ell}+\sum_{i=1}^{\frac{1.2 / \pi}{\sqrt{8 / \ell}}-1} \sqrt{8 / \ell} \cdot 2^{-i^{2}}\right) \\
& \lesssim \frac{1}{s_{1} \sqrt{\ell}} .
\end{aligned}
$$

We prove the lower bound:

$$
\begin{aligned}
\int_{-\frac{1.2}{\pi s_{1}}}^{\frac{1.2}{\pi s_{1}}}\left(\operatorname{sinc}\left(s_{1} t\right)\right)^{\ell} \mathrm{d} t & =\frac{2}{s_{1}}\left(\int_{0}^{\sqrt{8 / \ell}}(\operatorname{sinc} t)^{\ell} \mathrm{d} t+\int_{\sqrt{8 / \ell}}^{1.2 / \pi}(\operatorname{sinc}(t))^{\ell} \mathrm{d} t\right) \\
& \geq \frac{2}{s_{1}}\left(\int_{0}^{\sqrt{8 / \ell}}\left(1-\frac{\pi^{2} t^{2}}{6}\right)^{\ell} \mathrm{d} t\right) \\
& \gtrsim \frac{1}{s_{1} \sqrt{\ell}} .
\end{aligned}
$$

We bound the integration outside $\left[-\frac{1.2}{\pi s_{1}}, \frac{1.2}{\pi s_{1}}\right]$ from the last property of the sinc function.

Claim 12.12.2. $\int_{\frac{1.2}{\pi s_{1}}}^{+\infty}\left(\operatorname{sinc}\left(s_{1} t\right)\right)^{\ell} \mathrm{d} t=O\left(1.2^{-\ell}\right)$.

From these two claims, we have the existence of $s_{0}$.

Claim 12.12.3. There exists a universal constant $C_{0}$ and $s_{0}=C_{0} s_{1} \sqrt{\ell}$ such that $H_{1}(0)=1$.

Proof. Because $\ell$ is a large even number, $\int_{\text {rect }_{1-2 / s_{1}}} \operatorname{sinc}\left(s_{1} t\right)^{\ell} \mathrm{d} t \approx \frac{1}{s_{1} \sqrt{\ell}}$ from all discussion above.

We show several useful properties about the Filter functions $\left(H_{1}(t), \widehat{H_{1}}(f)\right)$.
Lemma 12.12.4. Given $s_{0}, s_{1}, s_{2}$, $\ell$, where $\frac{s_{2}}{2}+\frac{1}{s_{1}} \leq 1 / 2$ and $s_{0}=C_{0} s_{1} \sqrt{\ell}$ for some constant


Figure 12.5: The light red area represents $\int_{\left(1 / 2-2 / s_{1}\right)}^{\left(1 / 2-2 / s_{1}\right)-t} s_{0} \cdot \operatorname{sinc}\left(s_{1}(\tau)\right)^{\ell} \mathrm{d} \tau$ and the light green area represents $\int_{1 / 2-2 / s_{1}}^{1 / 2-2 / s_{1}+t} s_{0} \cdot \operatorname{sinc}\left(s_{1}(\tau)\right)^{\ell} \mathrm{d} \tau$.
$C_{0}$. The filter function $\left(H_{1}(t), \widehat{H}_{1}(f)\right)\left[s_{0}, s_{1}, s_{2}, \ell\right]$ has the following properties,
Property I: $\quad H_{1}(t) \in\left[1-\frac{s_{0}}{s_{1}} \cdot \frac{2 \pi^{-\ell}}{\ell-1}, 1\right], \quad$ if $|t| \leq \frac{s_{2}}{2}-\frac{1}{s_{1}}$.
Property II : $\quad H_{1}(t) \in[0,1], \quad$ if $\frac{s_{2}}{2}-\frac{1}{s_{1}} \leq|t| \leq \frac{1}{2}$
Property III : $\quad H_{1}(t) \leq s_{0} s_{2}\left(\left(s_{1}|t|-s_{1}+2\right)^{2}+1\right)^{-\ell}, \forall|t|>\frac{1}{2}$
Property IV : $\operatorname{supp}\left(\widehat{H}_{1}(f)\right) \subseteq\left[-\frac{s_{1} \ell}{2}, \frac{s_{1} \ell}{2}\right]$

Proof of Property I. First, $H_{1}(0)=1$ follows by definition of $s_{0}$, then we can prove the upper bound for $H_{1}(t)$ by showing for any $t>0, H_{1}(0)-H_{1}(t)>0$ always holds,

By definition of sinc function, we know that $\operatorname{sinc}\left(s_{1} t\right)^{\ell}$ reaches 0 at all the points $\left\{\left.\frac{1}{s_{1}}+i \frac{2}{s_{1}} \right\rvert\, i \in \mathbb{N}\right\}$. By definition of $s_{1}$, we know that $\frac{1}{s_{1}} \ll \frac{1}{2}-\frac{1}{s_{1}}$. For any $t>0$,

$$
\begin{aligned}
& H_{1}(0)-H_{1}(t) \\
= & \int s_{0} \cdot \operatorname{sinc}\left(s_{1}(\tau)\right)^{\ell} \cdot \operatorname{rect}_{1-2 / s_{1}}(0-\tau) \mathrm{d} \tau-\int s_{0} \cdot \operatorname{sinc}\left(s_{1}(\tau)\right)^{\ell} \cdot \operatorname{rect}_{1-2 / s_{1}}(t-\tau) \mathrm{d} \tau \\
= & \int_{-\left(1 / 2-2 / s_{1}\right)}^{1 / 2-2 / s_{1}} s_{0} \cdot \operatorname{sinc}\left(s_{1}(\tau)\right)^{\ell} \mathrm{d} \tau-\int_{-\left(1 / 2-2 / s_{1}\right)+t}^{1 / 2-2 / s_{1}+t} s_{0} \cdot \operatorname{sinc}\left(s_{1}(\tau)\right)^{\ell} \mathrm{d} \tau \\
= & \int_{-\left(1 / 2-2 / s_{1}\right)}^{-\left(1 / 2-2 / s_{1}\right)+t} s_{0} \cdot \operatorname{sinc}\left(s_{1}(\tau)\right)^{\ell} \mathrm{d} \tau-\int_{1 / 2-2 / s_{1}}^{1 / 2-2 / s_{1}+t} s_{0} \cdot \operatorname{sinc}\left(s_{1}(\tau)\right)^{\ell} \mathrm{d} \tau \\
\geq & 0
\end{aligned}
$$

where the third step can be observed in Figure 12.5, and the last inequality follows by choosing $s_{1}$ to be an integer. Thus, we prove an upper bound for $H_{1}(t)$. Third, we show the lower bound for $H_{1}(t)$,

$$
\begin{aligned}
H_{1}(t) & =\int_{-\infty}^{+\infty} s_{0} \cdot \operatorname{sinc}\left(s_{1} \tau\right)^{\cdot \ell} \operatorname{rect}_{s_{2}}(t-\tau) \mathrm{d} \tau \\
& =\int_{t-\frac{s_{2}}{2}}^{t+\frac{s_{2}}{2}} s_{0} \cdot \operatorname{sinc}\left(s_{1} \tau\right)^{\cdot \ell} \mathrm{d} \tau \\
& =1-\underbrace{\int_{t+\frac{s_{2}}{2}}^{+\infty} s_{0} \cdot \operatorname{sinc}\left(s_{1} \tau\right)^{\cdot \ell} \mathrm{d} \tau}_{A}-\underbrace{\int_{-\infty}^{t-\frac{s_{2}}{2}} s_{0} \cdot \operatorname{sinc}\left(s_{1} \tau\right)^{\cdot \ell} \mathrm{d} \tau}_{B}
\end{aligned}
$$

Thus, as long as we can upper bound the term $A$ and $B$, then we will have a lower bound
for the $H_{1}(t)$, for any $|t| \leq \frac{s_{2}}{2}-\frac{1}{s_{1}}$.

$$
\begin{aligned}
A & =\int_{t+\frac{s_{2}}{2}}^{+\infty} s_{0} \cdot \operatorname{sinc}\left(s_{1} \tau\right)^{\cdot \ell} \mathrm{d} \tau \\
& \leq \int_{\frac{1}{s_{1}}}^{+\infty} s_{0} \cdot \operatorname{sinc}\left(s_{1} \tau\right)^{\cdot \ell} \mathrm{d} \tau \\
& \leq \int_{\frac{1}{s_{1}}}^{+\infty} s_{0} \cdot\left(s_{1} \pi \tau\right)^{-\ell} \mathrm{d} \tau \\
& =s_{0} \cdot\left(s_{1} \pi\right)^{-\ell} \frac{1}{\ell-1}\left(1 / s_{1}\right)^{-\ell+1} \\
& =\frac{s_{0}}{s_{1}} \cdot(\pi)^{-\ell} \frac{1}{\ell-1}
\end{aligned}
$$

Similarly, we can bound the term $B$ in the same way.

Proof of Property II. In the proof of Property I, we already show that $\forall t, H_{1}(t) \leq 1$. Thus, the upper bound of Property II is also holding. The lower bound follows by both $\operatorname{sinc}\left(s_{1} t\right)^{\cdot \ell}$ and $\operatorname{rect}_{s_{2}}(t)$ are always nonnegative, thus the convolution of these two functions has to be nonnegative.

Proof of Property III. Let's prove the case when $t>1$, since $H_{1}(t)$ is symmetric, then the case $t<-1$ will also hold. By definition of $\left(H_{1}(t), \widehat{H}_{1}(f)\right)$, we have

$$
\begin{aligned}
H_{1}(t) & =s_{0} \cdot \int_{-\infty}^{+\infty} \operatorname{sinc}\left(s_{1}(t-\tau)\right)^{\cdot \ell} \operatorname{rect}_{s_{2}}(\tau) \mathrm{d} \tau \\
& =s_{0} \cdot \int_{-\frac{s_{2}}{2}}^{\frac{s_{2}}{2}} \operatorname{sinc}\left(s_{1}(t-\tau)\right)^{\cdot \ell} \mathrm{d} \tau \\
& =s_{0} \cdot \int_{-\frac{s_{2}}{2}}^{\frac{s_{2}}{2}} \operatorname{sinc}\left(s_{1}(\tau-t)\right)^{\cdot \ell} \mathrm{d} \tau
\end{aligned}
$$

We'd like to choose a middle point $\tau_{0}$, and then separated the interval into two parts, one is $\left[-\frac{s_{2}}{2}, \tau_{0}\right]$ and the other is $\left[\tau_{0}, \frac{s_{2}}{2}\right]$. To choose a reasonable $\tau_{0}$, we need to use the following simple facts,

$$
\begin{aligned}
& \left(\frac{\sin (x)}{x}\right)^{\ell} \leq x^{-\ell} \text { if } x \geq 1.2 \\
& \left(\frac{\sin (x)}{x}\right)^{\ell} \leq\left(1-\frac{x^{2}}{8}\right)^{\ell} \text { if } x<1.2
\end{aligned}
$$

Thus, $\left|\pi s_{1}\left(\tau_{0}-t\right)\right|=1.2$, which implies that $\tau_{0}^{+}=t+\frac{1.2}{\pi s_{1}}$ or $\tau_{0}^{-}=t-\frac{1.2}{\pi s_{1}}$. By relationship between $s_{1}$ and $s_{2}$, we know $\tau_{0}^{-}>\frac{1}{2}-\frac{1.2}{\pi s_{1}}>\frac{1}{2}-\frac{1}{s_{1}} \geq \frac{s_{2}}{2}$. Thus, we can use the case $x<1.2$ to upper bound the $H_{1}(t)$,

$$
\begin{aligned}
& H_{1}(t) \\
\leq & s_{0} \cdot s_{2} \cdot \max _{\tau \in\left[-s_{2} / 2, s_{2} / 2\right]} \operatorname{sinc}\left(s_{1}(\tau-t)\right)^{\cdot \ell} \\
\leq s_{0} \cdot s_{2} \max _{\tau \in\left[-s_{2} / 2, s_{2} / 2\right]}\left(1-\frac{\left(s_{1} \pi(\tau-t)\right)^{2}}{8}\right)^{\ell} & \\
=s_{0} \cdot s_{2}\left(1-\frac{\left(s_{1} \pi\left(\frac{s_{2}}{2}-t\right)\right)^{2}}{8}\right)^{\ell} & \\
\leq s_{0} \cdot s_{2} \cdot\left(e^{\left.-\frac{\left(s _ { 1 } \pi \left(\frac{\left.\left.s_{2}-t\right)\right)^{2}}{8}\right.\right.}{8}\right)^{\ell}}\right. & \text { by } 1-x \leq e^{-x} \\
\leq s_{0} \cdot s_{2} \cdot\left(e^{\left.\left(s_{1}\left(t-s_{2} / 2\right)\right)^{2}\right)^{-\ell}}\right. & \text { by } 1<\pi^{2} / 8 \\
\leq s_{0} \cdot s_{2} \cdot\left(1+\left(s_{1}\left(t-s_{2} / 2\right)\right)^{2}\right)^{-\ell} & \text { by } 1+x \leq e^{x} \\
\leq s_{0} \cdot\left(1+\left(s_{1}\left(t-s_{2} / 2\right)\right)^{2}\right)^{-\ell} & \text { by } s_{1} \leq 1 .
\end{aligned}
$$

Thus, we complete the proof.

Proof of Property IV. Because of the support of $\operatorname{rect}_{s_{1}}(f)$ is $s_{1}$, then the support of
$\left(\operatorname{rect}_{s_{1}}(f)\right)^{* \ell}=s_{1} \ell$. Since $\widehat{H}_{1}(f)$ is defined to be the $\left(\operatorname{rect}_{s_{1}}(f)\right)^{* \ell}$ multiplied by $\operatorname{sinc}\left(f s_{2}\right)$, thus $\operatorname{supp}\left(\widehat{H}_{1}(f)\right) \subseteq\left[-\frac{s_{1} \ell}{2}, \frac{s_{1} \ell}{2}\right]$.

Definition 12.12.2. Given any $0<s_{3}<1,0<\delta<1$, we define $(H(t), \widehat{H}(f))$ to be the filter function $\left(H_{1}(t), \widehat{H_{1}}(f)\right)$ by doing the following operations

- Setting $\ell=\Theta(k \log (k / \delta))$,
- Setting $s_{2}=1-\frac{2}{s_{1}}$,
- Shrinking by a factor $s_{3}$ in time domain,

$$
\begin{align*}
H(t) & =H_{1}\left(t / s_{3}\right)  \tag{12.29}\\
\widehat{H}(f) & =s_{3} \widehat{H_{1}}\left(s_{3} f\right) \tag{12.30}
\end{align*}
$$

We call the "heavy cluster" around a frequency $f_{0}$ to be the support of $\delta_{f_{0}}(f) * \widehat{H}(f)$ in the frequency domain and use

$$
\begin{equation*}
\Delta_{h}=|\operatorname{supp}(\widehat{H}(f))|=\frac{s_{1} \cdot \ell}{s_{3}} \tag{12.31}
\end{equation*}
$$

to denote the width of the cluster.

We show several useful properties about the Filter functions $(H(t), \widehat{H}(f))$.
Lemma 12.6.4. Given $s_{0}, s_{1}, 0<s_{3}<1, \ell>1,0<\delta<1$, where $\ell=\Theta(k \log (k / \delta))$. The filter
function $(H(t), \widehat{H}(f))$ has the following properties,
Property I: $\quad H(t) \in[1-\delta, 1]$, when $|t| \leq\left(\frac{1}{2}-\frac{2}{s_{1}}\right) s_{3}$.
Property II : $H(t) \in[0,1]$, when $\left(\frac{1}{2}-\frac{2}{s_{1}}\right) s_{3} \leq|t| \leq \frac{1}{2} s_{3}$.
Property III : $H(t) \leq s_{0} \cdot\left(s_{1}\left(\frac{|t|}{s_{3}}-\frac{1}{2}\right)+2\right)^{-\ell}, \forall|t|>\frac{1}{2} s_{3}$.
Property IV: $\operatorname{supp}(\widehat{H}(f)) \subseteq\left[-\frac{s_{1} \ell}{2 s_{3}}, \frac{s_{1} \ell}{2 s_{3}}\right]$.

For any exact $k$-Fourier-sparse signal $x^{*}(t)$, we shift the interval from $[0, T]$ to $[-1 / 2,1 / 2]$ and consider $x^{*}(t)$ for $t \in[-1 / 2,1 / 2]$ to be our observation, which is also $x^{*}(t) \cdot \operatorname{rect}_{1}(t)$.

$$
\begin{array}{ll}
\text { Property V : } & \int_{-\infty}^{+\infty}\left|x^{*}(t) \cdot H(t) \cdot\left(1-\operatorname{rect}_{1}(t)\right)\right|^{2} \mathrm{~d} t<\delta \int_{-\infty}^{+\infty}\left|x^{*}(t) \cdot \operatorname{rect}_{1}(t)\right|^{2} \mathrm{~d} t \\
\text { Property VI : } & \int_{-\infty}^{+\infty}\left|x^{*}(t) \cdot H(t) \cdot \operatorname{rect}_{1}(t)\right|^{2} \mathrm{~d} t \in[1-\epsilon, 1] \cdot \int_{-\infty}^{+\infty}\left|x^{*}(t) \cdot \operatorname{rect}_{1}(t)\right|^{2} \mathrm{~d} t .
\end{array}
$$

for arbitrarily small constant $\epsilon$.

The Property I, II, III and IV follow by filter function $H(t), \widehat{H}(f)$ inheriting $H_{1}(t), \widehat{H}_{1}(f)$.

Proof of Property $V . \quad \forall t \notin[-1 / 2,1 / 2]$, we have,

$$
\begin{array}{rlr} 
& \left|x^{*}(t) \cdot H(t)\right|^{2} & \\
\leq & \left|x^{*}(t)\right|^{2} \cdot|H(t)|^{2} & \text { by Property III of } H_{1}(t) \\
\leq & \left|x^{*}(t)\right|^{2} \cdot\left(s_{1}(|t|-1 / 2)+2\right)^{-\ell} & \text { by Lemma } 12.5 .5 \\
\leq & k^{7} \cdot(2 k t)^{2.5 k} \cdot \int_{-\infty}^{+\infty}\left|x^{*}(t) \cdot \operatorname{rect}_{1}(t)\right|^{2} \mathrm{~d} t \cdot\left(s_{1}\left(\frac{|t|}{s_{3}}-\frac{1}{2}\right)+2\right)^{-\ell} & \\
\leq & t^{O(k \log k)} \cdot \int_{-\infty}^{+\infty}\left|x^{*}(t) \cdot \operatorname{rect}_{1}(t)\right|^{2} \mathrm{~d} t \cdot\left(s_{1}\left(\frac{|t|}{s_{3}}-\frac{1}{2}\right)+2\right)^{-\ell / 2} . & \tag{12.32}
\end{array}
$$



Figure 12.6: $\widehat{H \cdot x^{*}}(f)$ and $H \cdot x^{*}(t)$.

Thus taking the integral finishes the proof because $\ell \gtrsim k \log (k / \delta)$.

Proof of Property VI. First, because of for any $t,\left|H_{1}(t)\right| \leq 1$, thus we prove the upper bound for LHS,

$$
\int_{-\infty}^{+\infty}\left|x^{*}(t) \cdot H(t) \cdot \operatorname{rect}_{1}(t)\right|^{2} \mathrm{~d} t \leq \int_{-\infty}^{+\infty}\left|x^{*}(t) \cdot 1 \cdot \operatorname{rect}_{1}(t)\right|^{2} \mathrm{~d} t .
$$

Second, as mentioned early, we need to prove the general case when $s_{3}=1-1 / \operatorname{poly}(k)$. Define interval $S=\left[-s_{3}\left(\frac{1}{2}-\frac{1}{s_{1}}\right), s_{3}\left(\frac{1}{2}-\frac{1}{s_{1}}\right)\right]$, by definition, $S \subset[-1 / 2,1 / 2]$. Then define
$\bar{S}=[-1 / 2,1 / 2] \backslash S$, which is $\left[-1 / 2,-s_{3}\left(\frac{1}{2}-\frac{1}{s_{1}}\right)\right) \cup\left(s_{3}\left(\frac{1}{2}-\frac{1}{s_{1}}\right), 1 / 2\right]$. By Property I, we have

$$
\begin{equation*}
\int_{S}\left|x^{*}(t) \cdot H(t)\right|^{2} \mathrm{~d} t \geq(1-\delta)^{2} \int_{S}\left|x^{*}(t)\right|^{2} \mathrm{~d} t \tag{12.33}
\end{equation*}
$$

Then we can show

$$
\begin{align*}
& \int_{\bar{S}}\left|x^{*}(t)\right|^{2} \mathrm{~d} t \\
\leq & |\bar{S}| \cdot \max _{t \in[-1 / 2,1 / 2]}\left|x^{*}(t)\right|^{2} \\
\leq & \left(1-s_{3}\left(1-\frac{2}{s_{1}}\right)\right) \cdot \widetilde{O}\left(k^{4}\right) \int_{-\frac{1}{2}}^{\frac{1}{2}}\left|x^{*}(t)\right|^{2} \mathrm{~d} t
\end{align*} \quad \text { by Lemma 12.5.1 } \quad \begin{array}{lr}
\lesssim & \int_{-\frac{1}{2}}^{\frac{1}{2}}\left|x^{*}(t)\right|^{2} \mathrm{~d} t
\end{array} \quad \text { by } \min \left(\frac{1}{1-s_{3}}, s_{1}\right) \geq \widetilde{O}\left(k^{4}\right) \text {. }
$$

Combining Equations (12.33) and (12.34) gives a lower bound for LHS,

$$
\begin{array}{rlr} 
& \int_{-\infty}^{+\infty}\left|x^{*}(t) \cdot H(t) \cdot \operatorname{rect}_{1}(t)\right|^{2} \mathrm{~d} t & \\
\geq & \int_{S}\left|x^{*}(t) H(t)\right|^{2} \mathrm{~d} t & \\
\geq & (1-2 \delta) \int_{S}\left|x^{*}(t)\right|^{2} \mathrm{~d} t & \\
\geq & (1-2 \delta) \int_{S \cup \bar{S}}\left|x^{*}(t)\right|^{2} \mathrm{~d} t-(1-2 \delta) \int_{\bar{S}}\left|x^{*}(t)\right|^{2} \mathrm{~d} t & \\
\geq & (1-2 \delta) \int_{S \cup \bar{S}}\left|x^{*}(t)\right|^{2} \mathrm{~d} t-(1-2 \delta) \epsilon \int_{S \cup \bar{S}}\left|x^{*}(t)\right|^{2} \mathrm{~d} t & \text { by Equation (12.33) } \\
\geq & (1-2 \delta-\epsilon) \int_{-\frac{1}{2}}^{\frac{1}{2}}\left|x^{*}(t)\right|^{2} \mathrm{~d} t & \\
\geq & (1-2 \epsilon) \int_{-\infty}^{+\infty}\left|x^{*}(t) \cdot \operatorname{rect}_{1}(t)\right|^{2} \mathrm{~d} t & \text { by } \quad \text { by } \epsilon \gg \delta
\end{array}
$$



Figure 12.7: Property VI of filter function $H(t)$, first two figures of four figures. The light green area represents RHS(without scalar) of Property VI of filter $H$.

Remark 12.12.1. To match $(H(t), \widehat{H}(f))$ on $[-1 / 2,1 / 2]$ with signal $x(t)$ on $[0, T]$, we will scale the time domain from $[-1 / 2,1 / 2]$ to $[-T / 2, T / 2]$ and shift it to $[0, T]$. For example, the rectangle function in Property V and VI will be replaced by $\operatorname{rect}_{T}(t-T / 2)$. For the parameters $s_{0}, s_{1}, s_{3}, \delta, \ell$ in the definition of $H$, we always treat them as numbers. We assume $T$ has seconds as unit and $\Delta_{h}$ has Hz as unit. For example, in time domain, the Property I becomes that given $T>0$,

$$
H(t) \in[1-\delta, 1] \text { if }\left|t-\frac{T}{2}\right| \leq\left(\frac{1}{2}-\frac{1}{s_{1}}\right) s_{3} \cdot T
$$



Figure 12.8: Property VI of filter function $H(t)$, last two figures of four figures. The light red area represents LHS of Property VI of filter $H$, the light yellow area represents the difference. Property VI says the light yellow area is only a small constant fraction of the light green area.

In frequency domain, the Property IV becomes

$$
\begin{equation*}
\operatorname{supp}(\widehat{H}(f)) \subseteq\left[-\frac{\Delta_{h}}{2}, \frac{\Delta_{h}}{2}\right], \text { where } \Delta_{h}=\frac{s_{1} \ell}{s_{3} T} \tag{12.35}
\end{equation*}
$$

Lemma 12.12.5. Let $H(t)$ denote the function defined in Definition 12.12.2. For any $t \in\left[-\frac{1}{2}, \frac{1}{2}\right]$, there exists an algorithm that takes $O\left(s_{1}+\ell \log \left(s_{1}\right)+\log (1 / \epsilon)\right)$ time to output a
value $\widetilde{H}(t)$ such that

$$
(1-\epsilon) H(t) \leq \widetilde{H}(t) \leq(1+\epsilon) H(t)
$$

Proof. We will show that using a low degree polynomial with sufficiently large degree is able to approximate the sinc function. By definition of filter function,

$$
\begin{aligned}
H(t) & =s_{0} \cdot \int_{-\infty}^{+\infty} \operatorname{sinc}\left(s_{1} \tau\right)^{\ell \ell} \operatorname{rect}_{s_{2}}(t-\tau) \mathrm{d} \tau \\
& =s_{0} \cdot \int_{t-\frac{s_{2}}{2}}^{t+\frac{s_{2}}{2}}\left(\frac{\sin \left(\pi s_{1} \tau\right)}{\pi s_{1} \tau}\right)^{\ell} \mathrm{d} \tau \\
& =\frac{s_{0}}{\pi s_{1}} \int_{\left(t-\frac{s_{2}}{2}\right) \pi s_{1}}^{\left(t+\frac{s_{2}}{2}\right) \pi s_{1}}\left(\frac{\sin (\tau)}{\tau}\right)^{\ell} \mathrm{d} \tau \\
& =\frac{s_{0}}{\pi s_{1}} \int_{\left(t-\frac{s_{2}}{2}\right) \pi s_{1}}^{\left(t+\frac{s_{2}}{2}\right) \pi s_{1}}\left(\sum_{i=0}^{\infty}(-1)^{i} \frac{\tau^{2 i}}{(2 i+1)!}\right)^{\ell} \mathrm{d} \tau \quad \text { by Taylor expansion } \\
& =\frac{s_{0}}{\pi s_{1}} \int_{\left(t-\frac{s_{2}}{2}\right) \pi s_{1}}^{\left(t+\frac{s_{2}}{2}\right) \pi s_{1}}(A+B)^{\ell} \mathrm{d} \tau
\end{aligned}
$$

where the last step follows by setting $A=\sum_{i=0}^{d}(-1)^{i} \frac{\tau^{2 i}}{(2 i+1)!}$, and $B=\sum_{i=d+1}^{\infty}(-1)^{i} \frac{\tau^{2 i}}{(2 i+1)!}$.
Denote $I^{+}=\left(t+\frac{s_{2}}{2}\right) \pi s_{1}$ and $I^{-}=\left(t-\frac{s_{2}}{2}\right) \pi s_{1}$. Because of $t \in[-1 / 2,1 / 2]$, then $\max \left(\left|I^{+}\right|,\left|I^{-}\right|\right)=O\left(s_{1}\right)$. The goal is to show that

$$
(1-\epsilon) \int_{I^{-}}^{I^{+}}(A+B)^{\ell} \mathrm{d} \tau \leq \int_{I^{-}}^{I^{+}} A^{\ell} \mathrm{d} \tau \leq(1+\epsilon) \int_{I^{-}}^{I^{+}}(A+B)^{\ell} \mathrm{d} \tau
$$

Let's prove an upper first,

$$
\begin{aligned}
& \int_{I^{-}}^{I^{+}}(A+B-B)^{\ell} \mathrm{d} \tau \\
= & \int_{I^{-}}^{I^{+}}(A+B)^{\ell} \mathrm{d} \tau+\sum_{j=1}^{\ell} \int_{I^{-}}^{I^{+}}\binom{\ell}{j}(A+B)^{\ell-j}(-B)^{j} \mathrm{~d} \tau \\
\leq & \int_{I^{-}}^{I^{+}}(A+B)^{\ell} \mathrm{d} \tau+\sum_{j=1}^{\ell} \int_{I^{-}}^{I^{+}}\binom{\ell}{j}|A+B|^{\ell-j}|B|^{j} \mathrm{~d} \tau \\
\leq & \\
\leq \int_{I^{-}}^{I^{+}}(A+B)^{\ell} \mathrm{d} \tau+\sum_{j=1}^{\ell} \int_{I^{-}}^{I^{+}}\binom{\ell}{j}|A+B|^{\ell-j} \mathrm{~d} \tau \cdot \max _{\tau \in\left[I^{-}, I^{+}\right]}|B|^{j} & \\
\leq & \int_{I^{-}}^{I^{+}}(A+B)^{\ell} \mathrm{d} \tau+\ell 2^{\ell} \cdot \max _{\tau \in\left[I^{-}, I^{+}\right]}|B| \\
\leq & \text { by }|H(t)| \leq 1 \text { and }|B|^{j} \leq|B| \\
\leq & \int_{I^{-}}^{I^{+}}(A+B)^{\ell} \mathrm{d} \tau+\epsilon \cdot\left(s_{1}\right)^{-\Theta(\ell)} \\
\leq & \text { by Claim } 12.12 .6
\end{aligned}
$$

where all the steps by setting $d \gtrsim s_{1}+\ell \log \left(s_{1}\right)+\log (1 / \epsilon)$. Similarly, we can prove a lower bound.

Claim 12.12.6. Let $B(\tau)=\sum_{i=d+1}^{+\infty}(-1) \frac{\tau^{2 i}}{(2 i+1)!}$, if $d \gtrsim \tau+\ell \log \left(s_{1}\right)+\log (1 / \epsilon)$ then $|B(\tau)| \leq$ $\epsilon\left(1 / s_{1}\right)^{O(\ell)}$.

Proof. We first show, for any $i \geq d+1$,

$$
\begin{array}{rr} 
& \frac{\tau^{2 i}}{(2 i+1)!} \\
\leq & \\
\hline e((2 i+1) / e)^{2 i+1} & \text { by } e(n / e)^{n} \leq n! \\
\leq 2^{-2 i} & \text { by } i \gtrsim \tau \\
\leq \epsilon\left(1 / s_{1}\right)^{O(\ell)} & \text { by } i \gtrsim \ell \log \left(s_{1}\right)+\log (1 / \epsilon)
\end{array}
$$

Second, we can show that

$$
\sum_{i=d+1}^{+\infty}(-1) \frac{\tau^{2 i}}{(2 i+1)!} \lesssim \frac{\tau^{2(d+1)}}{(2(d+1)+1)!} \leq \epsilon\left(1 / s_{1}\right)^{O(\ell)}
$$

Thus, we complete the proof.

Claim 12.12.7. $\min _{t \in[-1 / 2,1 / 2]}|H(t)| \geq\left(s_{1}\right)^{-\Omega(\ell)}$.

Proof. By the property of $H(t)$,

$$
\min _{\frac{1}{2} s_{3}<|t| \leq \frac{1}{2}} H(t)=\min _{|t| \leq \frac{1}{2}} H(t)
$$

Thus, it suffices to prove a lower bound on $H(t)$ for any $t$ such that $\frac{1}{2} s_{3}<|t| \leq \frac{1}{2}$. Because of symmetric property, we only need to prove a lower bound for one side. Let's consider $t \in\left[\frac{1}{2} s_{3}, 1 / 2\right]$,

$$
\begin{aligned}
H(t) & \geq \frac{s_{0}}{\pi s_{1}} \int_{\left(t-\frac{s_{2}}{2}\right) \pi s_{1}}^{\left(t+\frac{s_{2}}{2}\right) \pi s_{1}}\left(\frac{\sin (\tau)}{\tau}\right)^{\ell} \mathrm{d} \tau \\
& \geq \frac{s_{0}}{\pi s_{1}} \int_{\left(t+\frac{s_{2}}{4}\right) \pi s_{1}}^{\left(t+\frac{s_{2}}{2}\right) \pi s_{1}}\left(\frac{\sin (\tau)}{\tau}\right)^{\ell} \mathrm{d} \tau \\
& \geq \frac{s_{0}}{\pi s_{1}} \cdot \Theta\left(\left(t+\frac{s_{2}}{2}\right) s_{1}\right) \cdot \frac{1}{2} \cdot \pi \cdot \Theta\left(\left(t+\frac{s_{2}}{2}\right) \pi s_{1}\right)^{-\ell} \\
& \geq\left(s_{1}\right)^{-\Omega(\ell)}
\end{aligned}
$$

### 12.12.2 Analysis of filter function $(G(t), \widehat{G}(f))$

We construct $(G(t), \widehat{G}(f))$ in a similar way of $\left(H_{1}(t), \widehat{H_{1}}(f)\right)$ by switching the time domain and the frequency domain of $\left(H_{1}(t), \widehat{H_{1}}(f)\right)$ and modify the parameters for the permutation hashing $P_{\sigma, a, b}$.

Definition 12.12.3. Given $B>1, \delta>0, \alpha>0$, we construct $G(t), \widehat{G}(f)$ by doing the following operations,

- $s_{2}=\frac{\pi}{2 B}$,
- $s_{1}=\frac{B}{\alpha \pi}$,
- $\ell=l=\Theta(\log (k / \delta))$.

Then $G(t), \widehat{G}(f)$ becomes

$$
\begin{aligned}
G(t) & =b_{0} \cdot\left(\operatorname{rect}_{s_{1}}(t)\right)^{* l} \cdot \operatorname{sinc}\left(t s_{2}\right) \\
& =b_{0} \cdot\left(\operatorname{rect}_{\frac{B}{(\alpha \pi)}}(t)\right)^{* l} \cdot \operatorname{sinc}\left(t \frac{\pi}{2 B}\right), \\
\widehat{G}(f) & =b_{0} \cdot\left(\operatorname{sinc}\left(s_{1} f\right)\right)^{l} * \operatorname{rect}_{s_{2}}(f) \\
& =b_{0} \cdot\left(\operatorname{sinc}\left(\frac{B}{\alpha \pi} f\right)\right)^{l} * \operatorname{rect}_{\frac{\pi}{2 B}}(f) .
\end{aligned}
$$

where the scalar $b_{0}=\Theta\left(s_{1} \sqrt{l}\right)=\Theta(B \sqrt{l} / \alpha)$ satisfying $\widehat{G}(0)=1$.
Lemma 12.6.5. Given $B>1, \delta>0, \alpha>0$, we set $l=\Omega(\log (\delta / k))$. The filter function
$(G(t), \widehat{G}(f))[B, \delta, \alpha, l]$ satisfies the following properties,

$$
\begin{array}{ll}
\text { Property I: } & \widehat{G}(f) \in[1-\delta / k, 1], \text { if }|f| \leq(1-\alpha) \frac{2 \pi}{2 B} \\
\text { Property II : } & \widehat{G}(f) \in[0,1], \text { if }(1-\alpha) \frac{2 \pi}{2 B} \leq|f| \leq \frac{2 \pi}{2 B} \\
\text { Property III : } & \widehat{G}(f) \in[-\delta / k, \delta / k], \text { if }|f|>\frac{2 \pi}{2 B} . \\
\text { Property IV : } & \operatorname{supp}(G(t)) \subset\left[\frac{l}{2} \cdot \frac{-B}{\pi \alpha}, \frac{l}{2} \cdot \frac{B}{\pi \alpha}\right] . \\
\text { Property V : } & \max _{t}|G(t)| \lesssim \operatorname{poly}(B, l) .
\end{array}
$$

Proof. The first five Properties follows from Lemma 12.6.4 directly.


Figure 12.9: Parameters for $s_{1}, s_{3}$ and $\ell$.

### 12.12.3 Parameters setting for filters

One-cluster Recovery. In one-cluster, we donot need filter function $(G(t), \widehat{G}(f))$.
In section 12.12.1, by Equation (12.34) in the proof of Property VI of filter function $(H(t), \widehat{H}(f))$, we need $\min \left(\frac{1}{1-s_{3}}, s_{1}\right) \geq \widetilde{O}\left(k^{4}\right)$.

In section 12.12.1, by Equation (12.32) in the proof of Property V of filter function $(H(t), \widehat{H}(f))$, we set $\ell \gtrsim k \log (k / \delta)$.
$\Delta_{h}$ is determined by the parameters of filter $(H(t), \widehat{H}(f))$ in Equation (12.35): $\Delta_{h} \bar{\sim}$ $\frac{s_{1} \ell}{s_{3} T}$ in section 12.12.1. Combining the setting of $s_{1}, s_{3} \ell$, we should set $\Delta_{h} \geq \widetilde{O}\left(k^{5} \log (1 / \delta)\right) / T$.
$k$-cluster Recovery. Note that in the $k$-cluster recovery, we need to use filter function $(G(t), \widehat{G}(f))$. We choose $l=\log (k / \delta), \alpha \approx 1, B \approx k$, and $D=l / \alpha$.

By proof of Property II of $z$ in Lemma 12.7.18 from section 12.7.6, we need $T\left(1-s_{3}\right)>$
$\sigma B l$. By the same reason in one-cluster recovery, $1-s_{3} \leq \frac{1}{O\left(k^{4}\right)}$. Combining $T\left(1-s_{3}\right)>\sigma B l$ and $1-s_{3} \leq \frac{1}{\widehat{O}\left(k^{4}\right)}$, we obtain

$$
\begin{equation*}
\frac{T}{\widetilde{O}\left(k^{4}\right)}>\sigma B l \tag{12.36}
\end{equation*}
$$

Because in our algorithm, we will sample $\sigma$ from $\left[\frac{1}{B \Delta_{h}}, \frac{2}{B \Delta_{h}}\right]$. Thus, plugging $\sigma=\Theta\left(\frac{1}{B \Delta_{h}}\right)$ in Equation (12.36) we have

$$
\frac{T}{\widetilde{O}\left(k^{4}\right)}>\frac{l}{\Delta_{h}}
$$

which implies another lower bound for $\Delta_{h}$,

$$
\Delta_{h} \geq \widetilde{O}\left(k^{4}\right) l / T
$$

Combining the above bound with previous lower bound in one-cluster recovery, we get

$$
\Delta_{h} \geq \widetilde{O}\left(k^{4} \log (1 / \delta)\right) / T+\widetilde{O}\left(k^{5} \log (1 / \delta)\right) / T=\widetilde{O}\left(k^{5} \log (1 / \delta)\right) / T
$$

For $s_{1}$ and $\ell$, we still choose the same setting as before, $s_{1} \bar{\sim}\left(k^{4}\right)$ and $\ell \approx O(k \log (k / \delta))$.

### 12.12.4 Analysis of HashToBins

In this section, we explain the correctness of Procedure HashToBins in Algorithm 12.6. Before giving the proof of that algorithm, we show how to connect CFT, DTFT and DFT.

Lemma 12.12.8. For any signal $W: \mathbb{R} \rightarrow \mathbb{C}$, let $A: \mathbb{Z} \rightarrow \mathbb{C}$ and $B:[n] \rightarrow \mathbb{C}$ be defined as follows:

$$
A[i]=W(i), \forall i \in \mathbb{Z} \text { and } B[i]=\sum_{j \in \mathbb{Z}} A[i+j n], \forall i \in[n] .
$$

Then we consider the Fourier transform on $W, A$, and $B$ :

$$
\begin{aligned}
\mathrm{CFT} & \widehat{W}: \mathbb{R} \rightarrow \mathbb{C}, \\
\mathrm{DTFT} & \widehat{A}:[0,1] \rightarrow \mathbb{C} \\
\text { DFT } & \widehat{B}:[n] \rightarrow \mathbb{C}
\end{aligned}
$$

We have:

$$
\forall f \in[0,1), \widehat{A}(f)=\sum_{j \in \mathbb{Z}} \widehat{W}(f+j) ; \quad \forall i \in[n], \widehat{B}[i]=\sum_{j \in \mathbb{Z}} \widehat{W}(i / n+j)
$$

Proof. Recall that $\operatorname{Comb}_{s}(t)=\sum_{j \in \mathbb{Z}} \delta_{j s}(t)$. First, we show $\widehat{A}(f)=\sum_{j \in \mathbb{Z}} e^{2 \pi \mathrm{i} j f} A[j]$ equals
to $\sum_{j \in \mathbb{Z}} \widehat{W}(f+j)$ :

$$
\begin{array}{rlr}
\widehat{A}(f) & =\sum_{j \in \mathbb{Z}} e^{2 \pi \mathrm{i} j f} W[j] & \text { by } A[j]=W(j) \\
& =\int_{-\infty}^{+\infty} e^{2 \pi \mathrm{i} f j} W(j) \cdot \operatorname{Comb}_{1}(j) \mathrm{d} j & \\
& =W \cdot \operatorname{Comb}_{1}(f) & \\
& =\left(\widehat{W} * \widehat{\operatorname{Comb}}_{1}\right)(f) & \\
& =\sum_{j \in \mathbb{Z}} \widehat{W}(f+j) . \tag{12.37}
\end{array}
$$

Next, we prove that $\forall i \in[n], \widehat{B}[i]=\widehat{A}(i / n)$,

$$
\begin{array}{rlr}
\widehat{B}[i] & =\sum_{j=1}^{n} B[j] e^{\frac{2 \pi \mathrm{i}}{n} i j} & \text { by DFT } \\
& =\sum_{j=1}^{n}\left(\sum_{k \in \mathbb{Z}} A[j+k n]\right) e^{\frac{2 \pi \mathrm{i}}{n} i j} & \text { by } B[j]=\sum_{k \in \mathbb{Z}} A[j+k n] \\
& =\sum_{j=1}^{n} \sum_{k \in \mathbb{Z}} A[j+k n] e^{\frac{2 \pi \mathrm{i}}{n} i(j+k n)} & \text { by } e^{\frac{2 \pi \mathrm{i}}{n} \cdot i k n}=1 \\
& =\sum_{j \in \mathbb{Z}} A[j] e^{2 \pi \mathbf{i} j \frac{i}{n}}=\widehat{A}(i / n) & \text { by DTFT. }
\end{array}
$$

Combining Equation (12.38) and Equation (12.37), we obtain that $\widehat{B}[j]=\widehat{A}(j / n)=\sum_{i \in \mathbb{Z}} \widehat{W}(j / n+$ i) for all $j \in[n]$.

Claim 12.12.9. Let $u \in \mathbb{C}^{B}$ and $V \in \mathbb{C}^{B D}$ such that for any $j \in B, u[j]=\sum_{i \in[D]} V[j+(i-1) B]$. Then

$$
\widehat{u}[j]=\widehat{V}[j D], \forall j \in[B] .
$$

Proof. We prove it through the definition of the Fourier transform:

$$
\begin{array}{rlrl}
\widehat{V}[j D] & =\sum_{i=1}^{B D} V[i] \cdot e^{\frac{2 \pi \mathbf{i}}{B D} \cdot \cdot \cdot(j D)} & \text { by definition of DFT } \\
& =\sum_{i=1}^{B} \sum_{k=1}^{D} V[i+k B] e^{\frac{2 \pi \mathbf{i}}{B} \cdot(i+k B) \cdot j} & & \text { by replacing } i \text { by } i+k B \\
& =\sum_{i=1}^{B} e^{\frac{2 \pi \mathbf{i}}{B} \cdot j \cdot i} \sum_{k=1}^{D} V[i+(k-1) B] & & \text { by } e^{2 \pi \mathbf{i} j k}=1 \\
& =\sum_{i=1}^{B} e^{\frac{2 \pi \mathbf{i} \cdot}{B} \cdot j \cdot i} u[i]=\widehat{u}[j] & \text { by definition of DFT on } u
\end{array}
$$

We use Definition 12.6.1 and Lemma 12.6.1 to generalize Lemma 12.12.8,
Corollary 12.12.10. If for all $j \in[n], B[j]=\sum_{i \in \mathbb{Z}} W((j+i n) \sigma-\sigma a)$, then $\forall j \in[n]$,

$$
\widehat{B}[j]=\sum_{i \in \mathbb{Z}} \widehat{W}\left(\left(\frac{j}{n}+i\right) / \sigma\right) \cdot \frac{1}{\sigma} e^{-2 \pi \mathbf{i}\left(\frac{j}{n}+i\right) a}
$$

If for all $j \in[n], B[j]=\sum_{i \in \mathbb{Z}} W((j+i n) \sigma-\sigma a) e^{-2 \pi \mathbf{i} \sigma b(j+i n)}$, then $\forall j \in[n]$,

$$
\widehat{B}[j]=\sum_{i \in \mathbb{Z}} \widehat{W}\left(\left(\frac{j}{n}+i\right) / \sigma+b\right) \cdot \frac{1}{\sigma} e^{-2 \pi \mathbf{i}\left(\frac{j}{n}+i\right) a-2 \pi \mathbf{i} \sigma a b}
$$

Remark 12.12.2 (Samples of HashToBins). Procedure HashToBins in Algorithm 12.6 takes $B D$ samples in $x(t)$ :

$$
x(\sigma(1-a)), x(\sigma(2-a)), \cdots, x(\sigma(B D-a)) .
$$

To analyze our algorithm, we use filter function $(G(t), \widehat{G}(f))$ and $\operatorname{Comb}_{s}(t)=\sum_{j \in \mathbb{Z}} \delta_{s j}(t)$ to define the discretization of $G$.


Figure 12.10: $\widehat{G}_{\sigma, b}^{(j)}(f)$ where the top one is $j=0$ and the bottom one is $j=1, \mathcal{A}_{i, j}=$ $\left[\frac{1}{\sigma}\left(2 \pi\left(i+\frac{j}{B}\right)-\frac{2 \pi}{2 B}\right), \frac{1}{\sigma}\left(2 \pi\left(i+\frac{j}{B}\right)+\frac{2 \pi}{2 B}\right)\right], \mathcal{B}_{i, j}=\left[\frac{1}{\sigma}\left(2 \pi\left(i+\frac{j}{B}\right)-\frac{2 \pi(1-\alpha)}{2 B}\right), \frac{1}{\sigma}\left(2 \pi\left(i+\frac{j}{B}\right)+\frac{2 \pi(1-\alpha)}{2 B}\right)\right]$

Definition 12.12.4. Define the discretization of $G(t)$ and $\widehat{G}(f)$,

$$
\begin{aligned}
& G^{\text {dis }}(t)=G(t) \cdot \operatorname{Comb}_{s}(t) \\
& \widehat{G}^{\mathrm{dis}}(f)=\frac{1}{s}\left(\widehat{G} * \operatorname{Comb}_{1 / s}\right)(f) \\
& =\left(\widehat{G} * \mathrm{Comb}_{1}\right)(f)
\end{aligned}
$$

where $|\operatorname{supp}(G(t))|=\frac{l B}{\pi \alpha}, D=\frac{l}{\pi \alpha}, s=|\operatorname{supp}(G(t))| /(B D)=l /(\pi \alpha D)=1$.

Definition 12.6.6. $\forall \sigma>0, b$ and $j \in[B]$. Define,

$$
\begin{aligned}
G_{\sigma, b}^{(j)}(t) & =\frac{1}{\sigma} G(t / \sigma) e^{2 \pi \mathrm{i} t(j / B-\sigma b) / \sigma} \\
\widehat{G}_{\sigma, b}^{(j)}(f) & =\widehat{G}^{\mathrm{dis}}\left(\frac{j}{B}-\sigma f-\sigma b\right)=\sum_{i \in \mathbb{Z}} \widehat{G}\left(i+\frac{j}{B}-\sigma f-\sigma b\right)
\end{aligned}
$$

Lemma 12.6.7. Let $u \in \mathbb{C}^{B}$ be the result of HashToBins under permutation $P_{\sigma, a, b}$, and let $j \in[B]$. Define

$$
\widehat{z}=\widehat{x \cdot H} \cdot \widehat{G}_{\sigma, b}^{(j)}
$$

so

$$
z=(x \cdot H) * G_{\sigma, b}^{(j)} .
$$

Let vector $\widehat{u} \in \mathbb{C}^{B}$ denote the $B$-dimensional DFT of $u$, then $\forall j \in[B]$,

$$
\widehat{u}[j]=z_{\sigma a} .
$$

Proof. Recall $B$ is the number of hash bins. $B \cdot D$ is the number of samples in time signal. Let $W(t)=x \cdot H(t)$, define vector $y \in \mathbb{C}^{B D}$, then $\forall j \in[B D]$, define

$$
y[j]=W(\sigma(j-a)) e^{2 \pi \mathrm{i} \sigma b j}
$$

Recall $G(t)$ denote the $\operatorname{rect}_{B / \alpha}^{* l}(t) \cdot \operatorname{sinc}(t / B)$, then $|\operatorname{supp}(G(t))|=\frac{l B}{\alpha}$. Let vector $G^{\prime} \in \mathbb{C}^{B D}$ is the discretization of $G(t)$, where $G^{\prime}[i]=G(i)$. Then, $\forall j \in[B]$,

$$
u[j]=\sum_{i \in[D]} V[j+i B]
$$

where $V[j]=y[j] \cdot G^{\prime}[j]$ and $G^{\prime}[j]$ is the value at the $j$ th nonzero point of $G^{\text {dis }}(t)$. Applying Claim 12.12 .9 with the definition of $u[j]$ and $V[j+i B]$, gives $\widehat{u}[j]=\widehat{V}[j D], \forall j \in[B]$.

Because of $u$ is the result of $\operatorname{Hash} \operatorname{ToBins}\left(x \cdot H, P_{\sigma, a, b}, G\right)$ and $|\operatorname{supp}(G(t))|=B D$ (choosing $D=l / \alpha)$, then

$$
u[j]=\sum_{i \in \mathbb{Z}} W(\sigma(j+i B-a)) e^{-2 \pi \mathbf{i} \sigma b(j+i B)} G(j+i B)
$$

Then we define $G^{\prime \prime}(t)=G(t / \sigma+a) e^{-2 \pi \mathrm{i} b \sigma(t / \sigma+a)}$ and $Y(t)=W(t) \cdot G^{\prime \prime}(t)$, then immediately, we have

$$
\widehat{G}^{\prime \prime}(f)=\sigma \widehat{G}(\sigma(f-b)) e^{2 \pi \mathrm{i} a \sigma f} \quad \text { and } \quad \widehat{Y}(f)=\widehat{W}(f) * \widehat{G}^{\prime \prime}(f)
$$

Thus, we can rewrite $u[j]$ in the following sense,

$$
\begin{array}{rlr} 
& u[j] & \\
= & \sum_{i \in \mathbb{Z}} W(\sigma(j+i B-a)) e^{-2 \pi \mathbf{i} \sigma b(j+i B)} G(j+i B) & \\
= & \sum_{i \in \mathbb{Z}} W(\sigma(j+i B-a)) G^{\prime \prime}(\sigma(j+i B-a)) & \text { by } G^{\prime \prime}(t)=G(t / \sigma+a) e^{-2 \pi \mathrm{i} b \sigma(t / \sigma+a)} \\
= & \sum_{i \in \mathbb{Z}} Y(\sigma(j+i B-a)) & \quad \text { by } Y(t)=W(t) \cdot G^{\prime \prime}(t)
\end{array}
$$

Then

$$
\begin{array}{rlr} 
& \widehat{u}[j] & \\
= & \sum_{i \in \mathbb{Z}} \widehat{Y}\left(\left(\frac{j}{B}+i\right) / \sigma\right) \cdot \frac{1}{\sigma} \cdot e^{-2 \pi \mathbf{i}\left(\frac{j}{B}+i\right) a} & \text { by Corollary } 12.12 .10 \\
=\sum_{i \in \mathbb{Z}} \int_{-\infty}^{+\infty} \widehat{W}(s) \cdot \widehat{G}^{\prime \prime}\left(\frac{j / B+i}{\sigma}-s\right) \cdot \frac{1}{\sigma} \cdot e^{-2 \pi \mathbf{i} \cdot(j / B+i) a} \mathrm{~d} s & \text { by } \widehat{Y}(f)=\widehat{W}(f) * \widehat{G}^{\prime \prime}(f) \\
=\sum_{i \in \mathbb{Z}} \int_{-\infty}^{+\infty} \widehat{W}(s) \cdot \widehat{G}(j / B+i-\sigma s-\sigma b) \cdot e^{-2 \pi \mathbf{i} \cdot(-\sigma s) a} \mathrm{~d} s & \text { by } \widehat{G}^{\prime \prime}(f)=\sigma \widehat{G}(\sigma(f-b)) e^{2 \pi \mathbf{i} a \sigma f} \\
=\int_{-\infty}^{+\infty} \widehat{W}(s) \cdot \sum_{i \in \mathbb{Z}} \widehat{G}((j / B+i)-\sigma s-\sigma b) e^{2 \pi \mathbf{i} a \sigma s} \mathrm{~d} s & \\
=\int_{-\infty}^{+\infty} \widehat{W}(s) \cdot \widehat{G}^{\mathrm{dis}}\left(\frac{j}{B}-\sigma s-\sigma b\right) e^{-2 \pi \mathbf{i} a \sigma s} \mathrm{~d} s & \text { by } \widehat{G}^{\mathrm{dis}}(f)=\sum_{i \in \mathbb{Z}} \widehat{G}(f+i)
\end{array}
$$

By definition 12.12.4,

$$
\widehat{G}_{\sigma, b}^{(j)}=\widehat{G}^{\mathrm{dis}}\left(\frac{j}{B}-\sigma s-\sigma b\right)=\sum_{i \in \mathbb{Z}} \widehat{G}\left(i+\frac{j}{B}-\sigma s-\sigma b\right)
$$

By definition of $\widehat{z}$, we have

$$
\widehat{z}(s)=\widehat{x \cdot H}(s) \cdot \widehat{G}^{(j)}(s)=\widehat{W}(s) \cdot \widehat{G}^{(j)}(s)
$$

Then $\widehat{u}[j]$ is the $(a \sigma)^{t h}$ inverse Fourier coefficients of $\widehat{z}$, basically,

$$
\widehat{u}[j]=z_{a \sigma}=z(a \sigma)
$$

Thus, we can conclude first computing vector $u \in \mathbb{C}^{B}$. Getting vector $\widehat{u} \in \mathbb{C}^{B}$ by using the Discrete Fourier transform $\widehat{u}=\operatorname{DFT}(u)$. This procedure allows us to sample from time domain to implicitly access the time signal's Fourier transform $\widehat{z}$. If $z$ is one-cluster in frequency domain, then apply one-cluster recovery algorithm.

### 12.13 Algorithm

This section lists the pseudocode of our algorithms.

```
Algorithm 12.3 Get Empirical One Energy and Get Legal One Sample
    procedure GetEmpirical1Energy \((z, T, \Delta)\) - Claim 12.7.10
        \(R_{\text {est }} \leftarrow(T \Delta)^{2}\)
        for \(i=1 \rightarrow R_{\text {est }}\) do
            Choose \(\alpha_{i} \in[0, T]\) uniformly at random
            \(z_{\mathrm{emp}} \leftarrow z_{\mathrm{emp}}+\left|z\left(\alpha_{i}\right)\right|^{2}\)
        end for
        \(z_{\text {emp }} \leftarrow \sqrt{z_{\text {emp }} / R_{\text {est }}}\)
        return \(z_{\mathrm{emp}}\)
    end procedure
    procedure \(\operatorname{GetLegal} 1 \operatorname{Sample}\left(z, \Delta, T, \beta, z_{\mathrm{emp}}\right)\) - Lemma 12.7.1
        \(R_{\text {repeat }} \leftarrow(T \Delta)^{3}, S_{\text {heavy }} \leftarrow \emptyset\)
        for \(i=1 \rightarrow R_{\text {repeat }}\) do
            Choose \(\alpha_{i} \in[0, T]\) uniformly at random
            if \(\left|z\left(\alpha_{i}\right)\right| \geq 0.5 \cdot z_{\text {emp }}\) then
                \(S_{\text {heavy }} \leftarrow S_{\text {heavy }} \cup i\)
            end if
        end for
        for \(i \in S_{\text {heavy }}\) do
            \(w(i) \leftarrow\left|z\left(\alpha_{i}\right)\right|^{2}+\left|z\left(\alpha_{i}+\beta\right)\right|^{2}\)
        end for
        \(\alpha \leftarrow \alpha_{i}\) with probability \(w(i) / \sum_{j \in S_{\text {heavy }}} w(j)\) for \(i \in S_{\text {heavy }}\)
        return \(\alpha\)
    end procedure
```

```
Algorithm 12.4 Locate One Signal, Locate One Inner, and Frequency Recovery One Cluster
    procedure Locate1Signal \(\left(z, T, F, \Delta, z_{\mathrm{emp}}\right)\) - Lemma 12.7.14
        Set \(t \approx \log (F T), t^{\prime}=t / 4, D_{\max } \approx \log _{t^{\prime}}(F T), R_{\mathrm{loc}} \approx \log _{1 / c}(t c), L^{(1)}=2 F\)
        for \(i \in\left[D_{\text {max }}\right]\) do
            \(l \approx 2 F /\left(t^{\prime}\right)^{i-1} \Delta, s \approx c, \widehat{\beta}=\frac{t s}{2 \Delta l}\)
            if \(\widehat{\beta} \gtrsim T /(T \Delta)^{3 / 2}\) then
                break
            else
                \(L^{(i)} \leftarrow \operatorname{LOCATE} 1 \operatorname{InNER}\left(z, \Delta, T, \widehat{\beta}, z_{\mathrm{emp}}, L^{(i-1)}\right)\)
            end if
        end for
        return \(L^{(i)}\)
    end procedure
    procedure Locatelinner \(\left(z, \Delta, T, \widehat{\beta}, z_{\mathrm{emp}}, \widetilde{L}\right)\)
        Let \(v_{q} \leftarrow 0\) for \(q \in[t]\)
        while \(r=1 \rightarrow R_{\text {loc }}\) do
            Choose \(\beta \in\left[\frac{1}{2} \widehat{\beta}, \widehat{\beta}\right]\) uniformly at random
            \(\gamma \leftarrow \operatorname{GETLEGAL} 1 \operatorname{Sample}\left(z, \Delta, T, \beta, z_{\mathrm{emp}}\right)\)
            for \(i \in[m]\) do
                \(s_{i} \in[\beta(\widetilde{L}-\Delta l / 2), \beta(\widetilde{L}+\Delta l / 2)] \cap \mathbb{Z}_{+}, \theta_{i}=\frac{1}{2 \pi \sigma \beta}\left(\phi(x(\gamma) / x(\gamma+\beta))+2 \pi s_{i}\right)\)
                Let \(\theta_{i}\) belong to region \((q)\)
                Then add a vote to region \((q)\) and its two neighbors, i.e., region \((q-1)\) and
    region \((q+1)\)
        end for
        end while
        \(q_{j}^{*} \leftarrow\left\{q \left\lvert\, v_{q}>\frac{R_{\text {loc }}}{2}\right.\right\}\)
        return \(L \leftarrow\) center of region \(\left(q_{j}^{*}\right)\)
    end procedure
    procedure \(\operatorname{FrequencyRecovery1Cluster~}(z, T, F, \Delta)\) - Theorem 12.7.4
        \(z_{\text {emp }} \leftarrow \operatorname{GEtEmpirical1EnERGy}(z, T, \Delta)\)
        for \(r=1 \rightarrow O(k)\) do
            \(L_{r} \leftarrow \operatorname{Locate} 1 \operatorname{SignaL}\left(z, T, F, \Delta, z_{\mathrm{emp}}\right)\)
        end for
        return \(L^{*} \leftarrow \underset{r \in[O(k)]}{\operatorname{median}} L_{r}\)
    end procedure
```

```
Algorithm 12.5 Main Algorithm for One-cluster Recovery
    procedure CFT1Culster \((x, H, T, F)\) - Theorem 12.8.1
        \(\widetilde{f}_{0} \leftarrow \operatorname{FrequencyREcovery} 1 \operatorname{Cluster}(x, H, T, F)\)
        \(\widetilde{x} \leftarrow \operatorname{SignALRECOVERY} 1 \operatorname{ClUSTER}\left(\widetilde{f}_{0}, \operatorname{poly}(k) \Delta_{h}\right)\)
        return \(\widetilde{x}\)
    end procedure
    procedure GenerateIntervals \((d)\)
        \(n \leftarrow y_{0} \leftarrow i \leftarrow 0, m \leftarrow \Theta(d)\)
        while \(y_{i} \leq 1-\frac{9}{m^{2}}\) do
            \(y_{i+1} \leftarrow y_{i}+\frac{\sqrt{1-y_{i}^{2}}}{m}, I_{n+1} \leftarrow\left[y_{i}, y_{i+1}\right], I_{n+2} \leftarrow\left[-y_{i+1},-y_{i}\right]\)
            \(i \leftarrow i+1, n \leftarrow n+2\)
        end while
        \(I_{n+1} \leftarrow\left[y_{i}, 1\right], I_{n+2} \leftarrow\left[-y_{i},-1\right], n \leftarrow n+2\)
        return \(n, I\)
    end procedure
    procedure RobustPolynomialLearning \((x, d, T)\) - Theorem 12.1.2
        \((n, I) \leftarrow \operatorname{GenerateIntervals}(d)\)
        for \(j=1 \rightarrow n\) do
            \(w_{j} \leftarrow\left|I_{j}\right| / 2\)
            Choose \(t_{j}\) from \(I_{j}\) uniformly at random
            \(z_{j} \leftarrow x\left(T \cdot \frac{t_{j}+1}{2}\right)\)
        end for
        \(\widetilde{A}_{j, i} \leftarrow t_{j}^{i}\), for each \((j, i) \in[n] \times\{0,1, \cdots, d\}\)
        \(\alpha \leftarrow \operatorname{LINEARREGRESSIONW}(\widetilde{A}, \widetilde{b}=z, w)\)
        \(Q(t) \leftarrow \sum_{\widetilde{Q}}^{d}{ }_{i=0} \alpha_{i} t^{i}\)
        return \(\widetilde{Q}(t)=Q\left(T \cdot \frac{t+1}{2}\right)\)
    end procedure
    procedure RobustPolynomialLearning \({ }^{+}(x, d, T)\) - Theorem 12.4.5 - a.k.a. Sig-
    NALRECOVERY1Cluster
        \(R \leftarrow \Theta(d)\)
        \((n, I) \leftarrow\) GenerateIntervals \((d)\)
        \(w_{j} \leftarrow\left|I_{j}\right| / 2\), for each \(j \in[n]\)
        for \(i=1 \rightarrow R\) do
            \(Q_{i} \leftarrow \operatorname{RobustPolynomialLEARning}(x, d, T)\)
        end for
        Choose \(t_{j}\) from \(I_{j}\) uniformly at random, for each \(j \in[n]\)
        for \(i=1 \rightarrow R\) do
            \(Q_{i}\left(t_{1}\right), Q_{i}\left(t_{2}\right), \cdots, Q_{i}\left(t_{n}\right) \leftarrow \operatorname{MultipointEvaluation}\left(Q_{i},\left\{t_{1}, t_{2}, \cdots, t_{n}\right\}\right)\)
        end for 912
        \(\widetilde{Q}_{j} \leftarrow \underset{i \in[R]}{\operatorname{median}} Q_{i}\left(t_{j}\right)\), for each \(j \in[n]\)
        \(\widetilde{A}_{j, i} \leftarrow t_{j}^{i}\), for each \((j, i) \in[n] \times\{0,1, \cdots, d\}\)
        \(\alpha \leftarrow \operatorname{LinEARREGRESSIONW}(\widetilde{A}, \widetilde{b}=\widetilde{Q}, w)\)
        return \(Q(t) \leftarrow \sum_{i=0}^{d} \alpha_{i} t^{i}\)
    end procedure
```

```
Algorithm 12.6 Locate \(k\) Signal, Locate \(k\) Inner, Hash To Bins
    procedure LocateKSignal \(\left(x, H, G, T, \Delta, \sigma, b, z_{\mathrm{emp}}\right)\) - Clain 12.7.25
        Set \(t \approx \log (F T), t^{\prime}=t / 4, D_{\max } \approx \log _{t^{\prime}}(F T), R_{\mathrm{loc}} \approx \log _{1 / c}(t c), L^{(1)}=2 F\)
        for \(i \in\left[D_{\max }\right]\) do
            \(\Delta l \approx 2 F /\left(t^{\prime}\right)^{i-1}, s \approx c, \widehat{\beta}=\frac{t s}{2 \sigma \Delta l}\)
            if \(\sigma \widehat{\beta} \gtrsim T /(T \Delta)^{3 / 2}\) then
                break
            else
                \(L^{(i)} \leftarrow \operatorname{LocateKInNER}\left(x, H, G, T, \Delta, \sigma, b, z_{\mathrm{emp}} \widehat{\beta}, U, L^{(i-1)}\right)\)
            end if
        end for
        return \(L^{(i)}\)
    end procedure
    procedure LocateKinner \(\left(x, H, G, T, \Delta, \sigma, b, z_{\mathrm{emp}} \widehat{\beta}, U, \widetilde{L}\right)\)
        Let \(v_{j, q} \leftarrow 0\) for \((j, q) \in[B] \times[t]\)
        for \(r=1 \rightarrow R_{\text {loc }}\) do
            Choose \(\beta \in\left[\frac{1}{2} \widehat{\beta}, \widehat{\beta}\right]\) uniformly at random
            \(\widehat{u}, \widehat{u}^{\prime} \leftarrow \operatorname{GetLegalKSample}\left(x, H, G, T, \Delta, \sigma, \beta, z_{\mathrm{emp}}\right)\)
            for \(j \in[B]\) do
                for \(i \in[m]\) do
                \(\theta_{j, i}=\frac{1}{2 \pi \sigma \beta}\left(\phi\left(\widehat{u}[j] / \widehat{u^{\prime}}[j]\right)+2 \pi s_{i}\right), s_{i} \in\left[\sigma \beta\left(\widetilde{L}_{j}-\Delta l / 2\right), \sigma \beta\left(\widetilde{L}_{j}+\Delta l / 2\right)\right] \cap \mathbb{Z}_{+}\)
                \(f_{j, i}=\theta_{j, i}+b(\bmod F)\)
                suppose \(f_{j, i}\) belongs to region \((j, q)\),
                add a vote to both region \((j, q)\) and two neighbors nearby that region, e.g.
    \(\operatorname{region}(j, q-1)\) and \(\operatorname{region}(j, q+1)\)
                end for
            end for
        end for
        for \(j \in[B]\) do
            \(q_{j}^{*} \leftarrow\left\{q \left\lvert\, v_{j, q}>\frac{R_{\text {loc }}}{2}\right.\right\}\)
            \(L_{j} \leftarrow\) center of region \(\left(j, q_{j}^{*}\right)\)
        end for
        return \(L\)
    end procedure
    procedure \(\operatorname{HashToBins}\left(x, H, G, P_{\sigma, a, b}\right)\) - Lemma 12.6.7
        Compute \(u[j]=\sum_{i \in D} v[j+i B]\)
        \(\widehat{u} \leftarrow \operatorname{FFT}(u)\)
        return \(\widehat{u}\)
    end procedure
```

Algorithm 12.7 Get Empirical $k$ Energy, Get Legal $k$ Sample, and Onestage
procedure GetEmpiricalKEnergy $(x, H, G, T, \Delta, \sigma, b)$ - Claim 12.7.24
$R_{\text {est }} \leftarrow(T \Delta)^{2}$
for $i=1 \rightarrow R_{\text {est }}$ do
Choose $\alpha \in[0, T]$ uniformly at random
$\widehat{u} \leftarrow \operatorname{HashToBins}\left(x, H, G, P_{\sigma, \alpha, b}\right)$
for $j=1 \rightarrow B$ do
$z_{\text {emp }}^{j} \leftarrow z_{\text {emp }}^{j}+\left|\widehat{u}_{j}\right|^{2}$
end for
end for
for $j=1 \rightarrow B$ do
$z_{\mathrm{emp}}^{j} \leftarrow \sqrt{z_{\mathrm{emp}}^{j} / R_{\text {est }}}$
end for
return $z_{\text {emp }}$.
end procedure
procedure GetLegalKSample $\left(x, H, G, T, \Delta, \beta, z_{\mathrm{emp}}\right)$ - Lemma 12.7.23
$R_{\text {repeat }} \leftarrow(T \Delta)^{3}$.
$S_{\text {heavy }}^{j} \leftarrow \emptyset, \forall j \in[B]$
for $i=1 \rightarrow R_{\text {repeat }}$ do
Choose $\alpha \in[0, T]$ uniformly at random
$\widehat{u}^{i} \leftarrow \operatorname{Hash} \operatorname{ToBins}\left(x, H, G, P_{\sigma, \alpha, b}\right)$
$\widehat{u}^{\prime i} \leftarrow \operatorname{Hash} T o B i n s\left(x, H, G, P_{\sigma, \alpha+\beta, b}\right)$
for $j=1 \rightarrow B$ do
if $\left|\widehat{u}_{j}^{i}\right| \geq 0.5 \cdot z_{\text {emp }}^{j}$ then
$S_{\text {heavy }, j} \leftarrow S_{\text {heavy }}^{j} \cup i$
end if
end for
end for
for $j=1 \rightarrow B$ do
for $i \in S_{\text {heavy }}^{j}$ do
$w(i) \leftarrow\left|\widehat{u}_{j}^{i}\right|^{2}+\left|\widehat{u}_{j}^{\prime}\right|^{2}$
end for
$\left(\widehat{v}_{j}, \widehat{v}_{j}^{\prime}\right) \leftarrow\left(\widehat{u}_{j}^{i}, \widehat{u}_{j}^{i}\right)$ with probability $w(i) / \sum_{i^{\prime} \in S_{\text {heavy }}^{j}} w\left(i^{\prime}\right)$ for $i \in S_{\text {heavy }}^{j}$
end for
return $\widehat{v}, \widehat{v}^{\prime} \in \mathbb{C}^{B}$
end procedure
procedure OneStage $(x, H, G, \sigma, b)$ - Lemma 12.7.20
$z_{\mathrm{emp}} \leftarrow \operatorname{GetEmpiricalKEnergy}(x, H, G, T, \Delta, \sigma, b)$
$L \leftarrow \operatorname{LocateKSignal}\left(x, H, G, T, \Delta 9 \not 4 \not d b, z_{\mathrm{emp}}\right)$
end procedure

```
```

Algorithm 12.8 Main Algorithm for $k$-cluster Recovery
procedure CFTKCluster $(x, H, G, T, F)$
$\left\{\widetilde{f}_{1}, \cdots, \widetilde{f}_{l}\right\} \leftarrow \operatorname{FREQUENCyRECOVERYKCluster}(x, H, G, T, F)$
$\widetilde{x} \leftarrow \operatorname{SigNALRECOVERYKCLUSTER}^{+}\left(\widetilde{f}_{1}, \cdots, \widetilde{f}_{l}, \Delta=\operatorname{poly}(k, \log (1 / \delta)) / T, T\right)$
return $\widetilde{x}$ as our hypothesis
end procedure
procedure FrequencyRecoveryKCluster $(x, H, G)$ - Theorem 12.2.6
for $r \in[R]$ do
Choose $\sigma \in\left[\frac{1}{B \Delta_{h}}, \frac{2}{B \Delta_{h}}\right]$ uniformly at random
Choose $b \in\left[0, \frac{2 \pi\lfloor F / \Delta h\rfloor\rfloor}{(\sigma B)}\right]$ uniformly at random
$L_{r} \leftarrow \operatorname{OneStage}(x, H, G, \sigma, b)$
end for
$L^{*} \leftarrow \operatorname{MergedStages}\left(L_{1}, L_{2}, \cdots, L_{R}\right)$
end procedure
procedure SignalRECOVERYKCLUSTER $\left(\widetilde{f}_{1}, \cdots, \widetilde{f}_{l}, \Delta, T\right)$
$d \leftarrow 5 \pi\left((\Delta T)^{1.5}+k^{3} \log k+k \log 1 / \delta\right)$
$m \leftarrow O\left((k d)^{C_{3}} \cdot \log ^{C_{3}} d\right)$ for a constant $C_{3}=5$
for $j=1 \rightarrow m$ do
Sample $t_{j}$ from $[0, T]$ uniformly at random
$\widetilde{\sim}_{j, i_{1} \cdot l+i_{2}} \leftarrow t_{j}^{i_{1}} \cdot e^{2 \pi \mathbf{i} \tilde{f}_{i_{2}} t_{j}}$ for each $\left(i_{1}, i_{2}\right) \in\{0, \cdots, d\} \times[l]$
$\widetilde{b}_{j} \leftarrow x\left(t_{j}\right)$
end for
$\alpha \leftarrow \operatorname{LinEARREGRESSIon}(\widetilde{A}, \widetilde{b})$
return $\widetilde{x}(t) \leftarrow \sum_{i_{1}=0}^{d} \sum_{i_{2}=1}^{l} \alpha_{i_{1} \cdot l+i_{2}} t^{i_{1}} \cdot e^{2 \pi \mathbf{i} \widetilde{i}_{2} t}$
end procedure
procedure SignalRecoveryKCluster ${ }^{+}\left(\widetilde{f}_{1}, \cdots, \widetilde{f}_{l}, \Delta, T\right)$ — Theorem 12.9.1
$R \leftarrow \Theta(k)$
$d \leftarrow 5 \pi\left((\Delta T)^{1.5}+k^{3} \log k+k \log 1 / \delta\right)$
$m \leftarrow O\left((k d)^{C_{3}} \cdot \log ^{C_{3}} d\right)$ for a constant $C_{3}=5$
for $i=1 \rightarrow R$ do
$\widetilde{x}_{i}(t) \leftarrow \operatorname{SignalRECOVERYKClUSTER}\left(\widetilde{f}_{1}, \cdots, \widetilde{f}_{l}, \Delta, T\right)$
end for
for $j=1 \rightarrow m$ do
Sample $t_{j}$ from $[0, T]$ uniformly at random
$\widetilde{A}_{j, i_{1} \cdot l+i_{2}} \leftarrow t_{j}^{i_{1}} \cdot e^{2 \pi \mathbf{i} \tilde{f}_{i_{2}} t_{j}}$ for each $\left(i_{1}, i_{2}\right) \in\{0, \cdots, d\} \times[l]$
$\widetilde{b}_{j} \leftarrow \underset{i \in[R]}{\operatorname{median}} \widetilde{x}_{i}\left(t_{j}\right)$
end for $\quad \sim \sim 15$
$\alpha \leftarrow \operatorname{LinEARREGRESSION}(\widetilde{A}, \widetilde{b})$
return $\widetilde{x}(t) \leftarrow \sum_{i_{1}=0}^{d} \sum_{i_{2}=1}^{l} \alpha_{i_{1} \cdot l+i_{2}} t^{i_{1}} \cdot e^{2 \pi \mathrm{i} \tilde{\mathrm{f}}_{i_{2}} t}$
end procedure

```

\section*{Chapter 13}

\section*{High Dimensional Fourier Transform}

We consider the extensively studied problem of computing a \(k\)-sparse approximation to the \(d\)-dimensional Fourier transform of a length \(n\) signal. Our algorithm uses \(O(k \log k \log n)\) samples, is dimension-free, operates for any universe size, and achieves the strongest \(\ell_{\infty} / \ell_{2}\) guarantee, while running in time comparable to the Fast Fourier Transform. All previous algorithms proceed either via the Restricted Isometry Property or via filter functions. Our approach totally departs from the aforementioned techniques, and we believe is a fresh look to the sparse Fourier transform problem.

\subsection*{13.1 Introduction}

Initiated in discrete signal processing, compressed sensing/sparse recovery is an extensively studied branch of mathematics and algorithms, which postulates that a number of small linear measurements suffice to approximately reconstruct the best \(k\)-sparse approximation of a vector \(x \in \mathbb{C}^{n}\) [CT06, CRT06b, Don06]. The literature on the subject is enormous, and has found applications in imaging, astronomy, seismology etc. One of the initial papers in the field, due to Candes, Romberg and Tao [CRT06a], counts almost 15, 000 references.

Probably the most important subtopic is the sparse Fourier transform, where one desires to reconstruct a \(k\)-sparse vector from Fourier measurements. In other words, measurements are not allowed to be generic, but have to belong to the so-called Fourier ensemble. In Optics imaging [Goo05, Voe11] and Magnetic resonance imaging (MRI) [ASSN08], the physics [Rey89] of the underlying device restricts us to the Fourier ensemble, where the sparse Fourier problem becomes highly relevant. In fact, one of the initial motivations of Candes, Romberg and Tao came out due to the aforementioned applications. The number of samples plays a crucial role: they determine the amount of radiation a patient receives in CT scans, and taking fewer samples can reduce the amount of time the patient needs to stay in the machine. The framework has found its way in practical life-changing applications. Software includes the Compressed Sensing GRAB-VIBE, CS SPACE, CS SEMAC and CS TOF by Siemens [Sie], as well as Compressed Sense by Phillips [Phi]. Its incorporation in the MRI technology allows faster acquisition rates, depiction of dynamic processes or moving organs, as well as acceleration of MRI scanning up to a factor of 40 . On the webpage of SIEMENS Healthineers, for example, one can see the following, as well as numerous similar statements.

This allows bringing the advantages of Compressed Sensing GRASP-VIBE to daily clinical routine.
- Perform push-button, free-breathing liver dynamics.
- Overcome timing challenges in dynamic imaging and respiratory artifacts.
- Expand the patient population eligible for abdominal MRI.

The Fourier transform is in fact ubiquitous: image processing, audio processing, telecommunications, seismology, polynomial multiplication, SUBSET SUM and other textbook algorithms are a few of the examples where the Fast Fourier Transform finds applications. The Fast Fourier Transform by Cooley and Tukey [CT65] runs in \(O(n \log n)\) time, and has far-reaching applications in all of the aforementioned cases. It is thus expected that algorithms which exploit sparsity assumptions about the input, and can outperform FFT in applications are of high practical value. Generally, the two most important parameters one would like to optimize are the sample complexity, i.e. the numbers needed to obtain from the time domain, as well as time needed to approximate the Fourier Transform.

Two different lines of research exist for the problem: the one focuses solely on sample complexity, while the other tries to achieve sublinear time while keeping the sample complexity as low as possible. The first line of research operates via the renowned Restricted Isometry Property (RIP), which proceeds by taking random samples and solving a linear/convex program, or an iterative thresholding procedure [CT06, DDTS06, TG07, BD08, DM08, RV08, BD09b, BD09a, NT09b, NV09, GK09, BD10, NV10, Fou11, Bou14, HR16]. The analysis of the algorithms is performed in the following way, in two steps. The first step ensures
that, after sampling an appropriate number of points from the time domain, the inverse DFT matrix restricted on the rows indexed by those points acts as a near isometry on the space of \(k\)-sparse vectors. All of the state of the art results [CT06, RV08, Bou14, HR16] employ chaining arguments to make the analysis of this sampling procedure as tight as possible. The second part is how to exploit the aforementioned near-isometry property to find the best \(k\)-sparse approximation to the signal. There the approaches either follow an iterative procedure which gradually denoise the signal [BD08, NT09b, NV09], or perform \(\ell_{1}\) minimization [CT06], a method that promotes sparsity of solutions.

The second line of research tries to implement arbitrary linear measurements via sampling Fourier coefficients [GL89, Man92, KM93, GGI+ 02, AGS03, GMS05, Iwe08, Iwe10, HIKP12a, HIKP12b, LWC13, Iwe13, PR14, IKP14, IK14, Kap16, Kap17, CI17, BZI17, MZIC17, LN19] and use sparse functions (in the time domain) which behave like bandpass filters in the frequency domain. The seminal work of Kapralov [Kap17] achieves \(O(k \log n)\) samples and running time that is some log factors away from the sample complexity. This would be the end of the story, apart from the fact that this algorithm does not scale well with dimension, since it has an exponential dependence on \(d\). Indeed, in many applications, one is interested in higher dimensions, rather than the one-dimensional case. The main reason \({ }^{1}\) why this curse of dimensionality appears is due to the lack of dimension-independent ways to construct functions that approximate the \(\ell_{\infty}\) ball and are sufficiently sparse in the time domain. A very nice work of Kapralov, Velingker and Zandieh [KVZ19] tries to remedy that by combining the standard execution of FFT with careful aliasing, but their algorithm

\footnotetext{
\({ }^{1}\) But not the only one: pseudorandom permutations for sparse FT in high dimensions also incur an exponential loss, and it is not known whether this can be avoided.
}
works in a noiseless setting, and has a polynomial, rather than linear, dependence on \(k\); the running time is polynomial in \(k, \log n\) and the exponential dependence is avoided. It is an important and challenging question whether a robust and more efficient algorithm can be found.

We note that in many applications, such as MRI or computed tomography (CT), the main focus is the sample complexity; the algorithms that have found their way to industry are, to the best of our knowledge, not concerned with sublinear running time, but with the number of measurements, which determine the acquisition time, or in CT the radiation dose the patient receives. Lastly, we bring to the readers' attention the recent work on sparse Fourier transform in the continuous setting, see [Iwe10, Iwe13, Iwe13, \(\mathrm{BCG}^{+} 12, \mathrm{PS} 15\), CKPS16, \(\mathrm{AKM}^{+}\)19].

Our Contribution. We give a new algorithm for the sparse Fourier transform problem, which has \(O(k \log n \log k)\) sample complexity for any dimension, and achieves the \(\ell_{\infty} / \ell_{2}\) guarantee \({ }^{2}\), while running in time \(\widetilde{O}(n)\). The previous state of the art algorithm that achieved such a guarantee is the work of Indyk and Kapralov [IK14], which has \(2^{O(d \log d)} k \log n\) sample complexity; an exponentially worse dependence on \(d\). The work of [HR16] obtains \(O\left(k \log n \log ^{2} k\right)\) samples in any dimension, but has a much weaker guarantee, while their approach requires \(\Omega(k \log n \log k)\) samples in high dimensions [Rao19]. Moreover, the algorithm in [IK14] operates when the universe size in each dimension is a power of 2 , whereas there is no restriction in our work. To obtain our result, we introduce a set of new tech-

\footnotetext{
\({ }^{2}\) This is the strongest guarantee in the sparse recovery literature. See also the caption of the table in Section 1.2
}
niques, deviating from previous work, which used the Restricted Isometry Property and/or filter functions.

\subsection*{13.1.1 Preliminaries}

For any positive integer \(n\), we use \([n]\) to denote \(\{1,2, \cdots, n\}\).
We assume that the universe size \(n=p^{d}\) for any positive integer \(p\). Our algorithm facilitates \(n=\Pi_{j=1}^{d} p_{j}\) for any positive integers \(p_{1}, \ldots, p_{d}\), but we decide to present the case \(n=p^{d}\) for ease of exposition; the proof is exactly the same in the more general case. Let \(\omega=e^{2 \pi \mathbf{i} / p}\) where \(\mathbf{i}=\sqrt{-1}\). We will work with the normalized \(d\)-dimensional Fourier transform
\[
\widehat{x}_{f}=\frac{1}{\sqrt{n}} \sum_{t \in[p]^{d}} x_{t} \cdot \omega^{f^{\top} t}, \forall f \in[p]^{d}
\]
and the inverse Fourier transform is
\[
x_{t}=\frac{1}{\sqrt{n}} \sum_{f \in[p]^{d}} \widehat{x}_{f} \cdot \omega^{-f^{\top} t}, \forall t \in[p]^{d} .
\]

For any vector \(x\) and integer \(k\), we denote \(x_{-k}\) to be the vector obtained by zeroing out the largest (in absolute value) \(k\) coordinates from \(x\).

\subsection*{13.1.2 Our result}

Apart from being dimension-independent and working for any universe size, our algorithm satisfies \(\ell_{\infty} / \ell_{2}\), which is the strongest guarantee out of the standard guarantees considered in compressed sensing tasks. A guarantee \(G_{1}\) is stronger than guarantee \(G_{2}\) if for any \(k\)-sparse recovery algorithm that satisfies \(G_{1}\) we can obtain a \(\Omega(k)\)-sparse recovery
algorithm that satisfies \(G_{2}\). See also below for a comparison between \(\ell_{\infty} / \ell_{2}\) and \(\ell_{2} / \ell_{2}\), the second stronger guarantee.

Previous work is summarized in Table 13.1. Our result is the following.
Theorem 13.1.1 (main result, informal version). Let \(n=p^{d}\) where both \(p\) and \(d\) are positive integers. Let \(x \in \mathbb{C}^{[p]^{d}}\). Let \(k \in\{1, \ldots, n\}\). Assume that \(R^{*} \geq\|\widehat{x}\|_{\infty} /\left\|\widehat{x}_{-k}\right\|_{2}\) where \(\log R^{*}=\) \(O(\log n)\) (signal-to-noise ratio). There is an algorithm that takes \(O(k \log k \log n)\) samples from \(x\), runs in \(\widetilde{O}(n)\) time, and outputs a \(O(k)\)-sparse vector \(y\) such that
\[
\|\widehat{x}-y\|_{\infty} \leq \frac{1}{\sqrt{k}}\left\|\widehat{x}_{-k}\right\|_{2}
\]
holds with probability at least \(1-1 / \operatorname{poly}(n)\).

Comparison between \(\ell_{\infty} / \ell_{2}\) and \(\ell_{2} / \ell_{2}\left(\right.\) or \(\left.\ell_{2} / \ell_{1}\right)\). For the sake of argument, we will consider only the \(\ell_{2} / \ell_{2}\) guarantee which is stronger than \(\ell_{2} / \ell_{1}\). The \(\ell_{2} / \ell_{2}\) guarantee is the following: for \(\widehat{x} \in \mathbb{C}^{n}\) one should output a \(z\) such that \(\|\widehat{x}-z\|_{2} \leq C\left\|\widehat{x}_{-k}\right\|_{2}\), where \(C>1\) is the approximation factor. Consider \(C=1.1^{3}\), and think of the following signal: for a set \(S\) of size \(0.05 k\) we have \(\left|\widehat{x}_{i}\right|=\frac{2}{\sqrt{k}}\left\|\widehat{x}_{\bar{S}}\right\|_{2}\). Then the all zeros vectors is a valid solution for the \(\ell_{2} / \ell_{2}\) guarantee, since
\[
\|\overrightarrow{0}-\widehat{x}\|_{2}^{2}=\left\|\widehat{x}_{S}\right\|_{2}^{2}+\left\|\widehat{x}_{\bar{S}}\right\|_{2}^{2}=0.05 k \cdot \frac{4}{k}\left\|\widehat{x}_{\bar{S}}\right\|_{2}^{2}+\left\|\widehat{x}_{\bar{S}}\right\|_{2}^{2}=1.2\left\|\widehat{x}_{\bar{S}}\right\|_{2}^{2}<1.1^{2}\left\|\widehat{x}_{\bar{S}}\right\|_{2}^{2}
\]

\footnotetext{
\({ }^{3}\) This is the case with the RIP based approaches, which obtain \(\ell_{2} / \ell_{1}\). In fact many filter-based algorithms facilitate \((1+\epsilon)\) on the right hand side, with the number of measurements being multiplied by \(\epsilon^{-1}\). By enabling the same dependence on \(\epsilon^{-1}\) our algorithm facilitates a multiplicative \(\epsilon\) factor on right hand side of the \(\ell_{\infty} / \ell_{2}\), which makes it much stronger. Thus, a similar argument can go through.
}
\begin{tabular}{|l|l|l|l|l|l|}
\hline Reference & Samples & Time & Filter & RIP & Guarantee \\
\hline [GMS05] & \(k \log ^{O(d)} n\) & \(k \log ^{O(d)} n\) & Yes & No & \(\ell_{2} / \ell_{2}\) \\
\hline [CT06] & \(k \log ^{6} n\) & poly \((n)\) & No & Yes & \(\ell_{2} / \ell_{1}\) \\
\hline [RV08] & \(k \log ^{2} k \log (k \log n) \log n\) & \(\widetilde{O}(n)\) & No & Yes & \(\ell_{2} / \ell_{1}\) \\
\hline [HIKP12a] & \(k \log ^{d} n \log (n / k)\) & \(k \log ^{d} n \log (n / k)\) & Yes & No & \(\ell_{2} / \ell_{2}\) \\
\hline [CGV13] & \(k \log ^{3} k \log n\) & \(\widetilde{O}(n)\) & No & Yes & \(\ell_{2} / \ell_{1}\) \\
\hline [IK14] & \(2^{d \log d} k \log n\) & \(\widetilde{O}(n)\) & Yes & No & \(\ell_{\infty} / \ell_{2}\) \\
\hline [Bou14] & \(k \log k \log ^{2} n\) & \(\widetilde{O}(n)\) & No & Yes & \(\ell_{2} / \ell_{1}\) \\
\hline [HR16] & \(k \log ^{2} k \log n\) & \(\widetilde{O}(n)\) & No & Yes & \(\ell_{2} / \ell_{1}\) \\
\hline [Kap16] & \(2^{d^{2}} k \log n \log \log n\) & \(2^{d^{2} k \log ^{d+3} n}\) & Yes & No & \(\ell_{2} / \ell_{2}\) \\
\hline [KVZ19] & \(k^{3} \log ^{2} k \log { }^{2} n\) & \(k^{3} \log ^{2} k \log ^{2} n\) & Yes & Yes & Exactly \(k\)-sparse \\
\hline Theorem 13.1.1 & \(k \log k \log n\) & \(\widetilde{O}(n)\) & No & No & \(\ell_{\infty} / \ell_{2}\) \\
\hline
\end{tabular}

Table 13.1: \(n=p^{d}\). We ignore the \(O\) for simplicity. The \(\ell_{\infty} / \ell_{2}\) is the strongest possible guarantee, with \(\ell_{2} / \ell_{2}\) coming second, \(\ell_{2} / \ell_{1}\) third and exactly \(k\)-sparse being the less strong. We note that [CT06, RV08, CGV13, Bou14, HR16] obtain a uniform guarantee, i.e. with \(1-1 / \operatorname{poly}(n)\) they allow reconstruction of all vectors; \(\ell_{\infty} / \ell_{2}\) and \(\ell_{2} / \ell_{2}\) are impossible in the uniform case [CDD09]. We also note that [RV08, CGV13, Bou14, HR16] give improved analysis of the Restricted Isometry property; the algorithm is suggested and analyzed (modulo the RIP property) in [BD08]. The work in [HIKP12a] does not explicitly state the extension to the \(d\)-dimensional case, but can easily be inferred from the arguments. [HIKP12a, IK14, Kap16, KVZ19] work when the universe size in each dimension are powers of 2 . We also assume that the signal-to-noise ratio is bounded by a polynomial of \(n\), which is a standard assumption in the sparse Fourier transform literature [HIKP12a, IK14, Kap16, Kap17, LN19].

It is clear that since \(\overrightarrow{0}\) is a possible output, we may not recover any of the coordinates in \(S\), which is the set of "interesting" coordinates. On the other hand, the \(\ell_{\infty} / \ell_{2}\) guarantee does allow the recovery of every coordinate in \(S\). This is a difference of recovering all \(0.05 k\) versus 0 coordinates. From the above discussion, one can conclude in the case where there is too much noise, \(\ell_{2} / \ell_{2}\) becomes much weaker than \(\ell_{\infty} / \ell_{2}\), and can be even meaningless. Thus, \(\ell_{\infty} / \ell_{2}\) is highly desirable, whenever it is possible. The same exact argument holds for
\(\ell_{2} / \ell_{1}\).

\subsection*{13.1.3 Summary of previous Filter function based technique}

One of the two ways to perform Fourier sparse recovery is by trying to implement arbitrary linear measurements, with algorithms similar to the ubiquitous CountSketch [CCF02]. In the general setting CountSketch hashes every coordinate to one of the \(O(k)\) buckets, and repeats \(O(\log n)\) times with fresh randomness. Then, it is guaranteed that every heavy coordinate will be isolated, and the contribution from non-heavy elements is small. To implement this in the Fourier setting becomes a highly non-trivial task however: one gets access only to the time-domain but not the frequency domain. One natural way to do this is to exploit the convolution theorem and find a function which is sparse in the time domain and approximates the indicator of an interval (rectangular pulse) in the frequency domain; these functions are called (bandpass) filters. Appropriate filters were designed in [HIKP12a, HIKP12b]: they were very good approximations of the rectangular pulse, i.e. the contribution from elements outside the passband zone contributed only by \(1 / \operatorname{poly}(n)\) their mass. These filters had an additional \(\log n\) factor (in one dimension) in the sparsity of the time domain and they are sufficient for the purposes of [HIKP12a], but in high dimensions this factor becomes \(\log ^{d} n\). Filters based on the Dirichlet kernel give a better dependence in terms of sparsity and dimension (although still an exponential dependence on the latter), but the leak to subsequent buckets, i.e. coordinates outside the passband zone contribute a constant fraction of their mass, in contrast to the filter used in [HIKP12a]. Thus one should perform additional denoising, which is a non-trivial task. The seminal work of Indyk and Kapralov [IK14] was the first that showed how to perform sparse recovery with these
filters, and then Kapralov [Kap16, Kap17] extended this result to run in sublinear time. We note, that any filter-based approach with filters which approximate the \(\ell_{\infty}\) box, suffers from the curse of dimensionality. [KVZ19] devised an algorithm which avoids the curse of dimensionality by using careful aliasing, but it works in the noiseless case and has a cubic dependence on \(k\).

\subsection*{13.1.4 RIP property-based algorithms: a quick overview}

We say the matrix \(A \in \mathbb{C}^{m \times n}\) satisfies RIP (Restricted Isometry Property [CT05]) of order \(k\) if for all \(k\)-sparse vectors \(x \in \mathbb{C}^{n}\) we have \(\|A x\|_{2}^{2} \approx\|x\|_{2}^{2}\). A celebrated result of Candes and Tao [CT06] shows that Basis Pursuit ( \(\ell_{1}\) minimization) suffices for sparse recovery, as long as the samples from the time domain satisfy RIP. In [CT06] it was also proved using generic chaining that random sampling with oversampling factor \(O\left(\log ^{6} n\right)\) gives RIP property for any orthonormal matrix with bounded entries by \(1 / \sqrt{n}\). Then [RV08] improved the bound to \(O\left(k \cdot \log ^{2} k \cdot \log (k \log n) \cdot \log n\right)\) and [CGV13] improved it to \(O\left(k \cdot \log ^{3} k\right.\). \(\log n)\). Subsequent improvement by Bourgain [Bou14] has lead to \(O\left(k \log k \cdot \log ^{2} n\right)\) samples, improved by Haviv and Regev to \(O\left(k \log ^{2} k \cdot \log n\right)[H R 16]\). The fastest set of algorithms are iterative ones: for example Iterative Hard Thresholding [BD09a] or CoSaMP [NT09b] run \(O(\log n)\) iterations \({ }^{4}\) and each iteration takes \(\widetilde{O}(n)\) time.

We note the very recent lower bound of [Rao19]: a subsampled Fourier matrix that satisfies the RIP properties should have \(\Omega(k \log k \cdot d)\) rows \(^{5}\). This bound is particularly useful in high dimensions, since it deteriorates to a trivial bound in low dimensions. We still believe

\footnotetext{
\({ }^{4}\) To be precise, their running time is logarithmic in the signal-to-noise ratio, but we assumed throughout the paper that this quantity is polynomial in \(n\).
\({ }^{5}\) [BLLM19] independently gives a similar bound for \(d=\log n\).
}
though that a bound of \(\Omega(k \log k \log n)\) should hold in all dimensions. Thus, what remains is to obtain the \(\ell_{2} / \ell_{2}\) guarantee by giving a tighter analysis, and removing the one \(\log k\) factor to match the lower bound, but our algorithm already allows Fourier sparse recovery with these number of samples, even with a stronger guarantee.

\subsection*{13.1.5 Overview of our technique}

Let \(x \in \mathbb{C}^{[p]^{d}}\) denote our input signal in the time domain. In the following we assume the knowledge of \(\mu=\frac{1}{\sqrt{k}}\left\|\widehat{x}_{-k}\right\|_{2}\) and \(R^{*}\) which is an upper bound of \(\|\widehat{x}\|_{\infty} / \mu\), and bounded by poly \((n)\). These are standard assumption [HIKP12a, IK14, Kap16, Kap17, LN19] in the sparse Fourier transform literature. The bound on \(R^{*}\) is useful for bounding the running time (or the number of measurements in [HIKP12a]) and in any of [HIKP12a, IK14, Kap16, Kap17, LN19] a \(\log n\) can be substituted by \(\log R^{*}\) in the general case, which is also the case for our algorithm. We note that our algorithm will be correct with probability \(1-1 / \operatorname{poly}(n)\) whenever \(R^{*}<2^{n^{100}}\); this is fine for every reasonable application. We assumed the rounding errors in FFT computation to be negligible, similarly to Remark 3.4 in [IK14].

Consider the simplest scenario: \(d=1, p\) is a prime number and a 1 -sparse signal \(\widehat{x}\) which is 1 on some frequency \(f^{*}\). From a sample \(x_{t}\) in the time-domain what would be the most reasonable way to find \(f^{*}\) ? For every \(f \in[p]\) we would compute
\[
\sqrt{n} \omega^{f t} x_{t}=\sqrt{n} \omega^{f t} \cdot \frac{1}{\sqrt{n}} \sum_{f^{\prime} \in[p]} \omega^{-f^{\prime} t} \widehat{x}_{f^{\prime}}=\omega^{\left(f-f^{*}\right) t}
\]
and keep, for \(t \neq 0\), the frequency that gives a real number. Since \(\left(f-f^{*}\right) t\) will be zero only for \(f=f^{*}\), we are guaranteed correct recovery. In the noisy and multi-dimensional case or \(p\) is an arbitrary integer, however, this argument will not work, because of the presence of
contribution from other elements and the fact that \(\left(f-f^{*}\right)^{\top} t\) can be zero modulo \(p\) for other frequencies apart from \(f\). However, we can take a number of samples \(t\) and average \(\sqrt{n} \omega^{f^{\top} t}\), and hope that this will make the contribution from other frequencies small enough, so that we can infer whether \(f\) corresponds to a heavy coordinate or not. More specifically, we pick a list \(T\) of size \(O(k)\) uniformly at random from \([p]^{d}\) and compute
\[
\frac{\sqrt{n}}{|T|} \sum_{t \in T} \omega^{f^{\top} t} x_{t}
\]
for all frequencies \(f\). We show that if \(|T|=O(k)\) our estimator is good on average (and later we will maintain \(O(\log n)\) independent instances and take the median to make sure with probability \(1-1 / \operatorname{poly}(n)\) the estimators for all the frequencies are good), and in fact behaves like a crude filter, similarly to the ones used in [IK14], in the sense that every coordinate contributes a non-trivial amount to every other coordinate. However, these estimators do not suffer from the curse of dimensionality and our case is a little bit different, requiring a quite different handling. The main reason is that in contrast to the filters used in [IK14], there is not an easy way to formulate an isolation argument from heavy elements that would allow easy measurement re-use, like Definition 5.2 and Lemma 5.4 from [IK14]. Buckets induced by filter functions have a property of locality, since they correspond to approximate \(\ell_{\infty}\) boxes (with a polynomial decay outside of the box) in \([p]^{d}\) : the closer two buckets are the more contribute the elements of one into the other. Our estimators on the other side do not enjoy such a property. Thus, one has to go via a different route.

In what follows, we will discuss how to combine the above estimators with an iterative loop that performs denoising, i.e. removes the contribution of every heavy element to other heavy elements.

We first implement a procedure which takes \(O(k \log n)\) uniform random measurements from \(x\) and has the guarantee that for any \(\nu \geq \mu\) any \(y \in \mathbb{C}^{[p]^{d}}\) where \(\|\widehat{x}-y\|_{\infty} \leq 2 \nu\) and \(y\) is independent from the randomness of the measurements, the procedure outputs a \(O(k)\)-sparse \(z \in \mathbb{C}^{[p]^{d}}\) such that \(\|\widehat{x}-y-z\|_{\infty} \leq \nu\) with probability \(1-1 / \operatorname{poly}(n)\).

Lemma 13.1.2 (LinfinityReduce procedure/data structure, informal). Let \(\mu=\frac{1}{\sqrt{k}}\left\|\widehat{x}_{-k}\right\|_{2}\), and \(\nu \geq \mu\). Let \(\mathcal{T}^{(0)}\) be a list of \(O(k \log n)\) i.i.d. elements in \([p]^{d}\). Let \(S\) be top \(O(k)\) coordinates in \(\widehat{x}\). There is a procedure that takes \(\left\{x_{t}\right\}_{t \in \mathcal{T}}, y \in \mathbb{C}^{[p]^{d}}\) and \(\nu\) as input, runs in \(\widetilde{O}(n)\) time, and outputs \(z \in \mathbb{C}^{[p]^{d}}\) so that if \(\|\widehat{x}-y\|_{\infty} \leq 2 \nu, \operatorname{supp}(y) \subseteq S\) and \(y\) is independent from the randomness of \(\mathcal{T}^{(0)}\), then \(\|\widehat{x}-y-z\|_{\infty} \leq \nu\) and \(\operatorname{supp}(z) \subseteq S\) with probability \(1-1 / \operatorname{poly}(n)\) under the randomness of \(\mathcal{T}^{(0)}\).

Namely, we can take \(O(k \log n)\) measurements and run the procedure in Lemma 13.1.2 to reduce (the upper bound of) the \(\ell_{\infty}\) norm of the residual signal by half. We call the procedure in Lemma 13.1.2 LinfinityReduce procedure. More generally, we can take \(O(H \cdot k \log n)\) measurements and run the LinfinityReduce procedure \(H\) times to reduce the \(\ell_{\infty}\) norm of the residual signal to \(1 / 2^{H}\) of its original magnitude, with failure probability at most \(1 / \operatorname{poly}(n)\). This is because if \(\nu \geq 2^{H} \mu\) and \(\|\widehat{x}-y\|_{\infty} \leq \nu\), then we can proceed in \(H\) iterations where in the \(h\)-th iteration \((h \in[H])\) we can take \(O(k \log n)\) fresh measurements from \(x\) and run the LinfinityReduce procedure to make the \(\ell_{\infty}\) norm of the residual signal at most \(2^{-h} \nu\). Note that if we set \(H=\log R^{*}\), we have already obtained a recovery algorithm taking \(O\left(k \log n \log R^{*}\right)\) measurements, because we can drive down (the upper bound of ) the \(\ell_{\infty}\) norm of the residual signal from \(\|\widehat{x}\|_{\infty}\) to \(\mu\) in \(\log R^{*}\) iterations.

\subsection*{13.1.5.1 \(O(k \log n)\) measurements for \(k=O(\log n)\)}

We first discuss a measurement reuse idea that leads us to a sparse recovery algorithm (Algorithm 13.1) taking \(O(k \log n)\) measurements for \(k=O(\log n)\). We set \(H=5\), and let \(\mathcal{T}=\left\{\mathcal{T}^{(1)}, \ldots, \mathcal{T}^{(H)}\right\}\), where each \(\mathcal{T}^{(h)}\) is a list of \(O(k \log n)\) i.i.d. elements in \([p]^{d}\). Note that \(\mathcal{T}^{(1)}, \ldots, \mathcal{T}^{(H)}\) are independent. In our sparse Fourier recovery algorithm, we will measure \(x_{t}\) for all \(t \in \mathcal{T}\).

In a nutshell, our approach finely discretizes the space of possible trajectories the algorithm could evolve, and carefully argues about the correctness of the algorithm by avoiding the intractable union-bound over all trajectories.

Recovery algorithm. The recovery algorithm proceeds in \(\log R^{*}-H+1\) iterations, where each iteration (except the last iteration) the goal is to reduce the upper bound of \(\ell_{\infty}\) norm of the residual signal by half. Initially, the upper bound is \(R^{*}\). It is important to note that we use the same measurements \(\mathcal{T}=\left\{\mathcal{T}^{(1)}, \ldots, \mathcal{T}^{(H)}\right\}\) in all of these \(\log R^{*}-H+1\) iterations.

In the following, we will describe one iteration of the recovery algorithm. Let \(y \in \mathbb{C}^{[p]^{d}}\) denote the sparse vector recovered so far, and let the upper bound of \(\|\widehat{x}-y\|_{\infty}\) be \(2 \nu\). Running the LinfinityReduce procedure \(H\) times where in the \(h\)-th time we use measurements in \(\mathcal{T}^{(h)}\), we obtain a \(O(k)\)-sparse \(z\) such that with probability \(1-1 / \operatorname{poly}(n),\|\widehat{x}-y-z\|_{\infty} \leq\) \(2^{1-H} \nu \leq 0.1 \nu\) (we call such \(z\) a desirable output by the LinfinityReduce procedure). Instead of taking \(y+z\) as our newly recovered sparse signal, for each \(f \in \operatorname{supp}(y+z)\), we project \(y_{f}+z_{f}\) to the nearst points in \(\mathcal{G}_{0.6 \nu}:=\{0.6 \nu(x+y \mathbf{i}): x, y \in \mathbb{Z}\}\) and assign to \(y_{f}^{\prime}\), where \(y^{\prime}\) denotes our newly recovered sparse signal. For all \(f \notin \operatorname{supp}(y+z)\), we let \(y_{f}^{\prime}=0\).
```

Algorithm 13.1 Fourier Sparse Recovery by Projection, $O(k \log n)$ Measurements When
$k=O(\log n)$
procedure FourierSparseRecoveryByProjection $\left(x, n, k, \mu, R^{*}\right)$
Section 13.1.5.1
2: $\quad$ Require that $\mu=\frac{1}{\sqrt{k}}\left\|\widehat{x}_{-k}\right\|_{2}$ and $R^{*} \geq\|\widehat{x}\|_{\infty} / \mu$
$H \leftarrow 5, \nu \leftarrow \mu R^{*} / 2, y \leftarrow \overrightarrow{0} \quad \triangleright y \in \mathbb{C}^{[p]^{d}}$ refers to the sparse vector recovered so far
Let $\mathcal{T}=\left\{\mathcal{T}^{(1)}, \cdots, \mathcal{T}^{(H)}\right\}$ where each $\mathfrak{T}^{(h)}$ is a list of i.i.d. uniform samples in $[p]^{d}$
while true do
$\nu^{\prime} \leftarrow 2^{1-H} \nu$
Use $\left\{x_{t}\right\}_{t \in \mathcal{T}}$ to run the LinfinityReduce procedure (in Lemma 13.1.2) $H$ times
(use samples in $\mathcal{T}^{(h)}$ for each $h \in[H]$ ), and finally it finds $z$ so that $\|\widehat{x}-y-z\|_{\infty} \leq \nu^{\prime}$
if $\nu^{\prime} \leq \mu$ then return $y+z \quad \triangleright$ We found the solution
$y^{\prime} \leftarrow \overrightarrow{0}$
for $f \in \operatorname{supp}(y+z)$ do
$y_{f}^{\prime} \leftarrow \Pi_{0.6 \nu}\left(y_{f}+z_{f}\right) \quad \triangleright$ We want $\left\|\widehat{x}-y^{\prime}\right\|_{\infty} \leq \nu$ and the depend-
end for $\quad \triangleright$ ence between $y^{\prime}$ and $\mathcal{T}$ is under control
$y \leftarrow y^{\prime}, \nu \leftarrow \nu / 2$
end while
end procedure

```

To simplify our exposition, here we introduce some notations. We call \(\mathcal{G}_{0.6 \nu}\) a grid of side length \(0.6 \nu\), and we generalize the definition to any side length. Namely, for any \(r_{g}>0\), let grid \(\mathcal{G}_{r_{g}}:=\left\{r_{g}(x+y \mathbf{i}): x, y \in \mathbb{Z}\right\}\). Moreover, we define \(\Pi_{r_{g}}: \mathbb{C} \rightarrow \mathcal{G}_{r_{g}}\) to be the mapping that maps any element in \(\mathbb{C}\) to the nearest element in \(\mathcal{G}_{r_{g}}\). Now we can write \(y^{\prime}\) as
\[
y_{f}^{\prime}= \begin{cases}\Pi_{0.6 \nu}\left(y_{f}+z_{f}\right), & \text { if } f \in \operatorname{supp}(y+z) \\ 0, & \text { if } f \notin \operatorname{supp}(y+z)\end{cases}
\]

At the end of each iteration, we assign \(y^{\prime}\) to \(y\), and shrink \(\nu\) by half. In the last iteration, we will not compute \(y^{\prime}\), instead we output \(y+z\). We present the algorithm in Algorithm 13.1.

Analysis. We analyze \(y^{\prime}\) conditioned on the event that \(\|\widehat{x}-y-z\|_{\infty} \leq 0.1 \nu\) (i.e. \(z\) is a desirable output by the LinfinityReduce procedure, which happens with probability \(1-1 / \operatorname{poly}(n))\). We will prove that \(y^{\prime}\) has two desirable properties: \((1)\left\|\widehat{x}-y^{\prime}\right\|_{\infty} \leq \nu\); (2) the dependence between \(y^{\prime}\) and our measurements \(\mathfrak{T}\) is under control so that after taking \(y^{\prime}\) as newly recovered sparse signal, subsequent executions of the LINFINITYREDUCE procedure with measurements \(\mathfrak{T}\) still work with good probability. Property (1) follows from triangle inequality and the fact that \(\|\widehat{x}-(y+z)\|_{\infty} \leq 0.1 \nu\) and \(\left\|(y+z)-y^{\prime}\right\|_{\infty} \leq 0.6 \nu\). We now elaborate on property (2). We can prove that for any \(f \in[p]^{d}\),
\[
y_{f}^{\prime} \in\left\{\Pi_{0.6 \nu}\left(\widehat{x}_{f}+0.1 \nu(\alpha+\beta \mathbf{i})\right): \alpha, \beta \in\{-1,1\}\right\} .
\]

Let \(S\) denote top \(26 k\) coordinates (in absolute value) of \(\widehat{x}\). We can further prove that for any \(f \in \bar{S}, y_{f}^{\prime}=0\). Therefore, the total number of possible \(y^{\prime}\) is upper bounded by \(4^{|S|}=4^{O(k)}\). If \(k=O(\log n)\), we can afford union bounding all \(4^{O(k)}=\operatorname{poly}(n)\) possible \(y^{\prime}\), and prove that with probability \(1-1 / \operatorname{poly}(n)\) for all possible value of \(y^{\prime}\) if we take \(y^{\prime}\) as our newly recovered sparse signal then in the next iteration the LinfinityReduce procedure with measurements \(\mathcal{T}\) gives us a desirable output.

Sufficient event. More rigorously, we formulate the event that guarantees successful execution of Algorithm 13.1. Let \(\mathcal{E}_{1}\) be the event that for all \(O\left(\log R^{*}\right)\) possible values of \(\nu \in\left\{\mu \frac{R^{*}}{2}, \mu \frac{R^{*}}{4}, \ldots, \mu 2^{H-1}\right\}\), for all possible vector \(y\) where \(y_{f}=0\) for \(f \in \bar{S}\) and \(y_{f} \in\left\{\Pi_{0.6 \nu}\left(\widehat{x}_{f}+0.1 \nu(\alpha+\beta \mathbf{i})\right): \alpha, \beta \in\{-1,1\}\right\}\) for \(f \in S\) (we also need to include the case that \(y=\overrightarrow{0}\) for the success of the first iteration), running the LinfinityReduce procedure (in Lemma 13.1.2) \(H\) times (where in the \(h\)-th time measurements \(\left\{x_{t}\right\}_{t \in \mathcal{T}^{(h)}}\) are used
to reduce the error from \(2^{2-h} \nu\) to \(2^{1-h} \nu\) ) finally gives \(z\) so that \(\|\widehat{x}-y-z\|_{\infty} \leq 2^{1-H} \nu\). The randomness of \(\mathcal{E}_{1}\) comes from \(\mathcal{T}=\left\{\mathcal{T}^{(1)}, \ldots, \mathcal{T}^{(H)}\right\}\).

First, event \(\mathcal{E}_{1}\) happens with probability \(1-1 / \operatorname{poly}(n)\). This is because there are \(4^{O(k)} \log R^{*}\) possible combinations of \(\nu\) and \(y\) to union bound, and each has failure probability at most \(1 / \operatorname{poly}(n)\). For \(k=O(\log n)\), and any \(R^{*}<2^{n^{100}}\) this gives the desired result. Second, conditioned on event \(\mathcal{E}_{1}\) happens, Algorithm 13.1 gives correct output. This can be proved by a mathematical induction that in the \(t\)-th iteration of the while-true loop in Algorithm 13.1, \(\|\widehat{x}-y\|_{\infty} \leq 2^{-t} \mu R^{*}\).

\subsection*{13.1.5.2 \(O(k \log k \log n)\) measurements for arbitrary \(k\)}

Using random shift to reduce projection size. We remark that in the analysis of the previous recovery algorithm, if we can make sure that every \(y_{f}+z_{f}\) has only one possible outcome when projecting to the grid \(\mathcal{G}_{0.6 \nu}\), then we no longer need to union bound \(4^{O(k)}\) events. However, if \(\widehat{x}_{f}\) is very close to a grid point in \(\mathcal{G}_{0.6 \nu}\) (or \(\widehat{x}_{f} \in \mathcal{G}_{0.6 \nu}\) ), then no matter how close \(y_{f}+z_{f}\) and \(\widehat{x}_{f}\) are, \(\Pi_{0.6 \nu}\left(y_{f}+z_{f}\right)\) will have 4 possible values.

To address this, we introduce random shift, whose property is captured by Lemma 13.1.3. To simplify notation, for any \(r_{b}>0\) and \(c \in \mathbb{C}\) we define box \(\mathcal{B}_{\infty}\left(c, r_{b}\right):=\left\{c+r_{b}(x+y \mathbf{i})\right.\) : \(x, y \in[-1,1]\}\). For any \(S \subseteq \mathbb{C}\), let \(\Pi_{r_{g}}(S)=\left\{\Pi_{r_{g}}(c): c \in S\right\}\).

Lemma 13.1.3 (property of a randomly shifted box, informal). If we take a box of side length \(2 r_{b}\) and shift it randomly by an offset in \(\mathcal{B}_{\infty}\left(0, r_{s}\right)\) (or equivalently, \(\left[-r_{s}, r_{s}\right] \times\left[-r_{s}, r_{s}\right]\) ) where \(r_{s} \geq r_{b}\), and next we round every point inside that shifted box to the closest point in \(G_{r_{g}}\) where \(r_{g} \geq 2 r_{s}\), then with probability at least \(\left(1-r_{b} / r_{s}\right)^{2}\) everyone will be rounded to the same point.

In the following, we present a sparse Fourier recovery algorithm that incorporates the random shift idea. The algorithm takes \(O(k \log k \log n)\) measurements. We set \(H=O(\log k)\) and take measurements of \(\mathcal{T}=\left\{\mathcal{T}^{(1)}, \ldots, \mathcal{T}^{(H)}\right\}\), where \(\mathcal{T}^{(h)}\) is a list of \(O(k \log n)\) i.i.d elements in \([p]^{d}\). Note that \(\mathcal{T}^{(1)}, \ldots, \mathcal{T}^{(H)}\) are independent, and the choice of \(H\) is different from Section 13.1.5.1.

In a nutshell, our approach finely discretizes the space of possible trajectories the algorithm could evolve; in contrast to the case of \(k=O(\log n)\), the number of trajectories becomes much larger. For that, we perform random shifting after the samples are taken, such that the number of trajectories is pruned, and we need to argue for a much smaller collection of events. We note that we make the decoding algorithm be randomized: the randomness in previous algorithms was present only when taking samples, and the rest of the algorithm was deterministic. However, here we need randomness in both cases, and that helps us prune the number of possible trajectories. To the best of our knowledge, this is a novel argument and approach, and might be helpful for future progress in the field.

Recovery algorithm. Similar to the \(k=O(\log n)\) case (Section 13.1.5.1), we assume that we have already obtained a \(O(k)\)-sparse \(y \in \mathbb{C}^{[p]^{d}}\) such that \(\|\widehat{x}-y\|_{\infty} \leq 2 \nu\) and \(y\) is "almost" independent from \(\mathcal{T}\). We show how to obtain \(y^{\prime} \in \mathbb{C}^{[p]^{d}}\) such that \(\left\|\widehat{x}-y^{\prime}\right\|_{\infty} \leq \nu\) with probability \(1-1 / \operatorname{poly}(n)\) and \(y^{\prime}\) is "almost" independent from \(\mathcal{T}\). The main idea is we first run LinfinityReduce procedure \(H=O(\log k)\) times to get an \(O(k)\)-sparse \(z \in \mathbb{C}^{[p]^{d}}\) such that \(\|\widehat{x}-y-z\|_{\infty} \leq \frac{1}{2^{20} k} \nu\). Then we repeatedly sample a uniform random shift \(s \in \mathbb{C}\) (where \(\|s\|_{\infty} \leq 10^{-3} \nu\); here we consider complex numbers as 2D vectors) until for every \(f \in \operatorname{supp}(y+z)\), all the points (or complex numbers) of the form \(y_{f}+z_{f}+s+a+b \mathbf{i}\) where
```

Algorithm 13.2 Fourier Sparse Recovery by Random Shift and Projection (Informal Ver-
sion)
procedure FourierSparseRECOVERY $\left(x, n, k, \mu, R^{*}\right) \quad \triangleright$ Theorem 13.1.1, $n=p^{d}$
Require that $\mu=\frac{1}{\sqrt{k}}\left\|\widehat{x}_{-k}\right\|_{2}$ and $R^{*} \geq\|\widehat{x}\|_{\infty} / \mu$
$H \leftarrow O(\log k), \nu \leftarrow \mu R^{*} / 2, y \leftarrow \overrightarrow{0} \quad \triangleright y \in \mathbb{C}^{[p]^{d}}$ refers to the sparse vector recovered
so far
Let $\mathcal{T}=\left\{\mathcal{T}^{(1)}, \cdots, \mathcal{T}^{(H)}\right\}$ where each $\mathfrak{T}^{(h)}$ is a list of i.i.d. uniform samples in $[p]^{d}$
while true do
$\nu^{\prime} \leftarrow \frac{1}{2^{20} k} \nu$
Use $\left\{x_{t}\right\}_{t \in \mathcal{T}}$ to run the LinfinityReduce procedure (in Lemma 13.1.2) $H$ times
(use samples in $\mathcal{T}^{h}$ for each $h \in[H]$ ), and finally it finds $z$ so that $\|\widehat{x}-y-z\|_{\infty} \leq \nu^{\prime}$
if $\nu^{\prime} \leq \mu$ then return $y+z \quad \triangleright$ We found the solution
repeat
Pick $s \in \mathcal{B}_{\infty}\left(0,10^{-3} \nu\right)$ uniformly at random
until $\forall f \in \operatorname{supp}(y+z),\left|\Pi_{0.04 \nu}\left(\mathcal{B}_{\infty}\left(y_{f}+z_{f}+s, \nu^{\prime}\right)\right)\right|=1$
$y^{\prime} \leftarrow \overrightarrow{0}$
for $f \in \operatorname{supp}(y+z)$ do
$y_{f}^{\prime} \leftarrow \Pi_{0.04 \nu}\left(y_{f}+z_{f}+s\right) \quad \triangleright$ We want $\left\|\widehat{x}-y^{\prime}\right\|_{\infty} \leq \nu$ and the depend-
end for $\quad \triangleright$ ence between $y^{\prime}$ and $\mathcal{T}$ is under control
$y \leftarrow y^{\prime}, \nu \leftarrow \nu / 2$
end while
end procedure

```
\(a, b \in\left[-\frac{\nu}{2^{20} k}, \frac{\nu}{2^{20 k}}\right]\) round to the same grid point in \(\mathcal{G}_{0.04 \nu}\). Finally, for every \(f \in \operatorname{supp}(y+z)\), we assign \(\Pi_{0.04 \nu}\left(y_{f}+z_{f}+s\right)\) to \(y_{f}^{\prime}\); all remaining coordinates in \(y^{\prime}\) will be assigned to 0 . We present an informal version of our algorithm in Algorithm 13.2, and defer its formal version to the appendix.

Analysis. Now we analyze the above approach. First, we have the guarantee that \(\| \widehat{x}-\) \(y^{\prime} \|_{\infty} \leq \nu\). Moreover, note that by our choice of \(s\), for every \(f \in \operatorname{supp}(y+z), y_{f}+z_{f}+s\) and \(\widehat{x}_{f}+s\) round to the same grid point in \(\mathcal{G}_{0.04 \nu}\). Therefore, for the new vector \(y^{\prime}\) we
have recovered, we "hide" the randomness in \(\mathcal{T}\), and the randomness only leaks from failed attempts of the shifts. In the following, we show that each attempt of shift succeeds with probability \(\frac{1}{2}\).

We can restate the procedure of choosing \(s\) to be:
repeatedly sample \(s \sim \mathcal{B}_{\infty}\left(0,10^{-3} \nu\right)\), until for all \(f \in \operatorname{supp}(y+z),\left|\Pi_{0.04 \nu}\left(\mathcal{B}_{\infty}\left(y_{f}+z_{f}+s, \frac{\nu}{2^{20} k}\right)\right)\right|=1\).

Note that \(|\operatorname{supp}(y+z)|=O(k)\). Let us say that we can always guarantee that \(|\operatorname{supp}(y+z)| \leq\) 50k. By Lemma 13.1.3 where we let \(r_{b}=\frac{\nu}{2^{20} k}, r_{s}=10^{-3} \nu\) and \(r_{g}=0.04 \nu\), for \(f \in \operatorname{supp}(y+z)\),
\[
\operatorname{Pr}\left[\left|\Pi_{0.04 \nu}\left(\mathcal{B}_{\infty}\left(y_{f}+z_{f}+s, \frac{\nu}{2^{20} k}\right)\right)\right|=1\right] \geq\left(1-\frac{r_{b}}{r_{s}}\right)^{2} \geq 1-\frac{1}{100 k} .
\]

By a union bound over \(f \in \operatorname{supp}(y+z)\), the probability is at least \(\frac{1}{2}\) that for all \(f \in\) \(\operatorname{supp}(y+z),\left|\Pi_{0.04 \nu}\left(\mathcal{B}_{\infty}\left(y_{f}+z_{f}+s, \frac{\nu}{2^{20 k}}\right)\right)\right|=1\).

Therefore, with probability \(1-1 / \operatorname{poly}(n)\), we will only try \(O(\log n)\) shifts. We can apply a union bound over \(O(\log n)\) possible shifts, and prove that with probability \(1-1 / \operatorname{poly}(n)\) if taking \(y^{\prime}\) as our new \(y\), and shrinking \(\nu\) by half, the LinfinityReduce procedure will work as desired as if there is no dependence issue.

Sufficient event. Let \(S\) be top \(O(k)\) coordinates in \(\widehat{x}\). Let \(L=O\left(\log R^{*}\right)\) denote the number of iterations in Algorithm 13.2. For \(\ell \in[L]\), let \(\nu_{\ell}=2^{-\ell} \mu R^{*}\). For \(\ell \in[L-1]\), let \(s_{\ell}^{(a)}\) be the \(a\)-th uniform randomly sampled from \(\mathcal{B}_{\infty}\left(0,10^{-3} \nu_{\ell}\right)\) as appeared on Line 10 in Algorithm 13.2. For the sake of analysis, we assume that Algorithm 13.2 actually produces an infinite sequence of shifts \(s_{\ell}^{(1)}, s_{\ell}^{(2)}, \ldots\). We formulate the event that guarantees successful
execution of Algorithm 13.2. We define event \(\mathcal{E}_{2}\) to be all of the following events hold.
1. For all \(\ell \in[L-1]\), there exists \(a \in[10 \log n]\) so that for all \(f \in S\),
\[
\left|\Pi_{0.04 \nu_{\ell}}\left(\mathcal{B}_{\infty}\left(\widehat{x}_{f}+s_{\ell}^{(a)}, \frac{1}{100 k} \nu_{\ell}\right)\right)\right|=1 .
\]
2. For \(\ell=1\), if we run the LinfinityReduce procedure \(H\) times with \(y=\overrightarrow{0}\) and measurements in \(\mathcal{T}\), we get \(z\) such that \(\|\widehat{x}-z\|_{\infty} \leq 2^{1-H} \nu_{1}\) and \(\operatorname{supp}(z) \subseteq S\).
3. For all \(\ell \in\{2, \ldots, L\}\), for all \(a \in[10 \log n]\), if we run the LinfinityReduce procedure \(H\) times with \(y=\xi\) where
\[
\xi_{f}= \begin{cases}\Pi_{0.04 \nu_{\ell}}\left(\widehat{x}_{f}+s_{\ell-1}^{(a)}\right), & \text { if } f \in S ; \\ 0, & \text { if } f \in \bar{S}\end{cases}
\]
then we get \(z\) such that \(\|\widehat{x}-y-z\|_{\infty} \leq 2^{1-H} \nu_{\ell}\) and \(\operatorname{supp}(y+z) \subseteq S\).
We can prove that event \(\mathcal{E}_{2}\) happens with probability \(1-1 / \operatorname{poly}(n)\). Moreover, we can prove that conditioned on event \(\mathcal{E}_{2}\) Algorithm 13.2 gives correct output. We defer both proofs in the appendix.

\subsection*{13.2 Algorithm for \(d\)-dimensional Sparse Fourier Transform}

In this section, we will give a Fourier sparse recovery algorithm that takes \(O(k \log k \log n)\) measurements with " \(\ell_{\infty} / \ell_{2}\) " guarantee. We assume the knowledge of \(\mu=\frac{1}{\sqrt{k}}\left\|\widehat{x}_{-k}\right\|_{2}\). In fact, a constant factor approximation suffices, but we prefer to assume exact knowledge of it in order to simplify exposition. All of the arguments go through in the other case, with minor changes in constants. We also assume we know \(R^{*}\) so that \(R^{*} \geq\|\widehat{x}\|_{\infty} / \mu\). We assume that \(\log R^{*}=O(\log n)\). For larger \(\log R^{*}=O(\operatorname{poly}(n))\), our algorithm will still work, but the decoding time will be worse by a factor of \(\frac{\log R^{*}}{\log n}\). Note that our assumptions on \(\mu\) and \(R^{*}\) are standard. For example, [IK14] make the same assumption. We assume that we can measure the signal \(x\) in the time domain, and we want to recover the signal \(\widehat{x}\) in the frequency domain.

In our algorithm, we will use \(\mu\) as a threshold for noise, and we will perform \(\log R^{*}\) iterations, where in each iteration the upper bound of \(\ell_{\infty}\) norm of the residual signal (in the frequency domain) shrinks by half. In Section 13.2.1, we give some definitions that will be used in the algorithm. Then we present our new algorithm for \(d\)-dimension Fourier sparse recovery in Section 13.2.2. In Section 13.2.3, we prove the correctness of the proposed algorithm.

\subsection*{13.2.1 Notations}

For a subset of samples (or measurements) \(\left\{x_{t}\right\}_{t \in T}\) from the time domain, where \(T\) is a list of elements in \([p]^{d}\), we define \(\widehat{x}^{[T]}\) in Definition 13.2.1 as our estimation to \(\widehat{x}\).

Definition 13.2.1 (Fourier transform of a subset of samples). Let \(x \in \mathbb{C}^{[p]^{d}}\). For any \(T\)
which is a list of elements in \([p]^{d}\), for any \(f \in[p]^{d}\), we define
\[
\widehat{x}_{f}^{[T]}=\frac{\sqrt{n}}{|T|} \sum_{t \in T} \omega^{f^{\top} t} x_{t} .
\]

In order to reuse samples across different iterations where we drive down the upper bound of the residual signal by half, in each iteration after we obtain estimations to heavy hitters (or equivalently large coordinates), instead of subtracting the estimates directly, we need to "hide" the randomness leaked by the samples. We interpret each estimate (which is a complex number) as a point on a 2-dimension plane, and hide the randomness by rounding the estimate to the nearest grid point (where the side length of the grid is chosen to be a small constant fraction of the target \(\ell_{\infty}\) norm of the residual signal in the frequency domain), which we call "projection onto grid". In Definition 13.2.2, we formally define box and grid, and in Definition 13.2.3 we define projection to grid. We illustrate these two definitions in Figure 13.1.

Definition 13.2.2 (box and grid). For any \(c \in \mathbb{C}\) and \(r \geq 0\), we define box \(\mathcal{B}_{\infty}(c, r) \subseteq \mathbb{C}\) as
\[
\mathcal{B}_{\infty}(c, r)=\{c+x+y \mathbf{i}: x, y \in[-r, r]\} .
\]

Namely, if we consider complex numbers as points on 2D plane, box \(\mathcal{B}_{\infty}(c, r)\) refers to \(\ell_{\infty}\) ball with radius \(r\) centered at \(c\).

For any \(r>0\), we define grid \(\mathcal{G}_{r} \subseteq \mathbb{C}\) as
\[
\mathcal{G}_{r}=\{x r+y r \mathbf{i}: x, y \in \mathbb{Z}\} .
\]


Figure 13.1: Illustration of box \(\mathcal{B}_{\infty}(c, r)\) and grid \(\mathcal{G}_{r_{g}}\). Box \(\mathcal{B}_{\infty}(c, r)\) refers to all the points in the square centered at \(c\) with side length \(2 r\). Grid \(\mathcal{G}_{r_{g}}\) refers to all the solid round points, and the distance between origin \(O\) and \(A_{0}\) is \(r_{g}\). Note that the dashed lines are decision boundaries of the projection \(\Pi_{r_{g}}\), and all the points inside a minimum cell separated by the dashed lines are mapped (by \(\Pi_{r_{g}}\) ) to the same grid point in \(\mathcal{G}_{r_{g}}\) (which is the center of the cell). We have \(\Pi_{r_{g}}(c)=A_{1}\) and \(\Pi_{r_{g}}\left(\mathcal{B}_{\infty}(c, r)\right)=\left\{A_{1}, A_{2}, A_{3}, A_{4}\right\}\).

Definition 13.2.3 (projection onto grid). For any \(r>0\), we define \(\Pi_{r}\) to be a maping from \(\mathbb{C}\) to \(\mathcal{G}_{r}\), so that for any \(c \in \mathbb{C}\),
\[
\Pi_{r}(c)=\arg \min _{c^{\prime} \in \mathcal{S}_{r}}\left|c-c^{\prime}\right|,
\]
where we break the tie by choosing the one with minimum \(\left|c^{\prime}\right|\). As a natural generalization, For \(C \subseteq \mathbb{C}\), we define
\[
\Pi_{r}(C)=\left\{\Pi_{r}(c): c \in C\right\} .
\]

\subsection*{13.2.2 Algorithm}

We present our new sparse Fourier recovery algorithm in Algorithm 13.3. Its auxiliary function LinfinityReduce is in Algorithm 13.4. Important constants are summarized in Table 13.2.

In Algorithm 13.3, we define "bucket size" \(B=O(k)\) and number of repetitions \(R=O(\log n)\). For each \(r \in[R]\), we choose \(\mathcal{T}_{r}\) to be a list of \(B\) independent and uniformly random elements in \([p]^{d}\). We will measure \(x_{t}\) for all \(t \in \cup_{r \in[R]} \mathcal{T}_{r}\), and use LinfinityREDUCE in Algorithm 13.4 to locate and estimate all the "heavy hitters" of the residual signal so that if we substract them then the \(\ell_{\infty}\) norm of the new residual signal shrinks by half. The input to LinfinityReduce is a signal \(x \in \mathbb{C}^{[p]^{d}}\) in the time domain (but we can only get access to \(x_{t}\) where \(t \in \cup_{r \in[R]} \mathcal{T}_{r}\) ), a sparse vector \(y \in \mathbb{C}^{[p]^{d}}\) in the frequency domain that we have recovered so far, and \(\nu \geq \mu\) such that \(\|\widehat{x}-y\|_{\infty} \leq 2 \nu\) where we will refer \(\widehat{x}-y\) as the currect residual signal (in the frequency domain). It is guaranteed that \(\operatorname{LinfinityReduce}\left(x, n, y,\left\{\mathcal{T}_{r}\right\}_{r=1}^{R}, \nu\right)\) returns a \(O(k)\)-sparse \(z\) so that \(\|\widehat{x}-y-z\| \leq \nu\) with probability \(1-1 / \operatorname{poly}(n)\).

Algorithm 13.3 in total maintains \(H=O(\log k)\) independent copies of such errorreduce data structures, where in the \(h\)-th copy it measures \(\mathcal{T}^{(h)}=\left\{\mathcal{T}_{r}^{(h)}\right\}_{r \in[R]}\) for \(h \in[H]\). We denote \(\mathcal{T}=\left\{\mathcal{T}^{(h)}\right\}_{r \in[R]}\). If \(\log R^{*} \leq H\), then we can simply use different \(\mathcal{T}^{(h)}\) in different iterations. In that case \(L=1\) and \(H=\log R^{*}\) in Algorithm 13.3. We will get \(z^{(1)}\) on Line 20 such that \(\left\|\widehat{x}-y^{(0)}-z^{(1)}\right\|_{\infty} \leq \mu\) (we will prove in the analysis this holds with probability \(1-1 / \operatorname{poly}(n))\) where \(y^{(0)}=0\), and return \(z^{(1)}+y^{(0)}\) on Line 22 .

If \(\log R^{*}>H\), we have to reuse the samples. We proceed in \(L\) iterations (in the loop between Line 14 and Line 33 in Algorithm 13.3), where \(L=\log R^{*}-H+1\). For \(\ell \in[L]\),


Figure 13.2: Illustration of the behavior of Line 16 to Line 20 in Algorithm 13.3. For any \(f \in\) \([p]^{d}\), we draw box \(\mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}+z_{f}, 2^{1-h} \nu_{\ell}\right)\) after \(h\) iterations of the for loop between Line 17 and Line 19 in Algorithm 13.3, where \(h \in\{0,1, \ldots, H\}\). Conditioned on LinfinityReduce is correct, for every \(h \in\{0,1, \ldots, H\}\), after \(h\)-th iteration we have \(\widehat{x}_{f} \in \mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}+z_{f}, 2^{1-h} \nu_{\ell}\right)\). When \(h=0\), i.e. before the loop between Line 17 and Line 19 starts, we know that \(\widehat{x}_{f} \in \mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}, 2 \nu_{\ell}\right)\) as depicted by (a). After each iteration in \(h\), the radius of the box shrinks by half (and its center might change). Finally after \(H\) iterations, as depicted by (c), we obtain \(z^{(\ell-1)}\) such that \(\widehat{x}_{f} \in \mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}, 2^{1-H} \nu_{\ell}\right)\).
as defined in Line \(15, \nu_{\ell}=2^{-\ell} \mu R^{*}\) refers to the target \(\ell_{\infty}\) of the residual signal in the \(\ell\)-th iteration (namely, for \(\ell \in[L-1]\) we want to obtain \(y^{(\ell)}\) so that \(\left.\left\|\widehat{x}-y^{(\ell)}\right\|_{\infty} \leq \nu_{\ell}\right)\). In the \(\ell\)-th iteration where \(\ell \in[L]\), by using the samples in \(\mathcal{T}=\left\{\mathcal{T}^{(h)}\right\}_{h \in H}\) (Line 16 to Line 20), the algorithm tries to get \(z^{(\ell)}\) so that \(\left\|\widehat{x}-y^{(\ell-1)}-z^{(\ell)}\right\|_{\infty} \leq 2^{1-H} \nu_{\ell}\). The intuition on the behavior of Line 16 to Line 20 is depicted in Figure 13.2.

If \(\ell=L\) the algorithm will return \(y^{(L-1)}+z^{(L)}\) as in Line 22 ; otherwise, the algorithm will try to compute \(y^{(\ell)}\) based on \(y^{(\ell-1)}+z^{(\ell)}\). In Line 25 to Line 28 , the algorithm repeatedly samples a uniform random shift \(s_{\ell} \in \mathcal{B}_{\infty}\left(0, \alpha \nu_{\ell}\right)\) (where \(\alpha \in(0,1)\) is a small constant chosen in Table 13.2) until the shift is good, where shift \(s_{\ell}\) is good if and only if for each \(f \in \operatorname{supp}\left(y^{(\ell-1)}+z^{(\ell)}\right)\), all the points in \(\mathcal{B}_{\infty}\left(y^{(\ell-1)}+z^{(\ell)}+s_{\ell}, 2^{1-H} \nu_{\ell}\right)\) (i.e. the box obtained by


Figure 13.3: Illustration of the iteration between Line 25 and Line 28 in Algorithm 13.3. The round solid points represent grid points in \(\mathcal{G}_{\beta \nu}\), and the dashed lines represent decision boundaries of \(\Pi_{\beta \nu_{\ell}}\). In this example we have \(\left|\operatorname{supp}\left(y^{(\ell-1)}+z^{(\ell)}\right)\right|=3\), and the dotted squares represent boxes \(\mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}, 2^{1-H} \nu_{\ell}\right)\) for \(f \in \operatorname{supp}\left(y^{(\ell-1)}+z^{(\ell)}\right)\). The algorithm repeatedly samples a random shift \(s \sim \mathcal{B}_{\infty}\left(0, \alpha \nu_{\ell}\right)\), until all the shifted boxes \(\left\{\mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}, 2^{1-H} \nu_{\ell}\right)+\right.\) \(s\}_{f \in \operatorname{supp}\left(y^{(\ell-1)}+z^{(\ell)}\right)}\) do not intersect with the dashed lines (i.e. decision boundaries of \(\Pi_{\beta \nu_{\ell}}\) ). In the figure, we color a shifted box in green if it does not intersect with dashed lines, and color in red otherwise. After a series of failed attempts from (a) to (b), we finally have a successful attempt in (c).
applying shift \(s_{\ell}\) to the box \(\left.\mathcal{B}_{\infty}\left(y^{(\ell-1)}+z^{(\ell)}, 2^{1-H} \nu_{\ell}\right)\right)\) project to the same grid point in \(\mathcal{G}_{\beta \nu_{\ell}}\). We depict the process of obtaining the shift \(s_{\ell}\) in Figure 13.3. It is crucial to note that if the shift \(s_{\ell}\) is good and the vector \(z^{(\ell)}\) we get is desirable (namely \(\left\|\widehat{x}-y^{(\ell-1)}-z^{(\ell)}\right\|_{\infty} \leq 2^{1-H} \nu_{\ell}\) ), then for each \(f \in \operatorname{supp}\left(y^{(\ell-1)}+z^{(\ell)}\right), \Pi_{\beta \nu_{\ell}}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}+s_{\ell}\right)=\Pi_{\beta \nu_{\ell}}\left(\widehat{x}_{f}+s_{\ell}\right)\).

On Line 31, we assign \(\Pi_{\beta \nu_{\ell}}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}+s_{\ell}\right)\) to \(y_{f}^{(\ell)}\). Because \(\beta\) is a small constant, we still have the guarantee that \(\left\|\widehat{x}-y^{(\ell)}\right\|_{\infty} \leq \nu_{\ell}\). Moreover, by assigning \(\Pi_{\beta \nu_{\ell}}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}+s_{\ell}\right)=\) \(\Pi_{\beta \nu_{\ell}}\left(\widehat{x}_{f}+s_{\ell}\right)\) to \(y_{f}^{(\ell)}\), we "hide" the randomness in \(\mathcal{T}=\left\{\mathcal{T}^{(h)}\right\}_{h \in[H]}\). Now the randomness in \(\mathfrak{T}\) only leaks from failed attempts of the shifts. For analysis purpose, we maintain a counter \(a_{\ell}\) for \(\ell \in[L-1]\) recording the number of attempts until we have sampled a good
one. By our choice of parameters, we can prove that with high probability \(a_{\ell} \leq 10 \log n\) for each \(\ell \in[L-1]\). Thus intuitively the leaked randomness is under control, and we can formally apply a union bound to prove that with good probability all possible invocations of LinfinityReduce by our FourierSparseRecovery produce desirable output.

\subsection*{13.2.3 Analysis}

In order to analyze the algorithm, let \(S \subseteq[n]\) be top \(C_{S} k\) coordinates of \(\widehat{x}\) where \(C_{S}=26\), and let \(\bar{S}=[n] \backslash S\). In order to analyze the performance of LinfinityREDUCE in Algorithm 13.4, we need the following definition.

Definition 13.2.4 (uniform sample). We say \(t\) is sampled from \([p]^{d}\) uniformly at random if for each \(i \in[d]\), we independently sample \(t_{i}\) from \([p]\) uniformly at random. We use \(t \sim[p]^{d}\) to denote it.

Fact 13.2.1. Let \(\omega=e^{2 \pi \mathbf{i} / p}\) where \(p\) is any positive integer. For a fixed \(f \in[p]^{d} \backslash\{\overrightarrow{0}\}\), \(\mathbb{E}_{t \sim[p]^{d}}\left[\omega^{f^{\top} t}\right]=0\).

Proof. Note that \(\mathbb{E}_{t \sim[p]^{d}}\left[\omega^{f^{\top} t}\right]=\prod_{i \in[d]} \mathbb{E}_{t_{i} \sim[p]}\left[\omega^{f_{i} t_{i}}\right]\) by the fact that \(t_{1}, \ldots, t_{d}\) are independent. Because \(f \neq \overrightarrow{0}\), there exists \(i \in[d]\) so that \(f_{i} \neq 0\). We have
\[
\begin{aligned}
\underset{t_{i} \sim[p]}{\mathbb{E}}\left[\omega^{f_{i} t_{i}}\right] & =\frac{1}{p} \sum_{j=0}^{p-1}\left(\omega^{f_{i}}\right)^{j} \\
& =\frac{1}{p} \cdot \frac{\left(\omega^{f_{i}}\right)^{0}\left(1-\left(\omega^{f_{i}}\right)^{p}\right)}{1-\omega^{f_{i}}} \\
& =0,
\end{aligned}
\]
where the second step follows from the sum of geometry series where \(\omega^{f_{i}} \neq 1\), adn the third step follow from \(\left(\omega^{f_{i}}\right)^{p}=e^{2 \pi \mathrm{i} f_{i}}=1\). Therefore, \(\mathbb{E}_{t \sim[p]^{d}}\left[\omega^{f^{\top} t}\right]=0\).
```

Algorithm 13.3 Fourier Sparse Recovery by Random Shift and Projection
procedure FourierSparseRecovery $\left(x, n, k, \mu, R^{*}\right) \quad \triangleright$ Theorem 13.2.9, $n=p^{d}$
Require that $\mu=\frac{1}{\sqrt{k}}\left\|\widehat{x}_{-k}\right\|_{2}$ and $R^{*} \geq\|\widehat{x}\|_{\infty} / \mu \quad \triangleright R^{*}$ is a power of 2
$B \leftarrow C_{B} \cdot k \quad \triangleright C_{B}$ is a constant defined in Table 13.2
$R \leftarrow C_{R} \cdot \log n \quad \triangleright C_{R}$ is a constant defined in Table 13.2
$H \leftarrow \min \left\{\log k+C_{H}, \log R^{*}\right\} \quad \triangleright C_{H}$ is a constant defined in Table 13.2
for $h=1 \rightarrow H$ do
for $r=1 \rightarrow R$ do
$\mathcal{T}_{r}^{(h)} \leftarrow$ a list of $B$ i.i.d elements in $[p]^{d}$
end for
$\mathcal{T}^{(h)} \leftarrow\left\{\mathfrak{T}_{r}^{(h)}\right\}_{r=1}^{R}$
end for
$\triangleright$ We will measure $x_{t}$ for $t \in \cup_{h \in[H], r \in[R]} \mathfrak{T}_{r}^{(h)}$
$y^{(0)} \leftarrow \overrightarrow{0} \quad \triangleright y^{(0)} \in \mathbb{C}^{n}$
$L \leftarrow \log R^{*}-H+1$
for $\ell=1 \rightarrow L$ do
$\nu_{\ell} \leftarrow 2^{-\ell} \mu R^{*} \triangleright$ Target $\ell_{\infty}$ of the residual signal in iteration $t$
$z \leftarrow \overrightarrow{0} \quad \triangleright z$ is a temporary variable used to compute $z^{(\ell)}$
for $h=1 \rightarrow H$ do
$z \leftarrow z+\operatorname{LiNFINITYREDUCE}\left(x, n, y^{(\ell-1)}+z, \mathfrak{T}^{(h)}, 2^{1-h} \nu_{\ell}\right)$
end for
$z^{(\ell)} \leftarrow z \quad \triangleright$ We want $\left\|\widehat{x}-y^{(\ell-1)}-z^{(\ell)}\right\|_{\infty} \leq 2^{1-H} \nu_{\ell}$
if $\ell=L$ then
return $y^{(L-1)}+z^{(L)}$
end if
$a \leftarrow 0 \quad \triangleright$ A temporary counter maintained for analysis purpose only
repeat
Pick $s_{\ell} \in \mathcal{B}_{\infty}\left(0, \alpha \nu_{\ell}\right)$ uniformly at random $\quad \triangleright \alpha \in(0,1)$ is a small constant
$a \leftarrow a+1 \quad \triangleright \beta$ in the next line is a small constant where $\alpha<\beta<0.1$
until $\forall f \in \operatorname{supp}\left(y^{(\ell-1)}+z^{(\ell)}\right),\left|\Pi_{\beta \nu_{\ell}}\left(\mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}+s_{\ell}, 2^{1-H} \nu_{\ell}\right)\right)\right|=1$
$a_{\ell} \leftarrow a$
for $f \in \operatorname{supp}\left(y^{(\ell-1)}+z^{(\ell)}\right)$ do
$y_{f}^{(\ell)} \leftarrow \Pi_{\beta \nu_{\ell}}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}+s_{\ell}\right) \quad \triangleright$ We want $\left\|\widehat{x}-y^{(\ell)}\right\|_{\infty} \leq \nu_{\ell}$
end for
end for
end procedure

```

We define measurement coefficient as follows:

Definition 13.2.5 (measurement coefficient). For any \(f \in[p]^{d}\) and any \(T\) which is a list of elements in \([p]^{d}\), we define
\[
c_{f}^{[T]}=\frac{1}{|T|} \sum_{t \in T} \omega^{f^{\top} t}
\]

By definition of \(c_{f}^{[T]}\) and \(d\)-dimension Fourier transform, we can decompose \(\widehat{x}_{f}^{[T]}\) as follows.

Lemma 13.2.2 (measurement decomposition). For any \(f \in[p]^{d}\) and any \(T\) which is a list of elements in \([p]^{d}\),
\[
\widehat{x}_{f}^{[T]}=\sum_{f^{\prime} \in[p]^{d}} c_{f-f^{\prime}}^{[T]} \widehat{x}_{f^{\prime}} .
\]

Proof. We have
\[
\begin{aligned}
\widehat{x}_{f}^{[T]} & =\frac{\sqrt{n}}{|T|} \sum_{t \in T} \omega^{f^{\top} t} x_{t} \\
& =\frac{\sqrt{n}}{|T|} \sum_{t \in T} \omega^{f^{\top} t} \frac{1}{\sqrt{n}} \sum_{f^{\prime} \in[p]^{d}} \omega^{-f^{\prime \top} t} \widehat{x}_{f^{\prime}} \\
& =\frac{\sqrt{n}}{|T|} \sum_{t \in T} \frac{1}{\sqrt{n}} \sum_{f^{\prime} \in[p]^{d}} \omega^{\left(f-f^{\prime}\right)^{\top} t} \widehat{x}_{f^{\prime}} \\
& =\sum_{f^{\prime} \in[p]^{d}}\left(\frac{1}{|T|} \sum_{t \in T} \omega^{\left(f-f^{\prime}\right)^{\top} t}\right) \widehat{x}_{f^{\prime}} \\
& =\sum_{f^{\prime} \in[p]^{d}} c_{f-f^{\prime}}^{[T]} \widehat{x}_{f^{\prime}},
\end{aligned}
\]
\begin{tabular}{|l|l|l|l|}
\hline Notation & Choice & Statement & Parameter \\
\hline\(C_{B}\) & \(10^{6}\) & Lemma 13.2.8 & \(B\) \\
\hline\(C_{R}\) & \(10^{3}\) & Lemma 13.2.7 & \(R\) \\
\hline\(C_{H}\) & 20 & Algorithm 13.3 & \(H\) \\
\hline\(\alpha\) & \(10^{-3}\) & Algorithm 13.3 Line 26 & shift range \\
\hline\(\beta\) & 0.04 & Algorithm 13.3 Line 28 & grid size \\
\hline\(C_{S}\) & 26 & Lemma 13.2.7, Lemma 13.2.8 & \(|S|\) \\
\hline
\end{tabular}

Table 13.2: Summary of important constants.
\begin{tabular}{|l|l|}
\hline Lemma & Meaning \\
\hline Lemma 13.2.2 & measurement decomposition \\
\hline Lemma 13.2.3 & properties of coefficient \\
\hline Lemma 13.2.4 & noise bound \\
\hline Lemma 13.2.5 & guarantee of LINFINITYREDUCE \\
\hline Lemma 13.2.6 & property of a randomly shifted box \\
\hline Lemma 13.2.7 & event E happens \\
\hline Lemma 13.2.8 & correctness of our algorithm \\
\hline
\end{tabular}

Table 13.3: Summary of Lemmas.
where the first step follow by the definition of \(\widehat{x}_{f}^{[T]}\) in Definition 13.2.1, second step follows by the definition of inverse \(d\)-dimensional Fourier transform (see Section 13.1.1), third and forth step follow by rearranging terms, last step follows by the definition of measurement coefficients \(c\) in Definition 13.2.5.

Let \(T\) be a list of i.i.d. samples from \([p]^{d}\), then the coeffcients \(c_{f}^{[T]}\) defined in Definition 13.2.5 have the following property.

Lemma 13.2.3 (properties of coeffcient \(c\) ). Let \(T\) be a list of \(B\) independent and uniform random elements in \([p]^{d}\). Then we have
1. \(c_{0}^{[T]}=1\).
2. For any \(f \in[p]^{d} \backslash\{0\}, \mathbb{E}_{T}\left[\left|c_{f}^{[T]}\right|^{2}\right]=\frac{1}{B}\).
3. For any \(f, f^{\prime} \in[p]^{d}, f \neq f^{\prime}, \mathbb{E}_{T}\left[c_{f}^{[T]} \cdot c_{f^{\prime}}^{[T]}\right]=0\).

Proof. Part 1. By definition of \(c_{0}^{[T]}\),
\[
c_{0}^{[T]}=\frac{1}{|T|} \sum_{t \in T} \omega^{0 \cdot t}=1
\]

Part 2. Let \(T=\left\{t_{1}, \ldots, t_{B}\right\}\), where \(t_{i}\) is independently and uniformly chosen from \([p]^{d}\). For any \(f \in[p]^{d} \backslash\{0\}\),
\[
\begin{aligned}
\underset{T}{\mathbb{E}}\left[\left|c_{f}^{[T]}\right|^{2}\right] & =\underset{T}{\mathbb{E}}\left[c_{f}^{[T]} \cdot \overline{c_{f}^{[T]}}\right] \\
& =\frac{1}{|T|^{2}} \underset{T}{\mathbb{E}}\left[\sum_{i, j \in[B]} \omega^{f^{\top}\left(t_{i}-t_{j}\right)}\right] \\
& =\frac{1}{|T|^{2}}\left(|T|+\underset{T}{\mathbb{E}}\left[\sum_{i, j \in[B], i \neq j} \omega^{f^{\top}\left(t_{i}-t_{j}\right)}\right]\right) \\
& =\frac{1}{|T|}+\frac{1}{|T|^{2}} \sum_{i, j \in[B], i \neq j} \underset{T}{\mathbb{E}}\left[\omega^{f^{\top}\left(t_{i}-t_{j}\right)}\right] \\
& =\frac{1}{|T|}-\frac{1}{|T|^{2}} \cdot 0 \\
& =\frac{1}{|T|}=\frac{1}{B},
\end{aligned}
\]
where the forth step follows by \(\mathbb{E}_{T}\left[\omega^{f^{\top}\left(t_{i}-t_{j}\right)}\right]=\mathbb{E}_{t \sim[p]}\left[\omega^{f^{\top} t}\right]=0\), in which \(\mathbb{E}_{T}\left[\omega^{f^{\top}\left(t_{i}-t_{j}\right)}\right]=\) \(\mathbb{E}_{t \sim[p]^{d}}\left[\omega^{f^{\top} t}\right]\) because \(i \neq j, t_{i}, t_{j}\) are independent and uniformly random distributed in \([p]^{d}\), \(t_{i}-t_{j} \sim[p]^{d} ; \mathbb{E}_{t \sim[p]^{d}}\left[\omega^{f^{\top} t}\right]=0\) follows by by Fact 13.2 .1 and \(f\) is not a zero vector.

Part 3. For any \(f, f^{\prime} \in[p]^{d}, f \neq f^{\prime}\),
\[
\begin{aligned}
\underset{T}{\mathbb{E}}\left[c_{f}^{[T]} \cdot \overline{c_{f^{\prime}}^{[T]}}\right] & =\frac{1}{|T|^{2}} \underset{T}{\mathbb{E}}\left[\sum_{i, j \in[B]} \omega^{f^{\top} t_{i}-f^{\prime} t^{\top}}\right] \\
& =\frac{1}{|T|^{2}}\left(\sum_{i, j \in[B], i \neq j} \underset{T}{\mathbb{E}}\left[\omega^{f^{\top} t_{i}-f^{\prime \top} t_{j}}\right]+\sum_{i \in[B]}^{\underset{T}{\mathbb{E}}}\left[\omega^{\left(f-f^{\prime}\right)^{\top} t_{i}}\right]\right) \\
& =\frac{1}{|T|^{2}}\left(\sum_{i, j \in[B], i \neq j} \underset{\left.t_{i} \sim \sim p\right]^{d}}{\mathbb{E}}\left[\omega^{f^{\top} t_{i}}\right] \underset{t_{j} \sim[p]^{d}}{\mathbb{E}}\left[\omega^{-f^{\prime \top} t_{j}}\right]+\sum_{i \in[B]} \underset{t_{i} \sim[p]^{d}}{\mathbb{E}}\left[\omega^{\left(f-f^{\prime}\right)^{\top} t_{i}}\right]\right) \\
& =0,
\end{aligned}
\]
where the second step follows from separating diagonal term and off-diagonal terms, the third step follows from \(t_{i}\) and \(t_{j}\) are independent, the last step follows from Fact 13.2.1 where \(f-f^{\prime} \neq \overrightarrow{0}\), and at least one of \(f\) and \(f^{\prime}\) is not \(\overrightarrow{0}\).

Let \(T\) be a list of independent and uniformly random elements from \([p]^{d}\). We are going to measure \(x_{t}\) for \(t \in T\), and take \(\widehat{x}_{f}^{[T]}\) (recall its definition in Definition 13.2.1) as estimate to \(\widehat{x}_{f}\). By Lemma 13.2.2, \(\widehat{x}_{f}^{[T]}=\sum_{f^{\prime} \in[p]^{d}} c_{f-f^{\prime}}^{[T]} \widehat{x}_{f^{\prime}}\). The following lemma bounds the contribution of coordinates from \(V\) where \(V \subseteq[p]^{d} \backslash\{f\}\), namely \(\left|\sum_{f^{\prime} \in V} c_{f-f^{\prime}}^{[T]} \widehat{x}_{f^{\prime}}\right|\). When analyzing the quality of \(\widehat{x}_{f}^{[T]}\) as an approximation to \(\widehat{x}_{f}\), we consider coordinates in \(V\) as noise, and we usually set \(V=[p]^{d} \backslash\{f\}\).

Lemma 13.2.4 (noise bound). For any \(f \in[p]^{d}\), \(T\) which is a list of \(B\) i.i.d. samples from \([p]^{d}\) and \(V \subseteq[n]\) such that \(f \notin V\),
\[
\underset{T}{\operatorname{Pr}}\left[\left|\sum_{f^{\prime} \in V} c_{f-f^{\prime}}^{[T]} \widehat{x}_{f^{\prime}}\right| \geq \frac{10}{\sqrt{B}}\left\|\widehat{x}_{V}\right\|_{2}\right] \leq \frac{1}{100} .
\]

Proof. First, we can prove that \(\mathbb{E}_{T}\left[\left|\sum_{f^{\prime} \in V} c_{f-f^{\prime}}^{[T]} \widehat{x}_{f^{\prime}}\right|^{2}\right]=\frac{1}{B}\left\|\widehat{x}_{V}\right\|_{2}^{2}\), because
\[
\begin{aligned}
\underset{T}{\mathbb{E}}\left[\left|\sum_{f^{\prime} \in V} c_{f-f^{\prime}}^{[T]} \widehat{x}_{f^{\prime}}\right|^{2}\right] & =\underset{T}{\mathbb{E}}\left[\sum_{f_{1}, f_{2} \in V}\left(c_{f-f_{1}}^{[T]} \widehat{x}_{f_{1}}\right) \overline{\left(c_{f-f_{2}}^{[T]} \widehat{x}_{f_{2}}\right)}\right] \\
& =\sum_{f_{1}, f_{2} \in V} \underset{T}{\mathbb{E}}\left[c_{f-f_{1}}^{[T]} \overline{c_{f-f_{2}}^{[T]}}\right] \widehat{x}_{f_{1}} \overline{\widehat{x}_{f_{2}}} \\
& =\sum_{f^{\prime} \in V} \underset{T}{\mathbb{E}}\left[\left|c_{f-f^{\prime}}^{[T]}\right|^{2}\right]\left|\widehat{x}_{f^{\prime}}\right|^{2} \\
& =\frac{1}{B}\left\|\widehat{x}_{V}\right\|_{2}^{2},
\end{aligned}
\]
where the third step follows from Lemma 13.2.3 that for \(f-f_{1} \neq f-f_{2}, \mathbb{E}_{T}\left[c_{f-f_{1}}^{[T]} \overline{c_{f-f_{2}}^{[T]}}\right]=0\), and the last step follows from \(\mathbb{E}_{T}\left[\left|c_{f-f^{\prime}}^{[T]}\right|^{2}\right]=1 / B\) in Lemma 13.2.3.

Then the lemma follows from Chebyshev Inequality and the fact that
\[
\underset{T}{\mathbb{V}}\left[\left|\sum_{f^{\prime} \in V} c_{f-f^{\prime}}^{[T]} \widehat{x}_{f^{\prime}}\right|\right] \leq \underset{T}{\mathbb{E}}\left[\left|\sum_{f^{\prime} \in V} c_{f-f^{\prime}}^{[T]} \widehat{x}_{f^{\prime}}\right|^{2}\right]=\frac{1}{B}\left\|\widehat{x}_{V}\right\|_{2}^{2}
\]

In the next lemma, we show the guarantee of LinfinityReduce in Algorithm 13.4.
Lemma 13.2.5 (guarantee of LinfinityReduce in Algorithm 13.4). Let \(x \in \mathbb{C}^{[p]^{d}}\), and \(n=p^{d}\). Let \(R=C_{R} \log n\), and \(B=C_{B}\). Let \(C_{B} \geq 10^{6}\) and \(C_{R} \geq 10^{3}\). Let \(\mu=\frac{1}{\sqrt{k}}\left\|\widehat{x}_{-k}\right\|_{2}\), and \(\nu \geq \mu\). For \(r \in[R]\), let \(\mathcal{T}_{r}\) be a list of \(B\) i.i.d. elements in \([p]^{d}\). Let \(z \in \mathbb{C}^{n}\) denote the output of
\[
\operatorname{LinfinityREDUCE}\left(x, n, y,\left\{\mathcal{T}_{r}\right\}_{r=1}^{R}, \nu\right)
\]
```

Algorithm 13.4 Procedure for Reducing $\ell_{\infty}$-norm of the Residual Signal
procedure LinfinityREDUCE $\left(x, n, y,\left\{\mathcal{T}_{r}\right\}_{r=1}^{R}, \nu\right)$
$\triangleright$ Lemma 13.2.5
Require that $\|\widehat{x}-y\|_{\infty} \leq 2 \nu$
Let $w$ be inverse Fourier transform of $y \quad \triangleright$ We have $\widehat{w}=y$
for $r=1 \rightarrow R$ do
for $f=1 \rightarrow n$ do $\quad \triangleright$ Implemented by FFT which takes $O(n \log n)$ time
$u_{f, r} \leftarrow \frac{\sqrt{n}}{\left|\mathcal{T}_{r}\right|} \sum_{t \in \mathcal{J}_{r}} \omega^{f^{\top} t}\left(x_{t}-w_{t}\right) \quad \triangleright \omega=e^{2 \pi \mathbf{i} / p}, u_{f, r}=(\widehat{x-w})_{f}^{\left[\mathcal{T}_{r}\right]}$
end for
end for
for $f=1 \rightarrow n$ do
$\eta=\operatorname{median}_{r \in[R]}\left\{u_{f, r}\right\} \quad \triangleright$ Take the median coordinate-wise
if $|\eta| \leq \nu / 2$ then
$z_{f} \leftarrow \eta$
else
$z_{f} \leftarrow 0$
end if
end for
return $z \quad \triangleright$ Guarantee $\|\widehat{x}-y-z\|_{\infty} \leq \nu$
end procedure

```

Let \(S\) be top \(C_{S} k\) coordinates in \(\widehat{x}\), where \(C_{S}=26\). If \(\|\widehat{x}-y\|_{\infty} \leq 2 \nu, \operatorname{supp}(y) \subseteq S\) and \(y\) is independent from the randomness of \(\left\{\mathcal{T}_{r}\right\}_{r=1}^{R}\), then with probability \(1-1 / \operatorname{poly}(n)\) under the randomness of \(\left\{\mathcal{T}_{r}\right\}_{r=1}^{R},\|\widehat{x}-y-z\|_{\infty} \leq \nu\) and \(\operatorname{supp}(z) \subseteq S\). Moreover, the running time of LinfinityReduce is \(O\left(n \log ^{2} n\right)\).

Proof. Note that \(\forall f \in \bar{S}\),
\[
\left|\widehat{x}_{f}\right| \leq \sqrt{\frac{\left\|\widehat{x}_{-k}\right\|_{2}^{2}}{C_{S} k-k}}=\frac{1}{5} \mu
\]
where the last step follows from choice of \(C_{S}\).
Let \(w\) denote the inverse Fourier transform of \(y\). Note that on Line 6 in Algo-
rithm 13.4, for any \(f \in[p]^{d}\) and \(r \in[R]\),
\[
\begin{aligned}
u_{f, r} & =\frac{\sqrt{n}}{\left|\mathcal{T}_{r}\right|} \sum_{t \in \mathcal{I}_{r}} \omega^{f^{\top} t}\left(x_{t}-w_{t}\right) \\
& =(\widehat{x-w})_{f}^{\left[\mathcal{T}_{r}\right]} \\
& =\sum_{f^{\prime} \in[n]} c_{f-f^{\prime}}^{\left[\mathcal{T}_{r}\right]}\left(\widehat{x}_{f^{\prime}}-y_{f^{\prime}}\right),
\end{aligned}
\]
where the second step follows by the notation in Definition 13.2.1, and the third step follows by Lemma 13.2.2. Therefore,
\[
\begin{equation*}
\widehat{x}_{f}-y_{f}=u_{f, r}-\sum_{f^{\prime} \in[p p]^{d} \backslash\{f\}} c_{f-f^{\prime}}^{\left[\mathcal{T}_{r}\right]}\left(\widehat{x}_{f^{\prime}}-y_{f^{\prime}}\right), \tag{13.1}
\end{equation*}
\]

By Lemma 13.2.4,
\[
\begin{equation*}
\underset{\mathcal{T}_{r}}{\operatorname{Pr}}\left[\left|\sum_{f^{\prime} \in[p]^{d} \backslash\{f\}} c_{f-f^{\prime}}^{\left[\mathcal{T}_{r}\right]}\left(\widehat{x}_{f^{\prime}}-y_{f^{\prime}}\right)\right| \geq \frac{10}{\sqrt{B}}\left\|(\widehat{x}-y)_{[p]^{d} \backslash\{f\}}\right\|_{2}\right] \leq \frac{1}{100} . \tag{13.2}
\end{equation*}
\]

We have
\[
\begin{align*}
\frac{10}{\sqrt{B}}\left\|(\widehat{x}-y)_{[p]]^{d} \backslash\{f\}}\right\|_{2} & \leq \frac{10}{\sqrt{B}}\left(\left\|(\widehat{x}-y)_{S \backslash\{f\}}\right\|_{2}+\left\|(\widehat{x}-y)_{\bar{S} \backslash\{f\}}\right\|_{2}\right) \\
& \leq \frac{10}{\sqrt{B}}\left(\|\widehat{x}-y\|_{\infty} \cdot \sqrt{|S|}+\left\|\widehat{x}_{\bar{S} \backslash\{f\}}\right\|_{2}\right) \\
& \leq \frac{10}{\sqrt{B}}(2 \nu \cdot \sqrt{26 k}+\sqrt{k} \mu) \\
& \leq \frac{1}{100 \sqrt{k}}(2 \nu \cdot \sqrt{26 k}+\sqrt{k} \mu) \\
& <0.12 \nu \tag{13.3}
\end{align*}
\]
where the first step following by triangle inequality, the second step follows by the assumption that \(\operatorname{supp}(y) \subseteq S\), the forth step follows by \(C_{B} \geq 10^{6}\), the last step follows by \(\mu \leq \nu\).

Therefore,
\[
\begin{aligned}
\underset{\mathcal{T}_{r}}{\operatorname{Pr}}\left[\left|u_{f, r}-\left(\widehat{x}_{f}-y_{f}\right)\right| \leq 0.12 \nu\right] & =\underset{\mathcal{T}_{r}}{\operatorname{Pr}}\left[\left|\sum_{f^{\prime} \in[p]^{\prime} d\{\{f\}} c_{f-f^{\prime}}^{\left[\mathcal{T}_{r}\right]}\left(\widehat{x}_{f^{\prime}}-y_{f^{\prime}}\right)\right| \leq 0.12 \nu\right] \\
& =1-\underset{\mathcal{T}_{r}}{\operatorname{Pr}}\left[\left|\sum_{f^{\prime} \in[p]^{d} \backslash\{f\}} c_{f-f^{\prime}}^{\left[\mathcal{T}_{r}\right]}\left(\widehat{x}_{f^{\prime}}-y_{f^{\prime}}\right)\right|>0.12 \nu\right] \\
& \geq 1-\underset{\mathcal{T}_{r}}{\operatorname{Pr}}\left[\left|\sum_{f^{\prime} \in[p]^{\prime} \backslash\{f\}} c_{f-f^{\prime}}^{\left[\mathcal{T}_{r}\right]}\left(\widehat{x}_{f^{\prime}}-y_{f^{\prime}}\right)\right| \geq \frac{10}{\sqrt{B}}\left\|(\widehat{x}-y)_{[p]^{d} \backslash\{f\}}\right\|_{2}\right] \\
& \geq 1-\frac{1}{100},
\end{aligned}
\]
where the first step follows by (13.1), the third step follows by (13.3), and the last step follows by (13.2).

Thus we have
\[
\underset{\mathcal{T}_{r}}{\operatorname{Pr}}\left[u_{f, r} \in \mathcal{B}_{\infty}\left(\widehat{x}_{f}-y_{f}, 0.12 \nu\right)\right] \geq \underset{\mathcal{J}_{r}}{\operatorname{Pr} r}\left[\left|u_{f, r}-\left(\widehat{x}_{f}-y_{f}\right)\right| \leq 0.12 \nu\right] \geq 1-\frac{1}{100}
\]

Let \(\eta_{f}=\operatorname{median}_{r \in[R]} u_{f, r}\) as on Line 10 in Algorithm 13.4. By Chernoff bound, with probability \(1-1 / \operatorname{poly}(n)\), more than \(\frac{1}{2} R\) elements in \(\left\{u_{f, r}\right\}_{r=1}^{R}\) are contained in box \(\mathcal{B}_{\infty}\left(\widehat{x}_{f}-y_{f}, 0.12 \nu\right)\), so that \(\eta_{f} \in \mathcal{B}_{\infty}\left(\widehat{x}_{f}-y_{f}, 0.12 \nu\right)\).

Therefore, we have
\[
\operatorname{Pr}\left[\left|\eta_{f}-\left(\widehat{x}_{f}-y_{f}\right)\right| \leq 0.17 \nu\right] \geq \operatorname{Pr}\left[\left|\eta_{f}-\left(\widehat{x}_{f}-y_{f}\right)\right| \leq \sqrt{2} \cdot 0.12 \nu\right] \geq 1-1 / \operatorname{poly}(n)
\]

Let \(E\) be the event that for all \(f \in[p]^{d},\left|\eta_{f}-\left(\widehat{x}_{f}-y_{f}\right)\right| \leq 0.17 \nu\). By a union bound over \(f \in[p]^{d}\), event \(E\) happens with probability \(1-1 / \operatorname{poly}(n)\). In the rest of the proof, we condition on event \(E\).
(Case 1) For \(f \in \bar{S}\), note that
\[
\left|\eta_{f}\right| \leq 0.17 \nu+\left|\widehat{x}_{f}-y_{f}\right|=0.17 \nu+\left|\widehat{x}_{f}\right| \leq 0.17 \nu+0.2 \nu=0.37 \nu
\]

According to the if statement between Line 11 and Line 15 in Algorithm 13.4, \(z_{f}\) will be assigned to 0 . Thus \(\operatorname{supp}(z) \subseteq S\). In addition, \(\left|\widehat{x}_{f}-y_{f}-z_{f}\right|=\left|\widehat{x}_{f}\right| \leq \mu \leq \nu\).
(Case 2) For \(f \in S\), we have two cases. We prove that \(\left|\left(\widehat{x}_{f}-y_{f}\right)-z_{f}\right| \leq \nu\) for both cases.
(Case 2.1) \(\left|\eta_{f}\right| \leq 0.5 \nu . z_{f}\) is assigned as 0 . Because
\[
\left|\eta_{f}-\left(\widehat{x}_{f}-y_{f}\right)\right| \leq 0.17 \nu, \quad\left|\widehat{x}_{f}-y_{f}\right| \leq\left|\eta_{f}\right|+0.17 \nu \leq 0.67 \nu
\]

Therefore,
\[
\left|\left(\widehat{x}_{f}-y_{f}\right)-z_{f}\right| \leq 0.67 \nu \leq \nu
\]
(Case 2.2) \(\left|\eta_{f}\right|>0.5 \nu . z_{f}\) is assigned as \(\eta_{f}\). We have
\[
\left|\left(\widehat{x}_{f}-y_{f}\right)-z_{f}\right|=\left|\left(\widehat{x}_{f}-y_{f}\right)-\eta_{f}\right| \leq 0.17 \nu \leq \nu
\]

We thus have obtained that with probability \(1-1 / \operatorname{poly}(n),\|(\widehat{x}-y)-z\|_{\infty} \leq \nu\) and \(\operatorname{supp}(z) \subseteq S\).

The running time of LinfinityReduce is dominated by the loop between Line 4 and Line 8 , which takes \(O(R \cdot n \log n)=O\left(n \log ^{2} n\right)\) by FFT.

For a given box \(\mathcal{B}_{\infty}(c, r)\) and grid \(\mathcal{G}_{r_{g}}\), we say a shift \(s \in \mathbb{C}\) is good if after applying the shift, all the points in the shifted box \(\mathcal{B}_{\infty}(c, r)+s\) are mapped to the same point by


Figure 13.4: Illustration of good and bad shifts in Definition 13.2.6. In (a), the small square represents box \(\mathcal{B}_{\infty}\left(c, r_{b}\right)\), and the dashed lines represent the decision boundary of \(\Pi_{r_{g}}\). The arrows in (b) and (c) represent two different shifts, where the shift in (b) is an example of good shift, since the shifted box does not intersect with the decision boundaries of \(\Pi_{r_{g}}\), while the shift in (c) is an example of bad shift, since the shifted box intersects with the decision boundaries of \(\Pi_{r_{g}}\).
\(\Pi_{r_{g}}\) (recall that \(\Pi_{r_{g}}\) projects any point to the nearst grid point in \(\mathcal{G}_{r_{g}}\) ). We formulate the notation of a good shift in the following definition, and illustrate in Figure 13.4.

Definition 13.2 .6 (good shift). For any \(r_{g}, r_{b}\), and any \(c \in \mathbb{C}\), we say shift \(s \in \mathbb{C}\) is a good shift if
\[
\left|\Pi_{r_{g}}\left(\mathcal{B}_{\infty}\left(c, r_{b}\right)+s\right)\right|=1
\]

The following lemma intuitively states that if we take a box of radius \(r_{b}\) (or equivalently, side length \(2 r_{b}\) ) and shift it randomly by an offset in \(\mathcal{B}_{\infty}\left(0, r_{s}\right)\) (or equivalently, \(\left.\left[-r_{s}, r_{s}\right] \times\left[-r_{s}, r_{s}\right]\right)\) where \(r_{s} \geq r_{b}\), and next we round everyone inside that shifted box to the closest point in \(G_{r_{g}}\) where \(r_{g} \geq 2 r_{s}\), then with probability at least \(\left(1-r_{b} / r_{s}\right)^{2}\) everyone will be rounded to the same point. In other words, let \(s \sim \mathcal{B}_{\infty}\left(0, r_{s}\right)\), for box \(\mathcal{B}_{\infty}\left(c, r_{b}\right)\) and


Figure 13.5: Illustration of Lemma 13.2.6. In (a) the smallest square represents box \(\mathcal{B}_{\infty}\left(c, r_{b}\right)\), the medium-sized square represents \(\mathcal{B}_{\infty}\left(c, r_{s}\right)\), and the dashed lines represent decision boundaries of \(\Pi_{r_{g}}\). Note that for \(s \sim \mathcal{B}_{\infty}\left(0, r_{s}\right)\), the center of the shifted box \(s+\mathcal{B}_{\infty}\left(c, r_{b}\right)\) is \(s+c \sim \mathcal{B}_{\infty}\left(c, r_{s}\right)\). Shift \(s\) is good (recall in Definition 13.2.6) for box \(\mathcal{B}_{\infty}\left(c, r_{b}\right)\) and grid \(\mathcal{G}_{r_{g}}\) if and only if the distance between \(s+c\) and decision boundaries of \(\Pi_{r_{g}}\) is greater than \(r_{b}\). In (b), we draw in red the set of points which are within distance at most \(r_{b}\) to the decision boundaries of \(\Pi_{r_{g}}\). Then in (b) the red part inside \(\mathcal{B}_{\infty}\left(c, r_{s}\right)\) corresponds to bad shifts (plus \(c\) ), and the green part corresponds to good shifts (plus \(c\) ). Intuitively, the fraction of the green part is at least \(\left(1-r_{b} / r_{s}\right)^{2}\) because the vertical red strips can cover a width of at most \(2 r_{b}\) on the \(x\)-axis of \(\mathcal{B}_{\infty}\left(c, r_{s}\right)\) (whose side length is \(2 r_{s}\) ), and the horizontal red strips can cover a width of at most \(2 r_{b}\) on the \(y\)-axis.
grid \(\mathcal{G}_{r_{g}}, s\) is a good shift with probability at least \(\left(1-r_{b} / r_{s}\right)^{2}\). We illustrate the lemma in Figure 13.5.

Lemma 13.2.6 (property of a randomly shifted box). For any \(r_{g}, r_{s}\), \(r_{b}\) so that \(r_{g} / 2 \geq r_{s} \geq\) \(r_{b}>0\) and any \(c \in \mathbb{C}\), let \(s \in \mathbb{C}\) be uniform randomly chosen in \(\mathcal{B}_{\infty}\left(0, r_{s}\right)\), then
\[
\operatorname{Pr}_{s \sim \mathcal{B}_{\infty}\left(0, r_{s}\right)}\left[\left|\Pi_{r_{g}}\left(\mathcal{B}_{\infty}\left(c, r_{b}\right)+s\right)\right|=1\right] \geq\left(1-\frac{r_{b}}{r_{s}}\right)^{2}
\]
where we refer \(r_{g}, r_{s}, r_{b}\) as the radius of grid, shift and box respectively, and we use notation \(C+s\) to refer to \(\{c+s: c \in C\}\).

Proof. We consider complex numbers as points in 2D plane, where the real part is the coordinate on \(x\)-axis, and the imaginary part is the coordinate on \(y\)-axis. Note that the "decision boundary" of projection \(\Pi_{r_{g}}\) from \(\mathbb{C}\) onto grid \(\mathcal{G}_{r_{g}}\) consists of verticle lines of form \(x=\) \(\left(m+\frac{1}{2}\right) r_{g}\) and horizontal lines of form \(y=\left(m+\frac{1}{2}\right) r_{g}\), where \(m \in \mathbb{Z}\). \(\left|\Pi_{r_{g}}\left(\mathcal{B}_{\infty}\left(c, r_{b}\right)+s\right)\right|=1\) if and only if the shifted box \(\mathcal{B}_{\infty}\left(c, r_{b}\right)+s\) does not intersect with the "decision boundary".

Let \(s=s_{x}+s_{y} \mathbf{i}\) and \(c=c_{x}+c_{y} \mathbf{i}\). Then the shifted box does not intersect with the "decision boundary" if and only if both the interval
\[
\left[c_{x}-r_{b}+s_{x}, c_{x}+r_{b}+s_{x}\right] \text { and }\left[c_{y}-r_{b}+s_{y}, c_{y}+r_{b}+s_{y}\right]
\]
do not intersect with \(\left\{\left(m+\frac{1}{2}\right) r_{g}: m \in \mathbb{Z}\right\}\). The probability of each one is at least \(1-\frac{r_{b}}{r_{s}}\), and two events are independent. Therefore, we get the claimed result.

In the following, we define event \(\mathcal{E}\), which is a sufficient condition for the correctness of Algorithm 13.3. Event \(\mathcal{E}\) consists of three parts. Part 1 of \(\mathcal{E}\) is used to prove that \(a_{\ell} \leq 10 \log n\) for \(\ell \in[L-1]\) on Line 29 in Algorithm 13.3. Part 2 and Part 3 of \(\mathcal{E}\) are used to prove that Line 15 to Line 20 in Algorithm 13.3 give a desirable \(z^{(\ell)}\) for \(\ell \in[L]\).

Definition 13.2.7 (sufficient condition for the correctness of Algorithm 13.3). For input signal \(x \in \mathbb{C}^{n}\), let \(\mu=\frac{1}{\sqrt{k}}\left\|\widehat{x}_{-k}\right\|_{2}\) and \(R^{*}\) is an upper bound of \(\|\widehat{x}\|_{\infty} / \mu\). Let \(S\) be top \(C_{S} k\) coordinates in \(\widehat{x}\). Let \(H=\min \left\{\log k+C_{H}, \log R^{*}\right\}\), and \(L=\log R^{*}-H+1\). For \(\ell \in[L]\), let \(\nu_{\ell}=2^{-\ell} \mu R^{*}\). For \(\ell \in[L-1]\), let \(s_{\ell}^{(a)}\) be the \(a\)-th uniform randomly sampled from \(\mathcal{B}_{\infty}\left(0, \alpha \nu_{\ell}\right)\) as appeared on Line 26 in Algorithm 13.3 (i.e. \(s_{\ell}^{(1)}, \ldots, s_{\ell}^{\left(a_{\ell}\right)}\) are sampled, and \(s_{\ell}^{\left(a_{\ell}\right)}\) is the first that satisfies the condition on Line 28). For the sake of analysis, we assume that Algorithm 13.3 actually produces an infinite sequence of shifts \(s_{\ell}^{(1)}, s_{\ell}^{(2)}, \ldots\), and
chooses the smallest \(a_{\ell}\) so that \(s_{\ell}^{\left(a_{\ell}\right)}\) satisfies \(\forall f \in \operatorname{supp}\left(y^{(\ell-1)}+z^{(\ell)}\right), \mid \Pi_{\beta \nu_{\ell}}\left(\mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}+\right.\right.\) \(\left.\left.s_{\ell}^{\left(a_{\ell}\right)}, 2^{1-H} \nu_{\ell}\right)\right) \mid=1\) on Line 28 .

For \(\ell \in[L-1]\), we define random variable \(a_{\ell}^{\prime}\) to be the smallest \(a^{\prime}\) such that for all \(f \in S\),
\[
\left|\Pi_{\beta \nu_{\ell}}\left(\mathcal{B}_{\infty}\left(\widehat{x}_{f}+s_{\ell}^{\left(a^{\prime}\right)}, 2^{3-H} \nu_{\ell}\right)\right)\right|=1 .
\]

We define event \(\mathcal{E}\) to be all of the following events hold.
1. For all \(\ell \in[L-1], a_{\ell}^{\prime} \leq 10 \log n\).
2. For \(\ell=1\), if we execute Line 15 to Line 20 in Algorithm 13.3 with \(y^{(0)}=0\), we get \(z^{(1)}\) such that \(\left\|\widehat{x}-z^{(1)}\right\|_{\infty} \leq 2^{1-H} \nu_{1}\) and \(\operatorname{supp}\left(z^{(1)}\right) \subseteq S\).
3. For all \(\ell \in\{2, \ldots, L\}\), for all \(a \in[10 \log n]\), if we execute Line 15 to Line 20 in Algorithm 13.3 with \(y^{(\ell-1)}=\xi\) where
\[
\xi_{f}= \begin{cases}\Pi_{\beta \nu_{\ell}}\left(\widehat{x}_{f}+s_{\ell-1}^{(a)}\right), & \text { if } f \in S ; \\ 0, & \text { if } f \in \bar{S}\end{cases}
\]
then we get \(z^{(\ell)}\) such that \(\left\|\widehat{x}-y^{(\ell-1)}-z^{(\ell)}\right\|_{\infty} \leq 2^{1-H} \nu_{\ell}\) and \(\operatorname{supp}\left(y^{(\ell-1)}+z^{(\ell)}\right) \subseteq S\).

In the following, we will prove that for fixed \(x\), under the randomness of \(\left\{s_{\ell}^{(a)}\right\}_{\ell \in[L-1], a \in\{1, \ldots\}}\) and \(\mathcal{T}=\left\{\mathcal{T}^{(h)}\right\}_{h \in[H]}\), event \(\mathcal{E}\) (defined in Definition 13.2.7) happens with probability at least \(1-1 / \operatorname{poly}(n)\). Moreover, we will prove that event \(\mathcal{E}\) is a sufficient condition for the correctness of Algorithm 13.3. Namely, conditioned on event \(\mathcal{E}\), Algorithm 13.3 gives a desirable output.

Lemma 13.2.7 (event \(\mathcal{E}\) happens with high probability). Let \(\mathcal{E}\) in Definition 13.2.7. For any fixed \(x \in \mathbb{C}^{n}\), under the randomness of shifts \(\left\{s_{\ell}^{(a)}\right\}_{\ell \in[L-1], a \in\{1, \ldots\}}\) and \(\mathcal{T}=\left\{\mathcal{T}^{(h)}\right\}_{h \in[H]}\),
\[
\operatorname{Pr}[\varepsilon] \geq 1-1 / \operatorname{poly}(n)
\]

Proof. We bound the failure probability of each parts in event \(\mathcal{E}\) respectively as follows, and \(\operatorname{Pr}[\mathcal{E}] \geq 1-1 / \operatorname{poly}(n)\) follows by a union bound.

Part 1. If \(H=\log R^{*}\), then \(L=1\) and it is trivially true that "for all \(\ell \in[L-1], a_{\ell}^{\prime} \leq\) \(10 \log n "\). Otherwise, we have \(H=\log k+C_{H}\). By Lemma 13.2.6, for any \(\ell \in[L-1]\), for any \(f \in S\),
\[
\operatorname{Pr}_{s \sim \mathcal{B}_{\infty}\left(0, \alpha \nu_{\ell}\right)}\left[\left|\Pi_{\beta \nu_{\ell}}\left(\mathcal{B}_{\infty}\left(\widehat{x}_{f}, 2^{3-H} \nu_{\ell}\right)+s\right)\right|=1\right] \geq\left(1-\frac{2^{3-H} \nu_{\ell}}{\alpha \nu_{\ell}}\right)^{2}=\left(1-\frac{2^{3-H}}{\alpha}\right)^{2},
\]
where \(\left(1-2^{3-H} / \alpha\right)^{2} \geq 1-2^{4-C_{H}-\log k} / \alpha \geq 1-\frac{1}{100 k}\) by our choice of \(\alpha\) and \(C_{H}\) in Table 13.2.
For each \(\ell \in[L-1]\), by a union bound over all \(f\) in \(S\), the probability is at least \(1-\frac{C_{S} k}{100 k}=1-\frac{26 k}{100 k} \geq \frac{1}{2}\) that for all \(f \in S,\left|\Pi_{\beta \nu_{\ell}}\left(\mathcal{B}_{\infty}\left(\widehat{x}_{f}+s, 2^{3-H} \nu_{\ell}\right)\right)\right|=1\) where \(s \sim\) \(\mathcal{B}_{\infty}\left(0, \alpha \nu_{\ell}\right)\). Formally, we get
\[
\operatorname{Pr}_{s \sim \mathcal{B}}^{\operatorname{Pr}\left(0, \alpha \nu_{\ell}\right)}\left[\left|\Pi_{\beta \nu_{\ell}}\left(\mathcal{B}_{\infty}\left(\widehat{x}_{f}, 2^{3-H} \nu_{\ell}\right)+s\right)\right|=1, \forall f \in S\right] \geq 1 / 2
\]

Therefore, by definition of \(a_{\ell}^{\prime}\) in Definition 13.2.7,
\[
\operatorname{Pr}\left[a_{\ell}^{\prime} \leq 10 \log n\right] \geq 1-(1 / 2)^{10 \log n}=1-1 / n^{10}
\]

By a union bound over all \(\ell \in[L-1]\), the probability is at least \(1-L / n^{10}=1-1 / \operatorname{poly}(n)\) that for all \(\ell \in[L-1], a_{\ell}^{\prime} \leq 10 \log n\).

Part 2. By Lemma 13.2.5 and a union bound over all \(h \in[H]\), the failure probability is at most \(H / \operatorname{poly}(n)=1 / \operatorname{poly}(n)\), where \(H=O(\log k)\) and so \(H / \operatorname{poly}(n)\) is still \(1 / \operatorname{poly}(n)\).

Part 3. For each \(\ell \in\{2, \ldots, L\}\) and \(a \in[10 \log n]\), similar to the above argument, each has failure probability at most \(1 / \operatorname{poly}(n)\). By a union bound, the failure probability is
at most
\[
(L-1) \cdot(10 \log n) / \operatorname{poly}(n)=1 / \operatorname{poly}(n) .
\]

In the following lemma, we show that if event \(\mathcal{E}\) (defined in Definition 13.2.7) happens, then Algorithm 13.3 gives a desirable output.

Lemma 13.2.8 (correctness of Algorithm 13.3 conditioned on \(\mathcal{E}\) ). Let \(n=p^{d}\), and let \(k \in[n]\). Let \(x \in \mathbb{C}^{n}\) be input signal. Let \(\mu=\frac{1}{\sqrt{k}}\left\|\widehat{x}_{-k}\right\|_{2}\). Let \(R^{*} \geq\|\widehat{x}\|_{\infty} / \mu\) and \(R^{*}\) is a power of 2. Let \(H=\min \left\{\log k+C_{H}, \log R^{*}\right\}\). Let \(L=\log R^{*}-H+1\). For \(\ell \in\) \([L-1]\), let \(y^{(\ell)}\) be the vector obtained on Line 31 of Algorithm 13.3. For \(\ell \in[L]\), let \(z^{(\ell)}\) be the vector obtained on Line 20. Note that \(y^{(0)}=0\), and \(y^{(L-1)}+z^{(L)}\) is the output of FourierSparseRecovery \(\left(x, n, k, R^{*}, \mu\right)\) in Algorithm 13.3. Conditioned on the event \(\mathcal{E}\) (defined in Definition 13.2.7) happens, we have
\[
\left\|\widehat{x}-y^{(L-1)}-z^{(L)}\right\|_{\infty} \leq \frac{1}{\sqrt{k}}\left\|\widehat{x}_{-k}\right\|_{2} .
\]

Proof. We first discuss the case that \(H=\log R^{*}\). In that case, \(L=1\). Conditioned on the event \(\mathcal{E}\) (Part 2 of \(\mathcal{E}\) ), \(z^{(1)}\) obtained through Line 15 to Line 20 in Algorithm 13.3 satisfies \(\left\|\widehat{x}-z^{(1)}\right\|_{\infty} \leq 2^{1-H} \nu_{1}=2^{1-H}\left(2^{-1} \mu R^{*}\right)=\mu\).

In the rest of the proof, we discuss the case that \(H>\log R^{*}\). For \(\ell \in[L]\), let \(\nu_{\ell}=2^{-\ell} \mu R^{*}\). For \(\ell \in[L-1]\), let \(s_{\ell}^{\left(a_{\ell}\right)} \in \mathcal{B}_{\infty}\left(0, \alpha \nu_{\ell}\right)\) denote the first \(s_{\ell}^{(a)}\) on Line 26 in Algorithm 13.3 such that for all \(f \in \operatorname{supp}\left(y^{(\ell-1)}+z^{(\ell)}\right)\),
\[
\left|\Pi_{\beta \nu_{\ell}}\left(\mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}+s_{\ell}^{(a)}, 2^{1-H} \nu_{\ell}\right)\right)\right|=1
\]

For \(\ell \in[L-1]\), we define \(\xi^{(\ell)} \in \mathbb{C}^{[p]^{d}}\) as follows
\[
\xi_{f}^{(\ell)}= \begin{cases}\Pi_{\beta \nu_{\ell}}\left(\widehat{x}_{f}+s_{\ell}^{\left(a_{\ell}\right)}\right), & \text { if } f \in S ; \\ 0, & \text { if } f \in \bar{S} .\end{cases}
\]

We also define \(\xi^{(0)}=0, s_{0}^{(a)}=0\) for \(a \in\{1, \ldots\}\) and \(a_{0}=1\).
(Goal: Inductive Hypothesis) We are going to prove that conditioned on event \(\mathcal{E}\) (defined in Definition 13.2.7), for all \(\ell \in\{0, \ldots, L-1\}\),
\[
y^{(\ell)}=\xi^{(\ell)} \quad \text { and } \quad a_{\ell} \leq 10 \log n
\]
(Base case) Note that \(y^{(0)}=\xi^{(0)}=0\) and \(a_{0}=1 \leq 10 \log n\).
(Inductive step) We will prove that conditioned on event \(\mathcal{E}\), if \(y^{(\ell-1)}=\xi^{(\ell-1)}\) and \(a_{\ell-1} \leq 10 \log n\) for \(\ell \in[L-1]\), then \(y^{(\ell)}=\xi^{(\ell)}\) and \(a_{\ell} \leq 10 \log n\).
(Proving \(\left.a_{\ell} \leq 10 \log n\right)\) Conditioned on event \(\mathcal{E}\) (if \(L=1\) then from Part 2 of \(\mathcal{E}\), otherwise from Part 3 of \(\mathcal{E}\) and by the fact that \(\left.a_{\ell-1} \leq 10 \log n\right), z^{(\ell)}\) obtained through Line 15 to Line 20 in Algorithm 13.3 satisfies \(\left\|\widehat{x}-\xi^{(\ell-1)}-z^{(\ell)}\right\|_{\infty} \leq 2^{1-H} \nu_{\ell}\) and \(\operatorname{supp}\left(z^{(\ell)}\right) \subseteq S\). Namely, for all \(f \in[p]^{d}, \xi_{f}^{(\ell-1)}+z_{f}^{(\ell)} \in \mathcal{B}_{\infty}\left(\widehat{x}_{f}, 2^{1-H} \nu_{\ell}\right)\). Recall the definition of \(a_{\ell}^{\prime}\) in Definition 13.2.7. We can prove that \(a_{\ell} \leq a_{\ell}^{\prime}\) because if for all \(f \in S\),
\[
\left|\Pi_{\beta \nu_{\ell}}\left(\mathcal{B}_{\infty}\left(\widehat{x}_{f}+s_{\ell}^{\left(a_{\ell}^{\prime}\right)}, 2^{3-H} \nu_{\ell}\right)\right)\right|=1
\]
then for all \(f \in \operatorname{supp}\left(y^{(\ell-1)}+z^{(\ell)}\right)\),
\[
\left|\Pi_{\beta \nu_{\ell}}\left(\mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}+s_{\ell}^{\left(a_{\ell}^{\prime}\right)}, 2^{1-H} \nu_{\ell}\right)\right)\right|=1
\]
where
\[
\mathcal{B}_{\infty}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}+s_{\ell}^{\left(a_{\ell}^{\prime}\right)}, 2^{1-H} \nu_{\ell}\right) \subseteq \mathcal{B}_{\infty}\left(\widehat{x}_{f}+s_{\ell}^{\left(a_{\ell}^{\prime}\right)}, 2^{3-H} \nu_{\ell}\right)
\]
which follows by \(\xi_{f}^{(\ell-1)}+z_{f}^{(\ell)} \in \mathcal{B}_{\infty}\left(\widehat{x}_{f}, 2^{1-H} \nu_{\ell}\right)\). Therefore, conditioned on \(\mathcal{E}\) (Part 1 of \(\mathcal{E}\) ), \(a_{\ell} \leq a_{\ell}^{\prime} \leq 10 \log n\).
(Proving \(y_{f}^{(\ell)}=\xi_{f}^{(\ell)}\) ) For \(f \in[p]^{d}\), we will prove that \(y_{f}^{(\ell)}=\xi_{f}^{(\ell)}\) in two cases.
(Case 1) If \(f \in \operatorname{supp}\left(y^{(\ell-1)}+z^{(\ell)}\right) \subseteq S\). We have
\[
\begin{aligned}
y_{f}^{(\ell)} & =\Pi_{\beta \nu_{\ell}}\left(y_{f}^{(\ell-1)}+z_{f}^{(\ell)}+s_{\ell}^{\left(a_{\ell}\right)}\right) \\
& =\Pi_{\beta \nu_{\ell}}\left(\xi_{f}^{(\ell-1)}+z_{f}^{(\ell)}+s_{\ell}^{\left(a_{\ell}\right)}\right) .
\end{aligned}
\]

Because \(\xi_{f}^{(\ell-1)}+z_{f}^{(\ell)} \in \mathcal{B}_{\infty}\left(\widehat{x}_{f}, 2^{1-H} \nu_{\ell}\right)\), we have \(\xi_{f}^{(\ell-1)}+z_{f}^{(\ell)}+s_{\ell}^{\left(a_{\ell}\right)} \in \mathcal{B}_{\infty}\left(\widehat{x}_{f}+s_{\ell}^{\left(a_{\ell}\right)}, 2^{1-H} \nu_{\ell}\right)\). By the choice of \(s_{\ell}^{\left(a_{\ell}\right)}, \Pi_{\beta \nu_{\ell}}\left(\xi_{f}^{(\ell-1)}+z_{f}^{(\ell)}+s_{\ell}^{\left(a_{\ell}\right)}\right)=\Pi_{\beta \nu_{\ell}}\left(\widehat{x}_{f}+s_{\ell}^{\left(a_{\ell}\right)}\right)\). Thus \(y_{f}^{(\ell)}=\xi_{f}^{(\ell)}\).
(Case 2) If \(f \notin \operatorname{supp}\left(y^{(\ell-1)}+z^{(\ell)}\right)\). We have \(y_{f}^{(\ell)}=0\). Because \(\xi_{f}^{(\ell-1)}+z_{f}^{(\ell)} \in\) \(\mathcal{B}_{\infty}\left(\widehat{x}_{f}, 2^{1-H} \nu_{\ell}\right)\), we have \(\left|\widehat{x}_{f}\right|<2^{2-H} \nu_{\ell}<0.1 \beta \nu_{\ell}\) by our choice of \(H\). We can easily prove that \(\xi_{f}^{(\ell)}=0=y_{f}^{(\ell)}\) in the following two cases:
(Case 2.1) If \(f \in S\), we have \(\xi_{f}^{(\ell)}=\Pi_{\beta \nu_{\ell}}\left(\widehat{x}_{f}+s_{\ell}^{\left(a_{\ell}\right)}\right)=0\) because
\[
\left|\widehat{x}_{f}\right|+\left|s_{\ell}^{\left(a_{\ell}\right)}\right|<0.1 \beta \nu_{\ell}+2 \alpha \nu_{\ell}<0.5 \beta \nu_{\ell}
\]
(Case 2.2) If \(f \in \bar{S}, \xi_{f}^{(\ell)}=0\) by definition of \(\xi^{(\ell)}\).
Therefore, for all \(\ell \in[L-1], y^{(\ell)}=\xi^{(\ell)}\) and \(a_{\ell} \leq 10 \log n\). Again conditioned on event \(\mathcal{E}(\) Part 3 of \(\mathcal{E}), z^{(L)}\) obtained through Line 15 to Line 20 in Algorithm 13.3 satisfies
\[
\left\|\widehat{x}-y^{(L-1)}-z^{(L)}\right\|_{\infty} \leq 2^{1-H} \nu_{L}=2^{1-H}\left(2^{-\left(\log R^{*}-H+1\right)} \mu R^{*}\right)=\mu
\]

Therefore, \(y^{(L-1)}+z^{(L)}\) on Line 22 gives a desirable output.

Now we present our main theorem, which proves the correctness of Algorithm 13.3, and shows its sample complexity and time complexity.

Theorem 13.2.9 (main result, formal version). Let \(n=p^{d}\) where both \(p\) and \(d\) are positive integers. Let \(x \in \mathbb{C}^{[p]^{d}}\). Let \(k \in\{1, \ldots, n\}\). Assume we know \(\mu=\frac{1}{k}\left\|\widehat{x}_{-k}\right\|_{2}\) and \(R^{*} \geq\|\widehat{x}\|_{\infty} / \mu\) where \(\log R^{*}=O(\log n)\). There is an algorithm (Algorithm 13.3) that takes \(O(k \log k \log n)\) samples from \(x\), runs in \(O\left(n \log ^{3} n \log k\right)\) time, and outputs a \(O(k)\)-sparse vector \(y\) such that
\[
\|\widehat{x}-y\|_{\infty} \leq \frac{1}{\sqrt{k}} \min _{k-\text { sparse } x^{\prime}}\left\|\widehat{x}-x^{\prime}\right\|_{2}
\]
holds with probability at least \(1-1 / \operatorname{poly}(n)\).

Proof. The correctness of Algorithm 13.3 follows directly from Lemma 13.2.7 and Lemma 13.2.8. The number of samples from \(x\) is
\[
B \cdot R \cdot H=O(k \cdot \log n \cdot \log k)=O(k \log k \log n)
\]

Its running time is dominated by \(L \cdot H=O(\log k \log n)\) invocations of LinfinityReduce (in Algorithm 13.4). By Lemma 13.2.5, the running time of LinfinityReduce is \(O\left(n \log ^{2} n\right)\). Therefore, the running time of Algorithm 13.3 is
\[
O\left(L \cdot H \cdot n \log ^{2} n\right)=O\left(\log k \cdot \log n \cdot n \log ^{2} n\right)=O\left(n \log ^{3} n \log k\right)
\]

Chapter 14

\section*{Set Query}

\subsection*{14.1 Introduction}

This result presented in this section is based on a un-published manuscript, and is a joint work Vasileios Nakos and Zhengyu Wang [NSW18].

Our first result is on the set-query problem. We show that \(O(k / \epsilon)\) measurements, \(O(\log k)\) update time and \(O(k \log k)\) query time are possible. Our hash functions are only 2 -wise independent, and the amount of space to store the matrix \(\Phi\) is only \(O(\log k)\) words. The previous work of [Pri11] needed \(\Omega(n)\) time to store the matrix, since the columns had to be fully independent, something that makes the particular sketch not applicable in the streaming model. The problem of whether the matrix can be stored in less space was explicitly stated in that paper as an open problem:
"Our analysis assumes that the columns of \(A\) are fully independent. It would be valuable to reduce the independence needed, and hence the space required to store \(A\)."

Theorem 14.1.1 (Classical set query). There exists a randomized construction of a linear sketch \(\Phi \in \mathbb{R}^{m \times n}\) with \(m=O(k / \epsilon)\), such that given \(y=\Phi x\) and a set \(S \subset[n]\) of cardinality at most \(k\), we can find in \(O(k \log k)\) time a \(k\)-sparse vector such that
\[
\left\|x^{\prime}-x_{S}\right\|_{2}^{2} \leq \epsilon\left\|x_{[n] \backslash S}\right\|_{2}^{2}
\]
holds with probability 9/10. The time to update the linear sketch is \(O(\log k)\) and the space to store the matrix is \(O(\log k)\) words.

Our second main result concerns the ubiquitous problem of frequency estimation from Fourier measurements. For any vector \(x \in \mathbb{C}^{n}\), let \(\widehat{x} \in \mathbb{C}^{n}\) denote the discrete Fourier transform of \(x\). The only previous known result by the breakthrough work of [Kap17] achieves
\begin{tabular}{|l|l|l|l|l|}
\hline & Space & Measurements & Decoding Time & Update Time \\
\hline [Pri11] & \(n\) & \(k / \epsilon\) & \(k\) & 1 \\
\hline Theorem 14.1.1 & \(\log k\) & \(k / \epsilon\) & \(k \log k\) & \(\log k\) \\
\hline
\end{tabular}

Table 14.1: Results on set query. We ignore the " \(O\) " for simplicity. "Space" means the space to store matrix, which doesn't include number of measurements. We note that the decoding time of the set query algorithms do not depend on \(\epsilon\).
\begin{tabular}{|l|l|l|}
\hline & Measurements & Running Time \\
\hline [Kap17] & \(k / \epsilon\) & \(k / \epsilon \cdot \log ^{3.001} n\) \\
\hline Theorem 14.1.2 & \(k / \epsilon \cdot \log n\) & \(k / \epsilon \cdot \log n\) \\
\hline
\end{tabular}

Table 14.2: Results on Fourier set query. We ignore the " \(O\) " for simplicity. We assume signal is bounded by \(\operatorname{poly}(n)\). In Fourier set query, there is no update time concept.
sample complexity \(O(k / \epsilon)\) and running time \(O\left(\epsilon^{-1} k \log ^{2.001} n \log R^{*}\right)\) for \(\ell_{2} / \ell_{2}\) Fourier set query. Here, \(R^{*}\) is an upper bound on the \(\|\cdot\|_{\infty}\) norm of the vector, in most applications being considered poly \((n)\). We indicate how our approach immediately gives an algorithm with \(O\left(\epsilon^{-1} k \log n\right)\) running time. The result we get has no dependence on the running time on \(\log R^{*}\), but does not achieve the optimal sample complexity.

Theorem 14.1.2 (Fourier set query). Given a vector \(x \in \mathbb{C}^{n}\), for every \(\epsilon \in(0,1)\) and \(k \geq 1\), any \(S \subseteq[n],|S|=k\), there exists an algorithm that takes \(O\left(\epsilon^{-1} k \log (n / \delta)\right)\) samples, runs in \(O\left(\epsilon^{-1} k \log (n / \delta)\right)\) time, and outputs a vector \(x^{\prime} \in \mathbb{C}^{n}\) such that \(\left\|\left(x^{\prime}-\widehat{x}\right)_{S}\right\|_{2}^{2} \leq\) \(\epsilon\left\|\widehat{x}_{[n] \backslash S}\right\|_{2}^{2}+\delta\|\widehat{x}\|_{1}^{2}\) holds with probability at least \(9 / 10\).

Note that \(\delta\) dependence is standard in Fourier transform setting, see e.g. [HIKP12a, PS15, CKPS16].

\subsection*{14.2 Classical set query}

In this section, we present our set query data structure. We first give the construction of sketching matrix in Section 14.2.1. Then we show our new set query data structure in Section 14.2.2. We give an analysis of the data structure in Section 14.2.4, whose correctness relies on a technical lemma on iterative loop, which is proved in Section 14.2.3.

\subsection*{14.2.1 Construction of sketching matrix}

The sketching matrix used in our set query data structure consists of \(\log k\) parts, where the \(i\)-th part \((i \in[\log k])\) of the sketching matrix corresponds to a count sketch data structure that hashes \(n\) elements into \(B_{i}\) bins. We defer the choice of \(B_{i}\) to Section 14.2.2. We present the sketching matrix \(\Phi\) below.

Definition 14.2.1 (Sketching matrix \(\Phi\) ). For each \(i \in[\log k]\), we choose \(h_{i}:[n] \rightarrow\left[B_{i}\right]\) to be a pairwise independent hash function, and choose \(\sigma_{i}:[n] \rightarrow\{-1,+1\}\) to be a pairwise independent hash function.

For each \(i \in[\log k]\), we define Count-Sketch matrix \(\Phi^{(i)} \in \mathbb{R}^{B_{i} \times n}\) as follows
\[
\Phi_{j, l}^{(i)}= \begin{cases}\sigma_{i}(l) & \text { if } h_{i}(l)=j \\ 0 & \text { otherwise } .\end{cases}
\]

Let \(m=\sum_{i=1}^{\log k} B_{i}\). We define matrix \(\Phi \in \mathbb{R}^{m \times n}\) as follows
\[
\Phi=\left[\begin{array}{c}
\Phi^{(1)} \\
\Phi^{(2)} \\
\ldots \\
\Phi^{(\log k)}
\end{array}\right]
\]

\subsection*{14.2.2 Algorithm}

Let \(C\) be a sufficiently large constant, and \(\gamma\) be a sufficiently small constant. For \(i \in[\log k]\), let \(k_{i}=k \gamma^{i}\) and \(\epsilon_{i}=\epsilon(10 \gamma)^{i}\). Let \(B_{i}=C k_{i} / \epsilon_{i}\). The set query algorithm maintains the measurement \(y=\Phi x\).

On update operation, the algorithm follows directly through the definition of sketching matrix \(\Phi\). Namely, on \(\operatorname{UpdAtE}(i, \Delta)\), the algorithm finds the non-zero elements of the \(i\)-th column of \(\Phi\), and update the measurement \(y\) accordingly.

On query operation \(\operatorname{QUERY}(x, S)\), the recovery procedure uses \(x^{\prime}\) to record the current best approximation of \(x\), and initially set \(x^{\prime} \leftarrow 0\). The procedure consists of \(\log k\) iterations. In the \(i\)-th iteration, let \(S_{i} \subseteq S\) denote the set of coordinates in \(x\) that are yet to be recovered. Initially \(S_{1}=S\). Remind that sketching matrix \(\Phi\) in Definition 14.2 .1 consists of \(\log k\) parts, and the \(i\)-th part corresponds to a count sketch data structure with pair-wise independent hash function \(h_{i}\) and random sign function \(\sigma_{i}\). In the \(i\)-th iteration of the set query algorithm,
\[
T_{i} \leftarrow\left\{j \in S_{i} \mid h_{i}^{-1}(j) \cap S_{i}=\emptyset\right\} .
\]

Intuitively the hash function \(h_{i}\) maps \(n\) elements into \(B_{i}\) bins, for any \(j \in S_{i}\), we have \(j \in T_{i}\) if and only if \(j\) does not collide with any other elements in \(S_{i}\) (in that case, we say \(j\) is isolated).

The procedure also maintains \(y^{(i)}\), which is the measurement of the residual vector \(x-x^{\prime}\) (before the start of \(i\)-th iteration). After obtaining \(T_{i}\), the procedure uses \(y^{(i)}\) to estimate \(x_{j}\) for every \(j \in T_{i}\). Then the procedure updates \(S_{i}\) and \(x^{\prime}\) accordingly. Finally, the algorithm outputs \(x^{\prime}\). Our iterative set query algorithm is presented in Algorithm 14.1.
```

Algorithm 14.1 Iterative Set Query
procedure IterativeSetQuery $(\epsilon, k) \quad \triangleright$ Theorem 14.2.2
$C \leftarrow 20, \gamma \leftarrow 1 / 600$
for $i=1 \rightarrow \log k$ do
$k_{i} \leftarrow k \gamma^{i}, \epsilon_{i} \leftarrow \epsilon(10 \gamma)^{i}$
$B_{i} \leftarrow C k_{i} / \epsilon_{i}$
end for
$m \leftarrow \sum_{i=1}^{\log k} B_{i} \quad \triangleright m$ is the number of measurement
Construct matrix $\Phi$ according to Definition 14.2.1 $\triangleright \Phi \in \mathbb{R}^{m \times n}$
$y \leftarrow 0 \quad \triangleright y \in \mathbb{R}^{m}$
procedure $\operatorname{UPDATE}(i, \Delta)$
$y \leftarrow y+\Delta \cdot\left(\Phi \cdot e_{i}\right)$
end procedure
procedure $\operatorname{QUERY}(x, S)$
$S_{1} \leftarrow S$
$y^{(1)} \leftarrow y$
$x^{\prime} \leftarrow 0 \quad \triangleright x^{\prime} \in \mathbb{R}^{n}$ records an estimation of $x$
for $i=1 \rightarrow \log k$ do
Find $T_{i} \subseteq S_{i}$ such that for each $j \in T_{i}$, it is isolated from other coordinates in
$S_{i}$
$l \leftarrow \sum_{i^{\prime}=1}^{i-1} B_{i^{\prime}}$
For each $j \in T_{i}$, compute $\widehat{x}_{j} \leftarrow y_{l+h_{i}(j)}^{(i)}$
$S_{i+1} \leftarrow S_{i} \backslash T_{i}$
$y^{(i+1)} \leftarrow y^{(i)}-\Phi \widehat{x}_{T_{i}}$
$x^{\prime} \leftarrow x^{\prime}+\widehat{x}_{T_{i}}$
end for
return $x^{\prime}$
end procedure
end procedure

```

\subsection*{14.2.3 Iterative loop analysis}

In this subsection, we characterize and prove two properties of the iterative loop in Algorithm 14.1. The first property states that with good probability, as \(i\) increases, the size of \(S_{i}\) shrinks by a factor of at least \(1-\gamma\), and otherwise its size does not increase.

This property is guaranteed by identification property of count sketch and our choice of parameters. The second property states that with good probability, the estimation error in each round is small. This is guaranteed by estimation property of count sketch and our choice of parameters. We formalize both properties in Lemma 14.2.1 as follows.

Lemma 14.2.1. Given parameters \(\epsilon \in(0,1), k \geq 1\), for each \(i \in[\log k]\), we define
\[
k_{i}=k \gamma^{i}, \epsilon_{i}=\epsilon(10 \gamma)^{i}, B_{i}=C \cdot k_{i} / \epsilon_{i} .
\]

For each \(i \in[\log k]\), if
\[
\forall j \in[i-1],\left|S_{j+1}\right| \leq \gamma\left|S_{j}\right|,
\]
and
\[
\forall j \in[i-1], \sum_{t \in T_{j}}\left|x_{t}-\widehat{x}_{t}\right|^{2} \leq \epsilon \cdot 2^{-j}\left\|x_{\bar{S}}\right\|_{2}^{2}
\]

Then we have :
Property (I) : With probability at least \(1-\frac{1}{C \gamma}(10 \gamma)^{i}\),
\[
\left|S_{i+1}\right| \leq \gamma\left|S_{i}\right|,
\]
and otherwise \(\left|S_{i+1}\right| \leq\left|S_{i}\right|\).
Property (II) : With probability at least \(1-\frac{2}{C}(20 \gamma)^{i}\),
\[
\sum_{t \in T_{i}}\left|x_{t}-\widehat{x}_{t}\right|^{2} \leq \epsilon \cdot 2^{-i}\left\|x_{\bar{S}}\right\|_{2}^{2}
\]

Proof. Proof of Property (I).

For each \(i \in[\log k]\), for each \(j \in S_{i}\), we use \(Z_{i, j}\) to denote the boolean random variable where \(Z_{i, j}=1\) if there exists a \(j^{\prime} \in S_{i} \backslash\{j\}\) such that \(h_{i}(j)=h_{i}\left(j^{\prime}\right)\) and \(Z_{i, j}=0\) otherwise. By the randomness of hash function \(h_{i}\),
\[
\underset{h_{i}}{\operatorname{Pr}}\left[Z_{i, j}=1\right]=\underset{h_{i}}{\operatorname{Pr}}\left[j \text { is not isolated from } S_{i}\right] \leq\left|S_{i}\right| / B_{i} .
\]

By the statement of the lemma, we know that \(\left|S_{i}\right| \leq \gamma^{i} k\) and \(B_{i}=C k \gamma^{i} / \epsilon_{i}\), thus
\[
\underset{h_{i}}{\operatorname{Pr}}\left[Z_{i, j}=1\right] \leq \gamma^{i} k / B_{i} \leq \epsilon_{i} / C .
\]

We define \(Z_{i}=\sum_{j \in S_{i}} Z_{i, j}\). Next, we can show
\[
\underset{h_{i}}{\mathbb{E}}\left[Z_{i}\right]=\underset{h_{i}}{\mathbb{E}}\left[\mid\left\{j \in S_{i} \mid j \text { is not isolated from } S_{i}\right\} \mid\right]=\underset{h_{i}}{\mathbb{E}}\left[\sum_{j \in S_{i}} Z_{i, j}\right]=\sum_{j \in S_{i}} \underset{h_{i}}{\mathbb{E}}\left[Z_{i, j}\right] \leq\left|S_{i}\right| \epsilon_{i} / C,
\]
where the third step follows by linearity of expectation. Therefore,
\[
\underset{h_{i}}{\operatorname{Pr}}\left[Z_{i} \geq \gamma\left|S_{i}\right|\right] \leq \underset{h_{i}}{\mathbb{E}}\left[Z_{i}\right] /\left(\gamma\left|S_{i}\right|\right) \leq\left|S_{i}\right| \epsilon_{i} /\left(C \gamma\left|S_{i}\right|\right)=\epsilon_{i} /(C \gamma)
\]
where the first step follows by Markov's inequality, and second step follows by \(\mathbb{E}_{h_{i}}\left[Z_{i}\right] \leq\) \(\left|S_{i}\right| \epsilon_{i} / C\). This implies
\[
\operatorname{Pr}\left[\left|S_{i+1}\right| \geq \gamma\left|S_{i}\right|\right] \leq \epsilon(10 \gamma)^{i} /(C \gamma) \leq \frac{1}{C \gamma}(10 \gamma)^{i}
\]

\section*{Proof of Property (II).}

We first define random variable \(V_{i}\) to be
\[
V_{i}=\sum_{l \in\left(\cup_{\left.j \in S_{i} h_{i}^{-1}(j)\right) \backslash S_{i}}\left|x_{l}^{(i)}\right|^{2} . . . . ~ . ~\right.}
\]

That is, \(x^{(i)}\) is the residual signal in the \(i\)-th round, and \(V_{i}\) refers to the noise coordinates that are hashed to the same bins as coordinates in \(S_{i}\).

We can compute the expectation of \(V_{i}\),
\[
\begin{aligned}
\underset{h_{i}}{\mathbb{E}}\left[V_{i}\right] & =\underset{h_{i}}{\mathbb{E}}\left[\sum_{l \in\left(\cup_{j \in S_{i}}^{\left.h_{i}^{-1}(j)\right)} \backslash S_{i}\right.}\left|x_{l}^{(i)}\right|^{2}\right] \\
& \leq k_{i}\left\|x_{\overline{S_{i}}}^{(i)}\right\|_{2}^{2} / B_{i} \\
& =\frac{\epsilon_{i}}{C}\left\|x_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2} \\
& =\frac{(10 \gamma)^{i} \epsilon}{C}\left\|x_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}
\end{aligned}
\]
where the first step follows by the fact that for each \(l \in \overline{S_{i}}\), the probability that \(l \in\left(\cup_{j \in S_{i}} h_{i}^{-1}(j)\right) \backslash S_{i}\) under the randomness of \(h_{i}\) is \(\left|S_{i}\right| / B_{i} \leq k_{i} / B_{i}\), the second and third step follow by definitions of \(B_{i}\) and \(\epsilon_{i}\), respectively.

By Markov's inequality, we have
\[
\begin{equation*}
\underset{h_{i}}{\operatorname{Pr}}\left[V_{i} \geq \frac{1}{2} \cdot 10^{-i} \epsilon\left\|x_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}\right] \leq 2 \frac{(100 \gamma)^{i}}{C} . \tag{14.1}
\end{equation*}
\]

In the following, we argue conditioned on the event that all elements in \(T_{i}\) are isolated under \(h_{i}\) and \(V_{i} \leq 10^{-i} \epsilon\left\|x_{\bar{S}}^{(i)}\right\|_{2}^{2} / 2\).

Recall the definition of \(x^{(i)} \in \mathbb{R}^{n}\),
\[
\begin{equation*}
x^{(i)}=x-\sum_{t \in S \backslash S_{i}} \widehat{x}_{t}=x-\sum_{j=1}^{i} \sum_{t \in T_{j}} \widehat{x}_{t} . \tag{14.2}
\end{equation*}
\]

For any fixed \(t \in T_{i}\), let \(b=h_{i}(t)\). We have
\[
\begin{align*}
\underset{\sigma_{i}}{\mathbb{E}}\left[\left|x_{t}-\widehat{x}_{t}\right|^{2}\right] & =\underset{\sigma_{i}}{\mathbb{E}}\left[\left|x_{t}-\sigma_{i}(t) \cdot \sum_{l: h_{i}(l)=b} \sigma_{i}(l) x_{l}^{(i)}\right|^{2}\right] \\
& =\underset{\sigma_{i}}{\mathbb{E}}\left[\left|\sum_{l \in h_{i}^{-1}(b) \backslash\{t\}} \sigma_{i}(l) x_{l}^{(i)}\right|^{2}\right] \\
& =\sum_{l \in h_{i}^{-1}(b) \backslash\{t\}}\left|x_{l}^{(i)}\right|_{2}^{2} \tag{14.3}
\end{align*}
\]
where the first step follows by definition of \(\widehat{x}_{t}\) in Algorithm 14.1 and construction of sketching matrix \(\Phi\) in Definition 14.2.1, the third step follows by pairwise independence property of random sign function \(\sigma_{i}\).

We can upper bound \(\mathbb{E}_{\sigma_{i}}\left[\sum_{t \in T_{i}}\left|x_{t}-\widehat{x}_{t}\right|^{2}\right]\) in the following way,
\[
\begin{array}{rlr}
\underset{\sigma_{i}}{\mathbb{E}}\left[\sum_{t \in T_{i}}\left|x_{t}-\widehat{x}_{t}\right|^{2}\right] & =\sum_{t \in T_{i}} \sum_{l \in h_{i}^{-1}(t) \backslash\{t\}}\left|x_{l}^{(i)}\right|_{2}^{2} & \text { by Eq. (14.3) } \\
& \leq \sum_{l \in \cup_{t \in S_{i}} h_{i}^{-1}(t) \backslash\{t\}}\left|x_{l}^{(i)}\right|_{2}^{2} & \text { by definition of } T_{i} \\
& \leq \epsilon 10^{-i}\left\|x_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2} / 2 & \\
& \leq \epsilon 10^{-i}\left\|x_{\bar{S}}\right\|_{2}^{2} & \text { by Eq. (14.1) }  \tag{14.4}\\
& \text { by Eq. (14.5) }
\end{array}
\]

Second, we show how to upper bound \(\left\|x_{\overline{S_{i}}}^{(i)}\right\|_{2}^{2}\),
\[
\begin{array}{rlr}
\left\|x_{\overline{S_{i}}}^{(i)}\right\|_{2}^{2} & =\left\|x_{\bar{S} \cup S \backslash S_{i}}^{(i)}\right\|_{2}^{2} & \text { by } \overline{S_{i}}=\bar{S} \cup S \backslash S_{i} \\
& =\left\|x_{\bar{S}}^{(i)}\right\|_{2}^{2}+\left\|x_{S \backslash S_{i}}^{(i)}\right\|_{2}^{2} & \text { by } \bar{S} \cap\left(S \backslash S_{i}\right)=\emptyset \\
& =\left\|x_{\bar{S}}\right\|_{2}^{2}+\left\|x_{S \backslash S_{i}}^{(i)}\right\|_{2}^{2} & \text { by } x_{\bar{S}}^{(i)}=x_{\bar{S}} \\
& =\left\|x_{\bar{S}}\right\|_{2}^{2}+\sum_{j=1}^{i-1}\left\|x_{T_{j}}^{(i)}\right\|_{2}^{2} & \text { by } S \backslash S_{i}=T_{1} \cup T_{2} \cup \cdots \cup T_{i-1} \\
& =\left\|x_{\bar{S}}\right\|_{2}^{2}+\sum_{j=1}^{i-1} \sum_{t \in T_{j}}\left|x_{t}-\widehat{x}_{t}\right|^{2} & \text { by Eq. (14.2) } \\
& \leq\left\|x_{\bar{S}}\right\|_{2}^{2}+\sum_{j=1}^{i-1} \epsilon 2^{-j}\left\|x_{\bar{S}}\right\|_{2}^{2} & \text { by statement of Lemma } \\
& \leq 2\left\|x_{\bar{S}}\right\|_{2}^{2} . & \text { by geometric sum } \tag{14.5}
\end{array}
\]

Using Markov's inequality \(\operatorname{Pr}[X \geq a] \leq \mathbb{E}[X] / a\), we have
\[
\operatorname{Pr}\left[\sum_{t \in T_{i}}\left|x_{t}-\widehat{x}_{t}\right|^{2} \geq \epsilon 2^{-i}\left\|x_{\bar{S}}\right\|_{2}^{2}\right] \leq 5^{-i}
\]

Therefore, we complete the proof.

\subsection*{14.2.4 Main result}

In this section, we prove the correctness of our iterative set query algorithm, as well as justify the claimed space complexity, update time and query time. At the core of correctness is an analysis of the iterative loop, which is given as Lemma 14.2.1 in Section 14.2.3.

Theorem 14.2.2. Let \(m=O(k / \epsilon)\). We can give a randomized construction of a matrix \(\Phi \in \mathbb{R}^{m \times n}\) (Definition 14.2.1) such that for each set \(S \subseteq[n],|S| \leq k\), and for each vector
\(x \in \mathbb{R}^{n}\), there is an algorithm (Algorithm 14.1) that takes \(\Phi x\) as its input, and finds \(x^{\prime} \in \mathbb{R}^{k}\) such that
\[
\left\|x^{\prime}-x_{S}\right\|_{2}^{2} \leq \epsilon\left\|x_{[n] \backslash S}\right\|_{2}^{2}
\]
holds with probability \(9 / 10\). The space to store matrix \(\Phi\) is \(O(\log k)\) words. The columns of matrix \(\Phi\) are 2-wise independent. The update time is \(O(\log k)\). The decoding time is \(O(k \log k)\).

Proof. Our proof is based on properties of the iterative loop in Algorithm 14.1. The properties are formulated as Lemma 14.2.1 in Section 14.2.3. In summary, property (I) in Lemma 14.2 .1 shows that with good probability, as \(i\) increases, the size of \(S_{i}\) (i.e. the set of coordinates in \(S\) that yet to be recovered at the beginning of \(i\)-th round) shrinks by a constant factor, and otherwise its size does not increase. Property (II) in Lemma 14.2.1 shows that with good probability, the estimation error in each round is small.

Proof of success probability. It follows by recursively applying Lemma 14.2.1. By union bound, the failure probability is bounded by
\[
\sum_{i=1}^{\log k}\left(\frac{1}{C \gamma}(10 \gamma)^{i}+\frac{2}{C}(20 \gamma)^{i}\right)<1 / 10
\]
which follows from geometric sum and the choice of parameters.

\section*{Proof of estimation error.}
\[
\left\|x^{\prime}-x_{S}\right\|_{2}^{2} \leq \sum_{i=1}^{\log k} \sum_{t \in T_{i}}\left|x_{t}-\widehat{x}_{t}\right|^{2} \leq \sum_{i=1}^{\log k} \epsilon \cdot 2^{-i}\left\|x_{\bar{S}}\right\|_{2}^{2} \leq \epsilon\left\|x_{\bar{S}}\right\|_{2}^{2}
\]

Proof of space to store \(\Phi\). Since for each \(i \in[\log k], \Phi^{(i)}\) only uses 2 pairwise independent hash functions, the space to store matrix \(\Phi\) is \(O(\log k)\) words.

Proof of update time. Upon update \((j, \Delta)\), for each \(\Phi^{(i)}\), the algorithm only needs to update one measurement (located at \(\left.h_{i}(j)\right)\). Therefore, the update time is \(O(\log k)\). Since every element of the matrix can be computed in \(O(1)\) time.

Proof of query time. First of all, we note that the step in Line 18 can be performed in \(O\left(\left|S_{i}\right|\right)\) time upon query, since we can check every \(j \in S_{i}\), and mark the buckets that have more than one \(j\) hashed to them; in the next step we just keep only those coordinates \(j\) that are hashed to exactly one bucket. Recursively applying Lemma 14.2.1, we have with constant probability, such that for all \(i \in[\log k],\left|S_{i+1}\right| \leq \gamma\left|S_{i}\right|\). It implies that for all \(i \in[\log k]\), \(\left|S_{i}\right| \leq \gamma^{i-1}\left|S_{1}\right| \leq \gamma^{i-1} k\). Note that in iteration \(i \in[\log k]\), we need \(O\left(\left|S_{i}\right|\right)\) update time. Therefore, overall the iterations, the time of this part is
\[
\sum_{i=1}^{\log k} O\left(\left|S_{i}\right|\right) \leq O\left(\sum_{i=1}^{\log k} \gamma^{i-1} k\right)=O(k)
\]
where the last step follows by the choice of \(\gamma\).
We know that \(\|z\|_{0}=O(k)\) and \(\left|T_{i}\right|=O(k)\). Thus, in each iteration, the time of doing step \(y^{(i+1)} \leftarrow y^{(i)}-\Phi Z_{T_{i}}\) requires \(O(k)\) time. Thus the subtraction time over all iterations is \(O(k \log k)\), since there are \(O(\log k)\) iterations. Putting it together, we get \(O(k \log k)\) query time.

\subsection*{14.3 Fourier set query}

We present some definitions and backgrounds of Fourier transform in Section 14.3.1. Then we introduce spectrum permutations and filter functions in Section 14.3.2. They are used as hashing schemes in Fourier transform literatures. In Section 14.3.3, we introduce collision events. large offset events, and large noise events. We give a new algorithm and provide iterative loop analysis in Section 14.3.4. Finally, we present main result in Section 14.3.5.

\subsection*{14.3.1 Definitions and some backgrounds of Fourier transform}

We use \(\mathbf{i}\) to denote \(\sqrt{-1}\). Note that \(e^{\mathbf{i} \theta}=\cos (\theta)+\mathbf{i} \sin (\theta)\). For any complex number \(z \in \mathbb{C}\), we have \(z=a+\mathbf{i} b\), where \(a, b \in \mathbb{R}\). We define the complement of \(z\) as \(\bar{z}=a-\mathbf{i} b\). We define \(|z|=\sqrt{z \bar{z}}=\sqrt{a^{2}+b^{2}}\). For any complex vector \(x \in \mathbb{C}^{n}\), we use \(\operatorname{supp}(x)\) to denote the support of \(x\), and then \(\|x\|_{0}=|\operatorname{supp}(x)|\).

The discrete convolution of functions \(f\) and \(g\) is given by,
\[
(f * g)[n]=\sum_{m=-\infty}^{+\infty} f[m] g[n-m]
\]

For a complex vector \(x \in \mathbb{C}^{n}\), we use \(\widehat{x} \in \mathbb{C}^{n}\) to denote its Fourier spectrum,
\[
\widehat{x}_{i}=\frac{1}{\sqrt{n}} \sum_{j=1}^{n} e^{-2 \pi \mathbf{i} i j / n} x_{j}, \forall i \in[n] .
\]

Then the inverse transform is
\[
x_{j}=\frac{1}{\sqrt{n}} \sum_{i=1}^{n} e^{2 \pi \mathbf{i} i j / n} \widehat{x}_{i}, \forall j \in[n] .
\]

We define
\[
\operatorname{Err}(x, k)=\min _{k-\text { sparse } y}\|x-y\|_{2}
\]

\subsection*{14.3.2 Permutation and filter function}

We use the same (pseudorandom) spectrum permutation as [HIKP12a],

Definition 14.3.1. Suppose \(\sigma^{-1}\) exists \(\bmod n\). We define the permutation \(P_{\sigma, a, b}\) by
\[
\left(P_{\sigma, a, b} x\right)_{i}=x_{\sigma(i-a)} e^{-2 \pi \mathbf{i} \sigma b i / n}
\]

We also define \(\pi_{\sigma, b}=\sigma(i-b)(\bmod n)\). Then we have
Claim 14.3.1. \({\widehat{P_{\sigma, a, b} x}}_{\pi_{\sigma, b}(i)}=\widehat{x}_{i} e^{-2 \pi \mathrm{i} \sigma a i / n}\).
\(h_{\sigma, b}(i)=\operatorname{round}\left(\pi_{\sigma, b}(i) B / n\right)\) and \(o_{\sigma, b}(i)=\pi_{\sigma, b}(i)-h_{\sigma, b}(i) n / B\). We say \(h_{\sigma, b}(i)\) is the "bin" that frequency \(i\) is mapped into, and \(o_{\sigma, b}(i)\) is the "offset".

We use the same filter function as [HIKP12a, PS15, CKPS16],
Definition 14.3.2. Given parameters \(B \geq 1, \delta>0, \alpha>0\). We say that \(\left(G, \widehat{G}^{\prime}\right)=\) \(\left(G_{B, \delta, \alpha}, \widehat{G}_{B, \delta, \alpha}^{\prime}\right) \in \mathbb{R}^{n}\) is a filter function(Figure 14.1) if it satisfies the following properties
(I) \(|\operatorname{supp}(G)|=O\left(\alpha^{-1} B \log (n / \delta)\right)\),
(I) \(\widehat{G}_{i}^{\prime}=1\),
(II) \(\widehat{G}_{i}^{\prime}=0\),
\[
\text { if }|i| \leq(1-\alpha) n /(2 B)
\]
if \(|i| \geq n /(2 B)\),
(III) \(\widehat{G}_{i}^{\prime} \in[0,1], \quad\) for all \(i\),
(IV) \(\left\|\widehat{G}^{\prime}-\widehat{G}\right\|_{\infty}<\infty\).


Figure 14.1: Filter \(\widehat{G}^{\prime}\)

\subsection*{14.3.3 Collision event, large offset event, and large noise event}

We use three types of events defined in [HIKP12a] as basic building blocks for analyzing Fourier set query algorithms. For any \(i \in S\), we define three types of events associated with \(i\) and \(S\) and defined over the probability space induced by \(\sigma\) and \(b\) :

Definition 14.3.3 (Collision, large offset, large noise). "Collision" event \(E_{\text {coll }}(i)\) : holds iff \(h_{\sigma, b}(i) \in h_{\sigma, b}(S \backslash\{i\})\).
"Large offset" event \(E_{\text {off }}(i)\) : holds iff \(\left|o_{\sigma, b}(i)\right| \geq(1-\alpha) n /(2 B)\).
"Large noise" event \(E_{\text {noise }}(i)\) : holds iff
\[
\mathbb{E}\left[\left\|\widehat{x}_{h_{\sigma, b}^{-1}\left(h_{\sigma, b}(i)\right) \backslash S}^{\prime}\right\|_{2}^{2}\right] \geq \operatorname{Err}^{2}\left(\widehat{x}^{\prime}, k\right) /(\alpha B)
\]

Claim 14.3.2 (Claim 3.1 in [HIKP12a]). For any \(i \in S\), the event \(E_{\text {coll }}(i)\) holds with probability at most \(4|S| / B\).

Claim 14.3.3 (Claim 3.2 in [HIKP12a]). For any \(i \in S\), the event \(E_{\text {off }}(i)\) holds with probability at most \(\alpha\).

Claim 14.3.4 (Claim 4.1 in [HIKP12a]). For any \(i \in S, \operatorname{Pr}\left[E_{\text {noise }(i)}\right] \leq 4 \alpha\).

Lemma 14.3.5 (Lemma 4.2 in [HIKP12a]). Let \(a \in[n]\) uniformly at random, \(B\) divide \(n\), and the other parameters be arbitrary in
\[
\widehat{u}=\operatorname{HashToBins}\left(x, \widehat{z}, P_{\sigma, a, b}, B, \delta, \alpha\right) .
\]

Then for any \(i \in[n]\) with \(j=h_{\sigma, b}(i)\) and none of \(E_{\text {coll }}(i), E_{\text {off }}(i)\) or \(E_{\text {noise }}(i)\) holding,
\[
\mathbb{E}\left[\left|\widehat{u}_{j}-\widehat{x}_{i}^{\prime} e^{-\frac{2 \pi \mathrm{i}}{n} a \sigma i}\right|^{2}\right] \leq 2 \frac{\rho^{2}}{\alpha B} .
\]

Lemma 14.3.6. Suppose \(B\) divides \(n\). The output \(\widehat{u}\) of HAShToBins satisfies
\[
\left.\widehat{u}_{j}=\sum_{h_{\sigma, b}(i)=j} \widehat{(x-z}\right)_{i}\left(\widehat{G_{B, \delta, \alpha}^{\prime}}\right)_{-o_{\sigma, b}(i)} \omega^{a \sigma i} \pm \delta\|\widehat{x}\|_{1} .
\]

Let \(\zeta=\left|\left\{i \in \operatorname{supp}(\widehat{z}) \mid E_{\text {off }}(i)\right\}\right|\). The running time of HASHToBins is \(O\left(\frac{B}{\alpha} \log (n / \delta)+\right.\) \(\left.\|\widehat{z}\|_{0}+\zeta \log (n / \delta)\right)\).

\subsection*{14.3.4 Iterative loop analysis}

Iterative loop analysis for Fourier set query is more tricky than the classic set query, because in the Fourier case, hashing is not perfect, in the sense that by using spectrum permutation and filter function (as the counterpart of hashing techniques), one coordinate can non-trivially contribute to multiple bins. We give iterative loop induction in Lemma 14.3.7.

Lemma 14.3.7. Given parameters \(C \geq 1000, \gamma \leq 1 / 1000\). For any \(k \geq 1, \epsilon \in(0,1), R \geq 1\). For each \(i \in[R]\), we define
\[
k_{i}=k \gamma^{i}, \epsilon_{i}=\epsilon(10 \gamma)^{i}, \alpha_{i}=1 /\left(200 i^{3}\right), B_{i}=C \cdot k_{i} /\left(\alpha_{i}^{2} \epsilon_{i}\right) .
\]

For each \(i \in[R]\) :
If for all \(j \leq[i-1]\) we have
\[
\operatorname{supp}\left(\widehat{w}^{(j)}\right) \subseteq S_{j}, \quad\left|S_{j+1}\right| \leq k_{j+1}, \quad \widehat{z}^{(j+1)}=\widehat{z}^{(j)}+\widehat{w}^{(j)}, \quad \widehat{x}^{(j+1)}=\widehat{x}-\widehat{z}^{(j+1)},
\]
and
\[
\left\|\widehat{x}_{\bar{S}_{j+1}}^{(j+1)}\right\|_{2}^{2} \leq\left(1+\epsilon_{j}\right)\left\|\widehat{x}_{\bar{S}_{j}}^{(j)}\right\|_{2}^{2}+\epsilon_{j} \delta^{2} n\|\widehat{x}\|_{1}^{2} .
\]

Then, with probability \(1-10 \alpha_{i} / \gamma\), we have
\[
\operatorname{supp}\left(\widehat{w}^{(i)}\right) \subseteq S_{i}, \quad\left|S_{i+1}\right| \leq k_{i+1}, \quad \widehat{z}^{(i+1)}=\widehat{z}^{(i)}+\widehat{w}^{(i)}, \quad \widehat{x}^{(i+1)}=\widehat{x}-\widehat{z}^{(i+1)}
\]
and
\[
\left\|\widehat{x}_{\bar{S}_{i+1}}^{(i+1)}\right\|_{2}^{2} \leq\left(1+\epsilon_{i}\right)\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\epsilon_{i} \delta^{2} n\|\widehat{x}\|_{1}^{2} .
\]

Proof. We consider a particular step \(i\). We can condition on \(\left|S_{i}\right| \leq k_{i}\).
Collision. Using Claim 14.3.2, for any \(t \in S_{i}\), the event \(E_{\text {coll }}(t)\) holds with probability at most
\[
4\left|S_{i}\right| / B_{i} \leq \frac{4 k_{i}}{C k_{i} /\left(\alpha_{i}^{2} \epsilon_{i}\right)}=4 \alpha_{i}^{2} \epsilon_{i} / C \leq \alpha_{i} .
\]

It means
\[
\operatorname{Pr}_{\sigma, b}\left[E_{\text {coll }}(t)\right] \leq \alpha_{i} .
\]

Large offset. Using Claim 14.3.3, for any \(t \in S_{i}\), the event \(E_{\text {off }}(t)\) holds with probability at most \(\alpha_{i}\), i.e.
\[
\underset{\sigma, b}{\operatorname{Pr}}\left[E_{\text {off }}(t)\right] \leq \alpha_{i}
\]

Large noise. Using Claim 14.3.4, for any \(t \in S_{i}\),
\[
\underset{\sigma, b}{\operatorname{Pr}}\left[E_{\text {noise }}(t)\right] \leq 4 \alpha_{i}
\]

We say a coordinate \(t\) is "well isolated" if none of the above three events holding. By a union bound over the above three events, we have \(t\) is "well isolated" with probability at least \(1-6 \alpha_{i}\).

Therefore, each \(t \in S_{i}\) lies in \(T_{i}\) with probability at least \(1-6 \alpha_{i}\). Then by Markov's inequality, we have \(\left|S_{i} \backslash T_{i}\right| \leq \gamma k_{i}\) with probability \(1-6 \alpha_{i} / \gamma\). By definition \(S_{i+1}=S_{i} \backslash T_{i}\), then we know that
\[
\left|S_{i+1}\right|=\left|S_{i} \backslash T_{i}\right| \leq \gamma k_{i} \leq k_{i+1}
\]

Upper bound \(\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}\). By Lemma statement we have
\[
\begin{align*}
\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2} & \leq\left(1+\epsilon_{i}\right)\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\epsilon_{i} \delta^{2} n\|\widehat{x}\|_{1}^{2} \\
& \leq\left(1+\epsilon_{i}\right)\left(1+\epsilon_{i-1}\right)\left\|\widehat{x}_{\bar{S}_{i-1}}^{(i-1)}\right\|_{2}^{2}+\left(\left(1+\epsilon_{i}\right) \epsilon_{i-1}+\epsilon_{i}\right) \delta^{2} n\|\widehat{x}\|_{1}^{2} \\
& \leq \cdots \\
& \leq \prod_{j=1}^{i}\left(1+\epsilon_{j}\right)\left\|\widehat{x}_{\bar{S}}\right\|_{2}^{2}+\sum_{j=1}^{i} \epsilon_{j} \delta^{2} n\|\widehat{x}\|_{1}^{2} \prod_{l=j+1}^{i}\left(1+\epsilon_{l}\right) \\
& \leq 8\left(\left\|\widehat{x}_{\bar{S}}\right\|_{2}^{2}+\delta^{2} n\|\widehat{x}\|_{1}^{2}\right) \tag{14.6}
\end{align*}
\]
where the third step refers to recursively apply the second step, the last step follows by a geometric sum.

For a fixed \(t \in S_{i}\), let \(j=h_{\sigma, b}(t)\). By Lemma 3.3, we have
\[
\widehat{u}_{j}-\widehat{x}_{t}^{(i)} \omega^{a \sigma t}=\sum_{t^{\prime} \in T_{i}} \widehat{G}_{-o_{\sigma}\left(i^{\prime}\right)}^{\prime} \widehat{x}_{t^{\prime}}^{(i)} \omega^{a \sigma t^{\prime}} \pm \delta\|\widehat{x}\|_{1}
\]

For each \(t \in S_{i}\), we define set \(Q_{i, t}=h_{\sigma, b}^{-1}(j) \backslash\{t\}\). Let \(T_{i}\) be the set of coordinates \(t \in S_{i}\) such that \(Q_{i, t} \cap S_{i}=\emptyset\). Then it is easy to observe that
\[
\begin{aligned}
\sum_{t \in T_{i}}\left|\sum_{t^{\prime} \in Q_{i, t}} \widehat{G}_{-o_{\sigma}(t)}^{\prime} \widehat{x}_{t^{\prime}}^{(i)} \omega^{a \sigma t^{\prime}}\right|^{2} & =\sum_{t \in T}\left|\sum_{t^{\prime} \in Q_{i, t} \backslash S} \widehat{G}_{-o_{\sigma}\left(t^{\prime}\right)}^{\prime} \widehat{x}_{t^{\prime}}^{(i)} \omega^{a \sigma t^{\prime}}\right|^{2} \\
& \leq \sum_{t \in S}\left|\sum_{t^{\prime} \in Q_{i, t} \backslash S} \widehat{G}_{-o_{\sigma}\left(t^{\prime}\right)}^{\prime} \widehat{x}_{t^{\prime}}^{(i)} \omega^{a \sigma t^{\prime}}\right|^{2}
\end{aligned}
\]

We can calculate the expectation of \(\left\|\widehat{x}_{T_{i}}^{(i)}-\widehat{w}^{(i)}\right\|_{2}^{2}\),
\[
\begin{aligned}
\underset{\sigma, a, b}{\mathbb{E}}\left[\left\|\widehat{x}_{T_{i}}^{(i)}-\widehat{w}^{(i)}\right\|_{2}^{2}\right] & =\underset{\sigma, a, b}{\mathbb{E}}\left[\sum_{t \in T_{i}}\left\|\widehat{x}_{t}^{(i)}-\widehat{w}_{t}^{(i)}\right\|_{2}^{2}\right] \\
& \leq \sum_{i \in S} 2 \underset{\sigma, a, b}{\mathbb{E}}\left[\left|\sum_{i^{\prime} \in Q_{i} \backslash S} \widehat{G}_{-o_{\sigma}(i)}^{\prime} \omega^{a \sigma i}\right|^{2}\right]+\delta^{2}\|\widehat{x}\|_{1}^{2} \\
& \leq \sum_{t \in S_{i}}\left(\frac{1}{B_{i}}\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\delta^{2}\|\widehat{x}\|_{1}^{2}\right) \\
& \leq \frac{\left|S_{i}\right|}{B_{i}}\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\delta^{2}\left|S_{i}\right| \cdot\|\widehat{x}\|_{1}^{2} \\
& \leq \frac{\epsilon_{i} \alpha_{i}^{2}}{C}\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\delta^{2}\left|S_{i}\right| \cdot\|\widehat{x}\|_{1}^{2}
\end{aligned}
\]
where the last step follows from \(\left|S_{i}\right| \leq k_{i}\) and \(B_{i}=C \cdot k_{i} /\left(\alpha_{i}^{2} \epsilon_{i}\right)\).

Then, using Markov's inequality, we have,
\[
\operatorname{Pr}\left[\left\|\widehat{x}_{T_{i}}^{(i)}-\widehat{w}^{(i)}\right\|_{2}^{2} \geq \frac{\epsilon_{i} \alpha_{i}}{C}\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\delta^{2} \frac{\left|S_{i}\right|}{\alpha_{i}}\|\widehat{x}\|_{1}^{2}\right] \leq \alpha_{i} .
\]

Note that
\[
\begin{aligned}
\frac{\epsilon_{i} \alpha_{i}}{C}\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\delta^{2} \frac{\left|S_{i}\right|}{\alpha_{i}}\|\widehat{x}\|_{1}^{2} & \leq \frac{\epsilon_{i}}{C}\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\delta^{2} \frac{\left|S_{i}\right|}{\alpha_{i}}\|\widehat{x}\|_{1}^{2} \\
& \leq \frac{\epsilon_{i}}{C}\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\frac{\epsilon_{i}}{C} \delta^{2} B_{i}\|\widehat{x}\|_{1}^{2} \\
& \leq \frac{\epsilon_{i}}{C}\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\frac{\epsilon_{i}}{C} \delta^{2} n\|\widehat{x}\|_{1}^{2} \\
& \leq \frac{\epsilon_{i}}{20}\left(\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\delta^{2} n\|\widehat{x}\|_{1}^{2}\right)
\end{aligned}
\]
where the first step follows by \(\alpha_{i} \leq 1\), the second step follows by \(\left|S_{i}\right| \leq k_{i}=\epsilon_{i} B_{i} \alpha_{i}^{2} / C\), the third step follows by \(B_{i} \leq n\), the last step follows by \(C \geq 1000\).

Thus, we have
\[
\begin{equation*}
\operatorname{Pr}\left[\left\|\widehat{x}_{T_{i}}^{(i)}-\widehat{w}^{(i)}\right\|_{2}^{2} \leq \frac{\epsilon_{i}}{20}\left(\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\delta^{2} n\|\widehat{x}\|_{1}^{2}\right)\right] \geq 1-\alpha_{i} . \tag{14.7}
\end{equation*}
\]

Recall that \(\widehat{w}^{(i)}=\widehat{z}^{(i+1)}-\widehat{z}^{(i)}=\widehat{x}^{(i)}-\widehat{x}^{(i+1)}\). It is obvious that \(\operatorname{supp}\left(\widehat{w}^{(i)}\right) \subseteq T_{i}\).

Conditioning on all coordinates in \(T_{i}\) are well isolated and Eq. (14.7) holds, we have
\[
\begin{array}{rlr}
\left\|\widehat{x}_{\bar{S}_{i+1}}^{(i+1)}\right\|_{2}^{2} & =\left\|\left(\widehat{x}^{(i)}-\widehat{w}^{(i)}\right)_{\bar{S}_{i+1}}\right\|_{2}^{2} & \\
& =\left\|\widehat{x}_{\bar{S}_{i+1}}^{(i)}-\widehat{w}_{\bar{S}_{i+1}}^{(i)}\right\|_{2}^{2} & \\
& =\left\|\widehat{x}_{\bar{S}_{i+1}}^{(i)}-\widehat{w}^{(i)}\right\|_{2}^{2} & \\
& =\left\|\widehat{x}_{\bar{S}_{i} \cup T_{i}}^{(i)}-\widehat{w}^{(i)}\right\|_{2}^{2} & \text { by } S_{i}=T_{i} \cup S_{i+1} \\
& =\left\|\widehat{x}_{\bar{S}_{\bar{S}}}^{(i)}\right\|_{2}^{2}+\left\|\widehat{x}_{T_{i}}^{(i)}-\widehat{w}^{(i)}\right\|_{2}^{2} & \text { by } \bar{S}_{i} \cap T_{i}=\emptyset \\
& \leq\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\epsilon_{i}\left(\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\delta^{2} n\|\widehat{x}\|_{1}^{2}\right) & \text { by Eq. }(14.7) \\
& =\left(1+\epsilon_{i}\right)\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\epsilon_{i} \delta^{2} n\|\widehat{x}\|_{1}^{2} . & \text { by definition } \mu_{(i)}
\end{array}
\]

\subsection*{14.3.5 Main result}

In this subsection, we give the main result as the following theorem.
Theorem 14.3.8 (Main result). Given a vector \(x \in \mathbb{C}^{n}\), for every \(\epsilon \in(0,1)\) and \(k \geq 1\), any \(S \subseteq[n],|S|=k\), there exists an algorithm (Algorithm 14.2) that takes
\[
O\left(\epsilon^{-1} k \log (n / \delta)\right)
\]
samples, runs in
\[
O\left(\epsilon^{-1} k \log (n / \delta)\right)
\]
time, and outputs a vector \(x^{\prime} \in \mathbb{C}^{n}\) such that
\[
\left\|\left(x^{\prime}-\widehat{x}\right)_{S}\right\|_{2}^{2} \leq \epsilon\left\|\widehat{x}_{[n] \backslash S}\right\|_{2}^{2}+\delta\|\widehat{x}\|_{1}^{2}
\]
holds with probability at least 9/10.

Proof. Proof of Sample Complexity. The sample complexity of Estimation is
\[
\sum_{i=1}^{R}\left(B_{i} / \alpha_{i}\right) \log (n / \delta)=\epsilon^{-1} k \log (n / \delta)
\]

\section*{Proof of Success Probability.}

The failure probability is
\[
\sum_{i=1}^{R} 10 \alpha_{i} / \gamma<1 / 10
\]

\section*{Proof of Final Error}
\[
\begin{array}{rlr}
\left\|\widehat{x}_{S}-\widehat{z}^{(R+1)}\right\|_{2}^{2} & =\sum_{i=1}^{R}\left\|\widehat{x}_{T_{i}}^{(i)}-\widehat{w}^{(i)}\right\|_{2}^{2} & \\
& \leq \sum_{i=1}^{R} k_{i} \mu^{(i)} / 20 & \\
& \leq \sum_{i=1}^{R} \epsilon_{i}\left(\left\|\widehat{x}_{\bar{S}_{i}}^{(i)}\right\|_{2}^{2}+\delta^{2} n\|\widehat{x}\|_{1}^{2}\right) / 20 & \text { by Eq. (14.7) } \\
& \leq \sum_{i=1}^{R} \epsilon_{i} \cdot 10\left(\left\|\widehat{x}_{\bar{S}}\right\|_{2}^{2}+\delta^{2} n\|\widehat{x}\|_{1}^{2}\right) / 20 & \text { by Definition of } \mu^{(i)}  \tag{14.6}\\
& \leq \epsilon\left(\left\|\widehat{x}_{\bar{S}}\right\|_{2}^{2}+\delta^{2} n\|\widehat{x}\|_{1}^{2}\right) . & \text { by Eq. (14.6) }
\end{array}
\]
```

Algorithm 14.2 Fourier Set Query Algorithm
procedure FourierSetQuery $(x, S, \epsilon, k)$
$\triangleright$ Theorem 14.3.8
$\gamma \leftarrow 1 / 1000, C \leftarrow 1000, \widehat{z}^{(1)} \leftarrow 0, S_{1} \leftarrow S$
for $i=1 \rightarrow R$ do
$k_{i} \leftarrow k \gamma^{i}, \epsilon_{i} \leftarrow \epsilon(10 \gamma)^{i}, \alpha_{i} \leftarrow 1 /\left(100 i^{3}\right), B_{i} \leftarrow C \cdot k_{i} /\left(\alpha_{i}^{2} \epsilon_{i}\right)$
$\widehat{w}^{(i)}, T_{i} \leftarrow \operatorname{EstiamteValues}\left(x, \widehat{z}^{(i)}, S_{i}, B_{i}, \delta, \alpha_{i}\right) \quad \triangleright \widehat{w}^{(i)}$ is $\left|T_{i}\right|$-sparse
$S_{i+1} \leftarrow S_{i} \backslash T_{i}$
$\widehat{z}^{(i+1)} \leftarrow z^{(i)}+\widehat{w}^{(i)}$
end for
return $\widehat{z}^{(R+1)}$
end procedure
procedure EstimateValues $(x, \widehat{z}, S, B, \delta, \alpha) \quad \triangleright$ Lemma 14.3.7
Choose $a, b \in[n]$ uniformly at random
Choose $\sigma$ uniformly at random from the set of odd numbers in $[n]$
$\widehat{u} \leftarrow \operatorname{Hash} T o B i n s\left(x, \widehat{z}, P_{\sigma, a, b}, B, \delta, \alpha\right)$
$\widehat{w} \leftarrow 0, T \leftarrow \emptyset$
for $t \in S$ do
if $t$ is isolated from other coordinates of $S$ then $\quad \triangleright h_{\sigma, b}(t) \notin h_{\sigma, b}(S \backslash\{t\})$
if no large offset then
$\widehat{w}_{t} \leftarrow \widehat{u}_{h_{\sigma, b}(t)} e^{\frac{2 \pi \mathrm{i}}{n} \sigma a t}$
$T \leftarrow T \cup\{t\}$
end if
end if
end for
return $\widehat{w}, T$
end procedure
procedure $\operatorname{HAShToBins}\left(x, \widehat{z}, P_{\sigma, a, b}, B, \delta, \alpha\right)$
Compute $\widehat{y}_{j n / B}$ for $j \in[B]$, where $y=G_{B, \alpha, \delta} \cdot\left(P_{\sigma, a, b} x\right)$
Compute $\widehat{y}_{j n / B}^{\prime}=\widehat{y}_{j n / B}-\left(\widehat{G_{B, \alpha, \delta}^{\prime}} * \widehat{P_{\sigma, a, b} z}\right)_{j n / B}$
return $\widehat{u}_{j}=\widehat{y}_{j n / B}^{\prime}$
end procedure

```

\section*{Chapter 15}

\section*{Robust Sparse Recovery via M-estimators}

We consider the extensively studied problem of sparse recovery, where the error is measured with respect to \(M\)-estimators, such as the Huber or the Tukey loss functions. All previous works have studied the problem under the \(\ell_{p}\) norms. Thus, our work is the first to explore the more robust error measures of \(M\)-estimators, which are very popular in statistics. In specific, given linear measurements \(y=A x\), we want to output a \(k\)-sparse vector \(x^{\prime} \in \mathbb{R}^{n}\) such that
\[
\left\|x^{\prime}-x\right\|_{M} \leq(1+\epsilon) \min _{k-\text { sparse } y}\|y-x\|_{M}
\]
where \(\|x\|_{M}\) is corresponds to some \(M\)-estimator of \(x\). We show that the measurement complexity of different versions of the sparse recovery problem under these more sophisticated error measures does not change, igniting thus an interesting new line of research.

\subsection*{15.1 Introduction}

The mathematical framework of compressed sensing, or sparse recovery, aims to reconstruct an approximately \(k\)-sparse vector \(x \in \mathbb{R}^{n}\) from linear measurements \(y=A x\). This model can capture many significant problems, which appear in a variety of fields, such as Discrete Signal Processing[LDP07b], Machine Learning [LDW+18], Statistics [Rip04] etc. Depending on the application, the measurement matrix \(A\) is drawn from some random matrix ensemble, or can be designed without constraints.

What is important and interesting about this framework, is that it enables us to bypass the Shannon/Nyquist sampling theorem in the case of \(x\) being \(k\)-sparse or approximately \(k\)-sparse. The Shannon/Nyquist theorem tells us that in order to preserve information when uniformly sampling a signal, we must sample at least two times faster than its bandwidth, thus putting a lower bound on the number of measurements we need to obtain.

The applications of sparse recovery to Machine Learning are numerous, see [CDHB09, VGT12, DPFJ16, CMM11, SH11]. Moreover, deep learning approaches combined with sparse recovery techniques in [MPB15, ABEZ16, \(\mathrm{KLT}^{+} 16, \mathrm{LDW}^{+} 18\), ZLSD18, BJPD17, BPD18] are proposed, in order to more finely reconstruct signals that are approximately \(k\)-sparse, with focus on images. In [VGT12], the authors employ compressed sensing techniques to develop routines for brain mapping situations in neuroscience, where the goal is to solve a support-recovery type of problem for functional MRI (fMRI) data. In [DPFJ16], the authors extend compressed sensing to signals that are represented via a graphical model.

On a high-level, the reason why machine learning applications benefit from sparse recovery is because in such applications we aim to find a mapping between sensory data and
data semantics (the label). In order to do that, a standard approach is to exploit a set of features which captures important data characteristics, and is a closely associated to data labels. We thus transform the raw data into a feature space, where the similarity between data points corresponds to the similarity between features.

The work on compressed sensing is vast, initiated with the seminal paper of Candes and Tao [CRT06b]. The authors show that \(\ell_{1}\) minimization satisfies the so-called \(\ell_{2} / \ell_{1}\) guarantee, when the matrix \(A\) is chosen from a Gaussian ensemble. Other important guarantees, such as \(\ell_{1} / \ell_{1}\) and \(\ell_{2} / \ell_{2}\) have been extensively studied. The authors in [BIR08b, IR08b] showed that \(O\left(\epsilon^{-1} k \log (n / k)\right)\) measurements suffice for a uniform \((1+\epsilon) \ell_{1} / \ell_{1}\) guarantee. It is known that uniform reconstruction (i.e. deterministic algorithms) is not possible under \(\ell_{2} / \ell_{2}\), so one has to resort to randomized algorithm. In [PW11] Price and Woodruff give a variety of interesting results for the randomized cases of \(\ell_{2} / \ell_{2}\) and \(\ell_{1} / \ell_{1}\) sparse recovery, where they also show a difference between the exactly \(k\)-sparse and the \(O(k)\)-sparse case. More precisely, they show that if one is allowed to output a \(2 k\)-sparse vector, the dependence on \(\epsilon\) is strictly better than if one is allowed to output an exactly \(k\)-sparse vector. This fact, apart from being an interesting phenomenon, is useful in applications such as imaging, where we are not constrained to output an exactly \(k\)-spare vector. The state of the art for \(\ell_{2} / \ell_{2}\) is [GLPS10], where the authors show that \(O((k / \epsilon) \log (n / k))\) measurements are sufficient, and also give a sub-linear time recovery procedure.

All previous work was focused on \(\ell_{p} / \ell_{p}\), where the error is measured with respect to some \(\ell_{p}\) norm. In applications, sometimes these error measurements might be insufficient. For example, in the presence of outliers the strongest of the above \(\ell_{2} / \ell_{2}\) guarantees, behaves poorly. For that, the Huber norm is introduced [Hub64], which combines an \(\ell_{2}\) norm for
small \(x\) with an \(\ell_{1}\) norm for large \(x\). As mentioned in [CKY97, GS \({ }^{+} 99\) ] the Huber norm is of particular interest, because it is "recommended for almost all situations", and because it is "most robust" with respect to the presence of outliers and to the number of iterations than the \(\ell_{2}\) norm. The aforementioned facts, along with its smoothness properties, have made the Huber norm useful in Machine learning applications. Moreover, while some measures, such as \(\ell_{1}\), treat small residuals with the same importance as large residuals, it is often more appropriate to use Huber norm, since relatively small noise insensitive misfit measures can yield far more stable estimates in several cases, such as in geophysics applications [GS \({ }^{+} 99\) ].

Recently, [SK18] studies a formulation of robust sparse recovery, where the noise is sparse but adversarial.

Roadmap. Section 15.1 .1 provides a summary of our techniques. We present our main result in Section 15.1.2. We provide a complete proof for Huber set query and a short proof sketch for Huber sparse recovery in Section 15.2. We made a conclusion in Section 15.3. We deferred the complete proof for Huber, Tukey and other results into later sections.

Notations. We use \([n]\) to denote \(\{1,2, \cdots, n\}\). For a vector \(x \in \mathbb{R}^{n}\), we use \(\|x\|_{1}\) to denote its entry-wise \(\ell_{1}\) norm. We use \(\|x\|_{2}\) to denote its entry-wise \(\ell_{2}\) norm. We use \(\operatorname{supp}(x) \subseteq[n]\) to denote a set of non-zero indices. We use \(\|x\|_{0}\) to denote the number of non-zero entries in \(x . \quad\left(\|x\|_{0}=|\operatorname{supp}(x)|\right)\) Let \(x_{-k}\) denote the vector obtained by zeroing out largest \(k\) coordinates (in absolute) of \(x\). We will sometimes refer to \(x_{-k}\) as the tail of \(x\), whereas we will refer to \(x-x_{-k}\) as the head of \(x\).

\subsection*{15.1.1 Our techniques}

All of our results are based on the ubiquitous CountSketch data structure, which hashes every coordinate \(i \in[n]\) to one of \(O(k)\) buckets, and repeats \(O(\log n)\) times. This data structure can be used to solve the \(\ell_{p}\) heavy hitters problems, as well as the sparse recovery under the \(\ell_{2} / \ell_{2}\) guarantee [PW11]. It is not hard to see that \(\ell_{1} / \ell_{1}\) can also be achieved via the same data structure. Interestingly, our theorems indicate that the Huber and Tukey guarantees can also be obtained via CountSketch.

In specific, we use the CountSketch to obtain estimates for all \(i \in[n]\), and then form the set \(S\) of the \(k\) coordinates with the largest estimates, as in previous work. Now, it might be the case that some elements in the head are not included in \(S\), which means that they have been displaced by some other elements in the tail. In the case that the threshold \(\tau\) for the Huber function is "small" enough, i.e. less than every \(\left|x_{j}\right|\), for every \(j \in S\) and \(j \in \operatorname{head}(k)\), we can bound the error similarly to the \(\ell_{1} / \ell_{1}\) case. An analogous thing happens in the case that \(\tau\), is large enough, i.e. \(\tau \geq\|x\|_{\infty}\), where can bound the error similarly to the \(\ell_{2} / \ell_{2}\) case. Things are more complicated in the case that \(\tau\) falls somewhere in between \(S\) and head \((k)\), where our guarantee needs to bound the \(\ell_{1}\) norm on the left-hand side, by the \(\ell_{2}\) norm on the right-hand side. This at first glance seems to be something impossible to do, but a careful glance reveals an interesting property: in the case of \(\tau=1\), if we estimate every \(\left|x_{i}\right|\) up to \(\Delta\) and if \(\Delta\) is, then every displaced element in the head should should be \(1+\Delta=(1+\Theta(\Delta))^{2}\), so we can hope to employ a similar argument to the \(\ell_{2} / \ell_{2}\), as this property implicitly allows us to reduce the problem to the case where the right-hand side and the left-hand side are measured with respect to the \(\ell_{2}\) norm. If \(\Delta\) is large enough then intuitively the mass of coordinates displaced "does not matter". because they were small
compared to the noise, and we can charge them to the tail. Of course, a careful combination of arguments and techniques needs to be addressed for the analysis to work out.

We note that the algorithm we suggest is not novel per se, but we rather indicate that the classic CountSketch data structure gives also the stronger Huber guarantee.

\subsection*{15.1.2 Our results}

In the field of sparse recovery / compressive sensing, researchers usually consider \(\ell_{p}\) norms as error estimates. Two well-known problems, arise when setting \(p=2\), giving birth to \(\ell_{2} / \ell_{2}\) and \(\ell_{1} / \ell_{1}\) problems.
\[
\begin{aligned}
& \ell_{2} / \ell_{2}:\left\|x-x^{\prime}\right\|_{2} \leq(1+\epsilon)\left\|x_{-k}\right\|_{2}, \\
& \ell_{1} / \ell_{1}:\left\|x-x^{\prime}\right\|_{1} \leq(1+\epsilon)\left\|x_{-k}\right\|_{1},
\end{aligned}
\]
where \(x\) is the ground-truth vector, \(x_{-k}\) is the vector that is obtained after zeroing out the largest \(k\) coordinates (in absolute) of \(x, x^{\prime}\) is the \(k\)-sparse we output and \(\epsilon \in(0,1)\) is the accuracy parameter.

We consider the problem of sparse recovery, when the error is measured with respect to an \(M\)-estimators. \(M\)-estimators are specified by a function \(H: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}\) where \(H(x)=\) \(H(-x), H(0)=0\), and \(H()\) is non-decreasing \(|x|\). We use \(\|\cdot\|_{H}\) to denote a new "norm" (However, it might not be true norms in fact, because it doesn't satisfies triangle inequality). For a vector \(x \in \mathbb{R}^{n}\), we use \(\|x\|_{H}\) to denote \(\sum_{i=1}^{n} H\left(x_{i}\right)\).

Huber function [Hub64] and Tukey function [Tuk60] have been two well-known \(M\) estimators, which have been proven useful to researchers for more than 50 years. In this
work, we will mainly focus on these two estimators, which are also the most common ones. We first provide the definition of Huber function and Tukey function.

Definition 15.1.1 (Huber function). Given parameter \(\tau>0\),
\[
H(x)= \begin{cases}|x|^{2} / \tau & \text { if }|x| \leq \tau \\ |x| & \text { if }|x|>\tau\end{cases}
\]

Note that an alternative definition of Huber is \(|x|^{2} / 2 \tau\) if \(|x| \leq \tau\) and \(|x|-\tau / 2\) if \(|x|>\tau\). Our techniques work for both of those definitions. This is also being observed in some other problems, e.g. regression [CW15b] and low-rank approximation [CW15a].

We define Tukey function [Tuk60] as follows

Definition 15.1.2 (Tukey function). Given parameter \(\tau>0\),
\[
H(x)= \begin{cases}1-\left(1-(x / \tau)^{2}\right)^{3}, & \text { if }|x| \leq \tau \\ 1, & \text { if }|x|>\tau\end{cases}
\]

It is not hard to observe that both the Huber and Tukey function satisfy three basic properties: they vanish at zero, i.e., \(H(0)=0\), are symmetric and non-decreasing.

\subsection*{15.1.2.1 Set-query}

We start by presenting a problem called "set query" which plays a crucial role in solving sparse recovery problem. Set query problem is first formally defined by Price [Pri11] and recently got studied in the context of Sparse Fourier transform[Kap17].

Definition 15.1.3 (Set-query). Let \(A \in \mathbb{R}^{m \times n}\) be some implicit matrix we want to design. For any accuracy parameter \(\epsilon\) and set \(S\) with \(|S|=k\), our goal is to read \(A x\) and output


Figure 15.1: (a) Huber function (top) (b) Tukey function (bottom)
vector \(x^{\prime} \in \mathbb{R}^{k}\) such that
\[
\left\|x^{\prime}-x_{S}\right\|_{H} \leq \epsilon \cdot\left\|x_{[n] \backslash S}\right\|_{H},
\]
where \(\|x\|_{H}=\sum_{i} H\left(x_{i}\right)\) and \(H\) is some function.
We give some formal definitions of what we are able to achieve in terms of measurements, column sparsity and decoding time.

Definition 15.1.4 (Measurements, column sparsity, decoding time). There are several things we want to optimize,
1. The number of measurements, is the number of rows of \(A\).
2. Column sparsity is, for any \(u \cdot e_{i}\), the time we need to compute \(A \cdot u \cdot e_{i}\), where \(u \in \mathbb{R}\) and \(e_{i} \in \mathbb{R}^{n}\) has 1 in the \(i\)-th coordinate and 0 everywhere else.
3. Decoding time means once \(A x\) is given, the time we need to output \(x^{\prime}\).

Note that the decoding time is not necessarily larger than the number of measurements.

Our set-query results on \(M\)-estimators are given in the following theorem.

Theorem 15.1.1 (Set-query). Let function \(H: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}\) be either Huber function or Tukey function. There is an implicit matrix \(A \in \mathbb{R}^{m \times n}\) with \(m=O\left(\epsilon^{-1} k \log k\right)\) and \(O(\log k)\) column sparsity. For any \(S \subseteq[n]\) such that \(|S|=k\), there is an algorithm which given \(A x\), runs in \(O(k \log k)\) decoding time, and outputs a vector \(x^{\prime} \in \mathbb{R}^{k}\) such that
\[
\left\|x^{\prime}-x_{S}\right\|_{H} \leq \epsilon \cdot\left\|x_{[n] \backslash S}\right\|_{H}
\]
holds with probability 9/10.

Remark 15.1.1. Although the matrix has \(m=O\left(\epsilon^{-1} k \log k\right)\), but our algorithm doesn't need to use all the measurements for the set-query algorithm. The way our algorithm is designed implies that the decoding time is independent of accuracy parameter \(\epsilon\).

For a vector \(x \in \mathbb{C}^{n}\), we define its Fourier transform \(\widehat{x} \in \mathbb{C}^{n}\) where \(\forall i \in[n], \widehat{x}_{i}=\) \(\sum_{j=1}^{n} \frac{1}{\sqrt{n}} x_{j} e^{-2 \pi i \mathbf{i} j / n}\). Similarly, we define the inverse Fourier transform of \(\widehat{x} \in \mathbb{C}^{n}\) to be \(x\) where \(\forall j \in[n], x_{j}=\sum_{i=1}^{n} \frac{1}{\sqrt{n}} \widehat{x}_{i} e^{2 \pi \mathrm{i} i j / n}\). Now, we're ready to present our Fourier set query result,

Theorem 15.1.2 (Fourier set query). Given a vector \(x \in \mathbb{C}^{n}\), for every \(\epsilon \in(0,1)\) and \(k \geq 1\), any \(S \subseteq[n],|S| \leq k\), there exists an algorithm (Algorithm 15.3) that takes \(O\left(\epsilon^{-1} k \log k\right.\). \(\log (n / \delta))\) samples, runs in \(O\left(\epsilon^{-1} k \log k \cdot \log (n / \delta)\right)\) time, and outputs a vector \(x^{\prime} \in \mathbb{C}^{n}\) such that
\[
\left\|\left(x^{\prime}-\widehat{x}\right)_{S}\right\|_{H} \leq \epsilon\left\|\widehat{x}_{[n] \backslash S}\right\|_{H}+\delta\|\widehat{x}\|_{H}
\]
holds with probability at least \(1-1 / \operatorname{poly}(k)\).

We can think of \(\|\widehat{x}\|_{H}\) being bounded by \(\operatorname{poly}(n)\). Then we can set \(\delta=1 / \operatorname{poly}(n)\). This assumption, along with the additional \(\delta\) factor is common in Fourier literature, e.g. [HIKP12a, PS15, CKPS16].

\subsection*{15.1.2.2 Sparse Recovery}

We first present \(k\)-sparse recovery results under Huber/Tukey function. Our algorithm achieves \(O\left(\epsilon^{-2} k \log n\right)\) measurements which matches the state-of-the-art \(\ell_{2} / \ell_{2}\) algorithm (see [PW11]).

Theorem 15.1.3 (Main result, Sparse recovery, Optimal measurements). Let \(H\) denote either Huber function (see Def. 15.1.1) or Tukey function (see Def. 15.1.2). There exists a distribution of \(\Phi \in \mathbb{R}^{m \times n}\) with \(m=O\left(\epsilon^{-2} k \log n\right)\) and \(O(\log n)\) column sparsity, and a recovery algorithm, such that for each \(x \in \mathbb{R}^{n}, k \in[n]\) and \(\epsilon \in(0,1)\), the recovery algorithm takes \(\Phi x\) as its input, runs in \(O(n \log n)\) time, and outputs \(k\)-sparse vector \(x^{\prime} \in \mathbb{R}^{n}\) so that with probability \(1-1 / \operatorname{poly}(n)\),
\[
\left\|x^{\prime}-x\right\|_{H} \leq(1+\epsilon) \min _{k-\text { sparse } y}\|y-x\|_{H}
\]

\subsection*{15.2 Algorithm and Analysis}

We start by defining a well-known tool called Count-Sketch [CCF02].
Definition 15.2.1 (Count-Sketch [CCF02]). A count sketch data structure with bucket size \(B\) and replicates \(R\) corresponds to a measurement matrix \(\Phi \in R^{B R \times n}\). Matrix \(\Phi\) is a concatenation of \(\Phi^{(1)}, \ldots, \Phi^{(R)}\), where each \(\Phi^{(r)} \in \mathbb{R}^{B \times n}(r \in[R])\) is defined by a pariwise hash function \(h_{i}:[n] \rightarrow[B]\) and a pairwise sign function \(\sigma_{i}:[n] \rightarrow\{-1,1\}\), such that for
each \(j \in[B]\) and \(i \in[n]\),
\[
\Phi_{j, i}^{(r)}= \begin{cases}\sigma_{r}(i) & \text { if } h_{r}(i)=j \\ 0 & \text { otherwise }\end{cases}
\]

The data structure maintains the value of \(\Phi x\) and for each \(i \in[n]\), it gives
\[
x_{i}^{\prime}=\operatorname{median}\left(\sigma_{1}(i) \cdot\left(\Phi^{(1)} x\right)_{h_{1}(i)}, \ldots, \sigma_{R}(i) \cdot\left(\Phi^{(R)} x\right)_{h_{R}(i)}\right)
\]
as estimation to \(x_{i}\).

\subsection*{15.2.1 Set query algorithm for Huber}

Our set query algorithm for Huber loss function is presented in Algorithm 15.1. The algorithm uses count-sketch data structure with bucket number \(B=O\left(\epsilon^{-1} k\right)\) and replicate number \(R=O(\log k)\).
```

Algorithm 15.1 Set Query Algorithm
procedure $\operatorname{SetQuery}(n, k, \epsilon, S, x) \quad \triangleright$ Theorem 15.2.1
$B \leftarrow O\left(\epsilon^{-1} k\right)$
$R \leftarrow O(\log k)$
Construct count sketch with bucket size $B$ and replicates $R$ according to Defini-
tion 15.2.1
$y \leftarrow \Phi x \quad \triangleright$ In streaming setting we maintain $y$ when receiving updates to $x$
for $i \in S$ do
$x_{i}^{\prime} \leftarrow \operatorname{median}_{r \in[R]} \sigma_{r}(i) y_{h_{r}(i)}^{(r)} \quad \triangleright y=\left(y^{(1)}, \ldots, y^{(R)}\right)$
end for
return $x_{S}^{\prime}$
end procedure

```

Now we provide a correctness proof for the set-query algorithm under Huber function. For simplicity we let \(\tau=1\), but our proof extends to arbitrary \(\tau\). In order to express \(\left\|x_{[n] \backslash S}\right\|_{H}\)
in a handy form, we first partition the coordinates in \([n] \backslash S\) into \(P\) and \(Q\), where \(P\) contains small coordinates whose absolute values are less than \(\tau\) and \(Q\) contains those who are larger than \(\tau\). By guarantee of the count-sketch, we can bound the error of each estimate for coordinates in \(S\). We then bound the error in Huber loss function by discussing the absolute error is either less than or greater than \(\tau\).

Theorem 15.2.1 (Set-query, Huber case of Theorem 15.1.1). Let function \(H\) be defined as Definition 15.1.1. There is an algorithm (Algorithm 15.1) that takes \(O\left(\epsilon^{-1} k \log k\right)\) samples/times and outputs a vector \(x^{\prime} \in \mathbb{R}^{n}\) such that
\[
\left\|x_{S}^{\prime}-x_{S}\right\|_{H} \leq \epsilon \cdot\left\|x_{[n] \backslash S}\right\|_{H}
\]
holds with probability \(1-1 / \operatorname{poly}(k)\).

Proof. We choose \(B=\epsilon^{-1} k\). We define set \(P, Q \subset[n] \backslash S\) as follows
\[
P=\left\{i| | x_{i} \mid \leq 1, i \in[n] \backslash S\right\}, Q=\left\{i| | x_{i} \mid \geq 1, i \in[n] \backslash S\right\} .
\]

We define \(P_{B / 6}\) to be top \(B / 6\) coordinates in \(P\), and define \(P^{\prime}=P \backslash P_{B / 6}\). Similarly we define \(Q_{B / 6}\) to be top \(B / 6\) coordinates in \(Q\), and define \(Q^{\prime}=Q \backslash Q_{B / 6}\).

Let \(R=O(\log k)\). For each \(i \in S\), we estimate \(x_{i}\) as \(x_{i}^{\prime}=\operatorname{median}_{j \in[R]}\left\{y_{h_{j}(i)}^{(j)}\right\}\). We can prove that with probability \(1-1 / \operatorname{poly}(k)\),
\[
\left|x_{i}^{\prime}-x_{i}\right| \leq \sqrt{\left(\left\|x_{P^{\prime}}\right\|_{2}^{2}+\left\|x_{Q^{\prime}}\right\|_{2}^{2}\right) / B}
\]

Note that \(\left\|x_{\bar{S}}\right\|_{H}=\left\|x_{P}\right\|_{2}^{2}+\left\|x_{Q}\right\|_{1}\).

We prove the final result by considering two cases:
Case (1)
\[
\sqrt{\left(\left\|x_{P^{\prime}}\right\|_{2}^{2}+\left\|x_{Q^{\prime}}\right\|_{2}^{2}\right) / B} \leq 1
\]

Case (2)
\[
\sqrt{\left(\left\|x_{P^{\prime}}\right\|_{2}^{2}+\left\|x_{Q^{\prime}}\right\|_{2}^{2}\right) / B}>1
\]

Case (1) \(\sqrt{\left(\left\|x_{P^{\prime}}\right\|_{2}^{2}+\left\|x_{Q^{\prime}}\right\|_{2}^{2}\right) / B} \leq 1\).
In the following, we prove that \(\left\|x_{Q^{\prime}}\right\|_{2}^{2} \leq\left\|x_{Q^{\prime}}\right\|_{1}\). First, note that
\[
\left\|x_{Q^{\prime}}\right\|_{2} \leq \sqrt{\left\|x_{P^{\prime}}\right\|_{2}^{2}+\left\|x_{Q^{\prime}}\right\|_{2}^{2}} \leq \sqrt{B}
\]

By Lemma 15.5.2, we can also bound \(\left\|x_{Q^{\prime}}\right\|_{2}\) to be
\[
\left\|x_{Q^{\prime}}\right\|_{2} \leq \frac{1}{\sqrt{B}}\left\|x_{Q}\right\|_{1}
\]

Combining these two inequality on \(\left\|x_{Q^{\prime}}\right\|_{2}\), we get
\[
\left\|x_{Q^{\prime}}\right\|_{2}^{2} \leq \sqrt{B} \cdot \frac{1}{\sqrt{B}}\left\|x_{Q}\right\|_{1} \leq\left\|x_{Q}\right\|_{1} .
\]

By definition of Huber norm,
\[
\begin{aligned}
\left\|x_{S}^{\prime}-x_{S}\right\|_{H} & \leq \frac{k}{B}\left(\left\|x_{P^{\prime}}\right\|_{2}^{2}+\left\|x_{Q^{\prime}}\right\|_{2}^{2}\right) \\
& \leq \frac{k}{B}\left(\left\|x_{P}\right\|_{2}^{2}+\left\|x_{Q}\right\|_{1}\right) \\
& =\epsilon\left\|x_{\bar{S}}\right\|_{H} .
\end{aligned}
\]

Case (2) \(\sqrt{\frac{\left\|x_{P^{\prime}}\right\|_{2}^{2}+\left\|x_{Q^{\prime}}\right\|_{2}^{2}}{B}}>1\).


Figure 15.2: \(\left\|x_{S \cap \bar{T}}\right\|_{H}-\left\|x_{T \cap \bar{S}}\right\|_{H}\) is the displacement error

The error is \(\left\|x_{S}^{\prime}-x_{S}\right\|_{H} \leq \frac{k}{\sqrt{B}}\left(\left\|x_{P^{\prime}}\right\|_{2}+\left\|x_{Q^{\prime}}\right\|_{2}\right)\). We can bound \(\left\|x_{Q^{\prime}}\right\|_{2} \leq \frac{1}{\sqrt{B}}\left\|x_{Q}\right\|_{1}\) by Lemma 15.5.2. Now trying to bound \(\left\|x_{P^{\prime}}\right\|_{2}\). If \(\left\|x_{P^{\prime}}\right\|_{2} \leq\left\|x_{Q^{\prime}}\right\|_{2}\) we are done. Otherwise we have \(\left\|x_{P^{\prime}}\right\|_{2} \geq \sqrt{B / 2}\) and so \(\left\|x_{P^{\prime}}\right\|_{2} \leq \frac{\left\|x_{P^{\prime}}\right\|_{2}}{\sqrt{B / 2}}\left\|x_{P^{\prime}}\right\|_{2} \leq \sqrt{2 / B}\left\|x_{P}\right\|_{2}^{2}\). Either case we can finally bound \(\left\|x_{S}^{\prime}-x_{S}\right\|_{H}\) by \(2 \epsilon\left\|x_{\bar{S}}\right\|_{H}\).

\subsection*{15.2.2 Proof of Theorem 15.1.3}
```

Algorithm 15.2 Sparse Recovery Algorithm
procedure $\operatorname{SparseRecovery}(n, k, \epsilon, x)$
$B \leftarrow O\left(\epsilon^{-2} k\right)$
$R \leftarrow O(\log n)$
Construct count sketch with bucket size $B$ and replicates $R$ according to Defini-
tion 15.2.1
$y \leftarrow \Phi x \quad \triangleright$ In streaming setting we maintain $y$ when receiving updates to $x$
for $i=1 \rightarrow n$ do
$x_{i}^{\prime} \leftarrow \operatorname{median}_{r \in[R]} \sigma_{r}(i) y_{h_{r}(i)}^{(r)} \quad \triangleright y=\left(y^{(1)}, \ldots, y^{(R)}\right)$
end for
Let $T$ denote the set of coordinates that are $k$ largest (in absolute) of $x^{\prime}$
return $x_{T}^{\prime}$
end procedure

```

The high-level idea. Let \(S\) denote the top- \(k\) largest coordinates in absolute. Let
\(T\) denote the set coordinates the algorithm outputs. Then we can rewrite \(\left\|x-x_{T}^{\prime}\right\|_{H}\) as
\[
\begin{aligned}
& \left\|x-x_{T}^{\prime}\right\|_{H} \\
= & \left\|\left(x-x^{\prime}\right)_{T}\right\|_{H}+\left\|x_{\bar{T}}\right\|_{H} \\
= & \left\|\left(x-x^{\prime}\right)_{T}\right\|_{H}+\left(\left\|x_{\bar{T}}\right\|_{H}-\left\|x_{\bar{S}}\right\|_{H}\right)+\left\|x_{\bar{S}}\right\|_{H} \\
= & \underbrace{\left\|\left(x-x^{\prime}\right)_{T}\right\|_{H}}_{\text {set-query }}+\underbrace{\left(\left\|x_{S \cap \bar{T}}\right\|_{H}-\left\|x_{T \cap \bar{S}}\right\|_{H}\right)}_{\text {displacement }}+\left\|x_{\bar{S}}\right\|_{H} .
\end{aligned}
\]

Using Set-query, we know that \(\left\|\left(x-x^{\prime}\right)_{T}\right\|_{H} \leq(\epsilon / 2)\left\|x_{\bar{S}}\right\|_{H}\). Now the question is how to bound "displacement" by \((\epsilon / 2)\left\|x_{\bar{S}}\right\|_{H}\) and we explain the formal proof in next a few pages.

Proof. of Theorem 15.1.3. Our algorithm is presented in Algorithm 15.2. Let \(x_{-k}\) denotes the vector \(x\) with absolute largest \(k\) coordinates removed. We use count sketch to give each \(x_{i}\) (where \(i \in[n]\) ) an estimate. In particular, we choose the number of buckets \(B=O\left(\epsilon^{-2} k\right)\) and the number of duplicates \(R=O(\log n)\) for the count sketch data structure. Let \(x_{i}^{\prime}\) denote the estimate obtained by the count sketch for \(x_{i}\). Let \(T\) denote the largest \(k\) coordinates in absolute value in \(x^{\prime}\). The recovery algorithm outputs \(x_{T}^{\prime}\) as a \(k\)-sparse approximation to \(x\).

Note that the count sketch data structure takes \(B \cdot R=O\left(\epsilon^{-2} k \log n\right)\) measurements. Its update time is \(O(R)=O(\log n)\) time, and the time to obtain \(x^{\prime}\) is \(O(n \cdot R)=O(n \log n)\). Taking largest \(k\) coordinates from \(x^{\prime}\) requires \(O(n \log n)\) time. Thus the total query time is \(O(n \log n)\). In the following, we show that \(\left\|x-x_{T}^{\prime}\right\|_{H} \leq(1+\epsilon)\left\|x_{-k}\right\|_{H}\).

By Lemma 15.5.1, for each \(i \in[n]\), with probability \(1-1 / \operatorname{poly}(n)\),
\[
\begin{equation*}
\left(x_{i}-x_{i}^{\prime}\right)^{2} \leq \frac{\epsilon^{2}}{4 k}\left\|x_{-4 k / \epsilon^{2}}\right\|_{2}^{2} \leq \frac{\epsilon^{2}}{4 k}\left\|x_{-4 k}\right\|_{2}^{2} \tag{15.1}
\end{equation*}
\]

Let \(S\) denote the largest \(k\) coordinates in absolute value in \(x\). Recall that \(\bar{S}=[n] \backslash S\), \(\bar{T}=[n] \backslash T\), and \(x_{-k}=x_{\bar{S}}\). The recovery error is
\[
\begin{equation*}
\left\|x-x_{T}^{\prime}\right\|_{H}=\left\|\left(x-x^{\prime}\right)_{T}\right\|_{H}+\left\|x_{\bar{T}}\right\|_{H} . \tag{15.2}
\end{equation*}
\]

By point estimation guarantee given in Eq. (15.5) and results from Huber set query.
\[
\begin{aligned}
\left\|\left(x-x^{\prime}\right)_{T}\right\|_{H} & =\sum_{i \in T} H\left(x_{i}-x_{i}^{\prime}\right) \\
& \leq|T| \cdot H\left(\frac{\epsilon}{2 \sqrt{k}}\left\|x_{-4 k}\right\|_{2}\right) \\
& \leq \frac{\epsilon}{2}\left\|x_{-k}\right\|_{H} .
\end{aligned}
\]

Note that
\[
\begin{aligned}
\left\|x_{\bar{T}}\right\|_{H} & =\left(\left\|x_{\bar{T}}\right\|_{H}-\left\|x_{\bar{S}}\right\|_{H}\right)+\left\|x_{\bar{S}}\right\|_{H} \\
& =\left(\left\|x_{S \cap \bar{T}}\right\|_{H}-\left\|x_{\bar{S} \cap T}\right\|_{H}\right)+\left\|x_{-k}\right\|_{H}
\end{aligned}
\]
where the second step follows by canceling out the coordinates in \(\bar{S} \cap \bar{T}\). In the following, our goal is to show that \(\sum_{i \in S \cap \bar{T}} H\left(x_{i}\right)-\sum_{i \in \bar{S} \cap T} H\left(x_{i}\right) \leq \frac{\epsilon}{2}\left\|x_{-k}\right\|_{H}\).

If \(S \cap \bar{T}=\emptyset\), then \(S=T\) and \(\sum_{i \in S \cap \bar{T}} H\left(x_{i}\right)-\sum_{i \in \bar{S} \cap T} H\left(x_{i}\right)=0\). In the following, we assume that \(S \cap \bar{T} \neq \emptyset\). Let \(b=\max _{i \in S \cap \bar{T}}\left|x_{i}\right|\), and \(a=\min _{i \in \bar{S} \cap T}\left|x_{i}\right|\). Because Huber function \(H(t)\) is monotonically increasing as \(|t|\) increases,
\[
\left\|x_{S \cap \bar{T}}\right\|_{H}-\left\|x_{\bar{S} \cap T}\right\|_{H} \leq|\bar{S} \cap T| \cdot(H(b)-H(a)) .
\]

By Lemma 15.2.5, we have
\[
\begin{equation*}
b-a \leq \frac{\epsilon}{\sqrt{k}}\left\|x_{-4 k}\right\|_{2} \tag{15.3}
\end{equation*}
\]

Moreover, by definition that \(a=\min _{i \in \bar{S} \cap T}\left|x_{i}\right|\) and let \(r=|\bar{S} \cap T|\),
\[
\begin{equation*}
a \leq \frac{\left\|x_{\bar{S} \cap T}\right\|_{2}}{\sqrt{r}} \leq \frac{\left\|x_{-k}\right\|_{2}}{\sqrt{r}} . \tag{15.4}
\end{equation*}
\]

We are going to prove that \(H(b)-H(a) \leq \frac{\epsilon}{2} \frac{\left\|x_{-k}\right\|_{H}}{|\bar{S} \cap T|}\). We discuss in three cases: case (1) \(\tau \leq a\), case (2) \(a<\tau<b\), case (3) \(b \leq \tau\). We analyze these three cases in Claim 15.2.2, Claim 15.2.3 and Claim 15.2.4.

Claim 15.2.2. Let \(r=|\bar{S} \cap T|\). If \(\tau \leq a\), then \(H(b)-H(a) \leq \frac{\epsilon}{2} \frac{\left\|x_{-k}\right\|_{H}}{r}\).

Proof. For simplicity, we use \(r\) to denote \(|\bar{S} \cap T|\). By definition of \(T\), we know that \(1 \leq r \leq k\). Let \(P, Q\) be a partition of \(\bar{S}\) where
\[
P=\left\{i| | x_{i} \mid \leq \tau, i \in \bar{S}\right\}, Q=\left\{i| | x_{i} \mid>\tau, i \in \bar{S}\right\}
\]

By triangle inequality \(\left\|x_{-5 k}\right\|_{2} \leq\left\|x_{P}\right\|_{2}+\left\|\left(x_{Q}\right)_{-4 k}\right\|_{2}\). By definition of Huber function \(\left\|x_{-k}\right\|_{H}=\frac{1}{\tau}\left\|x_{P}\right\|_{2}^{2}+\left\|x_{Q}\right\|_{1}\). Because \(\tau \leq a, \bar{S} \cap T \subseteq Q\). Thus \(\left\|x_{Q}\right\|_{1} \geq\left\|x_{\bar{S} \cap T}\right\|_{1} \geq \tau|\bar{S} \cap T|=\) \(\tau \cdot r\). Therefore,
\[
\left\|x_{-k}\right\|_{H} \geq \frac{1}{\tau}\left\|x_{P}\right\|_{2}^{2}+\tau r \geq 2 \sqrt{\frac{1}{\tau}\left\|x_{P}\right\|_{2}^{2} \cdot \tau r}=2 \sqrt{r} \cdot\left\|x_{P}\right\|_{2}
\]

So we have \(\left\|x_{P}\right\|_{2} \leq \frac{1}{2 \sqrt{r}}\left\|x_{-k}\right\|_{H}\). By Lemma 15.5.2,
\[
\left\|\left(x_{Q}\right)_{-4 k}\right\|_{2} \leq \frac{1}{2 \sqrt{k}}\left\|x_{Q}\right\|_{1} \leq \frac{1}{2 \sqrt{r}}\left\|x_{Q}\right\|_{1} \leq \frac{1}{2 \sqrt{r}}\left\|x_{-k}\right\|_{H}
\]

Therefore,
\[
\begin{aligned}
H(b)-H(a) & =b-a \leq \frac{\epsilon}{\sqrt{r}}\left\|x_{-5 k}\right\|_{2} \\
& \leq \frac{\epsilon}{\sqrt{r}}\left(\left\|x_{P}\right\|_{2}+\left\|\left(x_{Q}\right)_{-4 k}\right\|_{2}\right) \\
& \leq \frac{\epsilon}{\sqrt{r}}\left(\frac{1}{2 \sqrt{r}}\left\|x_{-k}\right\|_{H}+\frac{1}{2 \sqrt{r}}\left\|x_{-k}\right\|_{H}\right) \\
& \leq \frac{\epsilon\left\|x_{-k}\right\|_{H}}{r} .
\end{aligned}
\]

Claim 15.2.3. If \(a<\tau<b\), then \(H(b)-H(a) \leq \frac{\epsilon}{2} \frac{\left\|x_{-k}\right\|_{H}}{|\bar{S} \cap T|}\).
Claim 15.2.4. If \(b \leq \tau\), then \(H(b)-H(a) \leq \frac{\epsilon}{2} \frac{\|x-k\|_{H}}{|S \cap T|}\).

We provide the proofs of the above two Claims in Section 15.7.

Lemma 15.2.5. Let vector \(u, v \in \mathbb{R}^{n}\) such that \(\|u-v\|_{\infty} \leq \Delta\). Let \(S\) be the largest \(k\) coordinates in \(u\), and let \(T\) be the largest \(k\) coordinates in \(v\). Let \(b=\max _{i \in S \cap \bar{T}} u_{i}\) and \(a=\min _{i \in \bar{S} \cap T} u_{i}\). Then \(b-a \leq 2 \Delta\).

Proof. Let \(\lambda=\max _{i \in \bar{T}} v_{i}\). We have
\[
b=\max _{i \in S \cap \bar{T}} u_{i} \leq \max _{i \in \bar{T}} u_{i} \leq \Delta+\max _{i \in \bar{T}} v_{i}=\Delta+\lambda .
\]

On the other hand,
\[
a=\min _{i \in \bar{S} \cap T} u_{i} \geq \min _{i \in T} u_{i} \geq-\Delta+\min _{i \in T} v_{i} \geq-\Delta+\lambda .
\]

Therefore,
\[
b-a \leq(\Delta+\lambda)-(-\Delta+\lambda) \leq 2 \Delta
\]

\subsection*{15.3 Conclusion}

In this work, we gave the first set of sparse recovery algorithms when the error is measured with respect to Huber and Tukey functions. Because of the usefulness of these Mestimators, we believe that this will be proved to be an important line of research, demanding new ideas and techniques, and offering fruitful results both in theory and in practice.

\subsection*{15.4 Preliminaries}

In this section, we give definitions of Huber function, reverse Huber function and Tukey function. All these functions have a parameter \(\tau>0\) such that the function is defined as one "simple" function on domain \(\{x:|x| \leq \tau\}\) and another "simple" function on domain \(\{x:|x|>\tau\}\). In particular, Huber function is defined to be a quadratic function for \(|x| \leq \tau\) and a linear function for \(|x|>\tau\). Inverse Huber function is opposite. Tukey function is defined approximately to be a quadratic function for \(|x| \leq \tau\), and a constant function for \(|x|>\tau\).

When measuring errors in terms of these functions, we are more robust to outliers. Because these functions in domain \(\{x:|x|>\tau\}\) grow slower than domain \(\{x:|x| \leq \tau\}\).

\subsection*{15.4.1 Huber estimator}

Given parameter \(\tau>0\), Huber function [Hub64] is defined as follows.
Definition 15.4.1 (Huber function).
\[
H(x)= \begin{cases}|x|^{2} / \tau & \text { if }|x| \leq \tau \\ |x| & \text { if }|x|>\tau\end{cases}
\]

In addition, as appeared in [SWZ18], one can also define reverse Huber function with parameter \(\tau\) as
\[
H(x)= \begin{cases}|x| & \text { if }|x| \leq \tau \\ |x|^{2} / \tau & \text { if }|x|>\tau\end{cases}
\]

\subsection*{15.4.2 Tukey estimator}

Tukey function [Tuk60] is defined as follows.

Definition 15.4.2 (Tukey function).
\[
H(x)= \begin{cases}1-\left(1-(x / \tau)^{2}\right)^{3}, & \text { if }|x| \leq \tau \\ 1, & \text { if }|x|>\tau\end{cases}
\]

It is easy to observe that for \(|x| \leq \tau\),
\[
\begin{aligned}
H(x) & =1-\left(1-3(x / \tau)^{2}+3(x / \tau)^{4}-(x / \tau)^{6}\right) \\
& =3(x / \tau)^{2}-3(x / \tau)^{4}+(x / \tau)^{6} \\
& =\frac{3}{4}(x / \tau)^{2}+\frac{9}{4}(x / \tau)^{2}-3(x / \tau)^{4}+(x / \tau)^{6} \\
& =\frac{3}{4}(x / \tau)^{2}+(x / \tau)^{2}\left(\frac{3}{2}-(x / \tau)^{2}\right)^{2} \\
& \geq \frac{3}{4}(x / \tau)^{2} .
\end{aligned}
\]

Similarly, we can also give an upper bound of \(H(x)\),
\[
\begin{aligned}
H(x) & =3(x / \tau)^{2}-3(x / \tau)^{4}+(x / \tau)^{6} \\
& \leq 3(x / \tau)^{2}+(x / \tau)^{6} \\
& \leq 4(x / \tau)^{2} .
\end{aligned}
\]

Therefore, we have the following fact, i.e. Tukey function is within a constant factor of the quadratic function for \(|x| \leq \tau\). Therefore, in our analysis in later sections, we approximate Tukey function as a quadratic function for \(|x| \leq \tau\).

Fact 15.4.1. If \(|x| \leq \tau, H(x) \in\left[\frac{3}{4}(x / \tau)^{2}, 4(x / \tau)^{2}\right]\); Otherwise \(H(x)=1\).

\subsection*{15.5 Set Query}

In this section, we present our set query results on reverse Huber function and Tukey function. The results on Huber function is presented in the main text. We summarize our results on set query in Table 15.1. Since we extensively apply the count-sketch data structure, we present as follows.

Definition 15.5.1 (Count Sketch). A count sketch data structure with bucket size \(B\) and replicates \(R\) corresponds to a measurement matrix \(\Phi \in R^{B R \times n}\). Matrix \(\Phi\) is a concatenation of \(\Phi^{(1)}, \ldots, \Phi^{(R)}\), where each \(\Phi^{(r)} \in \mathbb{R}^{B \times n}(r \in[R])\) is defined by a pariwise hash function \(h_{i}:[n] \rightarrow[B]\) and a pairwise sign function \(\sigma_{i}:[n] \rightarrow\{-1,1\}\), such that for each \(j \in[B]\) and \(i \in[n]\),
\[
\Phi_{j, i}^{(r)}= \begin{cases}\sigma_{r}(i) & \text { if } h_{r}(i)=j \\ 0 & \text { otherwise }\end{cases}
\]

The data structure maintains the value of \(\Phi x\) and for each \(i \in[n]\), it gives
\[
x_{i}^{\prime}=\operatorname{median}\left(\sigma_{1}(i) \cdot\left(\Phi^{(1)} x\right)_{h_{1}(i)}, \ldots, \sigma_{R}(i) \cdot\left(\Phi^{(R)} x\right)_{h_{R}(i)}\right)
\]
as esimation to \(x_{i}\).

Let \(x^{\prime} \in \mathbb{R}^{n}\) denote the approximation given by the count sketch data structure. We have the following guarantee for count sketch with bucket size \(B\) and replicates \(R\).

Lemma 15.5.1 ([CCF02]). Let \(x \in \mathbb{R}^{n}\), and let \(x^{\prime} \in \mathbb{R}^{n}\) be the output of a count sketch data structure (Definition 15.5.1) with bucket size \(B\) and replicate \(R\). For each \(i \in[n]\), with probability at least \(1-2^{-\Omega(R)}\),
\[
\left(x_{i}-x_{i}^{\prime}\right)^{2} \leq \frac{2}{B}\left\|x_{-B / 2}\right\|_{2}^{2}
\]

Table 15.1: Set Query
\begin{tabular}{|l|l|l|l|}
\hline & Measurements & Running time & Column sparsity \\
\hline Theorem 15.2.1, Huber & \(\epsilon^{-1} k \log k\) & \(k \log k\) & \(\log k\) \\
\hline Theorem 15.5.3, Reverse Huber & \(\epsilon^{-1} k \log k\) & \(k \log k\) & \(\log k\) \\
\hline Theorem 15.5.4, Tukey & \(\epsilon^{-1} k \log k\) & \(k \log k\) & \(\log k\) \\
\hline
\end{tabular}

The following is a standard lemma in compressive sensing literatures that bounds the \(\ell_{2}\) norm of a vector with largest (in absolute value) \(k\) coordinates removed in terms of \(\ell_{1}\) norm of its original vector. We also include its proof for completeness.

Lemma 15.5.2. For any \(x \in \mathbb{R}^{n}\) and any \(k \in[n],\left\|x_{-k}\right\|_{2} \leq \frac{1}{\sqrt{k}}\|x\|_{1}\), where \(x_{-k}\) denotes the vector obtained by zeroing out largest \(k\) coordinates (in absolute) of \(x\).

Proof. We partition the coordinates of \(x\) to be \(B_{1}, B_{2}, \ldots\) where \(B_{1}\) contains largest \(k\) coordinates, and \(B_{2}\) contains the next largest \(k\) coordinates, and so on. We have
\[
\begin{aligned}
\left\|x_{-k}\right\|_{2} & \leq \sum_{i=2}^{n / k}\left\|x_{B_{i}}\right\|_{2} \\
& \leq \sum_{i=2}^{n / k} \sqrt{k \cdot\left(\left\|x_{B_{i-1}}\right\|_{1} / k\right)^{2}} \\
& =\frac{1}{\sqrt{k}} \sum_{i=1}^{n / k-1}\left\|x_{B_{i}}\right\|_{1} \\
& \leq \frac{1}{k}\|x\|_{1} .
\end{aligned}
\]
where the second step follows by Cauchy-Schwarz inequality.

\subsection*{15.5.1 \(\epsilon\)-approximation for reverse Huber function}

We present our set query result for reverse Huber functions. The result is quite symmetric to our result for Huber function.

Theorem 15.5.3 (Reverse Huber set-query). Let function \(H\) be defined as Definition 15.4.1. There is an algorithm that takes \(O\left(\epsilon^{-1} k \log k\right)\) samples/times and outputs a vector \(x^{\prime} \in \mathbb{R}^{n}\) such that
\[
\left\|x_{S}^{\prime}-x_{S}\right\|_{H} \leq \epsilon \cdot\left\|x_{[n] \backslash S}\right\|_{H}
\]
holds with probability \(1-1 / \operatorname{poly}(k)\).

Proof. We choose \(B=O\left(\epsilon^{-1} k\right)\) and \(R=O(\log k)\), and use the count-sketch data structure with bucket size \(O(B)\) and replicate number \(O(R)\). Let \(x^{\prime} \in \mathbb{R}^{n}\) denote the estimate given by the count-sketch data structure. By Lemma 15.5.1, we have the guarantee that with probability \(1-1 / \operatorname{poly}(k)\),
\[
\left|x_{i}^{\prime}-x_{i}\right| \leq \sqrt{\frac{\left\|x_{-B}\right\|_{2}^{2}}{B}}
\]

We define set \(P, Q \subset[n] \backslash S\) as follows
\[
P=\left\{i| | x_{i} \mid \leq 1, i \in[n] \backslash S\right\}, Q=\left\{i| | x_{i} \mid \geq 1, i \in[n] \backslash S\right\} .
\]

We define \(P_{B / 6}\) to be top \(B / 6\) coordinates in \(P\), and define \(P^{\prime}=P \backslash P_{B / 6}\). Similarly we define \(Q_{B / 6}\) to be top \(B / 6\) coordinates in \(Q\), and define \(Q^{\prime}=Q \backslash Q_{B / 6}\). We have \(\sqrt{\frac{\left\|x_{-B}\right\|_{2}^{2}}{B}} \leq\) \(\sqrt{\frac{\left\|x_{P^{\prime}}\right\|_{2}^{2}+\left\|x_{Q^{\prime}}\right\|_{2}^{2}}{B}}\).

Note that
\[
\left\|x_{\bar{S}}\right\|_{H}=\left\|x_{P}\right\|_{1}+\left\|x_{Q}\right\|_{2}^{2}
\]

We prove the final result by considering two cases: case (1) \(\sqrt{\frac{\left\|x_{P^{\prime}}\right\|_{2}^{2}+\left\|x_{Q^{\prime}}\right\|_{2}^{2}}{B}} \leq 1\); Case (2) \(\sqrt{\frac{\left\|x_{P^{\prime}}\right\|_{2}^{2}+\left\|x_{Q^{\prime}}\right\|_{2}^{2}}{B}}>1\).

Case (1) \(\sqrt{\frac{\left\|x_{P^{\prime}}\right\|_{2}^{2}+\left\|x_{Q^{\prime}}\right\|_{2}^{2}}{B}} \leq 1\).
The error is
\[
\begin{aligned}
\left\|x_{S}^{\prime}-x_{S}\right\|_{H} & \leq k \sqrt{\frac{\left\|x_{P^{\prime}}\right\|_{2}^{2}+\left\|x_{Q^{\prime}}\right\|_{2}^{2}}{B}} \\
& \leq \frac{k}{\sqrt{B}}\left(\left\|x_{P^{\prime}}\right\|_{2}+\left\|x_{Q^{\prime}}\right\|_{2}\right) .
\end{aligned}
\]

By Lemma 15.5.2, we have \(\left\|x_{P^{\prime}}\right\|_{2} \leq \frac{1}{\sqrt{B}}\left\|x_{P}\right\|_{1}\). Now we are trying to upper bound \(\left\|x_{Q^{\prime}}\right\|_{2}\). If \(\left\|x_{Q^{\prime}}\right\|_{2} \leq\left\|x_{P^{\prime}}\right\|_{2}\) we are done. Otherwise we have \(\left\|x_{Q^{\prime}}\right\|_{2} \geq \sqrt{B / 2}\) and so \(\left\|x_{Q^{\prime}}\right\|_{2} \leq\) \(\frac{\left\|x_{Q^{\prime}}\right\|_{2}}{\sqrt{B / 2}}\left\|x_{Q^{\prime}}\right\|_{2} \leq \sqrt{2 / B}\left\|x_{Q}\right\|_{2}^{2}\). Either case we can finally bound \(\left\|x_{S}^{\prime}-x_{S}\right\|_{H}\) by \(2 \epsilon\left\|x_{\bar{S}}\right\|_{H}\).

Case (2) \(\sqrt{\frac{\left\|x_{P^{\prime}}\right\|_{2}^{2}+\left\|x_{Q^{\prime}}\right\|_{2}^{2}}{B}}>1\).
We have
\[
\begin{aligned}
\left\|x_{S}^{\prime}-x_{S}\right\|_{H} & \leq \frac{k}{B}\left(\left\|x_{P^{\prime}}\right\|_{2}^{2}+\left\|x_{Q^{\prime}}\right\|_{2}^{2}\right) \\
& \leq \frac{k}{B}\left(\left\|x_{P^{\prime}}\right\|_{1}+\left\|x_{Q^{\prime}}\right\|_{2}^{2}\right) \\
& \leq \frac{k}{B}\left(\left\|x_{P}\right\|_{1}+\left\|x_{Q}\right\|_{2}^{2}\right) \\
& =\epsilon\left\|x_{\bar{S}}\right\|_{H}
\end{aligned}
\]
where the second step follows by the fact that each coordinate in \(P^{\prime}\) is less than 1 .

\subsection*{15.5.2 \(\epsilon\)-approximation for Tukey function}

Theorem 15.5.4 (Tukey set-query). Let function \(H\) be defined as Definition 15.4.2. There is an algorithm that takes \(O\left(\epsilon^{-1} k \log k\right)\) samples, \(O(k \log k)\) decoding times, \(O(\log k)\) update time, and outputs a vector \(x^{\prime} \in \mathbb{R}^{n}\) such that
\[
\left\|x_{S}^{\prime}-x_{S}\right\|_{H} \leq \epsilon \cdot\left\|x_{[n] \backslash S}\right\|_{H}
\]
holds with probability 9/10.

Proof. Without loss of generality, we assume that \(\tau=1\). We approximate Tukey function to be \(H(x)=1\) for \(|x|>1\) and \(H(x)=x^{2}\) for \(|x| \leq 1\).

Let \(B=O\left(\epsilon^{-1} k\right)\) and \(R=O(\log k)\), and use the count-sketch data structure with bucket size \(O(B)\) and replicate number \(O(R)\). Let \(x^{\prime} \in \mathbb{R}^{n}\) denote the estimate given by the count-sketch data structure.

We define \(P, Q \subseteq[n] \backslash S\) as follows
\[
P=\left\{i| | x_{i} \mid \leq 1, i \in[n] \backslash S\right\}, Q=\left\{i| | x_{i} \mid>1, i \in[n] \backslash S\right\} .
\]

We consider the following two cases.
Case 1. \(|Q|>B / 4\).
We have
\[
\left\|x_{S}^{\prime}-x_{S}\right\|_{H} \leq k=\frac{k}{|Q|} \cdot|Q| \leq \frac{k}{B / 4}|Q| \leq \frac{4 k}{B}\left\|x_{-k}\right\|_{H} \leq \epsilon\left\|x_{-k}\right\|_{H}
\]

Case 2. \(|Q| \leq B / 4\).

By Lemma 15.5.1, for each \(i \in S\), with probability at least \(1-1 / \operatorname{poly}(k)\), we have
\[
\left|x_{i}^{\prime}-x_{i}\right| \leq \sqrt{\frac{\left\|x_{-2 B}\right\|_{2}^{2}}{2 B}} \leq \sqrt{\frac{\left\|x_{P}\right\|_{2}^{2}}{2 B}}
\]

Applying a union bound to coordinates in \(S\), we have with probability at least \(1-\operatorname{poly}(k)\),
\[
\left\|x_{S}^{\prime}-x_{S}\right\|_{H} \leq k H\left(\sqrt{\frac{\left\|x_{P}\right\|_{2}^{2}}{2 B}}\right)
\]

Then Case 2(a) \(\sqrt{\left\|x_{P}\right\|_{2}^{2} / 2 B} \geq 1\). We have
\[
\left\|x_{S}^{\prime}-x_{S}\right\|_{H} \leq k \leq \frac{k}{2 B}\left\|x_{P}\right\|_{2}^{2} \leq \frac{k}{2 B}\left\|x_{-k}\right\|_{H} \leq \epsilon\left\|x_{-k}\right\|_{H}
\]

Case 2(b) \(\sqrt{\left\|x_{P}\right\|_{2}^{2} / 2 B}<1\). We have
\[
\left\|x_{S}^{\prime}-x_{S}\right\|_{H} \leq \frac{k}{2 B}\left\|x_{P}\right\|_{2}^{2} \leq \frac{k}{2 B}\left\|x_{-k}\right\|_{H} \leq \epsilon\left\|x_{-k}\right\|_{H}
\]

\subsection*{15.6 Fourier set query}

In this section, we give set query algorithms from Fourier measurements, when the error is measured with respect to Huber norm and Tukey norm. All of the results in this section follow from the fact that we can simulate the count sketch data structure in Fourier case, by losing a factor of \(\log n\) on the number of measurements (we also \(\operatorname{lose}\) a \(\log n\) factor in update time and query time).

This simulation is standard (as in for example [HIKP12a]) by using (pseudo)random spectral permutation and filtering functions. Let \(B\) denote bucket size. Let \(P=\{i \in\) \([n]\left|\left|\widehat{x}_{i}\right|<\tau\right\}\) and \(Q=\left\{i \in[n]| | \widehat{x}_{i} \mid \geq \tau\right\}\). Let \(P^{\prime}\) denote the set after removing top \(O(k)\) coordinates from \(P\), and similarly let \(Q^{\prime}\) denote the set after removing top \(O(k)\) coordinates from \(Q\). We can guarantee that with high probability for each coordinate \(i \in S\), the error for our estimation to \(\widehat{x}_{i}\) is bounded by \(\sqrt{\frac{\left\|\widehat{x}_{P^{\prime}}\right\|_{2}^{2}+\left\|\widehat{x}_{Q^{\prime}}\right\|_{2}^{2}}{B}}\). Following the proof on set query results in Section 15.5, this yields to \(\ell_{H} / \ell_{H} \leq \epsilon\) guarantee for function \(H\) being Huber function or Tukey function.

Theorem 15.6.1 (Fourier point estimation guarantee). Given a vector \(x \in \mathbb{C}^{n}\), for every \(\epsilon \in(0,1)\) and \(k \geq 1\), any \(S \subseteq[n],|S| \leq k\), there exists an algorithm (Algorithm 15.3) that takes
\[
O\left(\epsilon^{-1} k \log k \cdot \log n\right)
\]
samples, runs in
\[
O\left(\epsilon^{-1} k \log k \cdot \log n\right)
\]
time, and outputs a vector \(x^{\prime} \in \mathbb{C}^{n}\) such that for each \(i \in S\),
\[
\left|x_{i}^{\prime}-\widehat{x}_{i}\right| \leq \sqrt{\frac{\left\|\widehat{x}_{P^{\prime}}\right\|_{2}^{2}+\left\|\widehat{x}_{Q^{\prime}}\right\|_{2}^{2}}{B}}
\]
holds with probability at least \(1-1 / \operatorname{poly}(k)\).

Before showing the proof to Theorem 15.6.1, we first present some basic definitions and basic building blocks for Fourier sparse recovery. Readers can refer to [HIKP12a] for more details.

For a complex vector \(x \in \mathbb{C}^{n}\), we use \(\widehat{x} \in \mathbb{C}^{n}\) to denote its Fourier spectrum,
\[
\widehat{x}_{i}=\frac{1}{\sqrt{n}} \sum_{j=1}^{n} e^{-2 \pi \mathbf{i} i j / n} x_{j}, \forall i \in[n]
\]

Then the inverse transform is
\[
x_{j}=\frac{1}{\sqrt{n}} \sum_{i=1}^{n} e^{2 \pi \mathrm{i} i j / n} \widehat{x}_{i}, \forall j \in[n] .
\]

The discrete convolution of functions \(f\) and \(g\) is given by,
\[
(f * g)[n]=\sum_{m=-\infty}^{+\infty} f[m] g[n-m]
\]

We use the same (pseudorandom) spectrum permutation as [HIKP12a],
Definition 15.6.1. Suppose \(\sigma^{-1}\) exists \(\bmod n\). We define the permutation \(P_{\sigma, a, b}\) by
\[
\left(P_{\sigma, a, b} x\right)_{i}=x_{\sigma(i-a)} e^{-2 \pi \mathbf{i} \sigma b i / n}
\]

We also define \(\pi_{\sigma, b}=\sigma(i-b)(\bmod n)\). Let \(h_{\sigma, b}(i)=\operatorname{round}\left(\pi_{\sigma, b}(i) B / n\right)\) and \(o_{\sigma, b}(i)=\) \(\pi_{\sigma, b}(i)-h_{\sigma, b}(i) n / B\). We say \(h_{\sigma, b}(i)\) is the "bin" that frequency \(i\) is mapped into, and \(o_{\sigma, b}(i)\) is the "offset".

We use the same filter function as [HIKP12a, PS15, CKPS16],

Definition 15.6.2. Given parameters \(B \geq 1, \delta>0, \alpha>0\). We say that \(\left(G, \widehat{G}^{\prime}\right)=\) \(\left(G_{B, \delta, \alpha}, \widehat{G}_{B, \delta, \alpha}^{\prime}\right) \in \mathbb{R}^{n}\) is a filter function if it satisfies the following properties
(I) \(|\operatorname{supp}(G)|=O\left(\alpha^{-1} B \log (n / \delta)\right)\),
(I) \(\widehat{G}_{i}^{\prime}=1\),
\[
\begin{array}{r}
\text { if }|i| \leq(1-\alpha) n /(2 B), \\
\quad \text { if }|i| \geq n /(2 B),
\end{array}
\]
(II) \(\widehat{G}_{i}^{\prime}=0\),
(III) \(\widehat{G}_{i}^{\prime} \in[0,1], \quad\) for all \(i\), (IV) \(\left\|\widehat{G}^{\prime}-\widehat{G}\right\|_{\infty}<\infty\).

We present our Fourier set query algorithm in Algorithm 15.3. We choose bucket size \(B=O(k / \epsilon)\) and number of repetitions \(R=O(\log k)\). For the filter function, we choose \(\delta=1 / \operatorname{poly}(n)\) so that the noise introduced by the filter is negligible. Let \(\alpha\) be a very small constant. We use procedure EstimateValues in \(r\)-th repetition where \(r \in[R]\) to obtain an unbiased estimator for \(\widehat{x}_{S}\) (up to error \(\delta\|\widehat{x}\|_{1}\) ), denoted by \(w^{(r)}\). Finally, for each \(i \in S\), we estimate \(\widehat{x}_{i}\) by \(x_{i}^{\prime} \leftarrow \operatorname{median}_{r \in[R]} w_{i}^{(r)}\). By taking the median of a set of complex numbers, we mean that we take the median of real part and imaginary part separately, and then add these two parts together.

We use three types of events defined in [HIKP12a] as basic building blocks for analyzing Fourier set query algorithms. For any \(i \in S\), we define three types of events associated with \(i\) and \(S\) and defined over the probability space induced by \(\sigma\) and \(b\) :

Definition 15.6.3 (Collision, large offset, large noise). "Collision" event \(E_{\text {coll }}(i)\) : holds iff \(h_{\sigma, b}(i) \in h_{\sigma, b}(S \backslash\{i\})\).
"Large offset" event \(E_{\text {off }}(i)\) : holds iff \(\left|o_{\sigma, b}(i)\right| \geq(1-\alpha) n /(2 B)\).
"Large noise" event \(E_{\text {noise }}(i)\) : holds iff
\[
\mathbb{E}\left[\left\|\widehat{x}_{h_{\sigma, b}^{-1}\left(h_{\sigma, b}(i)\right) \backslash S}^{\prime}\right\|_{2}^{2}\right] \geq \operatorname{Err}^{2}\left(\widehat{x}^{\prime}, k\right) /(\alpha B)
\]

Claim 15.6.2 (Claim 3.1 in [HIKP12a]). For any \(i \in S\), the event \(E_{\text {coll }}(i)\) holds with probability at most \(4|S| / B\).

Claim 15.6.3 (Claim 3.2 in [HIKP12a]). For any \(i \in S\), the event \(E_{\text {off }}(i)\) holds with probability at most \(\alpha\).

Claim 15.6.4 (Claim 4.1 in [HIKP12a]). For any \(i \in S, \operatorname{Pr}\left[E_{\text {noise }(i)}\right] \leq 4 \alpha\).

Lemma 15.6.5 (Lemma 4.2 in [HIKP12a]). Let \(a \in[n]\) uniformly at random, \(B\) divide \(n\), and the other parameters be arbitrary in
\[
\widehat{u}=\operatorname{HaShToBins}\left(x, P_{\sigma, a, b}, B, \delta, \alpha\right) .
\]

Then for any \(i \in[n]\) with \(j=h_{\sigma, b}(i)\) and none of \(E_{\text {coll }}(i), E_{\text {off }}(i)\) or \(E_{\text {noise }}(i)\) holding,
\[
\mathbb{E}\left[\left|\widehat{u}_{j}-\widehat{x}_{i}^{\prime} e^{-\frac{2 \pi \mathrm{i}}{n} a \sigma i}\right|^{2}\right]>2 \frac{\rho^{2}}{\alpha B}
\]

Lemma 15.6.6. Suppose \(B\) divides \(n\). The output \(\widehat{u}\) of HASHTOBINS satisfies
\[
\widehat{u}_{j}=\sum_{h_{\sigma, b}(i)=j} \widehat{x}_{i}\left(\widehat{G_{B, \delta, \alpha}^{\prime}}\right)_{-o_{\sigma, b}(i)} \omega^{a \sigma i} \pm \delta\|\widehat{x}\|_{1} .
\]

The running time of HashToBins is \(O\left(\frac{B}{\alpha} \log (n / \delta)\right)\).

Proof of Theorem 15.6.1. We give Fourier set query algorithm in Algorithm 15.3. The total number of measurements and decoding time is \(O(R) \cdot O\left(B \log \frac{n}{\alpha}\right)=O\left(\epsilon^{-1} k \log k \log n\right)\).

For each repetition \(r\), for each \(i \in S\), we can union bound the probability that none of three bad events (collision, large offset, large noise) happens by \(\frac{1}{10}\). Conditioned on no bad events happened, by Lemma 15.6.5, \(w_{i}^{(r)}\) is a good estimation for \(\widehat{x}_{i}\) in the sense that \(\mathbb{E}\left[\left|\widehat{x}_{i}-w_{i}^{(r)}\right|^{2}\right]=O\left(\frac{\left\|x_{P}\right\|_{2}^{2}+\left\|x_{Q}\right\|_{2}^{2}}{B}\right)\). By Markov inequality with only very small constant probability that \(\left|\widehat{x}_{i}-w_{i}^{(r)}\right| \geq \frac{1}{2} \sqrt{\frac{\left\|x_{P}\right\|_{2}^{2}+\left\|x_{Q}\right\|_{2}^{2}}{B}}\). By applying Chernoff bound, we can prove that with probability \(2^{-\Omega(R)},\left|\widehat{x}_{i}-\operatorname{median}_{r \in[R]} w_{i}^{(r)}\right| \geq \frac{1}{2} \sqrt{\frac{\left\|x_{P}\right\|_{2}^{2}+\left\|x_{Q}\right\|_{2}^{2}}{B}}\).

Corollary 15.6.7 (Fourier set query). Given a vector \(x \in \mathbb{C}^{n}\), for every \(\epsilon \in(0,1)\) and \(k \geq 1\), any \(S \subseteq[n],|S| \leq k\), there exists an algorithm (Algorithm 15.3) that takes
\[
O\left(\epsilon^{-1} k \log k \cdot \log n\right)
\]
samples, runs in
\[
O\left(\epsilon^{-1} k \log k \cdot \log n\right)
\]
time, and outputs a vector \(x^{\prime} \in \mathbb{C}^{n}\) such that
\[
\left\|\left(x^{\prime}-\widehat{x}\right)_{S}\right\|_{H} \leq \epsilon\left\|\widehat{x}_{\bar{S}}\right\|_{H},
\]
holds with probability at least \(1-1 / \operatorname{poly}(k)\). The corollary holds when function \(H\) is Huber function or Tukey function.
```

Algorithm 15.3 Fourier set query algorithm
procedure FourierCountSketch $(x, S, \epsilon, k) \quad \triangleright$ Theorem 15.6.1
$B \leftarrow O(k / \epsilon), R \leftarrow O(\log k), \delta \leftarrow 1 / \operatorname{poly}(n), \alpha \leftarrow 1 / 100$
for $r=1 \rightarrow R$ do
$\widehat{w}^{(r)} \leftarrow \operatorname{EstiamTEVALUES}(x, S, B, \delta, \alpha)$
end for
for $i \in S$ do
$x_{i}^{\prime} \leftarrow \operatorname{median}_{r \in[R]} \widehat{w}_{i}^{(r)}$
end for
return $x_{S}^{\prime}$
end procedure
procedure EstimateValues $(x, S, B, \delta, \alpha)$
Choose $a, b \in[n]$ uniformly at random
Choose $\sigma$ uniformly at random from the set of odd numbers in $[n]$
$\widehat{u} \leftarrow \operatorname{Hash} \operatorname{ToBins}\left(x, P_{\sigma, a, b}, B, \delta, \alpha\right)$
$\widehat{w} \leftarrow 0$
for $t \in S$ do
$\widehat{w}_{t} \leftarrow \widehat{u}_{h_{\sigma, b}(t)} e^{\frac{2 \pi i}{n} \sigma a t}$
end for
return $\widehat{w}$
end procedure
procedure $\operatorname{HAShToBins}\left(x, P_{\sigma, a, b}, B, \delta, \alpha\right)$
Compute $\widehat{y}_{j n / B}$ for $j \in[B]$, where $y=G_{B, \alpha, \delta} \cdot\left(P_{\sigma, a, b} x\right)$
return $\widehat{u}_{j}=\widehat{y}_{j n / B}$
end procedure

```

\subsection*{15.7 Sparse recovery}

\subsection*{15.7.1 \(k\)-sparse recovery for Huber function}

In this subsection, we present our \(k\)-sparse recovery results for Huber function.

Theorem 15.7.1. Let \(H\) denote Huber function defined in Definition 15.4.1. There exists a distribution of \(\Phi \in \mathbb{R}^{m \times n}\) with \(m=O\left(\epsilon^{-2} k \log n\right)\) and a recovery algorithm, such that for each \(x \in \mathbb{R}^{n}, k \in[n]\) and \(\epsilon \in(0,1)\), the recovery algorithm takes \(\Phi x\) as its input and outputs


Figure 15.3: \(\left\|x_{S \cap \bar{T}}\right\|_{H}-\left\|x_{T \cap \bar{S}}\right\|_{H}\) is the displacement error
\(k\)-sparse vector \(x^{\prime} \in \mathbb{R}^{n}\) so that with probability \(1-1 / \operatorname{poly}(n)\),
\[
\left\|x^{\prime}-x\right\|_{H} \leq(1+\epsilon) \min _{k-\text { sparse } y}\|y-x\|_{H} .
\]

The update time is \(O(\log n)\) and the recovery time is \(O(n \log n)\).

Proof. Let \(x_{-k}\) denotes the vector \(x\) with the largest \(k\) coordinates in magnitude zeroed out.
We use count sketch to give each \(x_{i}\) (where \(i \in[n]\) ) an estimate. In particular, we choose the number of buckets to be \(B=O\left(\epsilon^{-2} k\right)\) and the number of repetitions to be \(R=O(\log n)\). Let \(x_{i}^{\prime}\) denote the estimate obtained by the count sketch for \(x_{i}\). Let \(T\) denote the largest \(k\) coordinates in absolute value in \(x^{\prime}\). The recovery algorithm outputs \(x_{T}^{\prime}\) as a \(k\)-sparse approximation to \(x\).

Note that the count sketch data structure takes \(B \cdot R=O\left(\epsilon^{-2} k \log n\right)\) measurements. Its update time is \(O(R)=O(\log n)\) time, and the time to obtain \(x^{\prime}\) is \(O(n \cdot R)=O(n \log n)\). Taking largest \(k\) coordinates from \(x^{\prime}\) requires \(O(n \log n)\) time. Thus the total query time is \(O(n \log n)\). In the following, we show that \(\left\|x-x^{\prime}\right\|_{H} \leq(1+\epsilon)\left\|x_{-k}\right\|_{H}\).

By Lemma 15.5.1, for each \(i \in[n]\), with probability \(1-1 / \operatorname{poly}(n)\),
\[
\begin{equation*}
\left(x_{i}-x^{\prime}{ }_{i}\right)^{2} \leq \frac{\epsilon^{2}}{64 k}\left\|x_{-64 k / \epsilon^{2}}\right\|_{2}^{2} \leq \frac{\epsilon^{2}}{64 k}\left\|x_{-64 k}\right\|_{2}^{2} \tag{15.5}
\end{equation*}
\]

Let \(S\) denote the largest \(k\) coordinates of \(x\) in magnitude. The error is
\[
\begin{equation*}
\left\|x-x^{\prime}\right\|_{H}=\sum_{i \in T} H\left(x_{i}-x_{i}^{\prime}\right)+\sum_{i \in \bar{T}} H\left(x_{i}\right) . \tag{15.6}
\end{equation*}
\]

By point estimation guarantee given in Eq. (15.5) and results from Huber set query,
\[
\begin{aligned}
\sum_{i \in T} H\left(x_{i}-x_{i}^{\prime}\right) & \leq|T| \cdot H\left(\frac{\epsilon}{8 \sqrt{k}}\left\|x_{-64 k}\right\|_{2}\right) \\
& \leq \frac{\epsilon}{8}\left\|x_{-k}\right\|_{H}
\end{aligned}
\]

Note that
\[
\begin{aligned}
\sum_{i \in \bar{T}} H\left(x_{i}\right) & =\left(\sum_{i \in \bar{T}} H\left(x_{i}\right)-\sum_{i \in \bar{S}} H\left(x_{i}\right)\right)+\sum_{i \in \bar{S}} H\left(x_{i}\right) \\
& =\left(\sum_{i \in S \cap \bar{T}} H\left(x_{i}\right)-\sum_{i \in \bar{S} \cap T} H\left(x_{i}\right)\right)+\left\|x_{-k}\right\|_{H}
\end{aligned}
\]
where the second step follows by canceling out the coordinates in \(\bar{S} \cap \bar{T}\). In the following, our goal is to show that \(\sum_{i \in S \cap \bar{T}} H\left(x_{i}\right)-\sum_{i \in \bar{S} \cap T} H\left(x_{i}\right) \leq \frac{7}{8} \epsilon\left\|x_{-k}\right\|_{H}\).

If \(S \cap \bar{T}=\emptyset\), then \(S=T\) and \(\sum_{i \in S \cap \bar{T}} H\left(x_{i}\right)-\sum_{i \in \bar{S} \cap T} H\left(x_{i}\right)=0\). In the following, we assume that \(S \cap \bar{T} \neq \emptyset\). Let \(b=\max _{i \in S \cap \bar{T}}\left|x_{i}\right|\), and \(a=\min _{i \in \bar{S} \cap T}\left|x_{i}\right|\). Because Huber function \(H(t)\) is monotonically increasing as \(|t|\) increases,
\[
\sum_{i \in S \cap \bar{T}} H\left(x_{i}\right)-\sum_{i \in \bar{S} \cap T} H\left(x_{i}\right) \leq|\bar{S} \cap T| \cdot(H(b)-H(a)) .
\]

Our final goal is to prove that \(H(b)-H(a) \leq \frac{7 \epsilon\left\|x_{-k}\right\|_{H}}{8|\bar{S} \cap T|}\). We discuss in three cases: case (1) \(\tau \leq a\), case (2) \(a<\tau<b\), case (3) \(b \leq \tau\). We analyze these three cases in Claim 15.7.2, Claim 15.7.3 and Claim 15.7.4.

Before diving into case analysis, we first prove two inequalities: one is an upper bound of \(b-a\), and the other is an upper b ound of \(a\). They will be used repeatedly in the case analysis.

By Eq. (15.5), any \(i \in[n],\left|x_{i}-x_{i}^{\prime}\right| \leq \frac{\epsilon}{8 \sqrt{k}}\left\|x_{-64 k}\right\|_{2}\). That is, \(\left\|x-x^{\prime}\right\|_{\infty} \leq \Delta\) where \(\Delta=\frac{\epsilon}{8 \sqrt{k}}\left\|x_{-64 k}\right\|_{2}\). By Lemma 15.7.5,
\[
\begin{equation*}
b-a \leq \frac{\epsilon}{4 \sqrt{k}}\left\|x_{-64 k}\right\|_{2} . \tag{15.7}
\end{equation*}
\]

Moreover, by definition that \(a=\min _{i \in \bar{S} \cap T}\left|x_{i}\right|\),
\[
\begin{equation*}
a \leq \frac{\left\|x_{\bar{S} \cap T}\right\|_{2}}{\sqrt{|\bar{S} \cap T|}} \leq \frac{\left\|x_{-k}\right\|_{2}}{\sqrt{|\bar{S} \cap T|}} \tag{15.8}
\end{equation*}
\]

Claim 15.7.2. If \(\tau \leq a\), then \(H(b)-H(a) \leq \frac{7 \epsilon}{8} \frac{\left\|x_{-k}\right\|_{H}}{|\bar{S} \cap T|}\).

Proof. Let \(P, Q\) be a partition of \(\bar{S}\) where
\[
P=\left\{i| | x_{i} \mid \leq \tau, i \in \bar{S}\right\}, Q=\left\{i| | x_{i} \mid>\tau, i \in \bar{S}\right\} .
\]

Note that the coordinates in \(x_{-64 k}\), which equals to \(\bar{S}=P \cup Q\) with largest (in absolute value) \(63 k\) coordinates removed, are fully contained in \(P\) union \(Q\) with largest (in absolute value) \(63 k\) coordinates removed. By triangle inequality,
\[
\begin{equation*}
\left\|x_{-64 k}\right\|_{2} \leq\left\|x_{P}\right\|_{2}+\left\|\left(x_{Q}\right)_{-63 k}\right\|_{2} . \tag{15.9}
\end{equation*}
\]

In the following, we are going to bound both \(\left\|x_{P}\right\|_{2}\) and \(\left\|\left(x_{Q}\right)_{-63 k}\right\|_{2}\) in terms of \(\left\|x_{-k}\right\|_{H}\).

By definition of Huber function \(\left\|x_{-k}\right\|_{H}=\frac{1}{\tau}\left\|x_{P}\right\|_{2}^{2}+\left\|x_{Q}\right\|_{1}\). Because \(\tau \leq a, \bar{S} \cap T \subseteq Q\). Thus \(\left\|x_{Q}\right\|_{1} \geq\left\|x_{\bar{S} \cap T}\right\|_{1} \geq \tau|\bar{S} \cap T|\). Therefore,
\[
\begin{aligned}
\left\|x_{-k}\right\|_{H} & \geq \frac{1}{\tau}\left\|x_{P}\right\|_{2}^{2}+\tau|\bar{S} \cap T| \\
& \geq 2 \sqrt{\frac{1}{\tau}\left\|x_{P}\right\|_{2}^{2} \cdot \tau|\bar{S} \cap T|}=2 \sqrt{|\bar{S} \cap T|} \cdot\left\|x_{P}\right\|_{2}
\end{aligned}
\]

So we have
\[
\left\|x_{P}\right\|_{2} \leq \frac{1}{2 \sqrt{|\bar{S} \cap T|}}\left\|x_{-k}\right\|_{H}
\]

In order to bound \(\left\|\left(x_{Q}\right)_{-63 k}\right\|_{2}\), by Lemma 15.5.2,
\[
\begin{aligned}
\left\|\left(x_{Q}\right)_{-63 k}\right\|_{2} & \leq \frac{1}{\sqrt{63 k}}\left\|x_{Q}\right\|_{1} \leq \frac{1}{7 \sqrt{|\bar{S} \cap T|}}\left\|x_{Q}\right\|_{1} \\
& \leq \frac{1}{7 \sqrt{|\bar{S} \cap T|}}\left\|x_{-k}\right\|_{H}
\end{aligned}
\]

Therefore,
\[
\begin{aligned}
& H(b)-H(a) \\
= & b-a \\
\leq & \frac{\epsilon}{4 \sqrt{|\bar{S} \cap T|}}\left\|x_{-64 k}\right\|_{2} \\
\leq & \frac{\epsilon}{4 \sqrt{|\bar{S} \cap T|}}\left(\left\|x_{P}\right\|_{2}+\left\|\left(x_{Q}\right)_{-63 k}\right\|_{2}\right) \\
\leq & \frac{\epsilon}{4 \sqrt{|\bar{S} \cap T|}}\left(\frac{1}{2 \sqrt{|\bar{S} \cap T|}}\left\|x_{-k}\right\|_{H}+\frac{1}{7 \sqrt{|\bar{S} \cap T|}}\left\|x_{-k}\right\|_{H}\right) \\
\leq & \frac{\epsilon\left\|x_{-k}\right\|_{H}}{4|\bar{S} \cap T|}
\end{aligned}
\]
where the first line follows from definition of Huber function and \(a \geq \tau\), the second line follows by Eq. (15.7), the third follows by Eq. (15.9), the forth line follows by the bounds of \(\left\|x_{P}\right\|_{2}\) and \(\left\|\left(x_{Q}\right)_{-63 k}\right\|_{2}\) we just obtained.

Claim 15.7.3. If \(a<\tau<b\), then \(H(b)-H(a) \leq \frac{7 \epsilon}{8} \frac{\|x-k\|_{H}}{|\bar{S} \cap T|}\).

Proof. Let \(P=\left\{i| | x_{i} \mid \leq \tau, i \in \bar{S}\right\}, Q=\left\{i| | x_{i} \mid>\tau, i \in \bar{S}\right\}\). We discuss in two cases depending on whether the size of \(Q\) is greater than \(9 k\) or not.
(1) If \(|Q| \leq 9 k\). We first bound \(H(b)-H(a)\) as follows
\[
\begin{aligned}
& H(b)-H(a) \\
= & b-a^{2} / \tau \\
\leq & b^{2} / \tau-a^{2} / \tau \\
= & \frac{1}{\tau}(b-a)(b+a) \\
\leq & \frac{\epsilon\left\|x_{-10 k}\right\|_{2}}{4 \tau \sqrt{k}}\left(\frac{\epsilon\left\|x_{-10 k}\right\|_{2}}{4 \sqrt{k}}+2 a\right) \\
\leq & \frac{\epsilon\left\|x_{-10 k}\right\|_{2}}{4 \tau \sqrt{k}}\left(\frac{\epsilon\left\|x_{-10 k}\right\|_{2}}{4 \sqrt{k}}+2 \frac{\left\|x_{P}\right\|_{2}}{\sqrt{|\bar{S} \cap T|}}\right) \\
= & \frac{\epsilon^{2}\left\|x_{-10 k}\right\|_{2}^{2}}{16 \tau k}+\frac{\epsilon\left\|x_{-10 k}\right\|_{2}\left\|x_{P}\right\|_{2}}{2 \sqrt{k} \sqrt{|\bar{S} \cap T|}} \\
\leq & \frac{\epsilon}{\tau|\bar{S} \cap T|}\left(\frac{1}{16}\left\|x_{-10 k}\right\|_{2}^{2}+\frac{1}{2}\left\|x_{-10 k}\right\|_{2} \cdot\left\|x_{P}\right\|_{2}\right)
\end{aligned}
\]
where the forth step follows by Eq.(15.7), and the fifth step follows by Eq.(15.8).
Because \(|Q| \leq 9 k,\left\|x_{-10 k}\right\|_{2} \leq\left\|x_{P}\right\|_{2}\), and thus
\[
H(b)-H(a) \leq \frac{\epsilon}{\tau|\bar{S} \cap T|} \cdot \frac{9}{16}\left\|x_{P}\right\|_{2}^{2} \leq \frac{9 \epsilon\left\|x_{-k}\right\|_{H}}{16|\bar{S} \cap T|}
\]
(2) If \(|Q|>9 k\). First note that
\[
b-a^{2} / \tau \leq 2(b-a)
\]
because
\[
b+a^{2} / \tau \geq 2 \sqrt{b \cdot a^{2} / \tau} \geq 2 a
\]

Furthermore, \(b-a \leq \frac{\epsilon\left\|x_{-64}\right\|_{2}}{4 \sqrt{k}} \leq \frac{\epsilon\left\|x_{-10}\right\|_{2}}{4 \sqrt{k}}\), where
\[
\left\|x_{-10 k}\right\|_{2} \leq\left\|\left(x_{Q}\right)_{-9 k}\right\|_{2}+\left\|x_{P}\right\|_{2} \leq \frac{1}{\sqrt{9 k}}\left\|x_{Q}\right\|_{1}+\left\|x_{P}\right\|_{2}
\]

Also note that \(\left\|x_{-k}\right\|_{H}=\left\|x_{Q}\right\|_{1}+\frac{1}{\tau}\left\|x_{P}\right\|_{2}^{2} \geq 2 \sqrt{\left\|x_{Q}\right\|_{1} \cdot\left\|x_{P}\right\|_{2}^{2} / \tau} \geq 6 \sqrt{k}\left\|x_{P}\right\|_{2}\). Therefore,
\[
\begin{aligned}
H(b)-H(a) & =b-a^{2} / \tau \\
& \leq \frac{2 \epsilon}{4 \sqrt{k}}\left\|x_{-10}\right\|_{2} \\
& \leq \frac{2 \epsilon}{4 \sqrt{k}}\left(\frac{1}{\sqrt{9 k}}\left\|x_{-k}\right\|_{H}+\frac{1}{6 \sqrt{k}}\left\|x_{-k}\right\|_{H}\right) \\
& =\frac{\epsilon\left\|x_{-k}\right\|_{H}}{4 k} \leq \frac{\epsilon\left\|x_{-k}\right\|_{H}}{4|\bar{S} \cap T|} .
\end{aligned}
\]

Claim 15.7.4. If \(b \leq \tau\), then \(H(b)-H(a) \leq \frac{7 \in}{8} \frac{\left\|x_{-k}\right\|_{H}}{|\bar{S} \cap T|}\).

Proof. The analysis is similar to classical \(\ell_{2} / \ell_{2}\) sparse recovery case. We have
\[
\begin{aligned}
& H(b)-H(a) \\
= & \frac{b^{2}}{\tau}-\frac{a^{2}}{\tau} \\
= & \frac{1}{\tau}(b-a)(b+a) \\
= & \frac{1}{\tau}(b-a)(b-a+2 a) \\
\leq & \frac{1}{\tau} \cdot \frac{\epsilon}{4 \sqrt{k}}\left\|x_{-k}\right\|_{2} \cdot\left(\frac{\epsilon}{4 \sqrt{k}}\left\|x_{-k}\right\|_{2}+2 \frac{\left\|x_{-k}\right\|_{2}}{\sqrt{|\bar{S} \cap T|}}\right) \\
\leq & \frac{1}{\tau} \cdot \frac{\epsilon\left\|x_{-k}\right\|_{2}}{4 \sqrt{|\bar{S} \cap T|}} \cdot\left(\frac{\left\|x_{-k}\right\|_{2}}{4 \sqrt{|\bar{S} \cap T|}}+2 \frac{\left\|x_{-k}\right\|_{2}}{\sqrt{|\bar{S} \cap T|}}\right) \\
= & \frac{1}{\tau} \cdot \frac{9 \epsilon\left\|x_{-k}\right\|_{2}^{2}}{16|\bar{S} \cap T|} \\
= & \frac{9 \epsilon\left\|x_{-k}\right\|_{H}}{16|\bar{S} \cap T|}
\end{aligned}
\]
where the first step follows by definition of Huber function, the forth step follows by Eq. (15.7) and Eq.(15.8), the fifth step follows by \(|\bar{S} \cap T| \leq|T|=k\) and \(\epsilon \leq 1\), the last step follows because by definition of \(b=\max _{i \in S \cap \bar{T}}\left|x_{i}\right|\), for any \(j \in[n] \backslash S,\left|x_{j}\right| \leq \min _{i \in S}\left|x_{i}\right| \leq b \leq \tau\) and then by definition of Huber function \(\left\|x_{-k}\right\|_{H}=\frac{1}{\tau}\left\|x_{-k}\right\|_{2}^{2}\).

Lemma 15.7.5. Let vector \(u, v \in \mathbb{R}^{n}\) such that \(\|u-v\|_{\infty} \leq \Delta\). Let \(S\) be the largest \(k\) coordinates in \(u\), and let \(T\) be the largest \(k\) coordinates in \(v\). Let \(b=\max _{i \in S \cap \bar{T}} u_{i}\) and \(a=\min _{i \in \bar{S} \cap T} u_{i}\). Then \(b-a \leq 2 \Delta\).

Proof. Let \(\lambda=\max _{i \in \bar{T}} v_{i}\). We have
\[
b=\max _{i \in S \cap \bar{T}} u_{i} \leq \max _{i \in \bar{T}} u_{i} \leq \Delta+\max _{i \in \bar{T}} v_{i}=\Delta+\lambda .
\]

On the other hand,
\[
a=\min _{i \in \bar{S} \cap T} u_{i} \geq \min _{i \in T} u_{i} \geq-\Delta+\min _{i \in T} v_{i} \geq-\Delta+\lambda .
\]

Therefore
\[
b-a \leq(\Delta+\lambda)-(-\Delta+\lambda) \leq 2 \Delta
\]

\subsection*{15.7.2 \(k\)-sparse recovery for Tukey function}

In this subsection, we present our \(k\)-sparse recovery results for Tukey function. Similar to results for Huber function, we use count-sketch to estimate each coordinate and then find the top \(k\) ones. Moreover, our proof for Tukey function has a similar structure, where we first decompose estimation error into set query error and displacement error, and then discuss the displacement error in several cases.

Theorem 15.7.6. Let \(H\) denote Tukey function defined in Definition 15.4.2. There exists a distribution of \(\Phi \in \mathbb{R}^{m \times n}\) with \(m=O\left(\epsilon^{-2} k \log n\right)\) and a recovery algorithm, such that for each \(x \in \mathbb{R}^{n}, k \in[n]\) and \(\epsilon \in(0,1)\), the recovery algorithm takes \(\Phi x\) as its input and outputs \(k\)-sparse vector \(x^{\prime} \in \mathbb{R}^{n}\) so that with probability \(1-1 / \operatorname{poly}(n)\),
\[
\left\|x^{\prime}-x\right\|_{H} \leq(1+\epsilon) \min _{k-\text { sparse } y}\|y-x\|_{H}
\]

The update time is \(O(\log n)\) and the recovery time is \(O(n \log n)\).

Proof. Let \(x_{-k}\) denotes the vector \(x\) with absolute largest \(k\) coordinates removed.

We use count sketch to give each \(x_{i}\) (where \(i \in[n]\) ) an estimate. In particular, we choose the number of buckets \(B=O\left(\epsilon^{-2} k\right)\) and the number of duplicates \(R=O(\log n)\) for the count sketch data structure. Let \(x_{i}^{\prime}\) denote the estimate obtained by the count sketch for \(x_{i}\). Let \(T\) denote the largest \(k\) coordinates in absolute value in \(x^{\prime}\). The recovery algorithm outputs \(x_{T}^{\prime}\) as a \(k\)-sparse approximation to \(x\).

Note that the count sketch data structure takes \(B \cdot R=O\left(\epsilon^{-2} k \log n\right)\) measurements. Its update time is \(O(R)=O(\log n)\) time, and the time to obtain \(x^{\prime}\) is \(O(n \cdot R)=O(n \log n)\). Taking largest \(k\) coordinates from \(x^{\prime}\) requires \(O(n \log n)\) time. Thus the total query time is \(O(n \log n)\). In the following, we show that \(\left\|x-x_{T}^{\prime}\right\|_{H} \leq(1+\epsilon)\left\|x_{-k}\right\|_{H}\).

By Lemma 15.5.1, for each \(i \in[n]\), with probability \(1-1 / \operatorname{poly}(n)\),
\[
\left(x_{i}-x_{i}^{\prime}\right)^{2} \leq \frac{\epsilon^{2}}{10 k}\left\|x_{-10 k / \epsilon^{2}}\right\|_{2}^{2}
\]

Let \(S\) denote the largest \(k\) coordinates in absolute value in \(x\). The recovery error is
\[
\left\|x-x_{T}^{\prime}\right\|_{H}=\sum_{i \in T} H\left(x_{i}-x_{i}^{\prime}\right)+\sum_{i \in \bar{T}} H\left(x_{i}\right) .
\]

By point estimation guarantee and results from Tukey set query in Theorem 15.2.1,
\[
\sum_{i \in T} H\left(x_{i}-x_{i}^{\prime}\right) \leq|T| \cdot H\left(\frac{\epsilon}{\sqrt{10 k}}\left\|x_{-4 k}\right\|_{2}\right) \leq \frac{\epsilon}{3}\left\|x_{-k}\right\|_{H} .
\]

Note that
\[
\begin{aligned}
\sum_{i \in \bar{T}} H\left(x_{i}\right) & =\left(\sum_{i \in \bar{T}} H\left(x_{i}\right)-\sum_{i \in \bar{S}} H\left(x_{i}\right)\right)+\sum_{i \in \bar{S}} H\left(x_{i}\right) \\
& =\left(\sum_{i \in S \cap \bar{T}} H\left(x_{i}\right)-\sum_{i \in \bar{S} \cap T} H\left(x_{i}\right)\right)+\left\|x_{-k}\right\|_{H}
\end{aligned}
\]
where the second step follows by canceling out the coordinates in \(\bar{S} \cap \bar{T}\). In the following, our goal is to show that \(\sum_{i \in S \cap \bar{T}} H\left(x_{i}\right)-\sum_{i \in \bar{S} \cap T} H\left(x_{i}\right) \leq \frac{2}{3} \epsilon\left\|x_{-k}\right\|_{H}\).

If \(S \cap \bar{T}=\emptyset\), then \(S=T\) and \(\sum_{i \in S \cap \bar{T}} H\left(x_{i}\right)-\sum_{i \in \bar{S} \cap T} H\left(x_{i}\right)=0\). In the following, we assume that \(S \cap \bar{T} \neq \emptyset\). Let \(b=\max _{i \in S \cap \bar{T}}\left|x_{i}\right|\), and \(a=\min _{i \in \bar{S} \cap T}\left|x_{i}\right|\). Because Huber function \(H(t)\) is monotonically increasing as \(|t|\) increases,
\[
\sum_{i \in S \cap \bar{T}} H\left(x_{i}\right)-\sum_{i \in \bar{S} \cap T} H\left(x_{i}\right) \leq|\bar{S} \cap T| \cdot(H(b)-H(a)) .
\]

By Lemma 15.7.5
\[
b-a \leq \frac{\epsilon}{\sqrt{10 k}}\left\|x_{-10 k / \epsilon^{2}}\right\|_{2} .
\]

In the following, we are going to prove that \(H(b)-H(a) \leq \frac{\epsilon}{2} \frac{\left\|x_{-k}\right\|_{H}}{|\bar{S} \cap T|}\). We discuss in three cases: case (1) \(\tau \leq a\), case (2) \(a<\tau<b\), case (3) \(b \leq \tau\).

Case (1) \(\tau \leq a\).

In this case \(H(a)=H(b)=1\) and \(H(b)-H(a)=0\).
Case (2) \(a<\tau<b\).
Let \(P=\left\{i \in \bar{S}| | x_{i} \mid<\tau\right\}\), and \(Q=\left\{i \in \bar{S}| | x_{i} \mid \geq \tau\right\}\). We have
\[
\left\|x_{-k}\right\|_{H}=|Q|+\left\|x_{P}\right\|_{2}^{2}
\]
and
\[
\begin{aligned}
H(b)-H(a) & =1-a^{2}=(1-a)(1+a) \\
& \leq \frac{\epsilon\left\|x_{-10 k / \epsilon^{2}}\right\|_{2}}{\sqrt{k}} \cdot(1+a) \leq \frac{2 \epsilon\left\|x_{-10 k / \epsilon^{2}}\right\|_{2}}{\sqrt{k}} .
\end{aligned}
\]

If \(|Q|>\epsilon^{-2} k\), then we have \(H(b)-H(a) \leq 1\) while \(\left\|x_{-k}\right\|_{H} \geq|Q|>\epsilon^{-2} k\). Thus \(H(b)-\) \(H(a) \leq \frac{\epsilon}{k}\left\|x_{-k}\right\|_{H} \leq \frac{\epsilon}{|S \cap \bar{T}|}\left\|x_{-k}\right\|_{H}\).

If \(|Q| \leq \epsilon^{-2} k\) and \(\left\|x_{-10 k / \epsilon^{2}}\right\|_{2}>\sqrt{k}\), then \(\left\|x_{P}\right\|_{2} \geq\left\|x_{-10 k / \epsilon^{2}}\right\|_{2}>\sqrt{k}\), and so \(\left\|x_{-10 k / \epsilon^{2}}\right\|_{2} \leq \frac{\left\|x_{P}\right\|_{2}}{\sqrt{k}} \cdot\left\|x_{P}\right\|_{2}=\frac{\left\|x_{P}\right\|_{2}^{2}}{\sqrt{k}} \leq \frac{1}{\sqrt{k}}\left\|x_{-k}\right\|_{H}\). Thus \(H(b)-H(a) \leq \frac{2 \epsilon}{k}\left\|x_{-k}\right\|_{H} \leq\) \(\frac{2 \epsilon}{|S \cap \bar{T}|}\left\|x_{-k}\right\|_{H}\).

If \(|Q| \leq \epsilon^{-2} k\) and \(\left\|x_{-10 k / \epsilon^{2}}\right\|_{2} \leq \sqrt{k}\). We have \(\left\|x_{P}\right\|_{2} \geq\left\|x_{-10 k / \epsilon^{2}}\right\|_{2}\). By definition that \(a=\min _{i \in \bar{S} \cap T}\left|x_{i}\right|, a \leq \frac{\left\|x_{\bar{S} \cap T}\right\|_{2}}{\sqrt{|\bar{S} \cap T|}} \leq \frac{\left\|x_{P}\right\|_{2}}{\sqrt{|\bar{S} \cap T|}}\). Therefore,
\[
\begin{aligned}
H(b)-H(a) & \leq \frac{\epsilon\left\|x_{-10 k / \epsilon^{2}}\right\|_{2}}{\sqrt{k}} \cdot\left(\frac{\epsilon\left\|x_{-10 k / \epsilon^{2}}\right\|_{2}}{\sqrt{k}}+2 \frac{\left\|x_{P}\right\|_{2}}{\sqrt{|\bar{S} \cap T|}}\right) \\
& \leq \frac{\epsilon\left\|x_{P}\right\|_{2}}{\sqrt{|\bar{S} \cap T|}} \cdot 3 \frac{\left\|x_{P}\right\|_{2}}{\sqrt{|\bar{S} \cap T|}} \\
& \leq \frac{3 \epsilon}{|\bar{S} \cap T|}\left\|x_{P}\right\|_{2}^{2} \\
& \leq \frac{3 \epsilon}{|\bar{S} \cap T|}\left\|x_{-k}\right\|_{H} .
\end{aligned}
\]

Case (3) \(b \leq \tau\).

The analysis is similar to classical \(\ell_{2} / \ell_{2}\) sparse recovery case. We have
\[
\begin{aligned}
H(b)-H(a) & =\frac{b^{2}}{\tau}-\frac{a^{2}}{\tau} \\
& =\frac{1}{\tau}(b-a)(b+a) \\
& =\frac{1}{\tau}(b-a)(b-a+2 a) \\
& \leq \frac{1}{\tau} \cdot \frac{\epsilon}{\sqrt{k}}\left\|x_{-k}\right\|_{2} \cdot\left(\frac{\epsilon}{\sqrt{k}}\left\|x_{-k}\right\|_{2}+2 \frac{\left\|x_{-k}\right\|_{2}}{\sqrt{|\bar{S} \cap T|}}\right) \\
& \leq \frac{1}{\tau} \cdot \frac{\epsilon\left\|x_{-k}\right\|_{2}}{\sqrt{|\bar{S} \cap T|}} \cdot\left(\frac{\left\|x_{-k}\right\|_{2}}{\sqrt{|\bar{S} \cap T|}}+2 \frac{\left\|x_{-k}\right\|_{2}}{\sqrt{|\bar{S} \cap T|}}\right) \\
& =\frac{1}{\tau} \cdot \frac{3 \epsilon\left\|x_{-k}\right\|_{2}^{2}}{|\bar{S} \cap T|} \\
& =\frac{3 \epsilon\left\|x_{-k}\right\|_{H}}{|\bar{S} \cap T|},
\end{aligned}
\]
where the first step follows by definition of Huber function, the forth step follows by Eq. (15.7) and Eq.(15.8), the fifth step follows by \(|\bar{S} \cap T| \leq|T|=k\) and \(\epsilon \leq 1\), the last step follows because by definition of \(b=\max _{i \in S \cap \bar{T}}\left|x_{i}\right|\), for any \(j \in[n] \backslash S,\left|x_{j}\right| \leq \min _{i \in S}\left|x_{i}\right| \leq b \leq \tau\) and then by definition of Huber function \(\left\|x_{-k}\right\|_{H}=\frac{1}{\tau}\left\|x_{-k}\right\|_{2}^{2}\).

\section*{Part IV}

\section*{Matrix Factorizations}

\section*{Chapter 16}

\section*{Weighted Low Rank Approximation}

The classical low rank approximation problem is: given a matrix \(A\), find a rank- \(k\) matrix \(B\) such that the Frobenius norm of \(A-B\) is minimized. It can be solved efficiently using, for instance, the Singular Value Decomposition (SVD). If one allows randomization and approximation, it can be solved in time proportional to the number of non-zero entries of \(A\) with high probability.

Inspired by practical applications, we consider a weighted version of low rank approximation: for a non-negative weight matrix \(W\) we seek to minimize \(\sum_{i, j}\left(W_{i, j} \cdot\left(A_{i, j}-B_{i, j}\right)\right)^{2}\). The classical problem is a special case of this problem when all weights are 1. Weighted low rank approximation is known to be NP-hard, so we are interested in a meaningful parametrization that would allow efficient algorithms.

In this paper we present several efficient algorithms for the case of small \(k\) and under the assumption that the weight matrix \(W\) is of low rank, or has a small number of distinct columns. An important feature of our algorithms is that they do not assume anything about the matrix \(A\). We also obtain lower bounds that show that our algorithms are nearly optimal in these parameters. We give several applications in which these parameters are small. To the best of our knowledge, the present paper is the first to provide algorithms for the weighted low rank approximation problem with provable guarantees.

Perhaps even more importantly, our algorithms proceed via a new technique, which we call "guess the sketch". The technique turns out to be general enough to give solutions to several other fundamental problems: adversarial matrix completion, weighted non-negative matrix factorization and tensor completion.

\subsection*{16.1 Introduction}

Low rank approximation is arguably one of the most well-studied problems in randomized numerical linear algebra, with diverse applications to clustering [DFK \({ }^{+}\)04, FSS13, LBKW14, \(\mathrm{CEM}^{+}\)15], data mining [ \(\mathrm{AFK}^{+} 01\) ], distance matrix completion [Cha12], information retrieval [PRTV00], learning mixtures of distributions [AM05, KSV08], recommendation systems [DKR02], and web search [AFKM01, Kle99]. In practice one often has a low rank matrix which has been corrupted with noise of bounded norm, and low rank approximation allows one to approximately recover the original matrix. Low rank approximation may also help explain a dataset, revealing low dimensional structure in high dimensional data. Given a low rank approximation, one can store a matrix and compute a matrix-vector product much more efficiently by storing the corresponding factorization. It can also be used as a preprocessing step in applications, that is, by first projecting data onto a lower-dimensional subspace one preserves important properties of the input, but can now run subsequent algorithms in the lower-dimensional space. For example, it has been proposed to reduce the data dimension in Non-Negative Matrix Factorization [LS00] (more on this below), and Latent Dirichlet Allocation (LDA) [BNJ01].

The basic low rank approximation problem is: given an \(n \times n\) matrix \(A\), find a matrix \(\widehat{A}\) of rank at most \(k\) for which \(\|A-\widehat{A}\|_{F}\) is minimized, where for a matrix \(B\), \(\|B\|_{F}=\left(\sum_{i, j} B_{i, j}^{2}\right)^{1 / 2}\) is its Frobenius norm. This formulation intuitively corresponds to the matrix \(\widehat{A}\) capturing as much of the variance of \(A\) as possible. It is well-known that the optimal solution is given by \(A_{k}\), which if \(U \Sigma V^{\top}\) is the singular value decomposition (SVD) of \(A\), where \(U\) and \(V\) are orthogonal matrices and \(\Sigma\) is a non-negative diagonal matrix with \(\Sigma_{1,1} \geq \Sigma_{2,2} \geq \cdots \geq \Sigma_{n, n}\), then \(A_{k}=U \Sigma_{k} V^{\top}\), where \(\Sigma_{k}\) agrees with \(\Sigma\) on its first \(k\)
diagonal entries and is 0 otherwise. Although the SVD is computable in polynomial time, it is often acceptable to output a matrix \(\widehat{A}\) for which \(\|A-\widehat{A}\|_{F} \leq(1+\epsilon)\left\|A-A_{k}\right\|_{F}\) with high probability. In the latter case, much more efficient algorithms are known, and it is possible to compute such an \(\widehat{A}\) in \(n n z(A)+n \cdot \operatorname{poly}(k / \epsilon)\) time, where \(n n z(A)\) denotes the number of non-zero entries of \(A\) [CW13, MM13, NN13a]. We note that for typical applications \(k\) and \(1 / \epsilon\) are assumed to be much smaller than \(n\), e.g., in [Har14] they are treated as absolute constants.

Despite the large body of work on low rank approximation, the weighted case is not well understood. In this case one is given an \(n \times n\) matrix \(A\) and an \(n \times n\) matrix \(W\) with \(W_{i, j} \geq 0\), and one seeks to solve:
\[
\min _{\text {rank-k matrices } \widehat{A}}\|W \circ(A-\widehat{A})\|_{F}^{2}=\min _{\text {rank-k matrices }} \sum_{i, j} W_{i, j}^{2}\left(A_{i, j}-\widehat{A}_{i, j}\right)^{2} .
\]

The classical low rank approximation is a special case in which \(W_{i, j}=1\) for all \(i\) and \(j\). However, in general there may not be a good reason to weight all elements of the approximation error \(A-\widehat{A}\) equally, especially if one is given prior knowledge about the distribution of the errors. For example, suppose the columns of \(A\) each come from a low-dimensional subspace but one of the columns is then shifted by a fixed large vector so that its mean is different. One may first want to recenter the data by subtracting off the mean from each of the columns. While this is possible without weighted low rank approximation, suppose instead that each of the columns of \(A\) comes from a perturbation of columns in a low dimensional subspace but one of the columns has a much larger variance. Then if all weights were equal, it would be enough for \(\widehat{A}\) to fit this one single large variance column, which fails to capture the entire low-dimensional subspace. One way of fixing this is to reweight each entry of \(A\)
by the inverse of its variance. This is a common technique used in gene expression analysis, where the error model for microarray measurements provides entry-specific noise estimates, or when entries of \(A\) represent aggregates of many samples such as in word co-occurrence matrices and non-uniform weights are needed to appropriately capture any differences in the sample sizes; see [SJ03] for a discussion, and also the Wikipedia entry on weighted low rank approximation for a brief introduction.

While the extension of low rank approximation to the weighted case goes back to work of Young in 1940 [You40], its complexity is not well-understood, partly because the weighted case does not admit a solution via the SVD and may have many local minima [SJ03]. Early work by Shpak [Shp90] looked at gradient-based approaches while Lu et al. [LPW97, LA03] looked at alternating minimization methods. These were significantly sped up in practice by the work of Srebro and Jaakkola [SJ03], with success in various applications such as color image restoration [MES08], though there are no provable time bounds and in the worst case the running times could be exponential or worse. In fact, weighted low rank approximation is known to be NP-hard to approximate up to a \((1 \pm 1 / \operatorname{poly}(n))\) factor [GG11]. We note that this also follows from the fact that matrix completion, arguably one of the most important special cases of weighted low rank approximation in which case all weights are 0 or 1 , which we discuss more below, is also known to be NP-hard [Pee96, HMRW14]. Typically, though, assumptions such as incoherence and randomly sampled entries allow one to circumvent this hardness [CR09, LLR16]. There is some debate as to whether these assumptions are valid, for instance in [SW15] an argument is made why randomly missing entries may not hold for real-world datasets.

Many natural questions are left open from previous work. In particular the main
question as we see it is the following:
- For which weight matrices \(W\) is the problem tractable? More generally, is it possible to identify a natural parameter of \(W\) and to obtain parameterized complexity bounds in terms of that parameter?

\subsection*{16.1.1 Our Results}

In this work we provide an answer to the above questions by parameterizing the complexity of the problem in terms of the rank of the weight matrix \(W\) and the number of distinct columns of the weight matrix. Note that we make no assumptions about the input matrix \(A\). Let OPT denote the quantity \(\min _{\text {rank- } k \text { matrices }}\|W \circ(\widehat{A}-A)\|_{F}^{2}\). Our main theorems are the following. For a function \(f\), define \(\tilde{f}=f \cdot \operatorname{poly}(\log (f))\).

Theorem 16.1.1 (Algorithm for Weighted Low Rank Approximation). Let \(r\) be the rank of \(W\). There is an algorithm running in time \(n^{O\left(k^{2} r / \epsilon\right)}\) which outputs a factorization (into an \(n \times k\) and a \(k \times n\) matrix) of a rank-k matrix \(\widehat{A}\) for which
\[
\|W \circ(A-\widehat{A})\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}
\]
with probability at least 9/10.

We also have the following theorem.
Theorem 16.1.2 (Algorithm for Weighted Distinct Columns). Let \(r\) be the number of distinct columns of \(W\). For every \(\epsilon>0\) and for an arbitrarily small constant \(\gamma>0\), there is an algorithm running in time \(O\left((\operatorname{nnz}(A)+\mathrm{nnz}(W)) \cdot n^{\gamma}\right)+n \cdot 2^{\widetilde{O}\left(k^{2} r^{2} / \epsilon\right)}\) which outputs a factorization (into an \(n \times k\) and a \(k \times n\) matrix) of a rank-k matrix \(\widehat{A}\) for which
\[
\|W \circ(A-\widehat{A})\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}
\]
with probability at least 9/10.

Note that Theorem 16.1.2 does not make any assumptions on the number of distinct rows, which is important for the applications described below.

Let us point out two things here:
- Before only the case of \(W\) having rank 1 was known to be provably solvable in polynomial time by a direct reduction to the \(\mathrm{SVD}^{1}\);
- The result for at most \(r\) distinct columns implies that, with respect to this parametrization, the weighted low rank approximation problem is fixed-parameter tractable.

We complement the above positive results by proving a lower bound, which assumes the Exponential Time Hypothesis for the average case hardness of random 4-SAT. We state the assumption as follows.

Assumption 16.1.3 ("Random Exponential Time Hypothesis"). Let \(c>\ln 2\) be a constant. Consider a random 4-SAT formula on \(n\) variables in which each clause has 4 literals, and in which each of the \(16 n^{4}\) clauses is picked independently with probability \(c / n^{3}\). Then any algorithm which always outputs 1 when the random formula is satisfiable, and outputs 0 with probability at least \(1 / 2\) when the random formula is unsatisfiable, must run in \(2^{c^{\prime n}}\) time on some input, where \(c^{\prime}>0\) is an absolute constant.

We are not aware of prior work using this form of the Exponential Time Hypothesis, though both the Exponential Time Hypothesis and Feige's original assumption [Fei02] that there is no polynomial time algorithm for the problem in Assumption 16.1.3 are commonly

\footnotetext{
\({ }^{1}\) Namely, in that case, we can rewrite the problem as \(\min _{\text {rank- } k \text { matrices } \widehat{A}}\|D(\widehat{A}-A) E\|_{F}^{2}\) for diagonal matrices \(D\) and \(E\), from which we can replace \(A\) with \(D A E\) and solve the unweighted low rank approximation, obtaining an \(\widehat{A}\) for which we can then output \(D^{-1} \widehat{A} E^{-1}\) if the diagonal entries of \(D\) and \(E\) are non-zero. If they are zero we can first remove rows and columns from \(A\)
}
used. We do not know of any better algorithm for the problem in Assumption 16.1.3 and have consulted several experts \({ }^{2}\) about the assumption who do not know a counterexample to it. Instead of calling it "Random-ETH", some experts" suggest other names, e.g. average case ETH or Feige-ETH.

Theorem 16.1.4 (Weighted Low Rank Approximation Hardness). Let \(r\) be an upper bound on the number of distinct columns of \(W\). Under Assumption 16.1.3, there is an absolute constant \(\epsilon_{0} \in(0,1)\) for which any algorithm for solving the weighted low rank approximation algorithm with \(\epsilon \leq \epsilon_{0}\) and for any \(k \geq 1\), with constant probability, requires \(2^{\Omega(r)}\) time. Further, this holds even if \(W\) also only has \(r\) distinct rows.

Note that for constant \(k\) and \(\epsilon\), and \(r \geq C \log n\) for a constant \(C>0\), our upper bound assuming at most \(r\) distinct columns is \(2^{\widetilde{O}\left(r^{2}\right)}\), which nearly matches our lower bound in Theorem 16.1.4 of \(2^{\Omega(r)}\).

There are naturally arising applications in which the rank of the weight matrix or the number of distinct columns is small. Consider a matrix in which the rows correspond to users and the columns correspond to ratings of a movie, such as in the Netflix matrix. Further, suppose for each movie, there are \(r\) columns, indicating different aspects of the movie to be rated, such as acting, plot, sound effects, visual effects, etc. For a given user, one can look at the distribution of scores across movies along one of these aspects. These distributions may have different variances for that user and one can renormalize the scores by the reciprocal of their variance. In this case, the weight matrix consists of \(r\) distinct columns, one for each

\footnotetext{
\({ }^{2}\) Personal communication with Russell Impagliazzo and Ryan Williams.
\({ }^{3}\) Personal communication with Lijie Chen and Aviad Rubinstein.
}
aspect, each copied \(n / r\) times, where \(n / r\) is the total number of movies. Each entry of a column is a variance for a certain user for that aspect. This naturally generalizes to other applications for which the columns can be clustered into \(r\) groups, such as in stochastic block models or more general latent space models [Cha12]. Also, in some of these applications, one would want a low rank nonnegative factorization, and we remark that we can achieve this below.

Suppose now that \(A\) has constant rank \(k\). A consequence of Theorem 16.1.1, which we elaborate on more below, is that even if an adversary deletes up to \(O(1)\) entries in each column of \(A\) in an arbitrary way, one can still recover a matrix \(\widehat{A}\) of rank at most \(k\) which agrees with \(A\) on all of the remaining entries in \(\operatorname{poly}(n)\) time. This is a form of adversarial matrix completion, which was studied in [HKZ11, SW15]. In these works the authors had to make an incoherence assumption on the entries of \(A\), which may not always hold, e.g., in the presence of outliers. Without this assumption the prior results would not be able to recover such an \(\widehat{A}\) even if a single adversarially chosen entry was deleted in each column.

\subsection*{16.1.2 Other Results}

Our techniques can be applied to weighted non-negative factorization, a problem that has been extensively studied both in the unweighted (see, e.g., [AGKM12, Moi13]) and weighted cases (see, e.g., [GVS03, GZD10, KG12, ZZGM15, YC09]). In this problem, one seeks non-negative matrices \(U \in \mathbb{R}^{n \times k}\) and \(V \in \mathbb{R}^{k \times n}\) for which
\[
\forall i, j \text { such that } W_{i, j}>0, A_{i, j}=(U \cdot V)_{i, j},
\]
where the unweighted case corresponds to \(W_{i, j}=1\) for all \(i\) and \(j\). This problem naturally arises in applications such as topic modeling in which negative entries do not make sense. In a beautiful line of work [AGKM12, Moi13], an upper bound of \(n^{O\left(k^{2}\right)}\) was established for the unweighted case. By combining their techniques with ours, we obtain the first results for the weighted case when the weight matrix \(W\) has rank at most \(r\). We defer the details to the full version of the paper.

Finally, there is a simple bi-criteria approximation algorithm for weighted low rank approximation when the weight matrix has entries that are all either 0 or 1 . In this case it is possible to obtain a rank- \(r k\) approximation to \(A\) in \(n n z(A)+\mathrm{nnz}(W)+n \cdot \operatorname{poly}(r k / \epsilon)\) time with cost at most \((1+\epsilon)\) times the cost of the best rank- \(k\) approximation. Here \(r\) denotes the rank of the weight matrix. This has a better dependence on \(r\) and could be useful when one does not need to output a matrix of rank exactly \(k\). This is given in the full version of the paper.

\subsection*{16.1.3 Our Techniques}

To prove Theorem 16.1.1, we first prove a structural theorem about regression. Then our main new theme is what we call "guessing a sketched matrix".

Structural Regression Theorem Given any number \(t\) of multiple response regression instances, i.e., instances of the form \(\min _{X^{1}, X^{2}, \ldots, X^{t}}\left\|A^{i} X^{i}-B^{i}\right\|_{F}^{2}\), where \(A^{1}, \ldots, A^{t}\) and \(B^{1}, \ldots, B^{t}\) are matrices and each \(A^{i}\) has rank at most \(k\), if one is interested in simultaneously minimizing the sum of their costs, i.e., the objective function is \(\min _{X^{1}, X^{2}, \ldots, X^{t}} \sum_{i=1}^{t}\left\|A^{i} X^{i}-B^{i}\right\|_{F}^{2}\), then one can choose a Gaussian matrix \(S\) with \(O(k / \epsilon)\) rows and if one solves the problem \(\min _{Y^{1}, Y^{2}, \cdots, Y^{t}} \sum_{i=1}^{t}\left\|S A^{i} Y^{i}-S B^{i}\right\|_{F}^{2}\), then the minimizers \(Y^{1}, Y^{2}, \cdots, Y^{t}\) for this latter problem satisfy \(\sum_{i=1}^{t}\left\|A^{i} Y^{i}-B^{i}\right\|_{F}^{2} \leq(1+\epsilon) \min _{X^{1}, X^{2}, \cdots, X^{t}} \sum_{i=1}^{t}\left\|A^{i} X^{i}-B^{i}\right\|_{F}^{2}\) with high constant probability. Interestingly, the number of rows of \(S\) does not depend on the number \(t\) of regression instances, and is optimal already for \(t=1\) and when \(B^{1}\) only has a single column [CW09]. This also generalizes affine embeddings in [CW13] in which \(t=1\) and \(B^{1}\) may have multiple columns; we stress that the design matrices \(A^{i}\) may be different. The proof uses a novel observation that a Gaussian \(S\) is a subspace embedding on average over \(i\). This part uses a tail bound from [CD05] for the condition number for a Gaussian matrix. Having this, we bootstrap the approximation to \((1 \pm \epsilon)\) using the approximate matrix product property of a Gaussian matrix.

Guessing a Sketched Matrix Given our structural result, the main new theme of this paper is to "guess a sketched matrix". Suppose one is given an optimization problem of the
form
\[
p_{1}\left(x_{1}, \ldots, x_{v}\right) \geq 0, p_{2}\left(x_{1}, \ldots, x_{v}\right) \geq 0, \ldots, p_{m}\left(x_{1}, \ldots, x_{v}\right) \geq 0
\]
where each \(p_{i}\) is a polynomial of degree at most \(d\), the \(x_{1}, \ldots, x_{v}\) are indeterminates over the reals, and we are are interested in an assignment to \(x_{1}, \ldots, x_{v}\) which simultaneously satisfies these \(c\) polynomial inequalities. Then this can be solved in time \((m d)^{O(v)}\) using generic solvers [Ren92d, Ren92b, Ren92a, BPR96], and such techniques have been used in the context of database theory [EFW10], non-negative matrix factorization [AGKM12, Moi13], learning mixtures of Gaussians [LS15], computing approximate PSD factorizations [BDL16], and solving small-scale mixed-norm low rank approximation instances [CW15a]. There are two kinds of algorithms for semi-algebraic sets: the ones from [Ren92a, Ren92b, BPR96] are able to determine if a given sem-algebraic set is empty or not. The one from [Ren92d] is able to return a \(\delta\)-approximate solution to a given semi-algebraic formulae by paying an extra factor of \(\log \left(\frac{1}{\delta}\right)\) in the running time. For weighted low-rank approximation, there are two ways to output the matrices. One is using the algorithm from [Ren92d] directly. Another option is to perform binary search using the algorithm from [Ren92a, Ren92b, BPR96] for the entries of \(\widehat{A}_{i j}\) one by one.

Our main idea here is to use polynomial optimization for large-scale non-convex optimization by combining it directly with sketching. For example, suppose one is given the multiple response regression problem \(\min _{V}\|U V-A\|_{F}^{2}\) in which the number of columns of \(U\) is small. The twist though, is that both \(U\) and \(V\) are unknown! Then \(U V\) is just a low rank approximation to \(A\), and if we knew \(U\) we could solve the sketched optimization problem \(\min _{X}\|S U V-S A\|_{F}^{2}\), for a random oblivious sketching matrix \(S\), and our solution \(V\) would
be a good solution to the original problem. Since we do not know \(U\), we instead choose a random \(S\) and create variables for \(S \cdot U\), which is small, and also compute \(S \cdot A\), which we know. We then solve a regression problem for \(V\) in terms of the variables that we created for \(S \cdot U\). Given \(V\), we can then plug it into the original regression problem and solve for \(U\) in terms of \(V\) which is in turn in terms of our variables for \(S \cdot U\). Finally we can verify the solution by requiring that \(\|U V-A\|_{F}^{2}\) is small, which is now a system in a small number of variables. This verification step is essential because our \(S\) only has a probabilistic guarantee that it works for a fixed \(U\) with good probability, but crucially, we know there exists a \(U\) for which it works, and so by doing the verification step we will find such a \(U\). We note there are several issues with this approach which we discuss below.

While this may seem like an unnecessarily complicated way of doing standard low rank approximation, this idea proves crucial for weighted low rank approximation. In this case, using say, the rank constraint on the weight matrix \(W\), and using our structural result on multiple instances of regression, we are able to choose a single sketching matrix \(S\) and create variables for only \(r\) regression problems, \(S D_{W_{1}} U, \ldots, S D_{W_{r}} U\), where \(D_{W_{1}}, \ldots, D_{W_{r}}\) are diagonal matrices with independent columns of \(W\) on each diagonal, and \(U\) is a fixed optimal solution. We can then try to express all regression solutions \(\min _{X^{i}}\left\|D_{W_{i}} U X^{i}-D_{W_{i}} A^{i}\right\|_{2}^{2}\), for \(i=1, \ldots, n\) and where \(X^{i}\) and \(A^{i}\) are the \(i\)-th columns of \(X\) and \(A\) respectively, in terms of these variables, and hope to carry out the procedure above.

Dealing with Linear Dependencies At this point another obstacle arises which is that when the columns of a given \(D_{W_{i}} U\) are not linearly independent, there is no way to write down the pseudoinverse of \(D_{W_{i}} U\) in terms of the variables we have created. We also cannot afford
to create more than \(r \cdot \operatorname{poly}(k / \epsilon)\) variables, since our optimization procedure is exponential in this quantity, and so we cannot create new variables for each \(i=1, \ldots, n\). To get around this, we observe that there is a solution \(U \cdot V\) for which for all \(i\), the first \(\min \left(\left|\operatorname{supp}\left(D_{W_{i}}\right)\right|, k\right)\) columns of \(U\) are linearly independent, where \(\operatorname{supp}\left(D_{W_{i}}\right)\) denotes the set of non-zero entries of \(W_{i}\), and further the solution cost of \(U \cdot V\) is an arbitrarily small amount larger than that of the optimal cost. This follows by a simple perturbation argument applied to the optimal solution. This immediately gives an algorithm with additive error when we parameterize \(W\) by its rank.

In order to turn it into a relative error algorithm, we need a lower bound on the cost assuming that the cost is non-zero. The main idea here is that if we correctly guess which subsets of columns are linearly independent for the different matrices \(D_{W_{i}}\), then we can set up a non-negative polynomial system and provided this system has non-zero cost, we can apply known lower bounds on the cost of polynomial optimazation problems as a function of the degrees, number of variables, number of constraints, and coefficient sizes [Bas14]. While this is not an algorithmic procedure, since we cannot afford to make guess for each \(D_{W_{i}}\) without spending exponential in \(n\) time, it suffices for lower bounding the cost. Given such a lower bound, the above perturbation argument can then be used to argue that we achieve relative error.

While this leads to our time bound in the case in which we parameterize the weight matrix by its rank, in the case in which we parameterize by the number of distinct columns of \(W\), this is too slow. The issue is that we have at least \(n\) constraints to enforce, namely, that the first \(\min \left(\operatorname{supp}\left(W_{i}\right), k\right)\) columns of \(S D_{W_{i}} U\) are linearly independent. This would lead to a running time of \(n^{\operatorname{poly}(r k / \epsilon)}\) as opposed to the \(\operatorname{poly}(n) \cdot 2^{\operatorname{poly}(k r / \epsilon)}\) that we desire. We
notice though that when we have \(r\) distinct columns, there are only \(r\) constraints to enforce, one for each distinct column. These are "not equal" constraints but can be transformed to a single equality constraint of degree \(\operatorname{poly}(r k / \epsilon)\) by introducing a single auxiliary variable. While this ultimately enables us to write down all the entries of \(V\) using only a poly \((r k / \epsilon)\) number of variables, we are then faced with the task of writing down \(U\) in terms of such a \(V\). Here a priori we could have many distinct rows in \(W\), and may have \(n\) constraints to enforce. We observe though that since the entries of \(W\) are integers in \(\{1,2, \ldots, \operatorname{poly}(n)\}\), if we round them to the nearest power of \(1+\epsilon\), the solution cost changes only by a \((1+\epsilon)\) factor. Moreover, given that we have \(r\) distinct columns, for any row it is entirely specified on these \(r\) columns and after rounding, there are only \(O((\log n) / \epsilon)^{r}\) choices for entries on these \(r\) columns, which upper bounds the number of distinct rows. This ultimately allows us to write down \(U\) in terms of \(V\) with only \(O((\log n) / \epsilon)^{r}\) not equal constraints. This ultimately yields our improved running time when parameterizing \(W\) by its number of distinct columns.

Dealing with Rational Functions There is a subtle problem with the above arguments. When one solves for the \(i\)-th column \(V^{i}\) of \(V\) in terms of \(S D_{W_{i}} U\), the entries of \(V^{i}\) are rational functions rather than polynomials, and we cannot afford to clear the denominators of \(V^{i}\) for every \(i\) without blowing up the degree of the polynomials to \(\Omega(n)\), which would give a running time of \(n^{\operatorname{poly}(k r / \epsilon)}\). While this is not a problem when we parameterize the problem by the rank of \(W\), since we anyway spend this amount of time, this is a problem when we parameterize by the number of distinct columns of \(W\), in which we seek polynomial time even for super-constant \(r\) (and constant \(k / \epsilon\) ). Instead, we write \(V\) as \(V^{\prime} \cdot D\), where \(V^{\prime}\) has entries which are polynomials, and \(D\) is a diagonal matrix whose entries are \(\frac{1}{\operatorname{det}\left(U^{\top} D_{W_{i}} S^{\top} S D_{W_{i}} U\right)}\),
given by Cramer's rule. The entries of \(D\) are rational functions, and we would like to make them polynomials. Since \(W\) has at most \(r\) distinct columns, \(D\) has at most \(r\) distinct entries and we can create \(r\) new variables for the entries of \(D\). However, when we try to solve for \(U\) in terms of \(V\) we face the same problem again: we can write \(U\) as \(E \cdot U^{\prime}\), where \(U^{\prime}\) only has polynomial entries, but since we do not assume a small number of distinct rows of \(W\), it follows that \(E\) could have \(n\) distinct entries \(\frac{1}{\operatorname{det}\left(U^{\top} D_{W^{i}} S^{\top} S D_{W^{i}} U\right)}\) for \(i=1, \ldots, n\), where \(D_{W^{i}}\) is the diagonal matrix with the \(i\)-th row of \(W\) on the diagonal. As mentioned above, we fix this problem by rounding the entries of \(W\) to powers of \(1+\epsilon\), and then observing that the number of distinct rows of \(W\) can only be \(O((\log n) / \epsilon)^{r}\) after rounding.

Other Methods Our techniques for weighted non-negative matrix factorization and tensor completion largely follow these ideas as well, where we combine our "guessing a sketched matrix" approach with techniques of Arora, Ge, Kannan and Moitra [AGKM12] and Moitra [Moi13]. Our bi-criteria solution directly follows from known sketching results.

\subsection*{16.2 Preliminaries}

Notation Let \(\mathbb{R}\) denote the real numbers, and \(\mathbb{R}_{\geq 0}\) denote the nonnegative real numbers. Let \(\|A\|\) (and sometimes \(\|A\|_{2}\) ) denote the spectral norm of matrix \(A\). Let \(\|A\|_{F}^{2}=\sum_{i, j} A_{i, j}^{2}\) denote the Frobenius norm of \(A\). Let \(W \circ A\) denote the entry-wise product of matrices \(W\) and \(A\). Let \(\|A\|_{W}^{2}=\sum_{i, j} W_{i, j}^{2} A_{i, j}^{2}\) denote the weighted Frobenius norm of \(A\). Let nnz \((A)\) denote the number of nonzero entries of \(A\). Let \(\operatorname{det}(A)\) denote the determinant of a square matrix \(A\). Let \(A^{\top}\) denote the transpose of \(A\). Let \(A^{\dagger}\) denote the Moore-Penrose pseudoinverse of \(A\). Let \(A^{-1}\) denote the inverse of a full rank square matrix \(A\).

For the weight matrix \(W\), we always use \(W_{j}\) to denote the \(j\)-th column vector of \(W\), and \(W^{i}\) to denote the \(i\)-th row of \(W\). Let \(D_{W_{j}}\) denote the diagonal matrix with entries from the column vector \(W_{j}\) and \(D_{W^{i}}\) denote the diagonal matrix with entries from the row vector \(W^{i}\).

The following real algebraic geometry definitions are needed when proving a lower bound for the minimum nonzero cost of our problem. For a full discussion, we refer the reader to Bochnak et al. [BCR87]. Here we use the brief summary by Basu et al.[BPR05].

Definition 16.2.1 ([BPR05]). Let \(R\) be a real closed field.
Given \(x=\left(x_{1}, \cdots, x_{v}\right) \in R^{v}, r \in R, r>0\), we denote
\[
\begin{array}{ll}
B_{v}(x, r)=\left\{y \in R^{v} \mid\|y-x\|^{2}<r^{2}\right\} & \text { (open ball), } \\
\bar{B}_{v}(x, r)=\left\{y \in R^{v} \mid\|y-x\|^{2} \leq r^{2}\right\} & \text { (closed ball). }
\end{array}
\]

A set \(S \subset R^{v}\) is open if it is the union of open balls, i.e., if every point of \(U\) is contained in an open ball contained in \(U\).

A set \(S \subset R^{v}\) is closed if its complement is open. Clearly, the arbitrary union of open sets is open and the arbitrary intersection of closed sets is closed.

Semi-algebraic sets are defined by a finite number of polynomial inequalities and equalities.

A semi-algebraic set has a finite number of connected components, each of which is semi-algebraic. Here, we use the topological definition of a connected component, which is a maximal connected subset (ordered by inclusion), where connected means it cannot be divided into two disjoint nonempty closed sets.

A closed and bounded semi-algebraic set is compact.

A semi-algebraic set \(S \subset R^{v}\) is semi-algebraically connected if \(S\) is not the disjoint union of two non-empty semi-algebraic sets that are both closed in \(S\). Or, equivalently, \(S\) does not contain a non-empty semi-algebraic strict subset which is both open and closed in \(S\).

A semi-algebraically connected component of a semi-algebraic set \(S\) is a maximal semi-algebraically connected subset of \(S\).

Renegar [Ren92a, Ren92b] and Basu et al. [BPR96] independently provided an algorithm for the decision problem for the existential theory of the reals is to decide the truth or falsity of a sentence \(\left(x_{1}, \cdots, x_{v}\right) F\left(f_{1}, \cdots, f_{m}\right)\) where \(F\) is a quantifier-free Boolean formula with atoms of the form \(\operatorname{sign}\left(f_{i}\right)=\sigma\) with \(\sigma \in\{0,1,-1\}\). Note that this problem is equivalent to deciding if a given semi-algebraic set is empty or not. Here we formally state that theorem. For a full discussion of algorithms in real algebraic geometry, we refer reader to [BPR05] and [Bas14].

Theorem 16.2.1 (Decision Problem [Ren92a, Ren92b, BPR96]). Given a real polynomial system \(P\left(x_{1}, x_{2}, \cdots, x_{v}\right)\) having \(v\) variables and \(m\) polynomial constraints
\[
f_{i}\left(x_{1}, x_{2}, \cdots, x_{v}\right) \Delta_{i} 0, \forall i \in[m]
\]
where \(\Delta_{i}\) is any of the "standard relations": \(\{>, \geq,=, \neq, \leq,<\}\), let d denote the maximum degree of all the polynomial constraints and let \(H\) denote the maximum bitsize of the coefficients of all the polynomial constraints. Then in
\[
(m d)^{O(v)} \operatorname{poly}(H)
\]
time one can determine if there exists a solution to the polynomial system \(P\).

The key result we used for proving lower bound is the following bound on the minimum value attained by an integer polynomial restricted to a compact connected component of a basic closed semi-algebraic subset of \(\mathbb{R}^{v}\) defined by polynomials with integer coefficients in terms of the degrees and the bitsizes of the coefficients of the polynomials involved.

Theorem 16.2.2 (Jeronimo, Perrucci and Tsigaridas [JPT13]). Let \(T=\left\{x \in \mathbb{R}^{v} \mid f_{1}(x) \geq\right.\) \(\left.0, \cdots, f_{\ell}(x) \geq 0, f_{\ell+1}(x)=0, \cdots, f_{m}(x)=0\right\}\) be defined by polynomials \(f_{1}, \cdots, f_{m} \in\) \(\mathbb{Z}\left[x_{1}, \cdots, x_{v}\right]\) with \(n \geq 2\), degrees bounded by an even integer \(d\) and coefficients of absolute value at most \(H\), and let \(C\) be a compact connected component of \(T\). Let \(g \in \mathbb{Z}\left[x_{1}, \cdots, x_{v}\right]\) be a polynomial of degree at most \(d\) and coefficients of absolute value bounded by \(H\). Then, the minimum value that \(g\) takes over \(C\) is satisfies that if it is not zero, then its absolute value is greater than or equal to
\[
\left(2^{4-v / 2} \widetilde{H} d^{v}\right)^{-v 2^{v} d^{v}},
\]
where \(\widetilde{H}=\max \{H, 2 v+2 m\}\).

While the above theorem involves notions from topology, we shall apply it in an elementary way. Namely, in our setting \(T\) will be bounded and so every connected component, which is by definition closed, will also be bounded and therefore compact. As the connected components partition \(T\) the theorem will just be applied to give a global minimum value of \(g\) on \(T\) provided that it is non-zero.

\subsection*{16.3 Multiple Regression Sketch}

Theorem 16.3.1. Let \(A^{1}, \ldots, A^{m} \in \mathbb{R}^{n \times k}\) be \(m\) matrices of size \(n \times k\). Let \(b^{1}, \ldots, b^{m} \in \mathbb{R}^{n \times 1}\) be \(m\) column vectors of dimension \(n\).

For \(1 \leq i \leq m\) denote:
\[
x^{i}=\underset{x \in \mathbb{R}^{k \times 1}}{\operatorname{argmin}}\left\|A^{i} x-b^{i}\right\|_{2}^{2}
\]
the solution of the \(i\)-th regression problem.
Let \(S \in \mathbb{R}^{t \times n}\) be a random matrix with i.i.d. Gaussian entries with zero mean and standard deviation \(1 / \sqrt{t}\). For \(1 \leq i \leq m\) denote:
\[
y^{i}=\underset{y \in \mathbb{R}^{k \times 1}}{\operatorname{argmin}}\left\|S A^{i} y-S b^{i}\right\|_{2}^{2}
\]
the solution of the \(i\)-th regression problem in the sketch space.

We claim that for every \(0<\epsilon<1 / 2\) one can set \(t=O(k / \epsilon)\) such that:
\[
\sum_{i=1}^{m}\left\|A^{i} y^{i}-b^{i}\right\|_{2}^{2} \leq(1+\epsilon) \cdot \sum_{i=1}^{m}\left\|A^{i} x^{i}-b^{i}\right\|_{2}^{2} .
\]

The rest of this section is devoted to the proof of this theorem.
For \(1 \leq i \leq m\) we let \(D^{i} \geq 1\) denote the smallest number such that for every \(x \in \mathbb{R}^{k \times 1}\) and \(\lambda \in \mathbb{R}\) one has:
\[
\left\|S\left(A^{i} x+\lambda b^{i}\right)\right\|_{2}^{2} \in\left[\frac{1}{D^{i}} ; D^{i}\right] \cdot\left\|A^{i} x+\lambda b^{i}\right\|_{2}^{2}
\]

Claim 16.3.2. For every \(i\)
\[
\left\|A^{i} y^{i}-b^{i}\right\|_{2}^{2} \leq\left(D^{i}\right)^{2} \cdot\left\|A^{i} x^{i}-b^{i}\right\|_{2}^{2}
\]

Proof. This follows from the definition of \(D^{i}\).

Claim 16.3.3. One can set \(t=O(k / \epsilon)\) such that for every \(i\)
\[
\operatorname{Pr}_{S}\left[D^{i} \geq 1.01\right] \leq 2^{-\Omega(1 / \epsilon)} .
\]

Proof. This follows from Theorem 2.1 from [Woo14b].

Claim 16.3.4. One can set \(t=O(k / \epsilon)\) such that for every \(i\)
\[
\underset{S}{\mathbb{E}}\left[\left\|A^{i} y^{i}-b^{i}\right\|_{2}^{2}-\left\|A^{i} x^{i}-b^{i}\right\|_{2}^{2} \mid D^{i} \leq 1.01\right] \leq \epsilon \cdot\left\|A^{i} x^{i}-b^{i}\right\|_{2}^{2}
\]

Proof. This follows from the proofs of Theorem 2.8 and Theorem 2.16 from [Woo14b] (adapted to Gaussian matrices).

Claim 16.3.5. One can set \(t=O(k / \epsilon)\) such that for every \(i\)
\[
\underset{S}{\mathbb{E}}\left[\left(D^{i}\right)^{2} \mid D^{i} \geq 1.01\right]=O(1)
\]

Proof. One can see that \(D^{i}\) is polynomially related to the condition number of a random \(O(k / \epsilon) \times(k+1)\) matrix with i.i.d. Gaussian entries; indeed it corresponds to the maximum distortion of \(S\) applied to the vectors in the column span of an \(n \times(k+1)\) orthonormal matrix \(U\) whose columns span the space spanned by the columns of \(A^{i}\) together with \(b^{i}\). By rotational invariance, \(S \cdot U\) also has i.i.d. Gaussian entries. To understand the condition number one can invoke the main result from [CD05] which gives for all sufficiently large \(t\) : \(\operatorname{Pr}_{S}\left[D^{i} \geq t\right]=\frac{1}{t \Theta(k / \epsilon)}\). Thus,
\[
\underset{S}{\mathbb{E}}\left[\left(D^{i}\right)^{2} \mid D_{i} \geq 1.01\right] \leq O(1)+\int_{1.01}^{\infty} \frac{t^{2}}{t^{\Theta(k / \epsilon)}} d t=O(1)
\]

Having these Claims, let us complete the proof. We have for every \(1 \leq i \leq m\) :
\[
\begin{aligned}
& \underset{S}{\mathbb{E}}\left[\left\|A^{i} y^{i}-b^{i}\right\|_{2}^{2}-\left\|A^{i} x^{i}-b^{i}\right\|_{2}^{2}\right] \\
= & \operatorname{Pr}_{S}\left[D^{i} \geq 1.01\right] \cdot \underset{S}{\mathbb{E}}\left[\left\|A^{i} y^{i}-b^{i}\right\|_{2}^{2}-\left\|A^{i} x^{i}-b^{i}\right\|_{2}^{2} \mid D^{i} \geq 1.01\right] \\
+ & \operatorname{Pr}_{S}\left[D^{i} \leq 1.01\right] \cdot \underset{S}{\mathbb{E}}\left[\left\|A^{i} y^{i}-b^{i}\right\|_{2}^{2}-\left\|A^{i} x^{i}-b^{i}\right\|_{2}^{2} \mid D^{i} \leq 1.01\right] \\
\leq & 2^{-\Omega(1 / \epsilon)} \cdot \underset{S}{\mathbb{E}}\left[\left(D^{i}\right)^{2}-1 \mid D^{i} \geq 1.01\right] \cdot\left\|A^{i} x^{i}-b^{i}\right\|_{2}^{2} \\
+ & \operatorname{Pr}_{S}\left[D^{i} \leq 1.01\right] \cdot \underset{S}{\mathbb{E}}\left[\left\|A^{i} y^{i}-b^{i}\right\|_{2}^{2}-\left\|A^{i} x^{i}-b^{i}\right\|_{2}^{2} \mid D^{i} \leq 1.01\right] \\
\leq & 2^{-\Omega(1 / \epsilon)} \cdot \underset{S}{\mathbb{E}}\left[\left(D^{i}\right)^{2}-1 \mid D^{i} \geq 1.01\right] \cdot\left\|A^{i} x^{i}-b^{i}\right\|_{2}^{2} \\
+ & \epsilon \cdot\left\|A^{i} x^{i}-b^{i}\right\|_{2}^{2} \\
\leq & O(\epsilon) \cdot\left\|A^{i} x^{i}-b^{i}\right\|_{2}^{2},
\end{aligned}
\]
where the second step is by Claim 16.3.2 and Claim 16.3.3, the third step is by Claim 16.3.4, and the fourth step is by Claim 16.3.5.

Summing over \(i\) and applying the Markov's inequality, we are done.
While the above result is for Gaussian sketching matrices, one can also combine a Gaussian random matrix with a Count-Sketch matrix [CW13]. This way we are still getting \(O(k / \epsilon)\) rows, but now one can perform a matrix-vector multiplication in time proportional to the sparsity of the vector plus poly \((k \widetilde{r} / \epsilon)\), where \(\widetilde{r}\) is the dimension of the union of column spaces of \(A^{i}\) (which is at most \(k m\) in the worst case, but is much smaller for our applications).

\subsection*{16.4 Additive Approximation}

In this Section, to demonstrate the new technique, we prove the following theorem.

Theorem 16.4.1. Given \(A, W \in \mathbb{R}^{n}, 1 \leq k \leq n\) and \(0<\epsilon, \tau<0.1\) such that:
- \(\operatorname{rank}(W)=r\);
- all the non-zero entries of \(A\) and \(W\) are multiples of \(\delta>0\);
- all the entries of \(A\) and \(W\) are at most \(\Delta>0\) in absolute value,
one can output a number \(\Lambda\) in time \(n^{O\left(k^{2} r / \epsilon\right)} \cdot \log ^{O(1)}\left(\frac{\Delta}{\delta \tau}\right)\) such that \(\mathrm{OPT} \leq \Lambda \leq(1+\epsilon) \mathrm{OPT}+\tau\), where
\[
\mathrm{OPT}=\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times n}}\|(U V-A) \circ W\|_{F}^{2}
\]

We first assume that \(W\) has no zero entries; later, we will remove this assumption by being slightly more careful.

\section*{Lemma 16.4.2.}
\[
\mathrm{OPT} \leq \operatorname{poly}(n, \Delta)
\]

Proof. Set \(U\) and \(V\) to be the zero matrices.

Let us expand the objective function in two ways. On the one hand:
\[
\begin{equation*}
\|(U V-A) \circ W\|_{F}^{2}=\sum_{i=1}^{n}\left\|U^{i} V D_{W^{i}}-A^{i} D_{W^{i}}\right\|_{2}^{2} \tag{16.1}
\end{equation*}
\]

On the other hand:
\[
\begin{equation*}
\|(U V-A) \circ W\|_{F}^{2}=\sum_{j=1}^{n}\left\|D_{W_{j}} U V_{j}-D_{W_{j}} A_{j}\right\|_{2}^{2} \tag{16.2}
\end{equation*}
\]

We can sketch (16.1) and (16.2) using Gaussian matrices \(S^{\prime} \in \mathbb{R}^{n \times t}\) and \(S^{\prime \prime} \in \mathbb{R}^{t \times n}\) as follows:
\[
\sum_{i=1}^{n}\left\|U^{i} V D_{W^{i}} S^{\prime}-A^{i} D_{W^{i}} S^{\prime}\right\|_{2}^{2}
\]
and
\[
\sum_{j=1}^{n}\left\|S^{\prime \prime} D_{W_{j}} U V_{j}-S^{\prime \prime} D_{W_{j}} A_{j}\right\|_{2}^{2}
\]

Denote, for \(1 \leq i \leq n, P^{i}:=V D_{W^{i}} S^{\prime}\) and for \(1 \leq j \leq n\) denote \(Q_{j}:=S^{\prime \prime} D_{W_{j}} U\).
The crucial observation is that we can encode all \(P^{i}\) 's and \(Q_{j}\) 's using linear functions of \(2 k r t\) variables, since \(W\) has rank \(r\), and we can represent its rows/columns as linear combinations of \(r\) fixed rows/columns.

For fixed \(P^{i}\),s we define:
\[
\widehat{U}:=\underset{U \in \mathbb{R}^{n \times k}}{\operatorname{argmin}} \sum_{i=1}^{n}\left\|U^{i} P^{i}-A^{i} D_{W^{i}} S^{\prime}\right\|_{2}^{2} .
\]

Similarly, for fixed \(Q_{j}\) 's we define:
\[
\widehat{V}:=\underset{V \in \mathbb{R}^{k \times n}}{\operatorname{argmin}} \sum_{j=1}^{n}\left\|Q_{j} V_{j}-S^{\prime \prime} D_{W_{j}} A_{j}\right\|_{2}^{2}
\]

Let us use \(\|(\widehat{U} \widehat{V}-A) \circ W\|_{F}^{2}\) as a proxy for the objective function. We need to argue that we can, in fact, minimize the new objective function efficiently, and that it gives a good approximation to the original objective function, if \(t=\Theta(k / \epsilon)\), with high probability.

Optimization We assume that all \(P^{i}\) 's and \(Q_{j}\) 's have the maximum rank \(k\). Later we will show that, since \(W\) has no zero entries, this does not affect the quality of the solution found under this non-degeneracy constraint.

Assuming non-degeneracy of \(P^{i}\) 's and \(Q_{j}\) 's, we can express \(\widehat{U}\) and \(\widehat{V}\) as follows:
\[
\widehat{U}^{i}=A^{i} D_{W^{i}} S^{\prime}\left(P^{i}\right)^{\top}\left(P^{i}\left(P^{i}\right)^{\top}\right)^{-1}
\]
and
\[
\begin{equation*}
\widehat{V}_{j}=\left(Q_{j}^{\top} Q_{j}\right)^{-1} Q_{j}^{\top} S^{\prime \prime} D_{W_{j}} A_{j} \tag{16.3}
\end{equation*}
\]

Since the entries of \(P^{i}\) 's and \(Q_{j}\) 's are linear functions of \(2 k r t\) variables, we can represent the entries of \(\widehat{U}\) and \(\widehat{V}\) as rational functions over \(2 k r t\) variables and of degree \(O(k)\) (we use Cramer's formula for that).

Finally, we can represent \(\|(\widehat{U} \widehat{V}-A) \circ W\|_{F}^{2}\) as a rational function over \(2 k r t\) variables of degree \(O(k n)\).

We can minimize the objective function using the algorithm for checking the feasibility of a system of polynomial inequalities from Theorem 16.2.1, together with a binary search over the value of the objective function. Each iteration of the binary search takes time
\[
\left.(\# \text { degree of the polynomials })^{O(\# \text { variables })} \cdot \text { poly(input size }\right) .
\]

Since the degree is \(O(n k)\), the number of variables is \(O(r k t)=O\left(k^{2} r / \epsilon\right)\), and the input size is \(\operatorname{poly}(n, \log (\Delta / \delta))\), the running time of a single iteration of the binary search is
\[
n^{O\left(k^{2} r / \epsilon\right)} \cdot \log ^{O(1)}(\Delta / \delta)
\]

Finally, to perform the binary search, we need \(O(\log (n \Delta / \tau) / \epsilon)\) iterations to check for the existence of a solution with cost at most \((1+\epsilon) \mathrm{OPT}+\tau\) (assuming that sketching and non-degeneracy constraints on \(P^{i}\) 's and \(Q_{j}\) 's increase the cost to at most \((1+\epsilon) \mathrm{OPT}+\tau\), which we will show later), since, by Lemma 16.4.2, \(\mathrm{OPT} \leq \operatorname{poly}(n, \Delta)\). The overall running time is thus: \(n^{O\left(k^{2} r / \epsilon\right)} \cdot \log ^{O(1)}\left(\frac{\Delta}{\delta \tau}\right)\).

Near-optimality Here we show that one can set \(t=O(k / \epsilon)\) so that, with high probability,
\[
\begin{equation*}
\min _{P^{i}, Q_{j}}\|(\widehat{U} \widehat{V}-A) \circ W\|_{F}^{2} \leq(1+\epsilon) \cdot \min _{U, V}\|(U V-A) \circ W\|_{F}^{2}+\tau^{\prime} \tag{16.4}
\end{equation*}
\]
for every \(\tau^{\prime}>0\). This, together with the above discussion about the optimization procedure, concludes the analysis of the algorithm.

As such, (16.4) follows from Theorem 16.3.1. Indeed, if the optimal solution is \(U^{*} V^{*}\), then set \(Q_{j}:=S^{\prime \prime} D_{W_{j}} U^{*} . Q_{j}\) may be degenerate, but, since \(W\) has no zero entries, we can perturb \(U^{*}\) by an arbitrarily small amount to make \(Q_{j}\) non-degenerate (with probability one over \(\left.S^{\prime \prime}\right)\). Then, with high probability over \(S^{\prime \prime}\), we have
\[
\begin{equation*}
\left\|\left(U^{*} \widehat{V}-A\right) \circ W\right\|_{F}^{2} \leq(1+\epsilon) \cdot\left\|\left(U^{*} V^{*}-A\right) \circ W\right\|_{F}^{2}+\tau^{\prime} \tag{16.5}
\end{equation*}
\]
for an arbitrarily small \(\tau^{\prime}>0\). Similarly, we can set \(P^{i}=\widehat{V} D_{W^{i}} S^{\prime}\) (again, it can be degenerate, but the same argument as above for \(Q_{j}\) applies), which gives, with high probability over \(S^{\prime}\),
\[
\begin{equation*}
\|(\widehat{U} \widehat{V}-A) \circ W\|_{F}^{2} \leq(1+\epsilon) \cdot\left\|\left(U^{*} \widehat{V}-A\right) \circ W\right\|_{F}^{2}+\tau^{\prime} \tag{16.6}
\end{equation*}
\]
for an arbitrarily small \(\tau^{\prime}>0\).
Combining (16.5) and (16.6), we are done.

\subsection*{16.4.1 Handling Weight Matrices With Zero Entries}

Here we prove the version of Theorem 16.4.1 for the case when \(W\) is allowed to have zero entries. Let us first see what breaks in the previous argument.

What does not work anymore is that (after a small perturbation) \(Q_{j}=S^{\prime \prime} D_{W_{j}} U\) and \(P^{i}=V D_{W^{i}} S^{\prime}\) can be assumed to have the maximum possible rank \(k\). Nevertheless, we can assume that every \(Q_{j}\) has rank equal to
\[
t_{j}=\min \left(k, \text { the number of non-zero entries of } W_{j}\right)
\]

Moreover, we can assume that the first \(t_{j}\) columns of \(Q_{j}\) are linearly-independent. A similar argument applies to \(P^{i}\),s as well.

The above argument allows us to express \(\widehat{U}\) and \(\widehat{V}\) as before, but instead of \(Q_{j}\) we use the first \(t_{j}\) columns of \(Q_{j}\) (and, similarly for the \(P^{i}\), s ).

\subsection*{16.5 Multiplicative Approximation}

In order to get a genuine multiplicative \((1+\epsilon)\)-approximation, we need to lower bound OPT-provided that it is not equal to zero-which would allow us to set \(\tau \leq \epsilon \cdot\) OPT in the algorithm from the previous section.

We do this for the following optimization problem:
\[
\min _{\substack { U, V: \\
\begin{subarray}{c}{F}(\Delta / \delta)^{\operatorname{poly}(n){ U , V : \\
\begin{subarray} { c } { F } ( \Delta / \delta ) ^ { \operatorname { p o l y } ( n ) } }\end{subarray}}\|(U V-A) \circ W\|_{F}^{2}
\]

Note that we assume \(\|U V\|_{F}\) has an upper bound, as otherwise we cannot write down \(U\) and \(V\) using poly \((n)\) bits.

Using the approach outlined above, one can write down a rational function \(p\left(x_{1}, \ldots, x_{l}\right) / q\left(x_{1}, \ldots, x_{l}\right)\) such that:
- \(l=O\left(k^{2} r / \epsilon\right)\);
- for every \(x\) such that \(q(x) \neq 0\), one has \(p(x) / q(x) \geq\) OPT;
- for every \(\tau^{\prime}>0\), there exists \(x^{*}\) such that \(p\left(x^{*}\right) / q\left(x^{*}\right) \leq(1+\epsilon)\) OPT \(+\tau^{\prime}\);
- both \(p\) and \(q\) are homogeneous, and their degrees are \(O(k n)\);
- the coefficients of \(p\) and \(q\) are integers with absolute values at most \((\Delta / \delta)^{\operatorname{poly}(n)}\);
- \(q(x)=\prod_{i=1}^{2 n} g_{i}^{2}(x)\), where every \(g_{i}(x)\) is the determinant polynomial.

Only the fifth item needs an explanation. If sketch matrices \(S^{\prime}\) and \(S^{\prime \prime}\) were integer, then the last item would hold automatically (by scaling up all the coefficients). But, in
reality, \(S^{\prime}\) and \(S^{\prime \prime}\) are Gaussian matrices. Fortunately, one can show that it is possible to discretize them up to \(\pm 1 / \operatorname{poly}(n)\) so that the multiple regression theorem (Theorem 16.3.1) still goes through. Indeed, this just follows from the fact that discretization to \(\pm 1 / \operatorname{poly}(n)\) preserves condition number and subspace embeddings (since one argues about preserving lengths of unit vectors), and approximate matrix product properties used in that theorem.

Lower Bound We use the same way explained in section 16.4 to create variables, write down the system in a small number of variables, and also create some " \(\neq 0\) " constraints. It will generate \(2 n\) determinant polynomials, which are defined in the following way,
\[
\begin{aligned}
g_{i}(x) & =\operatorname{det}\left(\left(S D_{W_{i}} U\right)_{P_{i}}^{\top}\left(S D_{W_{i}} U\right)_{P_{i}}\right), \forall i \in[n] \\
g_{i+n}(x) & \left.=\operatorname{det}\left(\left(V D_{W^{i}} S\right)^{Q_{i}}\left(\left(V D_{W^{i}} S\right)^{Q_{i}}\right)\right)^{\top}\right), \forall i \in[n]
\end{aligned}
\]
where \(P_{i}\) and \(Q_{i}\) are maximal linearly indepdent subsets. Then we can write down the following optimization problem,
\[
\begin{array}{ll}
\min _{x \in \mathbb{R}^{l}} & p(x) / q(x) \\
\text { s.t. } & g_{i}^{2}(x) \neq 0, \forall i \in[2 n], \\
& q(x)=\prod_{i=1}^{2 n} g_{i}^{2}(x)
\end{array}
\]

To lower bound \(p(x) / q(x)\), let us introduce a new variable \(y\). Lower bounding \(p(x) / q(x)\) is equivalent to lower bounding \(p(x) y\) subject to \(q(x) y-1=0\). By assumption \(\|U\|_{F}^{2},\|V\|_{F}^{2} \leq(\Delta / \delta)^{\mathrm{poly}(n) 4}\), we know the upper bound of all the variables we created

\footnotetext{
\({ }^{4}\) A priori, we know only that \(\|U V\|_{F}^{2}\) is bounded. But we can get that each of the matrices is bounded by taking an optimal solution and orthonormalizing one of the matrices.
}
except for \(y\). In order to give an upper bound for \(y\), it suffices to show a lower bound for \(q(x)\). By definition, \(q(x)\) is the square of the product of all the determinant polynomials. Thus, for every determinant polynomial \(g_{i}(x)\), we need to show that if \(g_{i}^{2}(x)\) is nonzero, then it is at least something.

Define \(B\) to be an \(\ell_{2}\)-ball with bounded radius, e.g.,
\[
B=\left\{x \in \mathbb{R}^{l} \mid \sum_{i=1}^{l} x_{i}^{2} \leq(\Delta / \delta)^{\operatorname{poly}(n)}\right\}
\]

By [BPR05], we know that \(B\) is a closed and bounded semi-algebraic set. Thus \(B\) is also compact. Let \(x^{*}\) denote the optimal solution of the original problem \(\min _{x \in \mathbb{R}^{2}} g_{i}^{2}(x)\) when all variables are bounded. Because the radius of the ball \(B\) is large enough, \(x^{*} \in B\). Define \(T_{1}=\left\{x \in \mathbb{R}^{l} \mid g_{i}(x) \geq 0\right\}\) and let \(T=T_{1} \cap B\). By definition (see [BPR05]), \(T_{1}\) is a basic closed semi-algebraic set. Thus, the intersection of \(T_{1}\) and \(B\) is a semi-algebraic set with a finite number of connected components. Because \(B\) is compact and \(T_{1}\) is closed, each of these connected components is compact. There must exist one compact connected component \(C\) which contains the optimal soltuion \(x^{*}\). Applying Theorem 16.2.2 on system \(\left\{T, C, g_{i}^{2}(x)\right\}\), we conclude that if \(g_{i}^{2}(x)\) is not zero, then it is at least
\[
\begin{aligned}
\left((\Delta / \delta)^{\operatorname{poly}(n)}\right)^{-k^{O(l)}} & =(\Delta / \delta)^{-\operatorname{poly}(n) 2^{\tilde{O}(l)}} \\
& =(\Delta / \delta)^{-\operatorname{poly}(n)}
\end{aligned}
\]
which immediately gives us an upper bound for \(y\),
\[
y \leq\left((\Delta / \delta)^{\operatorname{poly}(n)}\right)^{n}=(\Delta / \delta)^{\operatorname{poly}(n)}
\]

Now, we are able to show a lower bound for \(p(x) y\). Define
\[
T_{1}=\left\{x \in \mathbb{R}^{l}, y \in \mathbb{R} \mid \prod_{i=1}^{m} g_{i}^{2}(x) y-1=0\right\}
\]

Define \(B\) to be a bounded ball over \(l+1\) variables,
\[
B=\left\{(x, y) \in \mathbb{R}^{l+1} \mid \sum_{i=1}^{l} x_{i}^{2}+y^{2} \leq(\Delta / \delta)^{\operatorname{poly}(n)}\right\} .
\]

By [BPR05], \(B\) is a closed and bounded semi-algebraic set. Thus \(B\) is also compact. Define \(T=T_{1} \cap B\). Let \(\left(x^{*}, y^{*}\right)\) denote the optimal solution of \(\min _{(x, y) \in T_{1}} p(x) y\), then \(\left(x^{*}, y^{*}\right)\) is also the optimal solution of \(\min _{(x, y) \in T} p(x) y\), because all variables are bounded and the radius of the ball is large enough.

By definition (see [BPR05]), \(T_{1}\) is a basic closed semi-algebraic set. Thus, the intersection of \(T_{1}\) and \(B\) is a semi-algebraic set with a finite number of connected components. Because \(B\) is compact and \(T_{1}\) is closed, each of the connected components is compact.

There must exist a compact connected component that contains the optimal solution \(\left(x^{*}, y^{*}\right)\). Let \(C\) denote that component. Applying Theorem 16.2.2 on system \(\{T, C, p(x) y\}\), where the number of constraints is bounded by \(O(1)\), the maximum coefficient of absolute value is bounded by \((\Delta / \delta)^{\operatorname{poly}(n)}\), the maximum degree is bounded by \(O(n k)\), the number of variables is bounded by \(l=O\left(r k^{2} / \epsilon\right)\), we conclude that if the minimum cost is not zero, then it is at least
\[
\begin{equation*}
\left((\Delta / \delta)^{\operatorname{poly}(n)}\right)^{-n^{O\left(r k^{2} / \epsilon\right)}}=\exp \left(-n^{O\left(k^{2} r / \epsilon\right)} \log ^{O(1)}(\Delta / \delta)\right) . \tag{16.7}
\end{equation*}
\]

Hence, OPT is at least (16.7) as well.
Plugging \(\tau \ll \epsilon \cdot(16.7) \leq \epsilon\) OPT into the algorithm from the previous section with an additional constraint \(\|\widehat{U} \widehat{V}\|_{F}^{2} \leq(\Delta / \delta)^{\text {poly }(n)}\), we do binary search to narrow down the range of \(\left[\Lambda^{-}, \Lambda^{+}\right]\)until we reach \(\Lambda\). During the \(j\) th step of binary search, we use Theorem 16.2.1
to check if the following semi-algebraic set is empty or not,
\[
S=\left\{x \in \mathbb{R}^{l} \mid p(x) \leq \Lambda_{j}^{+} q(x), p(x) \geq \Lambda_{j}^{-} q(x), q(x) \neq 0\right\}
\]
where \(\Lambda_{1}^{-}\)is initialized to be \(\tau\) and \(\Lambda_{1}^{+}\)is initialized to be poly \((n, \Delta)\). We obtain the following theorem.

Theorem 16.5.1. Given \(A, W \in \mathbb{R}^{n}, 1 \leq k \leq n\) and \(0<\epsilon, \tau<0.1\) such that:
- \(\operatorname{rank}(W)=r\);
- all the non-zero entries of \(A\) and \(W\) are multiples of \(\delta>0\);
- all the entries of \(A\) and \(W\) are at most \(\Delta>0\) in absolute value,
one can output a number \(\Lambda\) in time \(n^{O\left(k^{2} r / \epsilon\right)} \cdot \log ^{O(1)} \frac{\Delta}{\delta \tau}\) such that \(\mathrm{OPT} \leq \Lambda \leq(1+\epsilon) \mathrm{OPT}\), where
\[
\mathrm{OPT}=\min _{\substack{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times n} \\\|U V\|_{F}^{2} \leq\left(\frac{\Delta}{\delta}\right)^{\operatorname{poly}(n)}}}\|(U V-A) \circ W\|_{F}^{2}
\]

\subsection*{16.6 Recovering the Solution Itself}

Here we show how to recover an approximate solution, not only the value of OPT.
The idea is to recover the entries of \(U\) and \(V\) one by one and use the algorithm from the previous section for the corresponding decision problem. We initialize the semi-algebraic set to be
\[
S=\left\{x \in \mathbb{R}^{l} \mid q(x) \neq 0, p(x) \leq \Lambda q(x)\right\}
\]

We start by recovering the first entry of \(U\). We perform the binary search to localize the entry within an interval of length \(\delta^{\prime}\), which takes poly \((n) \cdot \log \left(\frac{\Delta}{\delta \delta^{\prime}}\right)\) invocations of the decision algorithm. For each step of binary search, we use 16.2.1 to determine if the following semi-algebraic set \(S\) is empty or not,
\[
S \cap\left\{U_{1,1}(x) \geq \widehat{U}_{1,1}^{-}, U_{1,1}(x) \leq \widehat{U}_{1,1}^{+}\right\}
\]

After that, we declare the first entry of \(U\) to be any point in this interval. This changes the cost of the solution by at most \(\delta^{\prime} \cdot\left(\frac{\Delta}{\delta}\right)^{\operatorname{poly}(n)}\). Then, we add an equality constraint that fixes the entry of \(\widehat{U}\) to this value, and add a new constraint into \(S\) permanently, e.g. \(S \leftarrow S \cap\left\{U_{1,1}(x)=\widehat{U}_{1,1}\right\}\). Next, we repeat the same with the second entry of \(U\) and so on.

This allows us to recover a solution of cost at most \((1+\epsilon) \mathrm{OPT}+\tau\) in time
\[
n^{O\left(k^{2} r / \epsilon\right)} \cdot \log ^{O(1)}(\Delta /(\delta \tau)) .
\]

\subsection*{16.7 Adversarial Matrix Completion}

Here we prove the following theorem.
Theorem 16.7.1. Let \(B \in \mathbb{R}^{n \times n}\) be a rank- \(k\) matrix with entries that are multiples of \(\delta>0\) bounded by \(\Delta>0\).

Let \(1 \leq r \leq n\) be an integer parameter.
Let \(C\) be an \(n \times n\) matrix, where in every column there are at most \(r\) question marks and other entries are equal to the corresponding entries of \(B\).

Then, there is an algorithm that:
- receives \(C\) as an input;
- outputs a rank-k matrix that is \(\tau\)-close to \(C\) in Frobenius norm (restricted to the entries that are not replaced with a question mark);
- has running time \(n^{O\left(k^{2} r\right)} \cdot \log ^{O(1)}\left(\frac{\Delta}{\delta \tau}\right)\).

A naive attempt is to use the algorithm from the previous section with \(W\) equal to 0 for the missing entries and to 1 for the "surviving" entries. Unfortunately, setting \(W\) like this does not work, since \(W\) may not have small rank.

We fix it by proving the following lemma.
Lemma 16.7.2. For every set system \(z_{1}, z_{2}, \ldots, z_{n} \subseteq[n]\) with \(\left|z_{j}\right| \leq r\) there exists a rank\((2 r+1)\) integer matrix \(W\) such that, for every \(1 \leq i \leq n\), the entry in the \(i\)-th column and j-th row,
\[
W_{i}^{j} \begin{cases}=0 & \text { if } j \in z_{i} \\ >0 & \text { if } j \in[n] \backslash z_{i}\end{cases}
\]

Moreover, all the entries of \(W\) are bounded by \(n^{O(r)}\) in absolute value.

Recall that for the \(i\)-th column of \(W, z_{i}\) denotes the set of row indices that have zero entry. For the \(i\)-th column of \(W\), define a polynomial \(p_{i}(x)\) such that
\[
p_{i}(x)=\prod_{\ell \in z_{i}}(x-j)^{2}=\sum_{\ell=1}^{2\left|z_{i}\right|+1} a_{i \ell} x^{\ell-1}
\]
where all \(a_{i \ell}\) are integers. Since \(2\left|z_{i}\right|+1\) is at most \(2 r+1\), we can also think of \(p_{i}(x)\) as a degree \(2 r\) polynomial,
\[
p_{i}(x)=\sum_{\ell=1}^{2 r+1} a_{i \ell} x^{\ell-1}
\]

Then, we can use \(a_{i \ell}\) to create a basis \(T \in \mathbb{R}^{n \times(2 r+1)}\) which has rank at most \(2 r+1\). Let \(T_{\ell}^{i}\) denote the entry of the \(i\)-th row and \(\ell\)-th column, then set \(T_{\ell}^{i}=a_{i \ell}, \forall i \in[n], \forall \ell \in[2 r+1]\). Let \(T_{\ell}\) denote the \(\ell\)-th column of matrix \(T, \forall \ell \in[2 r+1]\). To guarantee a linear combination of those \(2 r+1\) columns always outputs a nonnegative vector, we can just choose coefficients \(1, x, x^{2}, \cdots, x^{2 r}\). Let \(S\) denote a set of all possible vectors formed by the following linear combination,
\[
T(x)=\sum_{\ell=1}^{2 r+1} x^{\ell-1} T_{\ell} \in \mathbb{R}^{n \times 1}
\]

Moreover, we have
\[
\begin{aligned}
T(x) & =\sum_{\ell=1}^{2 r+1} x^{\ell-1} T_{\ell}=\sum_{\ell=1}^{2 r+1} x^{\ell-1}\left[\begin{array}{llll}
a_{1 \ell} & a_{2 \ell} & \cdots & a_{n \ell}
\end{array}\right]^{\top} \\
& =\left[\begin{array}{llll}
\sum_{\ell=1}^{2 r+1} x^{\ell-1} a_{1 \ell} & \sum_{\ell=1}^{2 r+1} x^{\ell-1} a_{2 \ell} & \cdots & \sum_{\ell=1}^{2 r+1} x^{\ell-1} a_{n \ell}
\end{array}\right]^{\top} \\
& =\left[\begin{array}{llll}
p_{1}(x) & p_{2}(x) & \cdots & p_{n}(x)
\end{array}\right]^{\top},
\end{aligned}
\]
where the second equality follows by the definition of \(T_{\ell}\) and the last equality follows by the definition of \(p_{i}(x)\). Let \(T^{j}(x)\) denote the \(j\) th entry of column vector \(T(x) \in \mathbb{R}^{n \times 1}\). For any
column vector \(W_{i}\), we assign \(T(i)\) to it,
\[
W_{i} \leftarrow T(i)
\]
which has the following property: for every \(1 \leq i \leq n\), the entry at the \(i\)-th column and \(j\)-th row satisfies
\[
W_{i}^{j}= \begin{cases}T^{j}(i)=p_{i}(j)=0 & \text { if } j \in z_{i}, \\ T^{j}(i)=p_{i}(j)>0 & \text { if } j \in[n] \backslash z_{i} .\end{cases}
\]

Note that for any column vector \(W_{i}\), we know that \(W_{i} \in S\) and \(\operatorname{rank}(S)=2 r+1\). Thus, \(\operatorname{rank}(W)=2 r+1\).

\subsection*{16.8 Few distinct columns}

In this section we show how to improve the running time from \(n^{\operatorname{poly}(k, r, 1 / \epsilon)}\) to \(\operatorname{poly}(n)\). \(2^{\operatorname{poly}(k, r, 1 / \epsilon)}\) under the following assumptions: (1) \(\Delta=\operatorname{poly}(n) \delta\); and (2) \(\|U V\|_{F}^{2} \leq(\Delta / \delta)^{n^{\gamma}}\), for an arbitrarily small constant \(\gamma>0\). In Section 16.8.1, as a warmup we assume that \(W\) has \(r\) distinct columns and \(r\) distinct rows, while in Section 16.8.2 and 16.8.3 we give our main result assuming only that \(W\) has \(r\) distinct columns.

A crucial observation is that the term \(n^{\text {poly }(k, r, 1 / \epsilon)}\) shows up in the "rank- \(r\) " algorithm due to the fact that the degree of polynomials we optimize is \(\Omega(n)\). The reason for this is that entries of \(\widehat{U}\) and \(\widehat{V}\) are rational functions with \(\Omega(n)\) potentially different denominators. When we combine them in a single rational function that corresponds to \(\|(\widehat{U} \widehat{V}-A) \circ W\|_{F}^{2}\), we get a denominator of degree \(\Omega(n)\).

\subsection*{16.8.1 Few Distinct Rows and Columns}

In this subsection, as a warmup we assume that \(W\) has \(r\) distinct rows and \(r\) distinct columns. Then, we get rid of the dependence on \(n\) in the degree. Indeed, now we have only \(2 r\) distinct denominators (w.l.o.g., assume the first \(r\) columns are distinct and the first \(r\) rows are distinct),
\[
\begin{aligned}
g_{i}(x) & =\operatorname{det}\left(\left(S D_{W_{i}} U\right)_{P_{i}}^{\top}\left(S D_{W_{i}} U\right)_{P_{i}}\right), \forall i \in[r] \\
f_{i}(x) & \left.=\operatorname{det}\left(\left(V D_{W^{i}} S\right)^{Q_{i}}\left(\left(V D_{W^{i}} S\right)^{Q_{i}}\right)\right)^{\top}\right), \forall i \in[r]
\end{aligned}
\]
where \(P_{i}\) and \(Q_{i}\) are maximal linearly independent subsets. Then we can write down the following optimization problem,
\[
\begin{array}{cl}
\min _{x \in \mathbb{R}^{l}} & p(x) / q(x) \\
\text { s.t. } & g_{i}^{2}(x) \neq 0, f_{i}^{2}(x) \neq 0, \forall i \in[r], \\
& q(x)=\prod_{i=1}^{r} g_{i}^{2}(x) f_{i}^{2}(x)
\end{array}
\]
where \(q(x)\) has degree \(O(r k)\), the maximum coefficient in absolute value is \((\Delta / \delta)^{O\left(r k n^{\gamma}\right)}\), and the number of variables \(O\left(r k^{2} / \epsilon\right)\). Using the same argument as in the rank- \(r\) case and applying Theorem 16.2.2, we can achieve the following minimum nonzero cost: \((\Delta / \delta)^{-n^{\gamma} 2^{\widetilde{O}\left(r k^{2} / \epsilon\right)}}\). Now, using the approach described in section 16.6, we can find the solution in time
\[
(\mathrm{nnz}(A)+\operatorname{nnz}(W)) n^{\gamma}+n 2^{\widetilde{O}\left(r k^{2} / \epsilon\right)} \log ^{O(1)}(\Delta /(\delta \tau))
\]
within a multiplicative factor of \(1+\epsilon\) and additive factor of \(\tau\).
One can adjust the lower bound on OPT accordingly, and conclude that an algorithm for approximating OPT within a multiplicative factor of \(1+\epsilon\) can be done in time
\[
(\mathrm{nnz}(A)+\mathrm{nnz}(W)) n^{\gamma}+n 2^{\widetilde{O}\left(r k^{2} / \epsilon\right)} \log ^{O(1)}(\Delta / \delta)
\]
where
\[
\mathrm{OPT}=\min _{\substack{U, V \\\|U V\|_{F}^{2} \leq(\Delta / \delta)^{n^{\gamma}}}}\|(U V-A) \circ W\|_{F}^{2}
\]

\subsection*{16.8.2 Few Distinct Columns When OPT \(=0\)}

This section, we explain how to find the solution to the weighted low rank approximation problem when \(W\) has at most \(r\) distinct columns and \(\mathrm{OPT}=0\).

The key observation is that for any matrix \(A\) and \(W \in \mathbb{R}_{\geq 0}^{n \times n}\), if there exists a solution of an \(n \times k\) matrix \(U\) and a \(k \times n\) matrix \(V\) such that,
\[
\|W \circ(U V-A)\|_{F}^{2}=\mathrm{OPT},
\]
then there exists another matrix \(W^{\prime} \in\{0,1\}^{n \times n}\) such that
\[
\left\|W^{\prime} \circ(U V-A)\right\|_{F}^{2}=\mathrm{OPT}
\]
where \(W_{i, j}^{\prime}=0\) if \(W_{i, j}=0\) and \(W_{i, j}^{\prime}=1\) if \(W_{i, j}>0\).
The above observation states that modifying the weight matrix to be Boolean does not change the optimal cost. Since \(W\) has \(r\) distinct columns, now that it is Boolean it has at most \(2^{r}\) distinct rows. Indeed, each row of \(W\) is completely determined after fixing its values on the \(r\) distinct columns, and there are only \(2^{r}\) possibly fixings. W.l.o.g. we assume that the first \(r\) columns are distinct. Instead of having at most \(2 r\) distinct denominators as in Section 16.8.1, we have at most \(r+2^{r}\) distinct denominators. We create \(l\) variables for \(\left\{S D_{W_{1}} U, \cdots, S D_{W_{r}} U\right\}\). Then we can write down \(\widehat{V}\) in the following way,
\[
\widehat{V}_{j}=\left(S D_{W_{j}} U\right)^{\dagger} \cdot S D_{W_{j}} A_{j}=\left(\left(\left(S D_{W_{j}} U\right)_{P_{i}}\right)^{\top} \cdot\left(S D_{W_{j}} U\right)_{P_{i}}\right)^{-1} \cdot\left(\left(S D_{W_{j}} U\right)_{P_{i}}\right)^{\top} S D_{W_{j}} A_{j}
\]
where \(P_{i}\) denotes a subset of rows. For all \(D_{W_{j}} \mathrm{~s}\) in the group \(Z_{i}\), they share the same \(P_{i}\), where for any \(j \in Z_{i}, D_{W_{j}}=D_{W_{i}}\).

Thus to express \(\widehat{V}\), there are only \(r\) distinct denominators \(g_{i}(x)\) which are the determinants of \(\left(\left(S D_{W_{j}} U\right)_{P_{i}}\right)^{\top} \cdot\left(S D_{W_{j}} U\right)_{P_{i}}\). In order to remove these denominators, we can create a new variable \(x_{l+i}\) and add a new equality constaint \(g_{i}(x) x_{l+i}-1=0\). Therefore, we do not have any denominators in \(\widehat{V}\).
W.l.o.g., we assume that the first \(2^{r}\) rows of \(W\) are distinct. Using \(\widehat{V}\) we can write down \(\widehat{U}\) in the following way,
\[
\begin{aligned}
\widehat{U}^{j} & =A^{j} D_{W^{j}}\left(\widehat{V} D_{W^{j}}\right)^{\dagger} \\
& =A^{j} D_{W^{j}}\left(\left(\widehat{V} D_{W^{j}}\right)^{Q_{i}}\right)^{\top}\left(\left(\widehat{V} D_{W^{j}}\right)^{Q_{i}}\left(\left(\widehat{V} D_{W^{j}}\right)^{Q_{i}}\right)^{\top}\right)^{-1}
\end{aligned}
\]
where \(Q_{i}\) denotes a subset of columns, where all \(D_{W^{j}} \mathrm{~S}\) in the group \(Z_{i}^{\prime}\) can share the same \(Q_{i}\), and for any \(j \in Z_{i}^{\prime}, D_{W^{j}}=D_{W^{i}}\).

Thus, to express \(\widehat{U}\), there are only \(2^{r}\) distinct denominators \(f_{j}(x)\) which are the determinants of the matrices \(\left(\widehat{V} D_{W^{j}}\right)^{Q_{i}}\left(\left(\widehat{V} D_{W^{j}}\right)^{Q_{i}}\right)^{\top}\).

Finally, we can use a small number of variables to represent all the entries of \(\widehat{U}\) and \(\widehat{V}\). It allows us to write the following optmization problem,
\[
\begin{array}{rl}
\min _{x \in \mathbb{R}^{l+r}} & p(x) / q(x) \\
\text { s.t. } & g_{i}(x) x_{l+i}-1=0, \forall i \in[r] \\
& f_{j}^{2}(x) \neq 0, \forall j \in\left[2^{r}\right] \\
& q(x)=\prod_{j=1}^{2^{r}} f_{j}^{2}(x)
\end{array}
\]
where \(q(x)\) has degree \(O\left(2^{r} k^{2}\right)\), maximum coefficients bounded in absolute value by \((\Delta / \delta)^{O\left(2^{r} k n^{\gamma}\right)}\), \(l=O\left(r k^{2} / \epsilon\right)\) and the number of variables \(O\left(r k^{2} / \epsilon\right)\). Using the same argument as for the
rank- \(r\) case and applying Theorem 16.2.2, we have the following minimum nonzero cost: \((\Delta / \delta)^{-n \gamma 2^{\tilde{O}\left(r^{2} k^{2} / \epsilon\right)} .}\)

Using the approach described in section 16.6, we can find the solution achieving zero cost in time
\[
(\mathrm{nnz}(A)+\operatorname{nnz}(W)) n^{\gamma}+n 2^{\widetilde{O}\left(r^{2} k^{2} / \epsilon\right)} \log ^{O(1)}(\Delta / \delta)
\]

\subsection*{16.8.3 Few Distinct Columns When OPT \(\neq 0\)}

Lower Bound Let \(U, V\) denote the optimal solution \(A, W\), which gives nonzero cost. We can modify \(W\) to a new matrix \(W^{\prime}\) in the following sense, \(W_{i, j}^{\prime}=\delta\) if \(W_{i, j} \neq 0\) and \(W_{i, j}^{\prime}=0\) if \(W_{i, j}=0\). Then we know that
\[
\left\|W^{\prime} \circ(U V-A)\right\|_{F}^{2} \leq\|W \circ(U V-A)\|_{F}^{2} \neq 0
\]

Note that if problem \(A, W^{\prime}\) has a zero cost solution, then problem \(A, W\) also has a zero cost solution, which contradicts our assumption in this section. Thus problem \(A, W^{\prime}\) does not have a zero cost solution. It follows from previous sections that the minimum nonzero cost
 of problem \(A, W^{\prime}\). Thus we have
\[
\|W \circ(U V-A)\|_{F}^{2} \geq\left\|W^{\prime} \circ(U V-A)\right\|_{F}^{2} \geq\left\|W^{\prime} \circ\left(U^{\prime} V^{\prime}-A\right)\right\|_{F}^{2}
\]


Algorithm For notational convenience, let \(\delta=1\). Then each entry of the input weight matrix \(W^{\prime}\) is in \(\{0,1,2, \cdots, \operatorname{poly}(n)\}\). For each entry \(W_{i, j}^{\prime}\), we round it to the smallest \((1+\epsilon)^{x}\) such that \(W_{i, j}^{\prime} \leq(1+\epsilon)^{x}\) where \(x\) is an integer. Because \(W^{\prime}\) is bounded, the total number of choices for the power \(x\) is \(O(\log (n) / \epsilon)\). Define \(W\) to be the matrix after rounding. Define OPT to be \(\min _{U, V}\left\|W^{\prime} \circ(U V-A)\right\|_{F}^{2}\). Then \(W\) has the following properties
1. \(W\) has \(r\) distinct columns,
2. \(W\) has \(R:=(\log (n) / \epsilon)^{O(r)}\) distinct rows,
3. \(\quad \mathrm{OPT} \leq \min _{U, V}\|W \circ(U V-A)\|_{F}^{2} \leq(1+\epsilon)^{2} \mathrm{OPT}\).

We prove the above three properties one by one.
The rounding is a deterministic procedure: if two values are the same in \(W^{\prime}\), then they are the same in \(W\). Hence, Property 1 holds.

To prove Property 2, take the \(r\) distinct columns \(i_{1}, \ldots, i_{r}\). Then every other column can be labeled \(j\) in \(\left\{i_{1}, \ldots, i_{r}\right\}\). If you fix the values on entries \(i_{1}, \ldots, i_{r}\) in a row, this fixes the values on every other column. So the number of distinct rows is the number of fixings to the values on \(i_{1}, \ldots, i_{r}\). Each entry has \(\log _{1+\epsilon} \operatorname{poly}(n)=O((\log n) / \epsilon)\) possibilities, so there are \(O((\log n) / \epsilon)^{r}\) distinct rows.

Because of the rounding procedure, each \(W_{i, j}^{\prime}\) satisifies that \(W_{i, j}^{\prime} \leq W_{i, j} \leq(1+\epsilon) W_{i, j}^{\prime}\), which implies Property 3.

We use the same approach as in Section 16.8.2 to create variables, write down the polynomial systems and add not equal constraints. Instead of having \(r+2^{r}\) distinct denominators, we have \(r+R\), where \(R=O((\log n) / \epsilon)^{r}\). We create \(l=O\left(r k^{2} / \epsilon\right)\) variables for \(\left\{S D_{W_{1}} U, \cdots, S D_{W_{r}} U\right\}\), then we can write down \(\widehat{V}\) with \(r\) distinct denominators \(g_{i}(x)\). Each \(g_{i}(x)\) is non-zero in an optimal solution using the perturbation argument in Section 16.4. We create new variables \(x_{i+l}\) to remove the denominators \(g_{i}(x)\). Then the entries of \(\widehat{V}\) are polynomials as opposed to rational functions. Using \(\widehat{V}\) we can express \(\widehat{U}\) with \(R=(\log (n) / \epsilon)^{O(r)}\) distinct denominators \(f_{i}(x)\), which are also non-zero by using the perturbation argument in 16.4, and using that \(W\) has at most this number of distinct rows. Finally we can write the
following optimization problem,
\[
\begin{array}{cl}
\min _{x \in \mathbb{R}^{l+r}} & p(x) / q(x) \\
\text { s.t. } & g_{i}(x) x_{l+i}-1=0, \forall i \in[r] \\
& f_{j}^{2}(x) \neq 0, \forall j \in[R] \\
& q(x)=\prod_{j=1}^{R} f_{j}^{2}(x)
\end{array}
\]

We then determine if there exists a solution to the above semi-algebraic set in time \(\left(k^{2} R\right)^{O\left(r k^{2} / \epsilon\right)}=\) \((\log (n) / \epsilon)^{O\left(r^{2} k^{2} / \epsilon\right)}\). Combining the binary search explained in section 16.5 and 16.6 with the lower bound we obtained, we can find the solution for the original problem in time
\[
(\mathrm{nnz}(A)+\operatorname{nnz}(W)) n^{\gamma}+n 2^{\widetilde{O}\left(r^{2} k^{2} / \epsilon\right)} \log ^{O(1)}(\Delta / \delta)
\]

Note that there is no \(\log \log n\) in the exponent \(2^{\widetilde{O}\left(r^{2} k^{2} / \epsilon\right)}\) since either \(r^{2} k^{2} / \epsilon=o(\log n / \log \log n)\), in which case this term is dominated by \(n^{1+\gamma}\), or \(\log \left(r^{2} k^{2} / \epsilon\right)=\Omega(\log \log n)\).

\subsection*{16.9 Hardness}

The Maximum Edge Biclique problem [AMS11] is defined as:
Input: An \(n\) by \(n\) bipartite graph \(G\).
Output: A \(k_{1}\) by \(k_{2}\) complete bipartite subgraph of \(G\).
Objective Function: Maximize \(k_{1} \cdot k_{2}\).
We use the Maximum Edge Biclique problem under the R4SAT assumption in [GL04], which extends the previous work done by Feige [Fei02] under the R3SAT assumption. That hardness result [GL04] shows under the R4SAT assumption there exist two constants \(\epsilon_{1}>\) \(\epsilon_{2}>0\) such that no efficient algorithm is able to distinguish between bipartite graphs \(G(U, V, E)\) with \(|U|=|V|=n\) which have a clique of size \(\geq(n / 16)^{2}\left(1+\epsilon_{1}\right)\) and those in which all bipartite cliques are of size \(\leq(n / 16)^{2}\left(1+\epsilon_{2}\right)\). Using the reduction of [GL04], one can show there exists a constant \(c\) such that for any instance of R4SAT with \(\widetilde{n}\) variables and \(c \widetilde{n}\) clauses, the corresponding bipartite graph \(G\) created in [GL04] has at least \(t n^{2}\) edges with large probability, for a constant \(t\), e.g. \(t=9 / 10\).

To construct a weighted low-rank appoximation problem from a given bipartite graph, for a given bipartite graph \(G(U, V, E)\), we generate the matrix \(A\) and \(W\) as in [GG11]: \(A_{i j}=1\) if edge \(\left(U_{i}, V_{j}\right) \in E, A_{i j}=0\) if edge \(\left(U_{i}, V_{j}\right) \notin E . W_{i j}=1\) if edge \(\left(U_{i}, V_{j}\right) \in E, W_{i j}=\operatorname{poly}(n)\) if edge \(\left(U_{i}, V_{j}\right) \notin E\). One can then show if there exists a biclique in \(G\) such the number of remaining edges is at most \(t n^{2}-(n / 16)^{2}\left(1+\epsilon_{1}\right)\), then the solution to min \(\|W \circ \widehat{A}-W \circ A\|_{F}^{2}\) has cost at most \(t n^{2}-(n / 16)^{2}\left(1+\epsilon_{1}\right)\). On the other hand, if there does not exist a biclique that has more than \((n / 16)^{2}\left(1+\epsilon_{2}\right)\) edges, which leads to the number of remaining edges being at least \(t n^{2}-(n / 16)^{2}\left(1+\epsilon_{2}\right)\), then any solution to \(\min \|W \circ \widehat{A}-W \circ A\|_{F}^{2}\) has cost at
least \(t n^{2}-(n / 16)^{2}\left(1+\epsilon_{2}\right)\).

\subsection*{16.10 Bicriteria Approximation}

\subsection*{16.10.1 Rank \(r\) and Binary Weights \(W\)}

This section, we explain how to get a polynomial time bicriteria algorithm when weighted matrix is boolean.

Lemma 16.10.1. Given matrix \(A \in \mathbb{R}^{n \times n}\) and rank-r weighted matrix \(W \in \mathbb{R}_{\geq 0}^{n \times n}\). We can find some rank-rk matrix \(B \in \mathbb{R}^{n \times n}\) such that
\[
\|B-W \circ A\|_{F}^{2} \leq(1+\epsilon) \min _{\operatorname{rank}\left(B^{\prime}\right)=r k}\left\|B^{\prime}-W \circ A\right\|_{F}^{2}
\]
in \(O(\operatorname{nnz}(A)+\operatorname{nnz}(W)+n \cdot \operatorname{poly}(r k / \epsilon))\) time.

Proof. It directly follows by the input sparsity low-rank approximation result in [CW13].
Lemma 16.10.2. For any matrix \(W, A \in \mathbb{R}^{n \times n}\) and \(\operatorname{rank}(W)=r\), we have
\[
\min _{\operatorname{rank}\left(B^{\prime}\right)=r k}\left\|B^{\prime}-W \circ A\right\|_{F}^{2} \leq \min _{\operatorname{rank}\left(C^{\prime}\right)=k}\left\|W \circ C^{\prime}-W \circ A\right\|_{F}^{2}
\]

Proof. For any rank- \(k\) matrix \(C^{\prime}\), matrix \(W \circ C^{\prime}\) has rank at most \(r k\). Then, for any rank- \(k\) matrix \(C^{\prime}\), there must exists a rank-rk matrix \(B^{\prime}\) such that
\[
\left\|B^{\prime}-W \circ A\right\|_{F}^{2} \leq\left\|W \circ C^{\prime}-W \circ A\right\|_{F}^{2}
\]

Theorem 16.10.3. Given two matrices \(W \in \mathbb{F}_{2}^{n \times n}\) and \(A \in \mathbb{R}^{n \times n}\), where \(\operatorname{rank}(W)=r\). Define \(A_{k}^{*}=\underset{\operatorname{rank}\left(A^{\prime}\right)=k}{\arg \min }\left\|A^{\prime}-A\right\|_{W}\). We can find a matrix \(B\) that has rank rk and such that
\[
\|B-A\|_{W}^{2} \leq(1+\epsilon)\left\|A_{k}^{*}-A\right\|_{W}^{2}
\]
in \(O(\operatorname{nnz}(A)+\operatorname{nnz}(W)+n \cdot \operatorname{poly}(k r / \epsilon))\) time.

Proof. Let matrix \(B\) denote the output of Lemma 16.10.1, let \(\Omega\) denote the support set of weighted matrix \(W, \Omega=\left\{(i, j) \mid W_{i, j}>0\right\}\).
\[
\begin{aligned}
& \|B-A\|_{W}^{2} \\
= & \|W \circ B-W \circ A\|_{F}^{2} \\
= & \sum_{(i, j) \in \Omega} W_{i j}^{2}\left(B_{i j}-A_{i j}\right)^{2}+\sum_{(i, j) \in[n]^{2} \backslash \Omega} W_{i j}^{2}\left(B_{i j}-A_{i j}\right)^{2} \\
= & \sum_{(i, j) \in \Omega} 1 \cdot\left(B_{i j}-A_{i j}\right)^{2}+\sum_{(i, j) \in[n]^{2} \backslash \Omega} 0 \cdot\left(B_{i j}-A_{i j}\right)^{2} \\
\leq & \sum_{(i, j) \in \Omega} 1 \cdot\left(B_{i j}-A_{i j}\right)^{2}+\sum_{(i, j) \in[n]^{2} \backslash \Omega} 1 \cdot\left(B_{i j}\right)^{2} \\
= & \|B-W \circ A\|_{F}^{2} \\
\leq & (1+\epsilon) \min _{\operatorname{rank}\left(B^{\prime}\right)=r k}\left\|B^{\prime}-W \circ A\right\|_{F}^{2} \text { by Lemma 16.10.1 } \\
\leq & \left(1+\epsilon \min _{\operatorname{rank}\left(C^{\prime}\right)=k}\left\|W \circ C^{\prime}-W \circ A\right\|_{F}^{2}\right. \text { by Lemma 16.10.2 }
\end{aligned}
\]

\subsection*{16.10.2 Modulo \(p\) Integer Weights \(W\)}

Theorem 16.10.4. Given two matrices \(W \in \mathbb{F}_{p}^{n \times n}\) and \(A \in \mathbb{R}^{n \times n}\), where \(\operatorname{rank}(W)=r\). Define \(A_{k}^{*}=\underset{\operatorname{rank}\left(A^{\prime}\right)=k}{\arg \min }\left\|A^{\prime}-A\right\|_{W}\). We can find a matrix \(A^{\prime}\) that has rank \(p^{r} \cdot k\) and satisfies the following
\[
\left\|A^{\prime}-A\right\|_{W}^{2} \leq(1+\epsilon)\left\|A_{k}^{*}-A\right\|_{W}^{2}
\]
in \(O\left(\operatorname{nnz}(A)+\operatorname{nnz}(W)+p^{r} \cdot n \cdot \operatorname{poly}(k / \epsilon)\right)\) time.

Proof. Since each entry of weight matrix \(W\) is chosen from \(\{0,1, \cdots, p-1\}\) and has rank \(r\). Thus, for any given \(r\) linearly independent basis, weight matrix \(W\) has at most \(p^{r}\) distinct columns. For any column vector \(A_{j}\) in \(A_{Z_{i}}\), its weighted column vector \(W_{j}\) is equal to \(W_{i}\). Then we can reduce the original problem to
\[
\begin{equation*}
\sum_{i=1}^{g} \min _{\operatorname{rank}\left(A_{i}^{\prime}\right)=k}\left\|D_{W_{i}} A_{i}^{\prime}-D_{W_{i}} A_{Z_{i}}\right\|_{F}^{2} \tag{16.8}
\end{equation*}
\]
where \(g \leq p^{r}, A_{Z_{i}} \in \mathbb{R}^{n \times\left|Z_{i}\right|}, \sum_{i=1}^{g}\left|Z_{i}\right|=n\). Since these \(g\) problems are independent, we can just solve them in parallel. Let \(\Omega_{i}\) denote the support set of rows of \(D_{W_{i}}\), ignoring the zero rows of \(D_{W_{i}} A_{Z_{i}}\), we can have a new matrix \(B_{i} \in \mathbb{R}^{\left|\Omega_{i}\right| \times\left|Z_{i}\right|}\). Consider the following low rank approximation problem
\[
\min _{\operatorname{rank}\left(B_{i}^{\prime}\right)=k}\left\|B_{i}^{\prime}-B_{i}\right\|_{F}^{2}
\]

By the input sparsity low-rank approximation result in [CW13], in time \(O\left(\mathrm{nnz}(B)+\max \left(\left|\Omega_{i}\right|,\left|Z_{i}\right|\right)\right.\). \(\operatorname{poly}(k / \epsilon)\) ), we can find some rank- \(k\) matrix \(B_{i}^{\prime} \in \mathbb{R}^{\left|\Omega_{i}\right| \times\left|Z_{i}\right|}\) such that
\[
\left\|B_{i}^{\prime}-B\right\|_{F}^{2} \leq(1+\epsilon) \min _{\operatorname{rank}\left(B_{i}^{*}\right)=k}\left\|B_{i}^{*}-B_{i}\right\|_{F}^{2}
\]

Since \(D_{W_{i}}\) is diagonal matrix, we can recover a rank- \(k\) matrix \(A_{i}^{\prime} \in \mathbb{R}^{n \times\left|Z_{i}\right|}\) by using matrix \(B_{i}^{\prime}\) and \(D_{W_{i}}\) in the follwoing way, for \(j\) th row of \(n \times\left|Z_{i}\right|\) matrix \(A_{i}^{\prime}\), if \(j\) th entry of \(W_{i}\) is zero,
then make \(j\) th row matrix \(A_{i}^{\prime}\) zero everywhere; otherwise, copy the corresponding row from \(B_{i}^{\prime}\) and rescale it by \(1 /\left(j\right.\) th entry of \(\left.W_{i}\right)\). Then the \(n \times\left|Z_{i}\right|\) matrix \(A_{i}^{\prime}\) we obtained has the following property
\[
\begin{aligned}
& \left\|D_{W_{i}} A_{i}^{\prime}-D_{W_{i}} A_{Z_{i}}\right\|_{F}^{2} \\
= & \left\|B_{i}^{\prime}-B_{i}\right\|_{F}^{2} \\
\leq & (1+\epsilon) \min _{\operatorname{rank}\left(B_{i}^{*}\right)=k}\left\|B_{i}^{*}-B_{i}\right\|_{F}^{2} \\
\leq & (1+\epsilon) \min _{\operatorname{rank}\left(A_{i}^{*}\right)=k}\left\|D_{W_{i}} A_{i}^{*}-D_{W_{i}} A_{Z_{i}}\right\|_{F}^{2} \text { by Lemma 16.10.2 }
\end{aligned}
\]

Since all the \(g\) subproblems are independent, then we can solve them separately and merge all the \(A_{i}^{\prime}\) to a matrix \(A^{\prime}\) that has rank- \(p^{r} k\). The entire procedure takes \(O(\mathrm{nnz}(A)+\mathrm{nnz}(W)+\) \(\left.p^{r} \cdot n \cdot \operatorname{poly}(k / \epsilon)\right)\).

\subsection*{16.10.3 \(r\) Distinct Columns Weights \(W\)}

Theorem 16.10.5. Given two matrices \(W \in \mathbb{R}_{\geq 0}^{n \times n}\) and \(A \in \mathbb{R}^{n \times n}\), \(W\) has \(r\) distinct columns. Define \(A_{k}^{*}=\underset{\operatorname{rank}\left(A^{\prime}\right)=k}{\arg \min }\left\|A^{\prime}-A\right\|_{W}\). We can find a matrix \(A^{\prime} \in \mathbb{R}^{n \times n}\) that has rank \(r \cdot k\) and satisfies the following
\[
\left\|A^{\prime}-A\right\|_{W}^{2} \leq(1+\epsilon)\left\|A_{k}^{*}-A\right\|_{W}^{2}
\]
in \(O(\mathrm{nnz}(A)+\mathrm{nnz}(W)+n r \cdot \operatorname{poly}(k / \epsilon))\) time.

Proof. In previous proof, we have \(p^{r}\) distinct columns in \(W\). But now, we only have \(r\) distinct columns in \(W\). We still use the same algorithm to solve each subproblem. Since the total numbe of subproblems is \(r\), thus the running time doesn't have exponential dependence for \(r\) any more.

\subsection*{16.11 Weighted Nonnegative Matrix Factorization Problem}

Arora, Ge, Moitra and Kannan [AGKM12] provided the first provably result for nonnegative matrix factorization. Later, Moitra [Mad13] improved it to almost optimal. This section explains the details of getting an algorithm for weighted nonnegative matrix factorization. We combine our new technique "guess the sketch" with previous work done by Moitra [Mad13] to write down a low degree polynomial system with some number of inequalities in terms of small number of variables. Next, we can use polynomial solver [Ren92a, Ren92b, Ren92c] to determine if there exists a nonnegative factorization of target matrix \(A\) under weighted matrix \(W\). The slover takes \((m d)^{O(v)}\) time to run, if \(m\) is the number of inequality constraints, \(d\) is the maximum degree of polynomials and \(v\) is the number of variables. We first show an algorithm that needs try \(2^{k^{O(r)}}\) guesses, the running time of which has double exponential dependence in \(r\). Then we expalin how to write down a slightly bigger polynomial systems without \(2^{k^{O(r)}}\) guesses, which reduces the double exponential of \(r\) down to single exponential dependence.

Notations Let \(W \in \mathbb{R}_{\geq 0}^{n \times n}\) denote the nonnegative weight matrix, let \(A\) denote the target \(n \times n\) matrix, our goal is to factorize \(A\) into two low-dimensional nonnegative matrices \(B \in \mathbb{R}_{\geq 0}^{n \times k}\) and \(C \in \mathbb{R}_{\geq 0}^{k \times n}\). Let \(A_{i}\) denote the \(i\) th column of matrix \(A\) and \(A^{j}\) denote the \(j\) th row of matrix \(A\). Let \(A \circ W\) denote the entry-wise product of two matrices \(A\) and \(W\), and let \(u \circ v\) denote the entry-wise product of two vectors \(A\) and \(W\). Let \(D_{\text {lev }}\) denote the leverage score matrix. Let \(D_{W_{i}}\) denote the diagonal matrix, where each entry on diagonal is from the vector \(W_{i}\), similarly for \(D_{W^{j}}\).

Definition 16.11.1. The nonnegative weighted \(\operatorname{rank} \operatorname{rank}^{+}(A)\) is defined to be the smallest
\(k\) such that \(A\) can be written as
\[
A_{i j}= \begin{cases}B^{i} C_{j}=(B C)_{i j} & \text { if } W_{i j}>0  \tag{16.9}\\ \text { arbitrary } & \text { otherwise }\end{cases}
\]
where \(B\) and \(C\) are nonnegative and have dimension \(n \times k\) and \(k \times n\) respectively, \(W\) has rank \(r\).

We use the notations \(\operatorname{aff}(A)\) and lexicographic ordering from [Mad13] directly, and extend the stable definition and addmissible definition into weighted version. The standard lexicographic ordering is usually for the subsets that have the same size, additionally, we assume the smaller size subset has early order than larger size subset.

Definition 16.11.2 (Affine Hull). Given \(n \times n\) matrix \(A\) and let \(S \subseteq[n]\). Define
\[
\operatorname{aff}\left(A_{S}\right)=\left\{\sum_{i \in S} \alpha_{i} A_{i} \mid \forall i, \alpha_{i} \geq 0\right\} .
\]

Definition 16.11.3 (Weighted Stable). \(\forall i \in[n]\), let \(S_{i} \subseteq[k]\) be the lexicographically first admissible subset ( of columns of \(B \in \mathbb{R}^{n \times k}\) ) for \(A_{i}\). \(\forall j \in[n]\), , let \(T^{j} \subseteq[k]\) be the lexicographically first admissible subset ( of rows of \(C \in \mathbb{R}^{k \times n}\) ) for \(A^{j}\). We call \(W \circ A=\) \(W \circ(B C)\) weighted stable if:
1. \(\forall i \in[n], C_{i}\) is supported in \(S_{i} \subseteq[k]\)
2. \(\forall j \in[n], B^{j}\) is supported in \(T^{j} \subseteq[k]\)

Definition 16.11.4 (Weighted Admissible). Given a weighted vector \(w \in \mathbb{R}_{\geq 0}^{n}\) and a vector \(v \in \mathbb{R}^{n}\), let \(S\) denote a subset of columns of \(A, S\) is weighted admissible if there exists some \(u \in \operatorname{aff}\left(A_{S}\right)\) such that \(w \circ v=w \circ u\).

Another interpretation of weighted admissible is thinking of we have \(*\) entry in vector \(v\) instead of thinking there is antoher weighted vector. Given a vector \(v \in\{\mathbb{R}, *\}^{n}\), let \(q \subseteq[n]\) denote the subset of coordinates that contain \(*\), we will call a subset \(S\) of columns of \(A\) admissible if there exists some \(u \in \operatorname{aff}\left(A_{S}\right)\) such that, \(\forall i \notin q, v_{i}=u_{i}\).

Lemma 16.11.1. Given matrices \(A \in \mathbb{R}^{n \times n}\) and \(W \in \mathbb{R}_{\geq 0}^{n \times n}\). If there exist some nonnegative inner-dimension \(k\) matrices \(B\) and \(C\) such that \(W \circ A=W \circ(B C)\), then there exist \(\widetilde{B}\) and \(\widetilde{C}\) such that:
1. \(W \circ A=W \circ(\widetilde{B} \widetilde{C}), \widetilde{B}\) and \(\widetilde{C}\) are nonnegative and have inner-dimension \(k\);
2. \(W \circ A=W \circ(\widetilde{B} \widetilde{C})\) is weighted stable.

Proof. In this proof, we modify the matrix \(A\) according to the zero position of matrix \(W\), \(A_{i j}=*\) if \(W_{i j}=0\). At the beginning, we have some \(B\) and \(C\) such that \(A\) and \(B C\) are entrywise consistent except the \(*\) positions. Then each update phase will alternately switching between \(B\)-updating phase and \(C\)-updating phase. In the \(C\)-updating phase, for each column vector \(A_{i} \in\{\mathbb{R}, *\}^{n \times 1}\), let \(S_{i}\) be the lexicographically first subset of columns of \(B\) that is admissible for \(A_{i}\). If \(S_{i} \subseteq[k]\) is lexicographically (strictly) earlier than the support of column vector \(C_{i} \in \mathbb{R}^{k \times 1}\), then we can find a column vector \(\widetilde{C}_{i} \in \mathbb{R}_{\geq 0}^{k \times 1}, \operatorname{supp}\left(\widetilde{C}_{i}\right) \subseteq S_{i}\) satisfies that \(A_{i}\) is entry-wise consistent with \(B \widetilde{C}_{i}\) except \(*\) position. Otherwise, we choose \(\widetilde{C}_{i}=C_{i}\). Since, we don't change anything for matrix \(B\) and \(C_{[n] \backslash i}\), then \(A\) is still entry-wise consistent with \(B \widetilde{C}\) except the * positions. The last step of this \(C\)-updating phase is updating \(C\) by \(\widetilde{C}\). Similarly, we can define the \(B\)-updating phase. During the overall updating, (1) the lexicographically order of support monotonically decrease, if \(B^{j}\) or \(C_{i}\) got updated. (2) the
matrices \(B\) and \(C\) always satisfy that \(A\) is entry-wise consistent with \(B C\) except \(*\) position, in other words. Thus, the overall updating phases will terminate with \(W \circ A=W \circ(B C)\).

\subsection*{16.11.1 Combining with Our Techniques}

Here we explain how to adapt our technique for solving the exact case to weighted nonnegative matrix factorization. (Note that there exists a low-rank matrix \(A^{\prime}\) to make the optimal cost 0, but it doesn't mean \(A\) is low-rank.) We first do leverage score sampling on the LHS of the problem, there exists some diagonal matrix \(D_{l e v}\) respect to \(\operatorname{span}\left(D_{W_{1}} B, D_{W_{2}} B, \cdots, D_{W_{n}} B\right)\) such that
\[
\min _{B, C} \sum_{i=1}^{n}\left\|D_{l e v} D_{W_{i}} B C_{i}\right\|_{2}^{2} \leq(1+\epsilon) \min _{B, C} \sum_{i=1}^{n}\left\|D_{W_{i}} B C_{i}\right\|_{2}^{2}
\]

If we know matrix \(B\), then we can compute the leverage sampling matrix deterministically and exactly. Here, we don't know the \(B\), thus, we have to either guess those nonzero probabilities or creat variables for nonzero probabilities. Let \(t\) denote the number of nonzero probabilities in leverage score sampling matrix, then

Claim 16.11.2. There are \(O\left(n^{t}\right)\) leverage score sampling matrix we need to guess.

Since \(D_{\text {lev }}\) is the leverage score sampling matrix that has \(t=\left(k r / \epsilon^{2}\right)\) nonzero probability on diagonal. Thus, the nonzero position of \(D_{\text {lev }}\) will determine the set of rows of \(B\). We create variables for the rows of \(B\) which are chosen by \(D_{l e v}\). There are \(t\) such rows, thus we need \(t r\) variables in total. To write down \(D_{l e v} D_{W_{i}} B\) by a small number of variables, we also need to create \(t\) variables for the nonzero probability of \(D_{\text {lev }}\). Thus,

Claim 16.11.3. \(\forall i \in[n], t \times k\) matrix \(D_{l e v} D_{W_{i}} B\) can expressed as \(O(t r)\) variables, where each entry of that matrix is single monomial that has degree at most 2.

In order to write down \(C_{i}\) by these \(O(t r)\) small number of variables, we need to group \(C_{i}\) by looking at the zero-pattern of \(D_{l e v} D_{W_{i}}\). Let \(g\) denote the number of different
zero-patterns, let \(z_{i}\) denote the set of indices belong to the same group,
\[
\min _{B, C} \sum_{i=1}^{n}\left\|D_{l e v} D_{W_{i}} B C_{i}\right\|_{2}^{2}=\min _{B, C} \sum_{\ell=1}^{g} \sum_{i \in z_{\ell}}\left\|D_{l e v} D_{W_{i}} B C_{i}\right\|_{2}^{2}
\]

By the similar result from previous section, there are only \(\binom{t}{r}=k^{r}\) different zeropatterns. Let \(D_{l e v} D_{W_{\ell}^{\prime}}\) denote the zero-pattern of group \(z_{\ell}, \forall \ell \in[g]\). Let rank \(\left(D_{l e v} D_{W_{\ell}^{\prime}}\right)=s_{\ell}\). For each group, we first guess \(s_{\ell}\) which is the size of maximal linearly independent of rows of \(D_{l e v} D_{W_{\ell}^{\prime}} B\), it has \(k\) possibilities. Second we need to guess the maximal linearly independent set of rows of \(D_{l e v} D_{W_{\ell}^{\prime}} B\), let \(U_{\ell}\) denote the subset of rows of \(D_{l e v} D_{W_{\ell}^{\prime}} B\). Then we define \(S_{1}, S_{2}, \cdots, S_{p}\) to be the full list of sets of \(s_{\ell}\) linearly independent columns of \(D_{l e v} D_{W_{\ell}^{\prime}} B\) in lexicographic order, where \(p \leq\binom{ k}{s_{\ell}} \leq 2^{k}\). By the definition of \(U_{\ell}\) and \(S_{i}\), we have the property that each \(s_{\ell} \times s_{\ell}\) submatrix \(\left(D_{l e v} D_{W_{\ell}^{\prime}} B\right)_{S_{i}}^{U_{\ell}}\) is full rank and invertible.

Claim 16.11.4. If \(D_{l e v} D_{W_{i}}\) and \(D_{l e v} D_{W_{j}}\) belong to the same zero pattern group \(z_{\ell}\), then they must have the same rank. There are \(k\) possibilities for the rank of \(D_{l e v} D_{W_{\ell}^{\prime}} B, \forall \ell \in\left[k^{r}\right]\). In total, there are \(O\left(k^{k^{r}}\right)\) guesses need to try.

Claim 16.11.5. If \(D_{l e v} D_{W_{i_{1}}}\) and \(D_{\text {lev }} D_{W_{i_{2}}}\) belong to the same zero pattern group \(z_{\ell}\), then they must have the same \(U_{\ell}\). There are \(r^{O(k)}\) possibilities for row subset \(U_{\ell}\) of \(D_{l e v} D_{W_{\ell}^{\prime}} B, \forall \ell \in\) [ \(k^{r}\) ]. In total, there are \(2^{k^{O(r)}}\) guesses need to try. Fix a guess of \(U_{\ell}\), there are at most \(2^{k}\) possibilities for column subset \(S_{j}\) of \(D_{l e v} D_{W_{\ell}^{\prime}} B, \forall \ell \in\left[k^{r}\right]\). In total, there are \(2^{k^{O(r)}}\) guesses need to try.

\subsection*{16.11.2 Few Entries Determine \(B\) and \(C\)}

We extend Lemma 2.7 in [Mad13] to the the following weighted version.

Lemma 16.11.6. For each column vector \(A_{i} \in \mathbb{R}^{n \times 1}\), let \(D_{l e v} D_{W_{i}}\) belong to the group \(z_{\ell}\), then among the set of vectors
\[
\mathcal{S}=\left\{\widehat{B}_{1}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}}, \widehat{B}_{2}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}}, \cdots, \widehat{B}_{p}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}}\right\}
\]
column vector \(C_{i} \in \mathbb{R}^{k \times 1}\) is the unique vector with lexicographically minimal support among all nonnegative vectors in the above set. Note that \(\left(D_{\text {lev }} D_{W_{i}} B\right)_{S_{j}}^{U_{\ell}} \in \mathbb{R}^{s_{\ell} \times s_{\ell}},\left(\left(D_{\text {lev }} D_{W_{i}} B\right)_{S_{j}}^{U_{\ell}}\right)^{\dagger} \in\) \(\mathbb{R}^{s_{\ell} \times s_{\ell}}, \widehat{B}_{j}^{i, \ell} \in \mathbb{R}^{k \times s_{\ell}}, \forall j \in[p]\),
\[
\widehat{B}_{j}^{i, \ell}=\underbrace{P_{\pi:[s e] \rightarrow[k]}}_{k \times s_{\ell}} \cdot\left(\left(D_{l e v} D_{W_{i}} B\right)_{S_{j}}^{U_{\ell}}\right)^{\dagger}, \forall i \in[n]
\]
\(P_{\pi:[s] \rightarrow[k]}\) denotes a permutation matrix that has exact one 1 in each column, and there are \(s\) out of \(k\) rows has exact one 1, everywhere else are all 0.

Proof. We already shown that the weighted stable \(C_{i}\) belongs to the set \(\mathcal{S}\). Thus, there must exist some \(j^{*} \in[p]\) such that \(C_{i}=\widehat{B}_{j^{*}}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}}\). Then, for any \(j \neq j^{*}\), consider the nonnegative vector \(\widehat{B}_{j}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}}=v\). We need to show that the support of \(v\) is lexicographically later than the support of \(C_{i}\).

First, we claim that if \(v \neq C_{i}\) then the support of \(C_{i}\) is not the same as the support of \(v\). Let's prove it by getting a contradiction. Suppose \(v \neq C_{i}\) and both of them have support \(R_{i}\). Indeed \(R_{i}\) must correspond to a linearly independent set of columns of \(D_{W_{i}} B\). By Claim 16.11.8, we obtain that \(\left(D_{W_{i}} B\left(v-C_{i}\right)\right)^{\Omega_{i}}=(\overrightarrow{0})^{\Omega_{i}}\). By \(v-C_{i} \neq \overrightarrow{0}\) and support of \(v-C_{i}\) contained in \(R_{i}\), we get \(\left(D_{W_{i}} B\left(v-C_{i}\right)\right)^{\Omega_{i}} \neq(\overrightarrow{0})^{\Omega_{i}}\). Thus, we get a contradiction.

So the support \(C_{i}\) and \(v\) are not identical and one of these must be lexicographically earlier. Suppose (for contradiction) that the support of \(v\) is earlier. By Claim 16.11.7, we have that the support of \(C_{i}\) is a weighted admissible set of columns of \(B\) for \(A_{i}\). This contradicts stability (because we could update \(C_{i}\) to \(v\) ), and so we can conclude that the support of \(C_{i}\) is lexicographically earlier.

It remains to prove Claim 16.11.7 and Claim 16.11.8, which are the weighted version of Lemma 2.8 and Lemma 2.9 in [Mad13].

Claim 16.11.7. For each column vector \(C_{i} \in \mathbb{R}^{k \times 1}\), it is contained in the set \(\mathcal{S}\).

Proof. Let \(R_{i}\) denote the support of column vector \(C_{i} \in \mathbb{R}^{k \times 1}\). Then \(R_{i}\) must correspond to a linearly independent set of columns of \(B\) - otherwise we could find a nonnegative \(\widetilde{C}_{i}\) whose support is a strict subset of \(R_{i}\) such that \(W_{i} \circ\left(B \widetilde{C}_{i}\right)=W_{i} \circ A_{i}\), but this would violate the condition of stability.

Suppose \(D_{l e v} D_{W_{i}}\) belongs to zero pattern \(z_{\ell}\), let \(U_{\ell}\) denote maximal linearly indepedent set of rows of \(D_{l e v} D_{W_{\ell}^{\prime}} B\) and \(s_{\ell}\) is the rank. Because the sets of linearly independent columns of \(D_{\text {lev }} D_{W_{i}} B \in \mathbb{R}^{n \times k}\) are a matroid, there is a set \(S_{j} \subset[k]\) of \(s_{\ell}\) linearly independent columns of \(D_{l e v} D_{W_{i}} B\) for which \(R_{i} \subset S_{j}\), hence
\[
\begin{aligned}
& \widehat{B}_{j}^{i, \ell}\left(D_{\text {lev }} D_{W_{i}} A_{i}\right)^{U_{\ell}} \\
= & \widehat{B}_{j}^{i, \ell}\left(D_{\text {lev }} D_{W_{i}} B C_{i}\right)^{U_{\ell}} \text { by } A_{i}=B C_{i} \\
= & \widehat{B}_{j}^{i, \ell}\left(D_{l e v} D_{W_{i}} B\right)^{U_{\ell}} C_{i} \\
= & P_{\pi:\left[s_{\ell}\right] \rightarrow[k]} \cdot\left(\left(D_{l e v} D_{W_{i}} B\right)_{S_{j}}^{U_{\ell}}\right)^{\dagger} \cdot\left(D_{l e v} D_{W_{i}} B\right)_{S_{j}}^{U_{\ell}} \cdot C_{i} \\
= & v
\end{aligned}
\]

Note that \(\widehat{B}_{j}^{i, \ell}\) is zero on rows outside the set \(S_{j}\) and support \(C_{i}\) is contained in \(S_{j}\), then we have \(C_{i}=v\)

Claim 16.11.8. For each column vector \(\widehat{B}_{j}^{i, \ell}\left(D_{\text {lev }} D_{W_{i}} A_{i}\right)^{U_{\ell}} \in \mathbb{R}^{k \times 1}\), we have
\[
W_{i} \circ\left(B \widehat{B}_{j}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}}\right)=W_{i} \circ A_{i}
\]
or equivalently
\[
D_{W_{i}}\left(B \widehat{B}_{j}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}}\right)=D_{W_{i}} A_{i}
\]

Proof. Let \(\Omega_{i}\) denote the support of \(W_{i}\). Define \(v=B \widehat{B}_{j}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}, \Omega_{i}} \in \mathbb{R}^{n \times 1}\). First, we prove that \(v^{U_{\ell}, \Omega_{i}}=A_{i}^{U_{\ell}, \Omega_{i}}\). Since \(\widehat{B}_{j}^{i, \ell} \in \mathbb{R}^{k \times\left|U_{\ell} \cap \Omega_{i}\right|}\) is zero on rows outside the set \(S_{j}\), we have
\[
\underbrace{B}_{n \times k} \underbrace{\widehat{B}_{j}^{i, \ell}}_{k \times\left|U \cap \Omega_{i}\right|}=B_{S_{j}}\left(\widehat{B}_{j}^{i, \ell}\right)^{S_{j}}=B_{S_{j}}\left(\left(D_{l e v} D_{W_{i}} B\right)_{S_{j}}^{U_{\ell}}\right)^{\dagger}
\]

Thus
\[
\begin{aligned}
v^{U_{\ell}, \Omega_{i}} & =\left(B \widehat{B}_{j}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}, \Omega_{i}}\right)^{U_{\ell}, \Omega_{i}} \\
& =\left(B \widehat{B}_{j}^{i, \ell}\right)^{U_{\ell}, \Omega_{i}}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}, \Omega_{i}} \\
& =B^{U_{\ell}, \Omega_{i}} \widehat{B}_{j}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}, \Omega_{i}} \\
& =B_{S_{j}}^{U_{\ell}, \Omega_{i}}\left(\left(D_{l e v} D_{W_{i}} B\right)_{S_{j}}^{U_{\ell}}\right)^{\dagger}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}, \Omega_{i}} \\
& =A_{i}^{U_{\ell}, \Omega_{i}}
\end{aligned}
\]

Second, we prove that \(v^{\overline{U_{\ell}}, \Omega_{i}}=A_{i}^{\overline{U_{\ell}}, \Omega_{i}}\). Since \(U_{\ell}\) is the maximal linearly independent set of rows of \(D_{\text {lev }} D_{W_{i}} B\) and \(D_{\text {lev }}\) preserves the rank of \(D_{W_{i}} B\). For any \(t \in \overline{U_{\ell}} \cap \Omega_{i}\), the row vector \(B^{t} \in \mathbb{R}^{1 \times k}\) can be written as a linear combination of rows in \(B\) in the set \(U_{\ell} \cap \Omega_{i}\) :
\[
B^{t}=\sum_{t^{\prime} \in U_{\ell} \cap \Omega_{i}} \alpha_{t, t^{\prime}} B^{t^{\prime}}
\]

Note that \(W_{i} \circ\left(B C_{i}\right)=W_{i} \circ A_{i}\) implies \(\left(B C_{i}\right)^{\Omega_{i}}=\left(A_{i}\right)^{\Omega_{i}}\). Then we have for \(t \in \overline{U_{\ell}} \cap \Omega_{i}\), \(A_{i}^{t}=B^{t} C_{i}=\sum_{t^{\prime} \in U_{\ell} \cap \Omega_{i}} \alpha_{t, t^{\prime}} B^{t^{\prime}} C_{i}=\sum_{t^{\prime} \in U_{\ell} \cap \Omega_{i}} \alpha_{t, t^{\prime}} A_{i}^{t^{\prime}}\) and hence:
\[
\begin{aligned}
v^{t} & =\left(B \widehat{B}_{j}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}, \Omega_{i}}\right)^{t} \\
& =B^{t} \widehat{B}_{j}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}, \Omega_{i}} \\
& =\sum_{t^{\prime} \in U_{\ell} \cap \Omega_{i}} \alpha_{t, t^{\prime}} B^{t^{\prime}} \widehat{B}_{j}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}, \Omega_{i}} \\
& =\sum_{t^{\prime} \in U_{\ell} \cap \Omega_{i}} \alpha_{t, t^{\prime}} v^{t^{\prime}} \\
& =\sum_{t^{\prime} \in U_{\ell} \cap \Omega_{i}} \alpha_{t, t^{\prime}} A_{i}^{t^{\prime}} \\
& =A_{i}^{t}
\end{aligned}
\]

We still need to show that leverage score sampling matrix doesn't destroy first lexicographically order property.

Claim 16.11.9. If \(\widetilde{C} \in \mathbb{R}^{k \times n}\) is the first lexicographically order solution of the original problem \(A=B C\) under \(W\), then it is also a solution to the problem \(D_{\text {lev }} A=D_{\text {lev }} B C\) under \(W\).

Proof. Subspace embedding preserves the equality. For each \(i \in[n], D_{W_{i}} A_{i}=D_{W_{i}} B \widetilde{C}_{i}\) implies that \(D_{l e v} D_{W_{i}} A_{i}=D_{l e v} D_{W_{i}} B \widetilde{C}_{i}\).

Claim 16.11.10. If \(C^{\prime} \in \mathbb{R}^{k \times n}\) is the first lexicographically order solution of the new problem \(D_{\text {lev }} A=D_{\text {lev }} B C\) under \(W\), then the \(C^{\prime}\) has to be equal to \(\widetilde{C}\).

Proof. Let's prove it by contradiction. Suppose the order of \(C^{\prime}\) is early than \(\widetilde{C}\). For each \(i \in[n]\), we know that \(D_{l e v} D_{W_{i}} A_{i}=D_{\text {lev }} D_{W_{i}} B C_{i}^{\prime}\). Since we also know that \(D_{l e v} D_{W_{i}} A_{i}=\) \(D_{l e v} D_{W_{i}} B \widetilde{C}_{i}\). Combing the above equation gives that \(D_{l e v} D_{W_{i}} B C_{i}^{\prime}=D_{\text {lev }} D_{W_{i}} B \widetilde{C}_{i}\), since \(C^{\prime} \neq \widetilde{C}\), then there must exist one \(i\) such that \(C_{i}^{\prime} \neq \widetilde{C}_{i}\), which means \(D_{W_{i}} B C_{i}^{\prime} \neq D_{W_{i}} B \widetilde{C}_{i}\).

By leverage score sampling, if \(D_{W_{i}} B C_{i}^{\prime} \neq D_{W_{i}} B \widetilde{C}_{i}\), then \(D_{\text {lev }} D_{W_{i}} B C_{i}^{\prime} \neq D_{\text {lev }} D_{W_{i}} B \widetilde{C}_{i}\)

\subsection*{16.11.3 A Semi-Algebraic Set with Double Exponential Depence in \(r\)}

Here, we explain the details of constraints should be added into the polynomials system solver.

Constraints 1 Define first \((\mathcal{S})\) to be the function that takes a collection of vectors as input and outputs the vector with lexicographically minimal support among all nonnegative vectors in \(\mathcal{S}\). For each \(k\)-dimensional column vector \(C_{i}, i \in[n]\), we need to add nonnegative constraints to it,
\[
C_{i} \geq \overrightarrow{0}, \forall i \in[n]
\]
where \(C_{i}=\operatorname{first}\left(\left\{\widehat{B}_{1}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}}, \widehat{B}_{2}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}}, \cdots, \widehat{B}_{p}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}}\right\}\right)\) and \(D_{l e v} D_{W_{i}}\) belong to the zero-pattern \(z_{\ell}\).

Similarly, we can choose a leverage score sampling matrix \(D_{l e v^{\prime}}\) on the RHS of the original problem and create variables for \(C^{\operatorname{supp}\left(D_{l e v^{\prime}}\right)}\). Thus, for each \(k\)-dimensional row vector \(B^{j}, j \in[n]\), we also need to add nonnegative constraints to it,
\[
B^{j} \geq \overrightarrow{0}, \forall j \in[n]
\]
where \(B^{j}=\operatorname{first}\left(\left\{\left(A^{j} D_{W^{j}} D_{l e v^{\prime}}\right)^{U_{\ell}} \widehat{C}_{1}^{j, \ell},\left(A^{j} D_{W^{j}} D_{l e v^{\prime}}\right)^{U_{\ell}} \widehat{C}_{2}^{j, \ell}, \cdots,\left(A^{j} D_{W^{j}} D_{l e v^{\prime}}\right)^{U_{\ell}} \widehat{C}_{p}^{j, \ell}\right\}\right)\) and \(D_{W^{j}} D_{l e v^{\prime}}\) belong to the zero-pattern \(z_{\ell}\).

Constraints 2 Since we know the representation of \(C_{i}, B^{j}\) and the exact value of \(A_{i}^{j}\), then we need to add \(n^{2}\) equality constraints,
\[
B^{j} C_{i}=A_{i}^{j}, \text { if } W_{i}^{j} \neq 0, \forall i \in[n], \forall j \in[n]
\]
where we need to clean up the denominator of \(B^{j}\) and \(C_{i}\) when writing down this constraint.

Constraints 3 We also need to write down the inequality constraints for the denominators of every \(\widehat{B}_{j}^{i, \ell}\), by definition, the denominator is from \(\left(\left(D_{l e v} D_{W_{i}} B\right)_{S_{j}}^{U_{\ell}}\right)^{\dagger}\). To guarantee the denominator is nonzero, we need to add this constraint,
the denominator of \(\left(\left(D_{\text {lev }} D_{W_{i}} B\right)_{S_{j}}^{U_{\ell}}\right)^{\dagger} \neq 0\).
where the number of such constraints is
\[
\# D_{W_{i}} \cdot \# S_{j} \leq n \cdot 2^{k}
\]
which is same for the other side \(C\).

Polynomial system Claim 2.12 in [Mad13] is able to write the explicit Boolean fucntion \(\mathbb{P}\) for the unweighted case of nonnegative matrix factorization problem. Note that \(\mathbb{P}\) is a function of sign constraints on polynomials. Here, we extend that the Boolean function to a weighted case of nonnegative matrix factorization.
\[
\wedge \wedge_{\ell \in[g]} \wedge_{\ell^{\prime} \in\left[g^{\prime}\right]}\left(\underset{i \in\left[z_{\ell}\right]}{\wedge} \wedge_{j \in\left[z_{\ell^{\prime}}\right]}\left(\underset{i_{1} \in\left[\binom{k}{s_{\ell}}\right]}{\vee} \underset{j_{1} \in\left[\binom{k}{\omega_{\ell^{\prime}}}\right]}{\vee}\left(f_{1} \wedge f_{2} \wedge f_{3} \wedge f_{4} \wedge f_{5} \wedge f_{6} \wedge f_{7}\right)\right)\right)
\]
where
\[
\begin{gathered}
f_{1}:\left(\left(A^{j} D_{W^{j}} D_{l e v^{\prime}}\right)_{V_{\ell^{\prime}}}{\widehat{C_{j}}}_{j, \ell^{\prime}} \cdot \widehat{B}_{i_{1}}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell},}-A_{j}^{i}\right)=0 \\
f_{2}:=\left(\widehat{B}_{i_{1}}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}}\right) \geq 0 \\
f_{3}:=\left(\left(A^{j} D_{W^{j}} D_{l e v^{\prime}}\right)_{V_{\ell^{\prime}}}{\widehat{j_{1}}}_{j, \ell^{\prime}}\right) \geq 0
\end{gathered}
\]
\[
\begin{gathered}
f_{4}:=\left(\widehat{B}_{i_{1}}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{\ell}} \text { is the first out of } \mathcal{S}^{i}\right) \\
f_{5}:=\left(\left(A^{j} D_{W^{j}} D_{l e v^{\prime}}\right)_{V_{\ell^{\prime}}} \widehat{C}_{j_{1}}^{j, \ell^{\prime}} \text { is the first out of } \mathcal{T}^{j}\right) \\
f_{6}:=\left(\text { denominator polynomial of } \widehat{B}_{i_{1}}^{i, \ell}\right) \neq 0 \\
f_{7}:=\left(\text { denominator polynomial of } \widehat{C}_{j_{1}}^{j, \ell^{\prime}}\right) \neq 0
\end{gathered}
\]
where the sign of the polynomial in each constraint can be recovered as a Boolean function of signs of degree at most poly \((k r)\).

Theorem 16.11.11. Given matrix \(A \in \mathbb{R}^{n \times n}\) and rank-r weighted matrix \(W \in \mathbb{R}_{\geq 0}^{n \times n}\). It takes \(n^{O\left(r k^{2}\right)} 2^{2^{O(r \log k)}}\) time to determine if there exists some matrices \(B \in \mathbb{R}_{\geq 0}^{n \times k}\) and \(C \in\) \(\mathbb{R}_{\geq 0}^{k \times n}\) such that \(W \circ(B C)=W \circ A\), which means existing a weighted nonnegative matrix factorization.

Proof. The final running time comes from two parts, (1) solving the polynomial system, it is known that the running time is
\[
(\text { degree } \cdot \# \text { constraints) })^{\# \text { variables }}
\]

The number of variables is \(O(t k)=O\left(r k^{2}\right)\). The degree is increasing by using Cramer rule to compute the inverse of full rank square matrix, the degree is at most poly \((k r)\). The number of constraints is at most \(n^{2} \cdot 2^{\text {poly }(k r)}\). Thus, running time of polynomial system solver is at most \(n^{O\left(r k^{2}\right)} 2^{\text {poly(kr)}}\). (2) The number of guesses for leverage score sampling matrix is \(n^{O(t)}\) and the number of guesses for each group's maximum linearly independent rows in \(B\) is \(2^{k^{O(r)}}\). Thus, we need to rerun the solver for \(n^{O(t)} 2^{k^{O(r)}}\) different polynomial systems. In summary, the entire running is \(n^{O\left(r k^{2}\right)} 2^{k^{O(r)}}\) by plugging \(t=O(k r)\).

\subsection*{16.11.4 A Semi-Algebraic Set with Single Exponential Dependence in \(r\)}

The intuition of improving double exponential dependence to single exponential dependence is, instead of guessing a possibility of rank and maximum linearly independent rows for each group at a time, we can write down all the possibilities into a slightly bigger polynomial system.

We define
\[
\begin{aligned}
& \mathcal{S}^{i, s_{\ell}, p_{s_{\ell}}}= \\
& \left\{\widehat{B}_{1,1}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{s_{\ell}, 1}}, \cdots, \widehat{B}_{1,\left(,{ }_{s_{\ell}}^{k}\right)}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{s_{\ell}, 1}}\right. \\
& \widehat{B}_{2,1}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{s_{\ell}, 2}}, \cdots, \widehat{B}_{\left.2,{ }_{s_{\ell}}^{k}\right)}^{i,}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{s_{\ell}, 2}}, \\
& \cdots, \\
& \left.\widehat{B}_{p_{s_{\ell}, 1}}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{s_{\ell}, p_{s}}}, \cdots, \widehat{B}_{p_{s_{\ell}},\left({ }_{s_{\ell}}^{k}\right)}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{s_{\ell}, p_{s_{\ell}}}}\right\}
\end{aligned}
\]
where \(z_{\ell}\) is corresponding to the zero-pattern of \(D_{l e v} D_{W_{i}}, s_{\ell}\) is the rank of \(D_{l e v} D_{W_{\ell}^{\prime}}\), \(U_{s_{\ell}, i}, \forall i \in\left[p_{s_{\ell}}\right]\) is a set of indices of \(s_{\ell}\) rows of submatrix \(D_{l e v} D_{W_{\ell}^{\prime}}, p_{s_{\ell}}=\binom{t}{s_{\ell}} \leq\binom{ t}{k}=O\left(r^{k}\right)\). Then we define
\[
\mathcal{S}^{i}=\bigcup_{s_{\ell}=1}^{k} \mathcal{S}^{i, s_{\ell}, p_{s_{\ell}}}
\]
where the number of entries in set \(S^{i}\) is at most \(k \cdot r^{k} \cdot 2^{k}\).
Then the Constraints 3 is slightly changing to the following,

Constraints 3' We also need to write down the inequality constraints for the denominators of every \(\widehat{B}_{i_{1}, i_{2}}^{i, \ell}\), by definition, the denominator is from \(\left(\left(D_{l e v} D_{W_{i}} B\right)_{S_{i_{2}}}^{U_{s_{i}, i_{1}}}\right)^{\dagger}\). To guarantee the
denominator is nonzero, we need to add this constraint,
\[
\text { the denominator of }\left(\left(D_{l e v} D_{W_{i}} B\right)_{S_{i_{2}}}^{U_{s_{i}, i_{1}}}\right)^{\dagger} \neq 0
\]
where the number of such constraints is
\[
\# D_{W_{i}} \cdot \# s_{\ell} \cdot \# i_{1} \cdot \# S_{i_{2}} \leq n \cdot k \cdot r^{O(k)} \cdot 2^{k}=n \cdot 2^{O(k \log r)}
\]
which is same for the other side \(C\).

Polynomial System Finally, we can write down the boolean function \(\mathbb{P}\) over the polynomial inequality constraints,
\(\wedge \wedge_{\ell \in[g]} \wedge \in\left(\underset{\ell^{\prime} \in\left[g^{\prime}\right]}{\wedge}\left(\underset{s_{\ell} \in[k]}{\vee} \underset{\omega_{\ell^{\prime}} \in[k]}{\vee} \underset{i_{1} \in\left[p_{s_{\ell}}\right]}{\vee} \underset{j_{1} \in\left[q_{\left.\omega_{\ell^{\prime}}\right]}\right]}{\vee}\left(\underset{i \in\left[z_{\ell}\right]}{\wedge} \wedge_{j \in\left[z_{\ell^{\prime}}\right]}^{\wedge}\left(\underset{i_{2} \in\left[\binom{k}{s_{\ell}}\right]}{\vee} \underset{j_{2} \in\left[\binom{k}{\omega_{\ell^{\prime}}}\right]}{\vee}\left(f_{1} \wedge f_{2} \wedge \cdots \wedge f_{6} \wedge f_{7}\right)\right)\right)\right)\right.\)
where
\[
\begin{gathered}
f_{1}:=\left(\left(A^{j} D_{W^{j}} D_{l e v^{\prime}}\right)_{V_{\omega_{\ell^{\prime}}, j_{1}}} \widehat{C}_{j_{1}, j_{2}}^{j, \ell^{\prime}} \cdot \widehat{B}_{i_{1}, i_{2}}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{s_{\ell}}, i_{1}}-A_{j}^{i}\right)=0 \\
f_{2}:=\left(\widehat{B}_{i_{1}, i_{2}}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{s_{\ell} \ell}, i_{1}}\right) \geq 0 \\
f_{3}:=\left(\left(A^{j} D_{W^{j}} D_{l e v^{\prime}}\right)_{V_{\omega_{\ell^{\prime}}, j_{1}}} \widehat{C}_{j_{1}, j_{2}}^{j, \ell^{\prime}}\right) \geq 0 \\
f_{4}:=\left(\widehat{B}_{i_{1}, i_{2}}^{i, \ell}\left(D_{l e v} D_{W_{i}} A_{i}\right)^{U_{s_{\ell},}, i_{1}} \text { is the first out of } \mathcal{S}^{i}\right) \\
f_{5}:=\left(\left(A^{j} D_{W^{j}} D_{l e v^{\prime}}\right)_{V_{\omega_{\ell^{\prime}, j_{1}}}} \widehat{C}_{j_{1}, j_{2}}^{j, \ell^{\prime}} \text { is the first out of } \mathcal{T}^{j}\right)
\end{gathered}
\]
\[
\begin{aligned}
& f_{6}:=\left(\text { denominator polynomial of } \widehat{B}_{i_{1}, i_{2}}^{i, \ell}\right) \neq 0 \\
& f_{7}:=\left(\text { denominator polynomial of } \widehat{C}_{j_{1}, j_{2}}^{j, \ell^{\prime}}\right) \neq 0
\end{aligned}
\]

Note that we only need to adding the functions \(f_{1}, f_{2}, f_{3}, f_{4}, f_{5}\) into the this big polynomial system when \(W_{j i} \neq 0\).

Theorem 16.11.12. Given matrix \(A \in \mathbb{R}^{n \times n}\) and rank-r weighted matrix \(W \in \mathbb{R}_{\geq 0}^{n \times n}\). It takes \(n^{O\left(r k^{2}\right)} 2^{\widetilde{O}\left(r k^{3}\right)}\) time to determine if there exists some matrices \(B \in \mathbb{R}_{\geq 0}^{n \times k}\) and \(C \in\) \(\mathbb{R}_{\geq 0}^{k \times n}\) such that \(W \circ(B C)=W \circ A\), which means existing a weighted nonnegative matrix factorization.
```

Algorithm 16.1 Bicriteria Approximation Algorithm - Section 16.10
procedure LowRankApproximation $(A, n, k, \epsilon)$
Use [CW13] to find rank- $k$ matrix $A^{\prime}$ such that:
$\left\|A^{\prime}-A\right\|_{F}^{2} \leq(1+\epsilon) \min _{\operatorname{rank}(B)=k}\|B-A\|_{F}^{2}$.
return $A^{\prime}$.
end procedure
procedure BicriteriaLowRankBooleanW $(A, W, n, r, k, \epsilon)$ - Section 16.10.1
Comment $\operatorname{rank}(W)=r$ and $W \in \mathbb{F}_{2}^{n \times n}$.
$B \leftarrow$ LowRankApproximation $(A \circ W, n, r k, \epsilon)$.
return $B$.
end procedure
procedure BicriteriaLowRankFiniteFieldW $(A, W, n, r, k, p, \epsilon)$ - Section 16.10.2
Comment $\operatorname{rank}(W)=r$ and $W \in \mathbb{F}_{p}^{n \times n}$.
$g \leftarrow \#$ distinct columns of $W$.
for $i=1 \rightarrow g$ do
Let $Z_{i}$ denote the subset of columns of $A$ such that, all the columns have the same
weight column vector.
$\Omega_{i} \leftarrow \operatorname{supp}\left(D_{W_{i}}\right)$.
$B_{i} \leftarrow\left(D_{W_{i}} A_{Z_{i}}\right)^{\Omega_{i}}$.
$B_{i}^{\prime} \leftarrow$ LowRankApproximation $\left(B_{i},\left(\left|\Omega_{i}\right|,\left|Z_{i}\right|\right), k, \epsilon\right)$.
end for
Reconstruct rank $k \cdot p^{r}$ matrix $B$ by merging all $B_{i}^{\prime}$.
return $B$.
end procedure
procedure BicriteriaRDistinctColumnsW $(A, W, n, r, k, \epsilon)$ - Section 16.10.3
Comment $W$ has $r$ distinct columns and $W \in \mathbb{R}_{\geq 0}^{n \times n}$
for $i=1 \rightarrow r$ do
Let $Z_{i}$ denote the subset of columns of $A$ such that, all the columns have the same
weight column vector.
$\Omega_{i} \leftarrow \operatorname{supp}\left(D_{W_{i}}\right)$.
$B_{i} \leftarrow\left(D_{W_{i}} A_{Z_{i}}\right)^{\Omega_{i}}$.
$B_{i}^{\prime} \leftarrow \operatorname{LowRankApproximation}\left(B_{i},\left(\left|\Omega_{i}\right|,\left|Z_{i}\right|\right), k, \epsilon\right)$.
end for
Reconstruct rank $r k$ matrix $B$ by merging all $B_{i}^{\prime}$.
return $B$
end procedure

```
```

Algorithm 16.2 Weighted Nonnegative Matrix Factorization
procedure WeightedNonnegativeMatrixFactorization $(A, W, n, r, k)$
$t \leftarrow O\left(k r / \epsilon^{2}\right)$
Guess leverage sampling matrix $D_{l e v}$ and $D_{l e v^{\prime}}$
Create $t$ variables for the nonzero probablities on diagonal of $D_{\text {lev }}$.
Create $t$ variables for the nonzero probablities on diagonal of $D_{l e v^{\prime}}$.
Create $t r$ variables for $r$ rows of $B$ based on the nonzero rows of $D_{\text {lev }}$.
Create $r t$ variables for $r$ columns of $C$ based on the nonzero columns of $D_{l e v^{\prime}}$.
for $i \in[n]$ do
Write down $D_{l e v} D_{W_{i}} B$ and $C D_{W^{i}} D_{l e v^{\prime}}$.
end for
$g \leftarrow$ \# zero-patterns of $\left\{D_{l e v} D_{W_{i}}\right\}_{i \in[n]}$
$g^{\prime} \leftarrow \#$ zero-patterns of $\left\{D_{W^{i}} D_{l e v^{\prime}}\right\}_{i \in[n]}$
for $\ell \in[g], \ell^{\prime} \in\left[g^{\prime}\right]$ do
Guess a rank $s_{\ell} \in[k]$ for zero-pattern $D_{\text {lev }} D_{W_{\ell}^{\prime}}$
Guess a maximal linearly independent set of rows $U_{\ell}$ for $D_{l e v} D_{W_{\ell}^{\prime}} B,\left|U_{\ell}\right|=s_{\ell}$
Guess a rank $s_{\ell^{\prime}} \in[k]$ for zero-pattern $D_{\text {lev }} D_{W_{1}^{\ell^{\prime}}}$
Guess a maximal linearly independent set of columns $V_{\ell^{\prime}}$ for $C D_{W_{\prime}^{\prime \prime}} D_{l e v^{\prime}},\left|U_{\ell}\right|=s_{\ell}$
end for
for $i \in[n], j \in[p]$ do
Write down $\widehat{B}_{j}^{i, \ell}$ in terms of $O(r t)$ variables created early
Write down $\widehat{C}_{j}^{i, \ell^{\prime}}$ in terms of $O(r t)$ variables created early
end for
Write down semi-algebraic set
Run an algorithm for deciding if the semi-algebraic set is non-empty.
end procedure

```
```

Algorithm 16.3 Weighted Nonnegative Matrix Factorization
procedure FasterWeightedNonnegativeMatrixFactorization $(A, W, n, r, k)$
$t \leftarrow O\left(k r / \epsilon^{2}\right)$
Guess leverage sampling matrix $D_{l e v}$ and $D_{l e v^{\prime}}$
Create $t$ variables for the nonzero probablities on diagonal of $D_{l e v}$.
Create $t$ variables for the nonzero probablities on diagonal of $D_{l e v^{\prime}}$.
Create $t r$ variables for $r$ rows of $B$ based on the nonzero rows of $D_{\text {lev }}$.
Create $r t$ variables for $r$ columns of $C$ based on the nonzero columns of $D_{l e v^{\prime}}$.
for $i \in[n]$ do
Write down $D_{l e v} D_{W_{i}} B$ and $C D_{W^{i}} D_{l e v^{\prime}}$.
end for
$g \leftarrow \#$ zero-patterns of $\left\{D_{l e v} D_{W_{i}}\right\}_{i \in[n]}$
$g^{\prime} \leftarrow \#$ zero-patterns of $\left\{D_{W^{i}} D_{l e v^{\prime}}\right\}_{i \in[n]}$
for $\ell \in[g]$ do
for $s_{\ell} \in[k]$ do
$\operatorname{rank} \leftarrow s_{\ell}, p_{s_{\ell}} \leftarrow \#$ size $s_{\ell}$ subset rows of $D_{\text {lev }} D_{W_{\ell}^{\prime}} B$
for $i_{1} \in\left[p_{s_{\ell}}\right]$ do
for $i \in\left[z_{\ell}\right]$ do
for $i_{2} \in\left[\binom{k}{s_{\ell}}\right]$ do
Write down $\widehat{B}_{i_{1}, i_{2}}^{i, \ell}$ in terms of $O(r t)$ variables created early, by
Cramer rule
end for
end for
end for
end for
end for
for $\ell^{\prime} \in\left[g^{\prime}\right]$ do
for $\omega_{\ell^{\prime}} \in[k]$ do
rank $\leftarrow \omega_{\ell^{\prime}}, q_{\omega_{\ell^{\prime}}} \leftarrow \#$ size $\omega_{\ell^{\prime}}$ subset columns of $C D_{W_{1}^{\ell^{\prime}}} D_{l e v^{\prime}}$
for $j_{1} \in\left[q_{\omega_{\ell}}\right]$ do
for $j \in\left[z_{\ell^{\prime}}\right]$ do
for $j_{2} \in\left[\binom{k}{\omega_{\ell^{\prime}}}\right]$ do
Write down $\widehat{C}_{j_{1}, j_{2}}^{j, \ell}$ in terms of $O(r t)$ variables created early, by
Cramer rule
end for
end for
end for
end for
end for1107
Write down semi-algebraic set
Run an algorithm for deciding if the semi-algebraic set is non-empty.
end procedure

```

\section*{Chapter 17}

\section*{Entry-wise L1 Low Rank Approximation}

We study the \(\ell_{1}\)-low rank approximation problem, where for a given \(n \times d\) matrix \(A\) and approximation factor \(\alpha \geq 1\), the goal is to output a rank- \(k\) matrix \(\widehat{A}\) for which
\[
\|A-\widehat{A}\|_{1} \leq \alpha \cdot \min _{\text {rank-k matrices } A^{\prime}}\left\|A-A^{\prime}\right\|_{1}
\]
where for an \(n \times d\) matrix \(C\), we let \(\|C\|_{1}=\sum_{i=1}^{n} \sum_{j=1}^{d}\left|C_{i, j}\right|\). This error measure is known to be more robust than the Frobenius norm in the presence of outliers and is indicated in models where Gaussian assumptions on the noise may not apply. The problem was shown to be NP-hard by Gillis and Vavasis and a number of heuristics have been proposed. It was asked in multiple places if there are any approximation algorithms.

We give the first provable approximation algorithms for \(\ell_{1}\)-low rank approximation, showing that it is possible to achieve approximation factor \(\alpha=(\log d) \cdot \operatorname{poly}(k)\) in \(n n z(A)+(n+d)\) poly \((k)\) time, where \(n n z(A)\) denotes the number of non-zero entries of \(A\). If \(k\) is constant, we further improve the approximation ratio to \(O(1)\) with a poly \((n d)\)-time algorithm. Under the Exponential Time Hypothesis, we show there is no poly(nd)-time algorithm achieving a \(\left(1+\frac{1}{\log ^{1+\gamma}(n d)}\right)\)-approximation, for \(\gamma>0\) an arbitrarily small constant, even when \(k=1\).

\subsection*{17.1 Introduction}

Two well-studied problems in numerical linear algebra are regression and low rank approximation. In regression, one is given an \(n \times d\) matrix \(A\), and an \(n \times 1\) vector \(b\), and one seeks an \(x \in \mathbb{R}^{d}\) which minimizes \(\|A x-b\|\) under some norm. For example, for least squares regression one minimizes \(\|A x-b\|_{2}\). In low rank approximation, one is given an \(n \times d\) matrix \(A\), and one seeks a rank- \(k\) matrix \(\widehat{A}\) which minimizes \(\|A-\widehat{A}\|\) under some norm. For example, in Frobenius norm low rank approximation, one minimizes \(\|A-\widehat{A}\|_{F}=\left(\sum_{i, j}\left(A_{i, j}-\widehat{A}_{i, j}\right)^{2}\right)^{1 / 2}\). Algorithms for regression are often used as subroutines for low rank approximation. Indeed, one of the main insights of [DMM06c, DMM06b, Sar06, DMM08, CW09] was to use results for generalized least squares regression for Frobenius norm low rank approximation. Algorithms for \(\ell_{1}\)-regression, in which one minimizes \(\|A x-b\|_{1}=\) \(\sum_{i}\left|(A x)_{i}-b_{i}\right|\), were also used [BD13, SW11] to fit a set of points to a hyperplane, which is a special case of entrywise \(\ell_{1}\)-low rank approximation, the more general problem being to find a rank- \(k\) matrix \(\widehat{A}\) minimizing \(\sum_{i, j}\left|A_{i, j}-\widehat{A}_{i, j}\right|\).

Randomization and approximation were introduced to significantly speed up algorithms for these problems, resulting in algorithms achieving relative error approximation with high probability. Such algorithms are based on sketching and sampling techniques; we refer to [Woo14b] for a survey. For least squares regression, a sequence of work [Sar06, CW13, MM13, NN13a, LMP13, BDN15, Coh16a] shows how to achieve algorithms running in \(n n z(A)+\operatorname{poly}(d)\) time. For Frobenius norm low rank approximation, using the advances for regression this resulted in \(n n z(A)+(n+d) \operatorname{poly}(k)\) time algorithms. For \(\ell_{1}\)-regression, sketching and sampling-based methods [Cla05, SW11, CDMI \({ }^{+}\)13, CW13, MM13, LMP13, WZ13, CW15b, CP15] led to an \(n n z(A)+\operatorname{poly}(d)\) time algorithm.

Just like Frobenius norm low rank approximation is the analogue of least squares regression, entrywise \(\ell_{1}\)-low rank approximation is the analogue of \(\ell_{1}\)-regression. Despite this analogy, no non-trivial upper bounds with provable guarantees are known for \(\ell_{1}\)-low rank approximation. Unlike Frobenius norm low rank approximation, which can be solved exactly using the singular value decomposition, no such algorithm or closed-form solution is known for \(\ell_{1}\)-low rank approximation. Moreover, the problem was recently shown to be NP-hard [GV15]. A major open question is whether there exist approximation algorithms, sketchingbased or otherwise, for \(\ell_{1}\)-low rank approximation. Indeed, the question of obtaining betters algorithms was posed in section 6 of [GV15], in [Exc13], and as the second part of open question 2 in [Woo14b], among other places. The earlier question of NP-hardness was posed in Section 1.4 of [KV09], for which the question of obtaining approximation algorithms is a natural followup. The goal of our work is to answer this question.

We now formally define the \(\ell_{1}\)-low rank approximation problem: we are given an \(n \times d\) matrix \(A\) and approximation factor \(\alpha \geq 1\), and we would like, with large constant probability, to output a rank- \(k\) matrix \(\widehat{A}\) for which
\[
\|A-\widehat{A}\|_{1} \leq \alpha \cdot \min _{\text {rank-k matrices } A^{\prime}}\left\|A-A^{\prime}\right\|_{1}
\]
where for an \(n \times d\) matrix \(C\), we let \(\|C\|_{1}=\sum_{i=1}^{n} \sum_{j=1}^{d}\left|C_{i, j}\right|\). This notion of low rank approximation has been proposed as a more robust alternative to Frobenius norm low rank approximation \(\left[K K 03, ~ K K 05\right.\), KLC \(^{+} 15\), Kwa08, ZLS \({ }^{+} 12\), BJ12, BD13, BDB13, MXZZ13, MKP13, MKP14, MKCP16, PK16], and is sometimes referred to as \(\ell_{1}\)-matrix factorization or robust PCA. \(\ell_{1}\)-low rank approximation gives improved results over Frobenius norm low rank approximation since outliers are less exaggerated, as one does not square their contribution
in the objective. The outlier values are often erroneous values that are far away from the nominal data, appear only a few times in the data matrix, and would not appear again under normal system operation. These works also argue \(\ell_{1}\)-low rank approximation can better handle missing data, is appropriate in noise models for which the noise is not Gaussian, e.g., it produces the maximum likelihood estimator for Laplacian noise [Gao08, \(\mathrm{KAC}^{+}\)08, VT01], and can be used in image processing to prevent image occlusion [YZD12].

To see that \(\ell_{1}\)-low rank approximation and Frobenius norm low rank approximation can give very different results, consider the \(n \times n\) matrix \(A=\left[\begin{array}{ll}n & 0 \\ 0 & B\end{array}\right]\), where \(B\) is any ( \(n-\) 1) \(\times(n-1)\) matrix with \(\|B\|_{F}<n\). The best rank-1 approximation with Frobenius norm error is given by \(\widehat{A}=n \cdot e_{1} e_{1}^{\top}\), where \(e_{1}\) is the first standard unit vector. Here \(\widehat{A}\) ignores all but the first row and column of \(A\), which may be undesirable in the case that this row and column represent an outlier. Note \(\|A-\widehat{A}\|_{1}=\|B\|_{1}\). If, for example, \(B\) is the all 1 s matrix, then \(\widehat{A}=[0,0 ; 0, B]\) is a rank- 1 approximation for which \(\|A-\widehat{A}\|_{1}=n\), and therefore this solution is a much better solution to the \(\ell_{1}\)-low rank approximation problem than \(n \cdot e_{1} e_{1}^{\top}\), for which \(\left\|A-n \cdot e_{1} e_{1}^{\top}\right\|_{1}=(n-1)^{2}\).

Despite the advantages of \(\ell_{1}\)-low rank approximation, its main disadvantage is its computationally intractability. It is not rotationally invariant and most tools for Frobenius low rank approximation do not apply. To the best of our knowledge, all previous works only provide heuristics. Using that for an \(n \times d\) matrix \(C,\|C\|_{F} \leq\|C\|_{1} \leq \sqrt{n d}\|C\|_{F}\), a Frobenius norm low rank approximation gives a \(\sqrt{n d}\) approximation for \(\ell_{1}\)-low rank approximation. A bit better is to use algorithms for low rank approximation with respect to the sum of distances, i.e., to find a rank- \(k\) matrix \(\widehat{A}\) minimizing \(\|A-\widehat{A}\|_{1,2}\), where for an \(n \times d\) matrix \(C,\|C\|_{1,2}=\sum_{i=1}^{n}\left\|C_{i}\right\|_{2}\), where \(C_{i}\) is the \(i\)-th row of \(C\). A sequence of work
[DV07, FMSW10, FL11, SV12a, CW15a] shows how to obtain an \(O(1)\)-approximation to this problem in \(\operatorname{nnz}(A)+(n+d)\) poly \((k)+\exp (k)\) time, and using that \(\|C\|_{1,2} \leq\|C\|_{1} \leq \sqrt{d}\|C\|_{1,2}\) results in an \(O(\sqrt{d})\)-approximation.

\subsection*{17.1.1 Our Results}

We give the first efficient algorithms for \(\ell_{1}\)-low rank approximation with provable approximation guarantees. By symmetry of the problem, we can assume \(d \leq n\). We first give an algorithm which runs in \(O(\operatorname{nnz}(A))+n \cdot \operatorname{poly}(k)\) time and solves the \(\ell_{1}\)-low rank approximation problem with approximation factor \((\log d) \cdot \operatorname{poly}(k)\). This is an exponential improvement over the previous approximation factor of \(O(\sqrt{d})\), provided \(k\) is not too large, and is polynomial time for every \(k\). Moreover, \(\operatorname{provided~} n n z(A) \geq n \cdot \operatorname{poly}(k)\), our time is optimal up to a constant factor as any relative error algorithm must spend nnz \((A)\) time. We also give a hard instance for our algorithm ruling out \(\frac{\log d}{k \log k}+k^{1 / 2-\gamma}\) approximation for arbitrarily small constant \(\gamma>0\), and hard instances for a general class of algorithms based on linear sketches, ruling out \(k^{1 / 2-\gamma}\) approximation.

Via a different algorithm, we show how to achieve an \(\widetilde{O}(k)\)-approximation factor in \(\operatorname{poly}(n) d^{\widetilde{O}(k)} 2^{\widetilde{O}\left(k^{2}\right)}\) time. This is useful for constant \(k\), for which it gives an \(O(1)\)-approximation in poly \((n)\) time, improving the \(O(\log d)\)-approximation for constant \(k\) of our earlier algorithm. The approximation ratio of this algorithm, although \(O(1)\) for constant \(k\), depends on \(k\). We also show one can find a rank- \(2 k\) matrix \(\widehat{A}\) in \(\operatorname{poly}(n)\) time for constant \(k\) for which \(\|A-\widehat{A}\|_{1} \leq C \min _{\text {rank-k matrices } A^{\prime}}\left\|A-A^{\prime}\right\|_{1}\), where \(C>1\) is an absolute constant independent of \(k\). We refer to this as a bicriteria algorithm. Finally, one can output a rank- \(k\) matrix \(\widehat{A}\), instead of a rank- \(2 k\) matrix \(\widehat{A}\), in \(\operatorname{poly}(n)\) time with the same absolute constant \(C\) approximation factor, under an additional assumption that the entries of \(\widehat{A}\) are integers in the range \(\{-b,-b+1, \ldots, b\}\) for an integer \(b \leq \operatorname{poly}(n)\). Unlike our previous algorithms, this very last algorithm has a bit complexity assumption, and runs in poly \((b)\) time instead of poly \((\log (b))\) time.

We also give a number of results for variants of \(\ell_{1}\)-low rank approximation which are studied for Frobenius norm low rank approximation; prior to our work nothing was known about these problems.

Column Subset Selection and CUR Decomposition: In the column subset selection problem, one seeks a small subset \(C\) of columns of \(A\) for which there is a matrix \(X\) for which \(\|C X-A\|\) is small, under some norm. The matrix \(C X\) provides a low rank approximation to \(A\) which is often more interpretable, since it stores actual columns of \(A\), preserves sparsity, etc. These have been extensively studied when the norm is the Frobenius or operator norm (see, e.g., [BMD09, DR10, BDM11] and the references therein). We initiate the study of this problem with respect to the \(\ell_{1}\)-norm. We first prove an existence result, namely, that there exist matrices \(A\) for which any subset \(C\) of \(\operatorname{poly}(k)\) columns satisfies \(\min _{X}\|C X-A\|_{1} \geq k^{1 / 2-\gamma} \cdot \min _{\text {rank-k matrices } A^{\prime}}\left\|A-A^{\prime}\right\|_{1}\), where \(\gamma>0\) is an arbitrarily small constant. This result is in stark contrast to the Frobenius norm for which for every matrix there exist \(O\left(\frac{k}{\epsilon}\right)\) columns for which the approximation factor is \(1+\epsilon\). We also show that our bound is nearly optimal in this regime, by showing for every matrix there exists a subset of \(O(k \log k)\) columns providing an \(O(\sqrt{k \log k})\)-approximation. One can find such columns in \(\operatorname{poly}(n) d^{O(k \log k)}\) time by enumerating and evaluating the cost of each subset. Although this is exponential in \(k\), we show it is possible to find \(O(k \log k)\) columns providing an \(O(\sqrt{k \log k} \log d)\)-approximation in polynomial time for every \(k\).

We extend these results to the CUR decomposition problem (see, e.g., [DMM08, BW14]), in which one seeks a factorization \(C U R\) for which \(C\) is a subset of columns of \(A, R\) is a subset of rows of \(A\), and \(\|C U R-A\|\) is as small as possible. In the case of Frobenius norm, one can choose \(O(k / \epsilon)\) columns and rows, have \(\operatorname{rank}(U)=k\), have \(\|C U R-A\|_{F}\) be at most
\((1+\epsilon)\) times the optimal cost, and find the factorization in \(n n z(A) \log n+n \cdot \operatorname{poly}((\log n) k / \epsilon)\) time [BW14]. Using our column subset selection results, we give an nnz \((A)+n \cdot \operatorname{poly}(k)\) time algorithm choosing \(O(k \log k)\) columns and rows, for which \(\operatorname{rank}(U)=k\), and for which \(\|C U R-A\|_{1}\) is poly \((k) \log d\) times the cost of any rank- \(k\) approximation to \(A\).

\subsection*{17.1.2 Technical Overview}

Initial Algorithm and Optimizations: Let \(A^{*}\) be a rank- \(k\) matrix for which \(\left\|A-A^{*}\right\|_{1}=\min _{\text {rank-k matrices } A^{\prime}}\left\|A-A^{\prime}\right\|_{1}\). Let \(A^{*}=U^{*} V^{*}\) be a factorization for which \(U^{*}\) is \(n \times k\) and \(V^{*}\) is \(k \times d\). Suppose we somehow knew \(U^{*}\) and consider the multi-response \(\ell_{1}\)-regression problem \(\min _{V}\left\|U^{*} V-A\right\|_{1}=\min _{V} \sum_{i=1}^{d}\left\|U^{*} V_{i}-A_{i}\right\|_{1}\), where \(V_{i}, A_{i}\) denote the \(i\)-th columns of \(V\) and \(A\), respectively. We could solve this with linear programming though this is not helpful for our argument here.

Instead, inspired by recent advances in sketching for linear algebra (see, e.g., [Woo14b] for a survey), we could choose a random matrix \(S\) and solve \(\min _{V}\left\|S U^{*} V-S A\right\|_{1}=\) \(\min _{V} \sum_{i=1}^{d}\left\|\left(S U^{*}\right) V_{i}-S A_{i}\right\|_{1}\). If \(V\) is an approximate minimizer of the latter problem, we could hope \(V\) is an approximate minimizer of the former problem. If also \(S\) has a small number \(t\) of rows, then we could instead solve \(\min _{V} \sum_{i=1}^{d}\left\|\left(S U^{*}\right) V_{i}-S A_{i}\right\|_{2}\), that is, minimize the sum of Euclidean norms rather than the sum of \(\ell_{1}\)-norms. Since \(t^{-1 / 2}\left\|\left(S U^{*}\right) V_{i}-S A_{i}\right\|_{1} \leq\) \(\left\|\left(S U^{*}\right) V_{i}-S A_{i}\right\|_{2} \leq\left\|\left(S U^{*}\right) V_{i}-S A_{i}\right\|_{1}\), we would obtain a \(\sqrt{t}\)-approximation to the problem \(\min _{V}\left\|S U^{*} V-S A\right\|_{1}\). A crucial observation is that the solution to \(\min _{V} \sum_{i=1}^{d} \|\left(S U^{*}\right) V_{i}-\) \(S A_{i} \|_{2}\) is given by \(V=\left(S U^{*}\right)^{\dagger} S A\), which implies that \(V\) is in the row span of \(S A\). If also \(S\) were oblivious to \(U^{*}\), then we could compute \(S A\) without ever knowing \(U^{*}\). Having a low-dimensional space containing a good solution in its span is our starting point.

For this to work, we need a distribution on oblivious matrices \(S\) with a small number of rows, for which an approximate minimizer \(V\) to \(\min _{V}\left\|S U^{*} V-S A\right\|_{1}\) is also an approximate minimizer to \(\min _{V}\left\|U^{*} V-A\right\|_{1}\). It is unknown if there exists a distribution on \(S\) with this property. What is known is that if \(S\) has \(O(d \log d)\) rows, then the Lewis weights (see, e.g., [CP15] and references therein) of the concatenated matrix \(\left[U^{*}, A\right]\) give a distribution for
which the optimal \(V\) for the latter problem is a \((1+\epsilon)\)-approximation to the former problem; see also earlier work on \(\ell_{1}\)-leverage scores [Cla05, \(\left.\mathrm{DDH}^{+} 09\right]\) which have poly \((d)\) rows and the same \((1+\epsilon)\)-approximation guarantee. Such distributions are not helpful here as (1) they are not oblivious, and (2) the number \(O(d \log d)\) of rows gives an \(O(\sqrt{d \log d})\) approximation factor, which is much larger than what we want.

There are a few oblivious distributions \(S\) which are useful for single-response \(\ell_{1-}{ }^{-}\) regression \(\min \left\|U^{*} v-a\right\|_{1}\) for column vectors \(v, a \in \mathbb{R}^{k}\) [SW11, CDMI \({ }^{+} 13\), WZ13]. In particular, if \(S\) is an \(O(k \log k) \times n\) matrix of i.i.d. Cauchy random variables, then the solution \(v\) to \(\min \left\|S U^{*} v-S a\right\|_{1}\) is an \(O(k \log k)\)-approximation to min \(\left\|U^{*} v-a\right\|_{1}\) [SW11]. The important property of Cauchy random variables is that if \(X\) and \(Y\) are independent Cauchy random variables, then \(\alpha X+\beta Y\) is distributed as a Cauchy random variable times \(|\alpha|+|\beta|\), for any scalars \(\alpha, \beta \in \mathbb{R}\). The \(O(k \log k)\) approximation arises because all possible regression solutions are in the column span of \(\left[U^{*}, a\right]\) which is \((k+1)\)-dimensional, and the sketch \(S\) gives an approximation factor of \(O(k \log k)\) to preserve every vector norm in this subspace. If we instead had a multi-response regression problem \(\min \left\|S U^{*} V^{*}-S A\right\|_{1}\) the dimension of the column span of \(\left[U^{*}, A\right]\) would be \(d+k\), and this approach would give an \(O(d \log d)\) approximation. Unlike Frobenius norm multi-response regression min \(\left\|S U^{*} V^{*}-S A\right\|_{F}\), which can be bounded if \(S\) is a subspace embedding for \(U^{*}\) and satisfies an approximate matrix product theorem [Sar06], there is no convenient linear-algebraic analogue for the \(\ell_{1}\)-norm.

We first note that since regression is a minimization problem, to obtain an \(O(\alpha)\) approximation by solving the sketched version of the problem, it suffices that (1) for the optimal \(V^{*}\), we have \(\left\|S U^{*} V^{*}-S A\right\|_{1} \leq O(\alpha)\left\|U^{*} V^{*}-A\right\|_{1}\), and (2) for all \(V\), we have \(\left\|S U^{*} V-S A\right\|_{1} \geq \Omega(1) \cdot\left\|U^{*} V-A\right\|_{1}\).

We show (1) holds for \(\alpha=O(\log d)\) and any number of rows of \(S\). Our analysis follows by truncating the Cauchy random variables \(\left(S U^{*} V_{j}^{*}-S A_{j}\right)_{i}\) for \(i \in[O(k \log k)]\) and \(j \in[d]\), so that their expectation exists, and applying linearity of expectation across the \(d\) columns. This is inspired from an argument of Indyk [Ind06] for embedding a vector into a lower-dimensional vector while preserving its \(\ell_{1}\)-norm; for single-response regression this is the statement that \(\left\|S U^{*} v^{*}-S a\right\|_{1}=\Theta(1)\left\|U^{*} v-a\right\|_{1}\), implied by [Ind06]. However, for multi-response regression we have to work entirely with expectations, rather than the tail bounds in [Ind06], since the Cauchy random variables \(\left(S U^{*} V_{j}-S A_{j}\right)_{i}\), while independent across \(i\), are dependent across \(j\). Moreover, our \(O(\log d)\)-approximation factor is not an artifact of our analysis - we show in Section 17.6 that there is an \(n \times d\) input matrix \(A\) for which with probability \(1-1 / \operatorname{poly}(k)\), there is no \(k\)-dimensional space in the span of \(S A\) achieving a \(\left(\frac{\log d}{t \log t}+k^{1 / 2-\gamma}\right)\)-approximation, for \(S\) a Cauchy matrix with \(t\) rows, where \(\gamma>0\) is an arbitrarily small constant. This shows \((k \log d)^{\Omega(1)}\)-inapproximability. Thus, the fact that we achieve \(O(\log d)\)-approximation instead of \(O(1)\) is fundamental for a matrix \(S\) of Cauchy random variables or any scaling of it.

While we cannot show (2), we instead show for all \(V,\left\|S U^{*} V-S A\right\|_{1} \geq \| U^{*} V-\) \(A\left\|_{1} / 2-O(\log d)\right\| U^{*} V^{*}-A \|_{1}\) if \(S\) has \(O(k \log k)\) rows. This suffices for regression, since the only matrices \(V\) for which the cost is much smaller in the sketch space are those providing an \(O(\log d)\) approximation in the original space. The guarantee follows from the triangle inequality: \(\left\|S U^{*} V-S A\right\|_{1} \geq\left\|S U^{*} V-S U^{*} V^{*}\right\|_{1}-\left\|S U^{*} V^{*}-S A\right\|_{1}\) and the fact that \(S\) is known to not contract any vector in the column span of \(U^{*}\) if \(S\) has \(O(k \log k)\) rows [SW11]. Because of this, we have \(\left\|S U^{*} V-S U^{*} V^{*}\right\|_{1}=\Omega(1)\left\|U^{*} V-U^{*} V^{*}\right\|_{1}=\Omega(1)\left(\left\|U^{*} V-A\right\|_{1}-\right.\) \(\left\|U^{*} V^{*}-A\right\|_{1}\) ), where we again use the triangle inequality. We also bound the additive term
\(\left\|S U^{*} V^{*}-S A\right\|_{1}\) by \(O(\log d)\left\|U^{*} V^{*}-A\right\|_{1}\) using (1) above.
Given that \(S A\) contains a good rank- \(k\) approximation in its row span, our algorithm with a slightly worse \(\operatorname{poly}(n)\) time and poly \((k \log (n))\)-approximation can be completely described here. Let \(S\) and \(T_{1}\) be independent \(O(k \log k) \times n\) matrices of i.i.d. Cauchy random variables, and let \(R\) and \(T_{2}\) be independent \(d \times O(k \log k)\) matrices of i.i.d. Cauchy random variables. Let
\[
X=\left(T_{1} A R\right)^{\dagger}\left(\left(T_{1} A R\right)\left(T_{1} A R\right)^{\dagger}\left(T_{1} A T_{2}\right)\left(S A T_{2}\right)\left(S A T_{2}\right)^{\dagger}\right)_{k}\left(S A T_{2}\right)^{\dagger}
\]
which is the rank- \(k\) matrix minimizing \(\left\|T_{1} A R X S A T_{2}-T_{1} A T_{2}\right\|_{F}\), where for a matrix \(C\), \(C_{k}\) is its best rank-k approximation in Frobenius norm. Output \(\widehat{A}=A R X S A\) as the solution to \(\ell_{1}\)-low rank approximation of \(A\). We show with constant probability that \(\widehat{A}\) is a poly \((k \log (n))\)-approximation.

To improve the approximation factor, after computing \(S A\), we \(\ell_{1}\)-project each of the rows of \(A\) onto \(S A\) using linear programming or fast algorithms for \(\ell_{1}\)-regression [CW13, MM13], obtaining an \(n \times d\) matrix \(B\) of \(\operatorname{rank} O(k \log k)\). We then apply the algorithm in the previous paragraph with \(A\) replaced by \(B\). This ultimately leads to a \(\log d \cdot \operatorname{poly}(k)\) approximation.

To improve the running time from \(\operatorname{poly}(n)\) to \(\operatorname{nnz}(A)+n \cdot \operatorname{poly}(k)\), we show a similar analysis holds for the sparse Cauchy matrices of [MM13]; see also the matrices in [WZ13].

CUR Decompositions: To obtain a CUR decomposition, we first find a \(\log d \cdot \operatorname{poly}(k)\) approximate rank- \(k\) approximation \(\widehat{A}\) as above. Let \(B_{1}\) be an \(n \times k\) matrix whose columns span those of \(\widehat{A}\), and consider the regression \(\min _{V}\left\|B_{1} V-A\right\|_{1}\). Unlike the problem \(\min _{V} \| U^{*} V-\)
\(A \|_{1}\) where \(U^{*}\) was unknown, we know \(B_{1}\) so can compute its Lewis weights efficiently, sample by them, and obtain a regression problem \(\min _{V}\left\|D_{1}\left(B_{1} V-A\right)\right\|_{1}\) where \(D_{1}\) is a sampling and rescaling matrix. Since
\[
\left\|D_{1}\left(B_{1} V-A\right)\right\|_{1} \leq\left\|D_{1}\left(B_{1} V-B_{1} V^{*}\right)\right\|_{1}+\left\|D_{1}\left(B_{1} V^{*}-A\right)\right\|_{1},
\]
where \(V^{*}=\operatorname{argmin}_{V}\left\|B_{1} V-A\right\|_{1}\), we can bound the first term by \(O\left(\left\|B_{1} V-B_{1} V^{*}\right\|_{1}\right)\) using that \(D_{1}\) is a subspace embedding if it has \(O(k \log k)\) rows, while the second term is \(O(1)\left\|B_{1} V^{*}-A\right\|_{1}\) by a Markov bound. Note that
\[
\left\|B_{1} V^{*}-A\right\|_{1} \leq(\log d) \cdot \operatorname{poly}(k) \min _{\text {rank-k matrices } A^{\prime}}\left\|A-A^{\prime}\right\|_{1}
\]

By switching to \(\ell_{2}\) as before, we see that \(\widehat{V}=\left(D_{1} B_{1}\right)^{\dagger} D_{1} A\) contains a \((\log d) \operatorname{poly}(k)\) approximation in its span. Here \(D_{1} A\) is an actual subset of rows of \(A\), as required in a CUR decomposition. Moreover the subset size is \(O(k \log k)\). We can sample by the Lewis weights of \(\widehat{V}\) to obtain a subset \(C\) of \(O(k \log k)\) rescaled columns of \(A\), together with a rank- \(k\) matrix \(U\) for which
\[
\|C U R-A\|_{1} \leq(\log d) \operatorname{poly}(k) \min _{\text {rank-k matrices } A^{\prime}}\left\|A-A^{\prime}\right\|_{1} .
\]

Algorithm for Small \(k\) : Our CUR decomposition shows how we might obtain an \(O(1)\) approximation for constant \(k\) in \(\operatorname{poly}(n)\) time. If we knew the Lewis weights of \(U^{*}\), an \(\alpha\) approximate solution to the problem \(\min _{V}\left\|D_{1}\left(U^{*} V-A\right)\right\|_{1}\) would be an \(O(\alpha)\)-approximate solution to the problem \(\min _{V}\left\|U^{*} V-A\right\|_{1}\), where \(D_{1}\) is a sampling and rescaling matrix of \(O(k \log k)\) rows of \(A\). Moreover, an \(O(\sqrt{k \log k})\)-approximate solution to \(\min _{V} \| D_{1}\left(U^{*} V-\right.\)
\(A) \|_{1}\) is given by \(V=\left(D_{1} U^{*}\right)^{\dagger} D_{1} A\), which implies the \(O(k \log k)\) rows of \(D_{1} A\) contain an \(O(\sqrt{k \log k})\)-approximation. For small \(k\), we can guess every subset of \(O(k \log k)\) rows of \(A\) in \(n^{O(k \log k)}\) time (if \(d \ll n\), by taking transposes at the beginning one can replace this with \(d^{O(k \log k)}\) time). For each guess, we set up the problem \(\min _{\text {rank-k } U}\left\|U\left(D_{1} A\right)-A\right\|_{1}\). If \(D_{2}\) is a sampling and rescaling matrix according to the Lewis weights of \(D_{1} A\), then by a similar triangle inequality argument as for our CUR decomposition, minimizing \(\left\|U\left(D_{1} A\right) D_{2}-A D_{2}\right\|_{1}\) gives an \(O(\sqrt{k \log k})\) approximation. By switching to \(\ell_{2}\), this implies there is an \(O(k \log k)\) approximation of the form \(A D_{2} W D_{1} A\), where \(W\) is an \(O\left(k \log ^{2} k\right) \times O(k \log k)\) matrix of rank \(k\). By setting up the problem \(\min _{\text {rank-k } W}\left\|A D_{2} W D_{1} A-A\right\|_{1}\), one can sample from Lewis weights on the left and right to reduce this to a problem independent of \(n\) and \(d\), after which one can use polynomial optimization to solve it in \(\exp (\operatorname{poly}(k))\) time. One of our guesses \(D_{1} A\) will be correct, and for this guess we obtain an \(\widetilde{O}(k)\)-approximation. For each guess we can compute its cost and take the best one found. This gives an \(O(1)\)-approximation for constant \(k\), removing the \(O(\log d)\)-factor from the approximation of our earlier algorithm.

Existential Results for Subset Selection: In our algorithm for small \(k\), the first step was to show there exist \(O(k \log k)\) rows of \(A\) which contain a rank- \(k\) space which is an \(O(\sqrt{k \log k})\)-approximation.

While for Frobenius norm one can find \(O(k)\) rows with an \(O(1)\)-approximation in their span, one of our main negative results for \(\ell_{1}\)-low rank approximation is that this is impossible, showing that the best approximation one can obtain with poly \((k)\) rows is \(k^{1 / 2-\gamma}\) for an arbitrarily small constant \(\gamma>0\). Our hard instance is an \(r \times(r+k)\) matrix \(A\) in which the first \(k\) columns are i.i.d. Gaussian, and the remaining \(r\) columns are an identity
matrix. Here, \(r\) can be twice the number of rows one is choosing. The optimal \(\ell_{1}\)-low rank approximation has cost at most \(r\), obtained by choosing the first \(k\) columns.

Let \(R \in \mathbb{R}^{r / 2 \times k}\) denote the first \(k\) entries of the \(r / 2\) chosen rows, and let \(y\) denote the first \(k\) entries of an unchosen row. For \(r / 2>k\), there exist many solutions \(x \in \mathbb{R}^{r / 2}\) for which \(x^{\top} R=y\). However, we can show the following tradeoff:
\[
\text { whenever }\left\|x^{\top} R-y\right\|_{1}<\frac{\sqrt{k}}{\operatorname{poly}(\log k)} \text {, then }\|x\|_{1}>\frac{\sqrt{k}}{\operatorname{poly}(\log k)} \text {. }
\]

Then no matter which linear combination \(x^{\top}\) of the rows of \(R\) one chooses to approximate \(y\) by, either one incurs a \(\frac{\sqrt{k}}{\operatorname{poly}(\log k)}\) cost on the first \(k\) coordinates, or since \(A\) contains an identity matrix, one incurs cost \(\|x\|_{1}>\frac{\sqrt{k}}{\text { poly }(\log k)}\) on the last \(r\) coordinates of \(x^{\top} R\).

To show the tradeoff, consider an \(x \in \mathbb{R}^{r / 2}\). We decompose \(x=x^{0}+\sum_{j \geq 1} x^{j}\), where \(x^{j}\) agrees with \(x\) on coordinates which have absolute value in the range \(\frac{1}{\sqrt{k} \log ^{c} k} \cdot\left[2^{-j}, 2^{-j+1}\right]\), and is zero otherwise. Here, \(c>0\) is a constant, and \(x^{0}\) denotes the restriction of \(x\) to all coordinates of absolute value at least \(\frac{1}{\sqrt{k} \log ^{c} k}\). Then \(\|x\|_{1}<\frac{\sqrt{k}}{\log ^{c} k}\), as otherwise we are done. Hence, \(x^{0}\) has small support. Thus, one can build a small net for all \(x^{0}\) vectors by choosing the support, then placing a net on it. For \(x^{j}\) for \(j>0\), the support sizes are increasing so the net size needed for all \(x^{j}\) vectors is larger. However, since \(x^{j}\) has all coordinates of roughly the same magnitude on its support, its \(\ell_{2}\)-norm is decreasing in \(j\). Since \(\left(x^{j}\right)^{\top} R \sim N\left(0,\left\|x^{j}\right\|_{2}^{2} I_{k}\right)\), this makes it much less likely that individual coordinates of \(\left(x^{j}\right)^{\top} R\) can be large. Since this probability goes down rapidly, we can afford to union bound over the larger net size. What we show is that for any sum of the form \(\sum_{j \geq 1} x^{j}\), at most \(\frac{k}{10}\) of its coordinates are at least \(\frac{1}{\log k}\) in magnitude.

For \(\left\|x^{\top} R-y\right\|_{1}\) to be at most \(\frac{\sqrt{k}}{\log ^{c} k}\), for at least \(\frac{k}{2}\) coordinates \(i\), we must have \(\left|\left(x^{\top} R-y\right)_{i}\right|<\frac{2}{\sqrt{k} \log ^{c} k}\). With probability \(1-2^{-\Omega(k)},\left|y_{i}\right| \geq \frac{1}{100}\) on at least \(\frac{2 k}{3}\) coordinates. From the previous paragraph, it follows there are at least \(\frac{k}{2}-\frac{k}{10}-\frac{k}{3}=\Omega(k)\) coordinates \(i\) of \(x\) for which (1) \(\left|\left(x^{\top} R-y\right)_{i}\right|<\frac{2}{\sqrt{k} \log ^{c} k}\), (2) \(\left|\sum_{j \geq 1} x_{i}^{j}\right|<\frac{1}{\log k}\), and (3) \(\left|y_{i}\right| \geq \frac{1}{100}\). On these \(i,\left(x^{0}\right)^{\top} R_{i}\) must be in an interval of width \(\frac{1}{\log k}\) at distance at least \(\frac{1}{100}\) from the origin. Since \(\left(x^{0}\right)^{\top} R \sim N\left(0,\left\|x^{0}\right\|_{2}^{2} I_{k}\right)\), for any value of \(\left\|x^{0}\right\|_{2}^{2}\) the probability this happens on \(\Omega(k)\) coordinates is at most \(2^{-\Theta(k)}\). Since the net size for \(x^{0}\) is small, we can union bound over every sequence \(x^{0}, x^{1}, \ldots\), coming from our nets.

Some care is needed to union bound over all possible subsets \(R\) of rows which can be chosen. We handle this by conditioning on a few events of \(A\) itself, which imply corresponding events for every subset of rows. These events are such that if \(R\) is the chosen set of half the rows, and \(S\) the remaining set of rows of \(A\), then the event that a constant fraction of rows in \(S\) are close to the row span of \(R\) is \(2^{-\Theta(k r)}\), which is small enough to union bound over all choices of \(R\).

Curiously, we also show there are some matrices \(A \in \mathbb{R}^{n \times d}\) for which any \(\ell_{1}\) rank- \(k\) approximation in the entire row span of \(A\) cannot achieve better than a \((2-\Theta(1 / d))\) approximation.

Bicriteria Algorithm: Our algorithm for small \(k\) gives an \(O(1)\)-approximation in poly \((n)\) time for constant \(k\), but the approximation factor depends on \(k\). We show how one can find a rank- \(2 k\) matrix \(\widehat{A}\) for which \(\|A-\widehat{A}\|_{1} \leq C\). OPT, where \(C\) is an absolute constant, and \(\mathrm{OPT}=\min _{\text {rank-k matrices } A^{\prime}}\left\|A-A^{\prime}\right\|_{1}\). We first find a rank- \(k\) matrix \(B^{1}\) for which \(\left\|A-B^{1}\right\|_{1} \leq p \cdot\) OPT for a factor \(1 \leq p \leq \operatorname{poly}(n)\). We can use any of our algorithms above
for this.
Next consider the problem \(\min _{V \in \mathbb{R}^{2 k \times d}}\left\|U^{*} V-\left(A-B^{1}\right)\right\|_{1}\), and let \(U^{*} V^{*}\) be a best \(\ell_{1}\)-low rank- \(2 k\) approximation to \(A-B^{1}\); we later explain why we look at this problem. We can assume \(V^{*}\) is an \(\ell_{1}\) well-conditioned basis [Cla05, \(\left.\mathrm{DDH}^{+} 09\right]\), since we can replace \(U^{*}\) with \(U^{*} R^{-1}\) and \(V^{*}\) with \(R V^{*}\) for any invertible linear transformation \(R\). For any vector \(x\) we then have \(\frac{\|x\|_{1}}{f} \leq\left\|x^{\top} V^{*}\right\|_{1} \leq e\|x\|_{1}\), where \(1 \leq e, f \leq \operatorname{poly}(k)\). This implies all entries of \(U^{*}\) are at most \(2 f\|A-B\|_{1}\), as otherwise one could replace \(U^{*}\) with \(0^{n \times 2 k}\) and reduce the cost. Also, any entry of \(U^{*}\) smaller than \(\frac{\|A-B\|_{1}}{100 \text { enkp }}\) can be replaced with 0 as this incurs additive error \(\frac{\mathrm{OPT}}{100}\). If we round the entries of \(U^{*}\) to integer multiples of \(\frac{\|A-B\|_{1}}{100 \text { enkp }}\), then we only have \(O(e n k p f)\) possibilities for each entry of \(U^{*}\), and still obtain an \(O(1)\)-approximation. We refer to the rounded \(U^{*}\) as \(U^{*}\), abusing notation.

Let \(D\) be a sampling and rescaling matrix with \(O(k \log k)\) non-zero diagonal entries, corresponding to sampling by the Lewis weights of \(U^{*}\). We do not know \(D\), but handle this below. By the triangle inequality, for any \(V\),
\[
\begin{aligned}
\left\|D\left(U^{*} V-\left(A-B^{1}\right)\right)\right\|_{1} & =\left\|D\left(U^{*} V-U^{*} V^{*}\right)\right\|_{1} \pm\left\|D\left(U^{*} V^{*}-\left(A-B^{1}\right)\right)\right\|_{1} \\
& =\Theta(1)\left\|U^{*} V-U^{*} V^{*}\right\|_{1} \pm O(1)\left\|U^{*} V^{*}-\left(A-B^{1}\right)\right\|_{1}
\end{aligned}
\]
where the Lewis weights give \(\left\|D\left(U^{*} V-U^{*} V^{*}\right)\right\|_{1}=\Theta(1)\left\|U^{*} V-U^{*} V^{*}\right\|_{1}\) and a Markov bound gives \(\left\|D\left(U^{*} V^{*}-\left(A-B^{1}\right)\right)\right\|_{1}=O(1)\left\|U^{*} V^{*}-\left(A-B^{1}\right)\right\|_{1}\). Thus, minimizing \(\| D U^{*} V-\) \(D\left(A-B^{1}\right) \|_{1}\) gives a fixed constant factor approximation to the problem \(\min _{V \in \mathbb{R}^{2 k \times d}} \| U^{*} V-\) \(\left(A-B^{1}\right) \|_{1}\). The non-zero diagonal entries of \(D\) can be assumed to be integers between 1 and \(n^{2}\).

We guess the entries of \(D U^{*}\) and note for each entry there are only \(O\left(\operatorname{enkpf} \log \left(n^{2}\right)\right)\)
possibilities. One of our guesses corresponds to Lewis weight sampling by \(U^{*}\). We solve for \(V\) and by the guarantees of Lewis weights, the row span of this \(V\) provides an \(O(1)\) approximation. We can find the corresponding \(U\) via linear programming. As mentioned above, we do not know \(D\), but can enumerate over all \(D\) and all possible \(D U^{*}\). The total time is \(n^{\text {poly (k) }}\).

After finding \(U\), which has \(2 k\) columns, we output the rank- \(3 k\) space formed by the column span of \(\left[U, B^{1}\right]\). By including the column span of \(B^{1}\), we ensure our original transformation of the problem \(\min _{V \in \mathbb{R}^{k \times d}}\left\|U^{*} \cdot V-A\right\|_{1}\) to the problem \(\min _{V \in \mathbb{R}^{2 k \times d}}\left\|U^{*} \cdot V-\left(A-B^{1}\right)\right\|_{1}\) is valid, since we can first use the column span of \(B^{1}\) to replace \(A\) with \(A-B^{1}\). Replacing \(A\) with \(A-B^{1}\) ultimately results in a rank- \(3 k\) output. Had we used \(A\) instead of \(A-B^{1}\) our output would have been rank \(k\) but would have additive error \(\frac{\|A\|_{1}}{\operatorname{poly}(k / \epsilon)}\). If we assume the entries of \(A\) are in \(\{-b,-b+1, \ldots, b\}\), then we can lower bound the cost \(\left\|U^{*} V-A\right\|_{1}\), given that it is non-zero, by \((n d b)^{-O(k)}\) (if it is zero then we output \(A\) ) using Lemma 4.1 in [CW09] and relating entrywise \(\ell_{1}\)-norm to Frobenius norm. We can go through the same arguments above with \(A-B\) replaced by \(A\) and our running time will now be \((n d b)^{\operatorname{poly}(k)}\).

Hard Instances for Cauchy Matrices and More General Sketches: We consider a \(d \times d\) matrix \(A=I_{d}+(\log d) e_{1}^{\top} e\), where \(e_{1}=(1,0, \ldots, 0)\) and \(e=(1,1, \ldots, 1)\) and \(I_{d}\) is the \(d \times d\) identity. For an \(O(k \log k) \times d\) matrix \(S\) of i.i.d. Cauchy random variables, \(S A=S+(\log d) S_{1}^{\top} e\), where \(S_{1}\) is the first column of \(S\). For a typical column of \(S A\), all entries are at most \(\operatorname{poly}(k) \log d\) in magnitude. Thus, in order to approximate the first row of \(A\), which is \((\log d) e\), by \(x^{\top} S A\) for an \(x \in \mathbb{R}^{k \log k}\), we need \(\|x\|_{1} \geq \frac{1}{\operatorname{poly}(k)}\). Also \(\left\|x^{\top} S\right\|_{1}=\Omega\left(\|x\|_{1} d \log d\right)\) with \(1-\exp (-k \log k)\) probability, for \(d\) large enough, so by a net
argument \(\|x\|_{1} \leq \operatorname{poly}(k)\) for all \(x\).

However, there are entries of \(S A\) that are very large, i.e., about one which is \(r=\) \(\Theta(d k \log k)\) in magnitude, and in general about \(2^{i}\) entries about \(r 2^{-i}\) in magnitude. These entries typically occur in columns \(C_{j}\) of \(S A\) for which all other entries in the column are bounded by poly \((k)\) in magnitude. Thus, \(\left|x^{\top} C_{j}\right| \approx r 2^{-i}\) for about \(2^{i}\) columns \(j\). For each such column, if \(r 2^{-i} \gg \log d\), then we incur cost \(\frac{r 2^{-i}}{\operatorname{poly}(k)}\) in approximating the first row of \(A\). In total the cost is \(\frac{r \log r}{\operatorname{poly}(k)}=\frac{d \log d}{\operatorname{poly}(k)}\), but the optimal cost is at most \(d\), giving a \(\frac{\log d}{\operatorname{poly}(k)}\) lower bound. We optimize this to a \(\frac{\log d}{k \log ^{2} k}\) lower bound.

When \(k\) is large this bound deteriorates, but we also show a \(k^{1 / 2-\gamma}\) lower bound for arbitrarily small constant \(\gamma>0\). This bound applies to any oblivious sketching matrix. The idea is similar to our row subset selection lower bound. Let \(A\) be as in our row subset selection lower bound, consider \(S A\), and write \(S=U \Sigma V^{\top}\) in its full SVD. Then \(S A\) is in the row span of the top \(O(k \log k)\) rows of \(V^{\top} A\), since \(\Sigma\) only has \(O(k \log k)\) non-zero singular values. Since the first \(k\) columns of \(A\) are rotationally invariant, \(V^{\top} A\) has first \(k\) columns i.i.d. Gaussian and remaining columns equal to \(V^{\top}\). Call the first \(O(k \log k)\) rows of \(V^{\top} A\) the matrix \(B\). We now try to approximate a row of \(A\) by a vector in the row span \(B\). There are two issues that make this setting different from row subset selection: (1) \(B\) no longer contains an identity submatrix, and (2) the rows of \(B\) depend on the rows of \(A\). We handle the first issue by building nets for subsets of coordinates of \(x^{\top} V^{\top}\) rather than \(x\) as before; since \(\left\|x^{\top} V^{\top}\right\|_{2}=\|x\|_{2}\) similar arguments can be applied. We handle the second issue by observing that if the number of rows of \(B\) is considerably smaller than that of \(A\), then the distribution of \(B\) had we replaced a random row of \(A\) with zeros would be statistically close to i.i.d. Gaussian. Hence, typical rows of \(A\) can be regarded as being independent of
\(B\).

\subsection*{17.1.3 Several Theorem Statements, an Algorithm, and a Roadmap}
```

Algorithm 17.1 Main Meta-Algorithm
procedure L1LowRAnkApprox $(A, n, d, k) \quad \triangleright$ Theorem 17.1.1
Choose sketching matrix $S$ (a Cauchy matrix or a sparse Cauchy matrix.)
Compute $S A$, form $C$ by $C^{i} \leftarrow \arg \min _{x}\left\|x S A-A^{i}\right\|_{1}$. Form $B=C \cdot S A$.
Choose sketching matrices $T_{1}, R, D, T_{2}$ (Cauchy matrices or sparse Cauchy matrices.)
Solve $\min _{X, Y}\left\|T_{1} B R X Y D B T_{2}-T_{1} B T_{2}\right\|_{F}$.
return $B R X, Y D B$.
end procedure

```

Theorem 17.1.1 (Informal Version of Theorem 17.4.6). Given \(A \in \mathbb{R}^{n \times d}\), there is an algorithm which in \(\mathrm{nnz}(A)+(n+d) \cdot \operatorname{poly}(k)\) time, outputs a (factorization of a) rank-k matrix \(A^{\prime}\) such that with probability \(9 / 10,\left\|A^{\prime}-A\right\|_{1} \leq(\log d) \operatorname{poly}(k) \min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|U V-A\|_{1}\).

Theorem 17.1.2 (Informal Version of Theorem 17.4.7). Given \(A \in \mathbb{R}^{n \times d}\), there is an algorithm that takes poly \((n) d^{\widetilde{O}(k)} 2^{\widetilde{O}\left(k^{2}\right)}\) time and outputs a rank-k matrix \(A^{\prime}\) such that, with probability 9/10, \(\left\|A^{\prime}-A\right\|_{1} \leq \widetilde{O}(k) \min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|U V-A\|_{1}\). In addition, \(A^{\prime}\) is a CUR decomposition.

Theorem 17.1.3 (Informal Version of Theorem 17.6.26). For any \(k \geq 1\), and any constant \(c \geq 1\), let \(n=k^{c}\). There exists a matrix \(A\) such that for any matrix \(A^{\prime}\) in the span of \(n / 2\) rows of \(A,\left\|A^{\prime}-A\right\|_{1}=\Omega\left(k^{0.5-\alpha}\right)_{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|U V-A\|_{1}\), where \(\alpha>0\) is an arbitrarily small constant.

Road map Section 17.2 introduces some notation and definitions. Section 17.3 includes several useful tools. We provide several \(\ell_{1}\)-low rank approximation algorithms in Section 17.4. Section 17.5 contains the no contraction and no dilation analysis for our main algorithm.

We provide our existential hardness results for Cauchy matrices, row subset selection and oblivious subspace embeddings in Section 17.6.

\subsection*{17.2 Notation}

Let \(\mathbb{N}_{+}\)denote the set of positive integers. For any \(n \in \mathbb{N}_{+}\), let \([n]\) denote the set \(\{1,2, \cdots, n\}\). For any \(p \in[1,2]\), the \(\ell_{p}\)-norm of a vector \(x \in \mathbb{R}^{d}\) is defined as
\[
\|x\|_{p}=\left(\sum_{i=1}^{d}\left|x_{i}\right|^{p}\right)^{1 / p}
\]

For any \(p \in[1,2)\), the \(\ell_{p}\)-norm of a matrix \(A \in \mathbb{R}^{n \times d}\) is defined as
\[
\|A\|_{p}=\left(\sum_{i=1}^{n} \sum_{j=1}^{d}\left|A_{i j}\right|^{p}\right)^{1 / p}
\]

Let \(\|A\|_{F}\) denote the Frobenius norm of matrix \(A\). Let \(\operatorname{nnz}(A)\) denote the number of nonzero entries of \(A\). Let \(\operatorname{det}(A)\) denote the determinant of a square matrix \(A\). Let \(A^{\top}\) denote the transpose of \(A\). Let \(A^{\dagger}\) denote the Moore-Penrose pseudoinverse of \(A\). Let \(A^{-1}\) denote the inverse of a full rank square matrix. We use \(A_{j}\) to denote the \(j^{\text {th }}\) column of \(A\), and \(A^{i}\) to denote the \(i^{\text {th }}\) row of \(A\). For an \(n \times d\) matrix \(A\), for \(S\) a subset of \([n]\) and \(T\) a subset of \([d]\), we let \(A^{S}\) denote the \(|S| \times d\) submatrix of \(A\) with rows indexed by \(S\), while \(A_{T}\) denotes the \(n \times|T|\) submatrix of \(A\) with columns indexed by \(T\), and \(A_{T}^{S}\) denote the \(|S| \times|T|\) submatrix \(A\) with rows in \(S\) and columns in \(T\).

For any function \(f\), we define \(\widetilde{O}(f)\) to be \(f \cdot \log ^{O(1)}(f)\). In addition to \(O(\cdot)\) notation, for two functions \(f, g\), we use the shorthand \(f \lesssim g\) (resp. \(\gtrsim\) ) to indicate that \(f \leq C g\) (resp. \(\geq\) ) for an absolute constant \(C\). We use \(f \approx g\) to mean \(c f \leq g \leq C f\) for constants \(c, C\). We use OPT to denote \(\min _{\text {rank }-k A_{k}}\left\|A_{k}-A\right\|_{1}\), unless otherwise specified.

\subsection*{17.3 Preliminaries}

\subsection*{17.3.1 Polynomial system verifier}

Renegar [Ren92a, Ren92b] and Basu et al. [BPR96] independently provided an algorithm for the decision problem for the existential theory of the reals, which is to decide the truth or falsity of a sentence \(\left(x_{1}, \cdots, x_{v}\right) F\left(f_{1}, \cdots, f_{m}\right)\) where \(F\) is a quantifier-free Boolean formula with atoms of the form \(\operatorname{sign}\left(f_{i}\right)=\sigma\) with \(\sigma \in\{0,1,-1\}\). Note that this problem is equivalent to deciding if a given semi-algebraic set is empty or not. Here we formally state that theorem. For a full discussion of algorithms in real algebraic geometry, we refer the reader to [BPR05] and [Bas14].

Theorem 17.3.1 (Decision Problem [Ren92a, Ren92b, BPR96]). Given a real polynomial system \(P\left(x_{1}, x_{2}, \cdots, x_{v}\right)\) having \(v\) variables and \(m\) polynomial constraints
\[
f_{i}\left(x_{1}, x_{2}, \cdots, x_{v}\right) \Delta_{i} 0, \forall i \in[m]
\]
where \(\Delta_{i}\) is any of the "standard relations": \(\{>, \geq,=, \neq, \leq,<\}\), let \(d\) denote the maximum degree of all the polynomial constraints and let \(H\) denote the maximum bitsize of the coefficients of all the polynomial constraints. Then in
\[
(m d)^{O(v)} \operatorname{poly}(H),
\]
time one can determine if there exists a solution to the polynomial system \(P\).

Recently, this technique has been used to solve a number of low-rank approximation and matrix factorization problems [AGKM12, Mad13, CW15a, BDL16, RSW16].

\subsection*{17.3.2 Cauchy and \(p\)-stable transform}

Definition 17.3.1 (Dense Cauchy transform). Let \(S=\sigma \cdot C \in \mathbb{R}^{m \times n}\) where \(\sigma\) is a scalar, and each entry of \(C \in \mathbb{R}^{m \times n}\) is chosen independently from the standard Cauchy distribution. For any matrix \(A \in \mathbb{R}^{n \times d}, S A\) can be computed in \(O(m \cdot \mathrm{nnz}(A))\) time.

Definition 17.3.2 (Sparse Cauchy transform). Let \(\Pi=\sigma \cdot S C \in \mathbb{R}^{m \times n}\), where \(\sigma\) is a scalar, \(S \in \mathbb{R}^{m \times n}\) has each column chosen independently and uniformly from the \(m\) standard basis vectors of \(\mathbb{R}^{m}\), and \(C \in \mathbb{R}^{n \times n}\) is a diagonal matrix with diagonals chosen independently from the standard Cauchy distribution. For any matrix \(A \in \mathbb{R}^{n \times d}, \Pi A\) can be computed in \(O(\operatorname{nnz}(A))\) time.

Definition 17.3.3 (Dense \(p\)-stable transform). Let \(p \in(1,2)\). Let \(S=\sigma \cdot C \in \mathbb{R}^{m \times n}\) where \(\sigma\) is a scalar, and each entry of \(C \in \mathbb{R}^{m \times n}\) is chosen independently from the standard \(p\)-stable distribution. For any matrix \(A \in \mathbb{R}^{n \times d}, S A\) can be computed in \(O(m \mathrm{nnz}(A))\) time.

Definition 17.3.4 (Sparse \(p\)-stable transform). Let \(p \in(1,2)\). Let \(\Pi=\sigma \cdot S C \in \mathbb{R}^{m \times n}\), where \(\sigma\) is a scalar, \(S \in \mathbb{R}^{m \times n}\) has each column chosen independently and uniformly from the \(m\) standard basis vectors of \(\mathbb{R}^{m}\), and \(C \in \mathbb{R}^{n \times n}\) is a diagonal matrix with diagonals chosen independently from the standard \(p\)-stable distribution. For any matrix \(A \in \mathbb{R}^{n \times d}, \Pi A\) can be computed in \(O(\mathrm{nnz}(A))\) time.

\subsection*{17.3.3 Lewis weights}

We follow the exposition of Lewis weights from [CP15].
Definition 17.3.5. For a matrix \(A\), let \(a_{i}\) denote \(i^{\text {th }}\) row of \(A\), and \(a_{i}\left(=\left(A^{i}\right)^{\top}\right)\) is a column vector. The statistical leverage score of a row \(a_{i}\) is
\[
\tau_{i}(A) \stackrel{\text { def }}{=} a_{i}^{\top}\left(A^{\top} A\right)^{-1} a_{i}=\left\|\left(A^{\top} A\right)^{-1 / 2} a_{i}\right\|_{2}^{2} .
\]

For a matrix \(A\) and norm \(p\), the \(\ell_{p}\) Lewis weights \(w\) are the unique weights such that for each row \(i\) we have
\[
w_{i}=\tau_{i}\left(W^{1 / 2-1 / p} A\right)
\]
or equivalently
\[
a_{i}^{\top}\left(A^{\top} W^{1-2 / p} A\right)^{-1} a_{i}=w_{i}^{2 / p}
\]

Lemma 17.3.2 (Lemma 2.4 of [CP15] and Lemma 7 of \(\left[\mathrm{CLM}^{+} 15\right]\) ). Given a matrix \(A \in\) \(\mathbb{R}^{n \times d}, n \geq d\), for any constant \(C>0,4>p \geq 1\), there is an algorithm which can compute \(C\)-approximate \(\ell_{p}\) Lewis weights for every row \(i\) of \(A\) in \(O\left(\left(n n z(A)+d^{\omega} \log d\right) \log n\right)\) time, where \(\omega<2.373\) is the matrix multiplication exponent/Str69, CW87, Wil12].

Lemma 17.3.3 (Theorem 7.1 of [CP15]). Given matrix \(A \in \mathbb{R}^{n \times d}(n \geq d)\) with \(\ell_{p}(4>p \geq\) 1) Lewis weights \(w\), for any set of sampling probabilities \(p_{i}, \sum_{i} p_{i}=N\),
\[
p_{i} \geq f(d, p) w_{i}
\]
if \(S \in \mathbb{R}^{N \times n}\) has each row chosen independently as the \(i^{\text {th }}\) standard basis vector, times \(1 / p_{i}^{1 / p}\), with probability \(p_{i} / N\), then with probability at least 0.999 ,
\[
\forall x \in \mathbb{R}^{d}, \frac{1}{2}\|A x\|_{p}^{p} \leq\|S A x\|_{p}^{p} \leq 2\|A x\|_{p}^{p}
\]

Furthermore, if \(p=1, N=O(d \log d)\). If \(1<p<2, N=O(d \log d \log \log d)\). If \(2 \leq p<4\), \(N=O\left(d^{p / 2} \log d\right)\).

Given a matrix \(A \in \mathbb{R}^{n \times d}(n \geq d)\), by Lemma 17.3.3 and Lemma 17.3.2, we are able to compute a sampling/rescaling matrix \(S\) in \(O\left(\left(n n z(A)+d^{\omega} \log d\right) \log n\right)\) with \(\widetilde{O}(d)\) nonzero entries such that
\[
\forall x \in \mathbb{R}^{d}, \frac{1}{2}\|A x\|_{p}^{p} \leq\|S A x\|_{p}^{p} \leq 2\|A x\|_{p}^{p}
\]

Sometimes, \(\operatorname{poly}(d)\) is much smaller than \(\log n\). In this case, we are able to compute the such sampling/rescaling matrix \(S\) in \(n\) poly \((d)\) time in the following way: basically we can run one of the input sparsity \(\ell_{p}\) embedding algorithm (see e.g. [MM13]) to compute a well conditioned basis \(U\) of column span of \(A\) in \(n\) poly \((d)\) time. By sampling according to the well conditioned basis (see e.g. [Cla05, \(\mathrm{DDH}^{+} 09\), Woo14b]), we can compute a sampling/rescaling matrix \(S_{1}\) such that \((1-\epsilon)\|A x\|_{p}^{p} \leq\left\|S_{1} A x\right\|_{p}^{p} \leq(1+\epsilon)\|A x\|_{p}^{p}\) where \(\epsilon \in(0,1)\) is an arbitrary constant. Notice that \(S_{1}\) has poly \((d)\) nonzero entries, thus \(S_{1} A\) has size poly \((d)\). Now, we apply Lewis weights sampling according to \(S_{1} A\), we can get a sampling/rescaling matrix \(S\) such that
\[
\forall x \in \mathbb{R}^{d},\left(1-\frac{1}{3}\right)\left\|S_{1} A x\right\|_{p}^{p} \leq\left\|S S_{1} A x\right\|_{p}^{p} \leq\left(1+\frac{1}{3}\right)\left\|S_{1} A x\right\|_{p}^{p}
\]

It means that
\[
\forall x \in \mathbb{R}^{d}, \frac{1}{2}\|A x\|_{p}^{p} \leq\left\|S S_{1} A x\right\|_{p}^{p} \leq 2\|A x\|_{p}^{p}
\]

Note that \(S S_{1}\) is still a sampling/rescaling matrix according to \(A\), and the number of non-zero entries is \(\widetilde{O}(d)\). And the total running time is thus \(n \operatorname{poly}(d)\).

\subsection*{17.3.4 Frobenius norm and \(\ell_{2}\) relaxation}

Theorem 17.3.4 (Generalized rank-constrained matrix approximations, Theorem 2 in [FT07]). Given matrices \(A \in \mathbb{R}^{n \times d}, B \in \mathbb{R}^{n \times p}\), and \(C \in \mathbb{R}^{q \times d}\), let the \(S V D\) of \(B\) be \(B=U_{B} \Sigma_{B} V_{B}^{\top}\) and the \(S V D\) of \(C\) be \(C=U_{C} \Sigma_{C} V_{C}^{\top}\). Then,
\[
B^{\dagger}\left(U_{B} U_{B}^{\top} A V_{C} C_{C}^{\top}\right)_{k} C^{\dagger}=\underset{\operatorname{rank}-k}{\arg \min } \| \in \mathbb{R}^{p \times q}\left(\|A-B X C\|_{F},\right.
\]
where \(\left(U_{B} U_{B}^{\top} A V_{C} V_{C}^{\top}\right)_{k} \in \mathbb{R}^{p \times q}\) is of rank at most \(k\) and denotes the best rank-k approximation to \(U_{B} U_{B}^{\top} A V_{C} V_{C}^{\top} \in \mathbb{R}^{p \times d}\) in Frobenius norm.

Claim 17.3.5 ( \(\ell_{2}\) relaxation of \(\ell_{p}\)-regression). Let \(p \in[1,2)\). For any \(A \in \mathbb{R}^{n \times d}\) and \(b \in \mathbb{R}^{n}\), define \(x^{*}=\underset{x \in \mathbb{R}^{d}}{\arg \min }\|A x-b\|_{p}\) and \(x^{\prime}=\underset{x \in \mathbb{R}^{d}}{\arg \min }\|A x-b\|_{2}\). Then,
\[
\left\|A x^{*}-b\right\|_{p} \leq\left\|A x^{\prime}-b\right\|_{p} \leq n^{1 / p-1 / 2} \cdot\left\|A x^{*}-b\right\|_{p}
\]

Proof. The lower bound trivially holds by definition; we will focus on proving the upper bound. Because \(A x-b\) is an \(n\)-dimensional vector, \(\forall x\),
\[
\begin{equation*}
\frac{1}{n^{1 / p-1 / 2}}\|A x-b\|_{p} \leq\|A x-b\|_{2} \leq\|A x-b\|_{p} \tag{17.1}
\end{equation*}
\]

Then,
\[
\begin{aligned}
& \left\|A x^{\prime}-b\right\|_{p} \\
\leq & \\
\leq & \\
\leq & \sqrt{n}\left\|A x^{\prime}-b\right\|_{2}-b \|_{2} \\
\leq & \text { by LHS of Equation }(17.1) \\
\leq & \text { by } x^{\prime}=\underset{x}{\arg \min }\|A x-b\|_{2}^{*}-b \|_{p}
\end{aligned} \quad \text { by RHS of Equation (17.1) }
\]

This completes the proof.

Claim 17.3.6 (Frobenius norm relaxation of \(\ell_{p}\)-low rank approximation). Let \(p \in[1,2)\) and for any matrix \(A \in \mathbb{R}^{n \times d}\), define \(A^{*}=\underset{\operatorname{rank}-k}{\arg \min } B \in \mathbb{R}^{n \times d}\|B-A\|_{p}\) and \(A^{\prime}=\underset{\operatorname{rank}-k}{\arg \min }\left\|\in \mathbb{R}^{n \times d}\right\| B-\) \(A \|_{F}\). Then
\[
\begin{equation*}
\left\|A^{*}-A\right\|_{p} \leq\left\|A^{\prime}-A\right\|_{p} \leq(n d)^{1 / p-1 / 2}\left\|A^{*}-A\right\|_{p} \tag{17.2}
\end{equation*}
\]

Proof. The lower bound of \(\left\|A^{\prime}-A\right\|_{p}\) trivially holds by definition. We show an upper bound of \(\left\|A^{\prime}-A\right\|_{p}\) in the rest of the proof. For any \(A^{\prime}-A \in \mathbb{R}^{n \times d}\), we have
\[
\begin{equation*}
\frac{1}{(n d)^{1 / p-1 / 2}}\left\|A^{\prime}-A\right\|_{p} \leq\left\|A^{\prime}-A\right\|_{F} \leq\left\|A^{\prime}-A\right\|_{p} \tag{17.3}
\end{equation*}
\]

Then,
\[
\begin{aligned}
& \left\|A^{\prime}-A\right\|_{p} \\
\leq & (n d)^{1 / p-1 / 2}\left\|A^{\prime}-A\right\|_{F} \\
\leq & \text { by LHS of Equation (17.3) } \\
\leq & (n d)^{1 / p-1 / 2}\left\|A^{*}-A\right\|_{F}
\end{aligned} \quad \text { by } A^{\prime}=\underset{\text { rank }-k B}{\arg \min }\|B-A\|_{p} .
\]

\subsection*{17.3.5 Converting entry-wise \(\ell_{1}\) and \(\ell_{p}\) objective functions into polynomials}

Claim 17.3.7 (Converting absolute value constraints into variables). Given \(m\) polynomials \(f_{1}(x), f_{2}(x), \cdots, f_{m}(x)\) where \(x \in \mathbb{R}^{v}\), solving the problem
\[
\begin{equation*}
\min _{x \in \mathbb{R}^{v}} \sum_{i=1}^{m}\left|f_{i}(x)\right| \tag{17.4}
\end{equation*}
\]
is equivalent to solving another minimization problem with \(O(m)\) extra constraints and \(m\) extra variables,
\[
\begin{aligned}
& \quad \min _{x \in \mathbb{R}^{v}, \sigma \in \mathbb{R}^{m}} \sum_{i=1}^{m} \sigma_{i} f_{i}(x) \\
& \text { s.t. } \\
& \sigma_{i}^{2}=1, \forall i \in[m] \\
& \\
& \quad f_{i}(x) \sigma_{i} \geq 0, \forall i \in[m] .
\end{aligned}
\]

Claim 17.3.8. (Handling \(\ell_{p}\) ) Given \(m\) polynomials \(f_{1}(x), f_{2}(x), \cdots, f_{m}(x)\) where \(x \in \mathbb{R}^{v}\) and \(p=a / b\) for positive integers \(a\) and \(b\), solving the problem
\[
\begin{equation*}
\min _{x \in \mathbb{R}^{v}} \sum_{i=1}^{m}\left|f_{i}(x)\right|^{p}, \tag{17.5}
\end{equation*}
\]
is equivalent to solving another minimization problem with \(O(m)\) extra constraints and \(O(m)\) extra variables,
\[
\begin{array}{ll} 
& \min _{x \in \mathbb{R}^{v}, \sigma \in \mathbb{R}^{m}} \sum_{i=1}^{m} y_{i} \\
\text { s.t. } & \sigma_{i}^{2}=1, \forall i \in[m] \\
& f_{i}(x) \sigma_{i} \geq 0, \forall i \in[m] \\
& \left(\sigma_{i} f_{i}(x)\right)^{a}=y_{i}^{b}, \forall i \in[m] \\
& y_{i} \geq 0, \forall i \in[m] .
\end{array}
\]

\subsection*{17.3.6 Converting entry-wise \(\ell_{1}\) objective function into a linear program}

Claim 17.3.9. Given any matrix \(A \in \mathbb{R}^{n \times d}\) and matrix \(B \in \mathbb{R}^{k \times d}\), the problem \(\min _{U \in \mathbb{R}^{n \times k}} \| U B\) \(A \|_{1}\) can be solved by solving the following linear program,
\[
\begin{aligned}
& \min _{U \in \mathbb{R}^{n \times k}, x \in \mathbb{R}^{n \times d}} \sum_{i=1}^{n} \sum_{j=1}^{m} x_{i, j} \\
& U_{i} B^{j}-A_{i, j} \leq x_{i, j}, \forall i \in[n], j \in[d] \\
& U_{i} B^{j}-A_{i, j} \geq-x_{i, j}, \forall i \in[n], j \in[d] \\
& x_{i, j} \geq 0, \forall i \in[n], j \in[d],
\end{aligned}
\]
where the number of constraints is \(O(n d)\) and the number of variables is \(O(n d)\).

\section*{17.4 \(\quad \ell_{1}\)-Low Rank Approximation}

This section presents our main \(\ell_{1}\)-low rank approximation algorithms. Section 17.4.1 provides our three existence results. Section 17.4.2 shows an input sparsity algorithm with \(\operatorname{poly}(k) \log ^{2} d \log n\)-approximation ratio. Section 17.4.3 improves the approximation ratio to \(\operatorname{poly}(k) \log d\). Section 17.4.4 explains how to obtain \(\widetilde{O}(k)\) approximation ratio. Section 17.4.5 improves the approximation ratio to \(O(1)\) by outputting a rank- \(3 k\) solution. Section 17.4.6 presents our algorithm for CUR decomposition. Section 17.4.7 includes some useful properties. Our \(\ell_{1}\)-low rank approximation algorithm for a rank- \(r\) (where \(k \leq r \leq(n, d)\) ) matrix is used as a black box (by setting \(r=\operatorname{poly}(k)\) ) in several other algorithms.

\subsection*{17.4.1 Existence results via dense Cauchy transforms, sparse Cauchy transforms, Lewis weights}

The goal of this section is to present the existence results in Corollary 17.4.2. We first provide some bicriteria algorithms in Theorem 17.4.1 which can be viewed as a "warmup". Then the proof of our bicriteria algorithm actually implies the existence results.

Theorem 17.4.1. Given matrix \(A \in \mathbb{R}^{n \times d}\), for any \(k \geq 1\), there exist bicriteria algorithms with running time \(T\) (specified below), which output two matrices \(U \in \mathbb{R}^{n \times m}, V \in \mathbb{R}^{m \times d}\) such that, with probability 9/10,
\[
\|U V-A\|_{1} \leq \alpha \min _{\mathrm{rank}-k}\left\|A_{k}-A\right\|_{1}
\]
(I). Using a dense Cauchy transform,
\(T=\operatorname{poly}(n, d, k), m=O(k \log k), \alpha=O(\sqrt{k \log k} \log d)\).
(II). Using a sparse Cauchy transform,
\(T=\operatorname{poly}(n, d, k), m=O\left(k^{5} \log ^{5} k\right), \alpha=O\left(k^{4.5} \log ^{4.5} k \log d\right)\).
(III). Sampling by Lewis weights,
\(T=(n d)^{\tilde{O}(k)}, m=O(k \log k), \alpha=O(\sqrt{k \log k})\).

The matrices in (I), (II), (III) here, are the same as those in (I), (II), (III), (IV) of Lemma 17.5.5. Thus, they have the properties shown in Section 17.5.2.

Proof. We define
\[
\mathrm{OPT}:=\min _{\operatorname{rank}-k A_{k}}\left\|A_{k}-A\right\|_{1} .
\]

We define \(U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times d}\) to be the optimal solution such that \(\left\|U^{*} V^{*}-A\right\|_{1}=\) OPT .
Part (I). Apply the dense Cauchy transform \(S \in \mathbb{R}^{m \times n}\) with \(m=O(k \log k)\) rows, and \(\beta=O(\log d)\).

Part (II). Apply the sparse Cauchy transform \(S\left(=\Pi \in \mathbb{R}^{m \times n}\right)\) with \(m=O\left(k^{5} \log ^{5} k\right)\) rows, and \(\beta=O(\sigma \log d)=O\left(k^{2} \log ^{2} k \log d\right)\).

Part (III). Use \(S\left(=D \in \mathbb{R}^{n \times k}\right)\) to denote an \(n \times n\) matrix which is a sampling and rescaling diagonal matrix according to the Lewis weights of matrix \(U^{*}\). It has \(m=O(k \log k)\) rows, and \(\beta=O(1)\). Sometimes we abuse notation, and should regard \(D\) as a matrix which has size \(m \times n\), where \(m=O(k \log k)\).

We can just replace \(M\) in Lemma 17.5.5 with \(U^{*} V^{*}-A\), replace \(U\) in Lemma 17.5.5 with \(U^{*}\), and replace \(c_{1} c_{2}\) with \(O(\beta)\). So, we can apply Lemma 17.5.5 for \(S\). Then we can plug it in Lemma 17.5.2, we have: with constant probability, for any \(c \geq 1\), for any \(V^{\prime} \in \mathbb{R}^{k \times d}\) which satisfies
\[
\begin{equation*}
\left\|S U^{*} V^{\prime}-S A\right\|_{1} \leq c \cdot \min _{V \in \mathbb{R}^{k \times d}}\left\|S U^{*} V-S A\right\|_{1}, \tag{17.6}
\end{equation*}
\]
it has
\[
\begin{equation*}
\left\|U^{*} V^{\prime}-A\right\|_{1} \leq c \cdot O(\beta)\left\|U^{*} V^{*}-A\right\|_{1} . \tag{17.7}
\end{equation*}
\]

Define \(\widehat{V}_{i}=\underset{V_{i} \in \mathbb{R}^{k}}{\arg \min }\left\|S U^{*} V_{i}-S A_{i}\right\|_{2}\) for each \(i \in[d]\). By using Claim 17.3.5 with
\(n=m\) and \(d=k\), it shows
\[
\begin{aligned}
\left\|S U^{*} \widehat{V}-S A\right\|_{1} & =\sum_{i=1}^{d}\left\|S U^{*} \widehat{V}_{i}-S A_{i}\right\|_{1} \\
& \leq \sum_{i=1}^{d} \sqrt{m}\left\|S U^{*} \widetilde{V}_{i}-S A_{i}\right\|_{1} \\
& =\sqrt{m} \min _{V \in \mathbb{R}^{k \times d}}\left\|S U^{*} V-S A\right\|_{1} .
\end{aligned}
\]
which means \(\widehat{V}\) is a \(\sqrt{m}\)-approximation solution to problem, \(\min _{V \in \mathbb{R}^{k x d}}\left\|S U^{*} V-S A\right\|_{1}\).
Now, let us look into Equation (17.6) and Equation (17.7), we can obtain that
\[
\left\|U^{*} \widehat{V}-A\right\|_{1} \leq \sqrt{m} O(\beta) \mathrm{OPT}
\]

Because \(\widehat{V}_{i}\) is the optimal solution of the \(\ell_{2}\) regression problem, we have
\[
\widehat{V}_{i}=\left(S U^{*}\right)^{\dagger} S A_{i} \in \mathbb{R}^{k}, \forall i \in[d] \text {, which means } \widehat{V}=\left(S U^{*}\right)^{\dagger} S A \in \mathbb{R}^{k \times d}
\]

Plugging \(\widehat{V}\) into original problem, we obtain
\[
\left\|U^{*}\left(S U^{*}\right)^{\dagger} \cdot S A-A\right\|_{1} \leq \sqrt{m} O(\beta) \mathrm{OPT}
\]

It means
\[
\begin{equation*}
\min _{\text {rank }-k \in \mathbb{R}^{n \times m}}\|X S A-A\|_{1} \leq \sqrt{m} O(\beta) \mathrm{OPT} \tag{17.8}
\end{equation*}
\]

If we ignore the constraint on the rank of \(X\), we can get a bicriteria solution:
For part (I), notice that \(X\) is an \(n \times m\) matrix which can be found by using a linear program, because matrices \(S A \in \mathbb{R}^{m \times d}\) and \(A \in \mathbb{R}^{n \times d}\) are known.

For part (II), notice that \(X\) is an \(n \times m\) matrix which can be found by using a linear program, because matrices \(S A \in \mathbb{R}^{m \times d}\) and \(A \in \mathbb{R}^{n \times d}\) are known.

For part (III), notice that \(X\) is an \(n \times m\) matrix which can be found by using a linear program, when the span of rows of \(D A \in \mathbb{R}^{m \times d}\) is known. We assume that \(D\) is known in all the above discussions. But \(D\) is actually unknown. So we need to try all the possible choices of the row span of \(D A\). Since \(D\) samples at most \(m=O(k \log k)\) rows of \(A\), then the total number of choices of selecting \(m\) rows from \(n\) rows is \(\binom{n}{m}=n^{O(k \log k)}\). This completes the proof.

Equation (17.8) in the proof of our bicriteria solution implies the following result,

Corollary 17.4.2. Given \(A \in \mathbb{R}^{n \times d}\), there exists a rank-k matrix \(A^{\prime} \in \mathbb{R}^{n \times d}\) such that \(A^{\prime} \in \operatorname{rowspan}\left(S^{\prime} A\right) \subseteq \operatorname{rowspan}(A)\) and \(\left\|A^{\prime}-A\right\|_{1} \leq \alpha \cdot \min _{\operatorname{rank}-k}\left\|A-A_{k}\right\|_{1}\), where \(S^{\prime} \in \mathbb{R}^{m \times n}\) is a sketching matrix. If \(S^{\prime}\)
(I). indicates the dense Cauchy transform, then \(\alpha=O(\sqrt{k \log k} \log d)\).
(II). indicates the sparse Cauchy transform, then \(\alpha=O\left(k^{4.5} \log ^{4.5} k \log d\right)\).
(III). indicates sampling by Lewis weights, then \(\alpha=O(\sqrt{k \log k})\).

Proof. Define \(\mathrm{OPT}=\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|U V-A\|_{1}\).
Proof of (I). Choose \(S\) to be a dense Cauchy transform matrix with \(m\) rows, then
\[
\min _{U \in \mathbb{R}^{n \times k}, Z \in \mathbb{R}^{k \times m}}\|U Z S A-A\|_{1} \leq O(\sqrt{m} \log d) \mathrm{OPT}
\]
where \(m=O(k \log k)\). Choosing \(A^{\prime}=U Z S A\) completes the proof.

Proof of (II). Choose \(\Pi=S D \in \mathbb{R}^{m \times n}\) where \(S \in \mathbb{R}^{m \times n}\) has each column chosen independently and uniformly from the \(m\) standard basis vectors of \(\mathbb{R}^{m}\), and where \(D\) is a diagonal matrix with diagonals chosen independently from the standard Cauchy distribution, then
\[
\min _{U \in \mathbb{R}^{n \times k}, Z \in \mathbb{R}^{k \times m}}\|U Z \Pi A-A\|_{1} \leq O(\sqrt{m} \sigma \log d) \mathrm{OPT}
\]
where \(m=O\left(k^{5} \log ^{5} k\right)\) and \(\sigma=O\left(k^{2} \log ^{2} k\right)\). Choosing \(A^{\prime}=U Z \Pi A\) completes the proof.

\section*{Proof of (III).}

Choose \(D\) to be the sampling and rescaling matrix corresponding to the Lewis weights of \(U^{*}\), and let it have \(m=O(k \log k)\) nonzero entries on the diagonal, then
\[
\min _{U \in \mathbb{R}^{n \times k}, Z \in \mathbb{R}^{k \times m}}\|U Z D A-A\|_{1} \leq O(\sqrt{m}) \mathrm{OPT}
\]

Choosing \(A^{\prime}=U Z D A\) completes the proof.
17.4.2 Input sparsity time, \(\operatorname{poly}(k, \log n, \log d)\)-approximation for an arbitrary matrix \(A\)

The algorithm described in this section is actually worse than the algorithm described in the next section. But this algorithm is easy to extend to the distributed and streaming settings (See full version of [SWZ17]).
```

Algorithm 17.2 Input Sparsity Time Algorithm
procedure L1LowRankApproxInputSparsity $(A, n, d, k) \quad \triangleright$ Theorem 17.4.3
Set $s \leftarrow r \leftarrow t_{1} \leftarrow \widetilde{O}\left(k^{5}\right), t_{2} \leftarrow \widetilde{O}(k)$.
Choose sparse Cauchy matrices $S \in \mathbb{R}^{s \times n}, R \in \mathbb{R}^{d \times r}, T_{1} \in \mathbb{R}^{t_{1} \times n}$.
Choose dense Cauchy matrices $T_{2} \in \mathbb{R}^{d \times t_{2}}$.
Compute $S \cdot A, A \cdot R$ and $T_{1} \cdot A \cdot T_{2}$.
Compute $X Y=\arg \min _{X, Y}\left\|T_{1} A R X Y S A T_{2}-T_{1} A T_{2}\right\|_{F}$.
return $A R X, Y S A$.
end procedure

```

Theorem 17.4.3. Given matrix \(A \in \mathbb{R}^{n \times d}\), for any \(k \geq 1\), there exists an algorithm which takes \(O(\operatorname{nnz}(A))+(n+d) \cdot \operatorname{poly}(k)\) time and outputs two matrices \(U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}\) such that
\[
\|U V-A\|_{1} \leq O\left(\operatorname{poly}(k) \log n \log ^{2} d\right) \min _{\operatorname{rank}-k A_{k}}\left\|A_{k}-A\right\|_{1}
\]
holds with probability 9/10.

Proof. Choose a Cauchy matrix \(S \in \mathbb{R}^{s \times n}\) (notice that \(S\) can be either a dense Cauchy transform matrix or a sparse Cauchy transform matrix). Using Corollary 17.4.2, we have
\[
\min _{U \in \mathbb{R}^{n \times k}, Z \in \mathbb{R}^{k \times s}}\|U Z S A-A\|_{1} \leq \alpha_{s} \mathrm{OPT}
\]
where \(\alpha_{s}\) is the approximation by using matrix \(S\). If \(S\) is a dense Cauchy transform matrix, then due to Part (I) of Corollary 17.4.2, \(\alpha_{s}=O(\sqrt{k \log k} \log d), s=O(k \log k)\), and computing \(S A\) takes \(O(s \mathrm{nnz}(A))\) time. If \(S\) is a sparse Cauchy transform matrix, then due to Part II of Corollary 17.4.2, \(\alpha_{s}=\widetilde{O}\left(k^{4.5} \log d\right), s=\widetilde{O}\left(k^{5}\right)\), and computing \(S A\) takes nnz \((A)\) time.

We define \(U^{*}, Z^{*}=\arg \min _{U, Z}\|U Z S A-A\|_{1}\). For the fixed \(Z^{*} \in \mathbb{R}^{k \times s}\), choose a Cauchy matrix \(R \in \mathbb{R}^{d \times r}\) (note that \(R\) can be either a dense Cauchy transform matrix or a sparse Cauchy transform matrix) and sketch on the right of \((U Z S A-A)\). If \(R\) is a dense Cauchy transform matrix, then \(\alpha_{r}=O(\log n), r=O(k \log k)\), computing \(A R\) takes \(O(r \cdot \operatorname{nnz}(A))\) time. If \(R\) is a sparse Cauchy transform matrix, then \(\alpha_{r}=\widetilde{O}\left(k^{2}\right) \log n\), \(r=\widetilde{O}\left(k^{5}\right)\), computing \(A R\) takes \(O(\mathrm{nnz}(A))\) time.

Define a row vector \(\widehat{U}^{j}=A^{j} R\left(\left(Z^{*} S A\right) R\right)^{\dagger} \in \mathbb{R}^{k}\). Then
\[
\forall j \in[n],\left\|\widehat{U}^{j} Z^{*} S A R-A^{j} R\right\|_{2}=\min _{x \in \mathbb{R}^{k}}\left\|x^{\top} Z^{*} S A R-A^{j} R\right\|_{2}
\]

Recall that \(r\) is the number of columns of \(R\). Due to Claim 17.3.5,
\[
\sum_{j=1}^{n}\left\|A^{j} R\left(\left(Z^{*} S A\right) R\right)^{\dagger} Z^{*} S A R-A^{j} R\right\|_{1} \leq O(\sqrt{r}) \sum_{j=1}^{n} \min _{U^{j} \in \mathbb{R}^{k}}\left\|U^{j} Z^{*} S A R-A^{j} R\right\|_{1}
\]
which is equivalent to
\[
\left\|A R\left(\left(Z^{*} S A\right) R\right)^{\dagger} Z^{*} S A R-A R\right\|_{1} \leq O(\sqrt{r}) \min _{U \in \mathbb{R}^{n \times k}}\left\|U Z^{*} S A R-A R\right\|_{1}
\]
where \(A R\) is an \(n \times r\) matrix and \(S A\) is an \(s \times d\) matrix.
Using Lemma 17.5.2, we obtain,
\[
\left\|A R\left(\left(Z^{*} S A\right) R\right)^{\dagger} Z^{*} S A-A\right\|_{1} \leq O\left(\sqrt{r} \alpha_{r}\right) \min _{U \in \mathbb{R}^{n \times k}}\left\|U Z^{*} S A-A\right\|_{1}
\]

We define \(X^{*} \in \mathbb{R}^{r \times k}, Y^{*} \in \mathbb{R}^{k \times s}\),
\[
X^{*}, Y^{*}=\underset{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}{\arg \min }\|A R X Y S A-A\|_{1} .
\]

Then,
\[
\begin{aligned}
\left\|A R X^{*} Y^{*} S A-A\right\|_{1} & \leq\left\|A R\left(\left(Z^{*} S A\right) R\right)^{\dagger} Z^{*} S A-A\right\|_{1} \\
& \leq O\left(\sqrt{r} \alpha_{r}\right) \min _{U \in \mathbb{R}^{n \times k}}\left\|U Z^{*} S A-A\right\|_{1} \\
& =O\left(\sqrt{r} \alpha_{r}\right) \min _{U \in \mathbb{R}^{n \times k}, Z \in \mathbb{R}^{k \times s}}\|U Z S A-A\|_{1} \\
& \leq O\left(\sqrt{r} \alpha_{r} \alpha_{s}\right) \mathrm{OPT} .
\end{aligned}
\]

It means that \(A R X^{*}, Y^{*} S A\) gives an \(O\left(\alpha_{r} \alpha_{s} \sqrt{r}\right)\)-approximation to the original problem.
Thus it suffices to use Lemma 17.4.4 to solve
\[
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\|A R X Y S A-A\|_{1}
\]
by losing an extra \(\operatorname{poly}(k) \log d\) factor in the approximation ratio.
By using a sparse Cauchy transform (for the place discussing the two options), combining the approximation ratios and running times all together, we can get \(\operatorname{poly}(k) \log (n) \log ^{2}(d)\) approximation ratio with \(O(n n z(A))+(n+d) \operatorname{poly}(k)\) running time. This completes the proof.

Lemma 17.4.4. Given matrices \(A \in \mathbb{R}^{n \times d}, S A \in \mathbb{R}^{s \times n}, R A \in \mathbb{R}^{r \times d}\) where \(S \in \mathbb{R}^{s \times n}\), \(R \in \mathbb{R}^{d \times r}\) with \(\min (n, d) \geq \max (r, s)\). For any \(1 \leq k \leq \min (r, s)\), there exists an algorithm that takes \(O(\mathrm{nnz}(A))+(n+d)\) poly \((s, r, k)\) time to output two matrices \(X^{\prime} \in \mathbb{R}^{r \times k}, Y^{\prime} \in \mathbb{R}^{k \times s}\) such that
\[
\left\|A R X^{\prime} \cdot Y^{\prime} S A-A\right\|_{1} \leq \operatorname{poly}(r, s) \log (d) \min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\|A R X Y S A-A\|_{1}
\]
holds with probability at least .999.

Proof. Choose sketching matrices \(T_{1} \in \mathbb{R}^{t_{1} \times n}\) to sketch on the left of ( \(A R X Y S A-A\) ) (note that \(S\) can be either a dense Cauchy transform matrix or a sparse Cauchy transform matrix). If \(T_{1}\) is a dense Cauchy transform matrix, then \(t_{1}=O(r \log r), \alpha_{t_{1}}=O(\log d)\), and computing \(T_{1} A\) takes \(O\left(t_{1} \cdot \operatorname{nnz}(A)\right)\) time. If \(T_{1}\) is a sparse Cauchy transform matrix, then \(t_{1}=\widetilde{O}\left(r^{5}\right)\), \(\alpha_{t_{1}}=\widetilde{O}\left(r^{2}\right) \log d\), and computing \(T_{1} A\) takes \(\mathrm{nnz}(A)\) time.

Choose dense Cauchy matrices \(T_{2}^{\top} \in \mathbb{R}^{t_{2} \times d}\) to sketch on the right of \(T_{1}(A R X Y S A-A)\) with \(t_{2}=O\left(\left(t_{1}+s\right) \log \left(t_{1}+s\right)\right)\). We get the following minimization problem,
\[
\begin{equation*}
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\left\|T_{1} A R X Y S A T_{2}-T_{1} A T_{2}\right\|_{1} . \tag{17.9}
\end{equation*}
\]

Define \(X^{\prime}, Y^{\prime}\) to be the optimal solution of
\[
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\left\|T_{1} A R X Y S A T_{2}-T_{1} A T_{2}\right\|_{F}
\]

Due to Claim 17.3.6,
\(\left\|T_{1} A R X^{\prime} Y^{\prime} S A T_{2}-T_{1} A T_{2}\right\|_{1} \leq \sqrt{t_{1} t_{2}} \min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\left\|T_{1} A R X Y S A T_{2}-T_{1} A T_{2}\right\|_{1}\).

Due to Lemma 17.5.4
\[
\left\|A R X^{\prime} Y^{\prime} S A-A\right\|_{1} \leq \sqrt{t_{1} t_{2}} \alpha_{t_{1}} \log t_{1} \min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\|A R X Y S A-A\|_{1}
\]

It remains to solve
\[
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\left\|T_{1} A R X Y S A T_{2}-T_{1} A T_{2}\right\|_{F}
\]

By using Theorem 17.3.4 and choosing \(T_{1}\) to be a sparse Cauchy transform matrix, we have that the optimal rank- \(k\) solution \(X^{\prime} Y^{\prime}\) is \(\left(T_{1} A R\right)^{\dagger}\left(U_{B} U_{B}^{\top}\left(T_{1} A T_{2}\right) V_{C} V_{C}^{\top}\right)_{k}\left(S A T_{2}\right)\) which can be computed in \(O(\operatorname{nnz}(A))+(n+d)\) poly \((s, r, k)\) time. Here, \(U_{B}\) are the left singular vectors of \(T_{1} A R . V_{C}\) are the right singular vectors of \(S A T_{2}\).

An alternative way of solving Equation (17.9) is using a polynomial system verifier. Note that a polynomial system verifier does not allow absolute value constraints. Using Claim 17.3.7, we are able to remove these absolute value constraints by introducing new constraints and variables. Thus, we can get a better approximation ratio but by spending exponential running time in \(k\). In the previous step, we should always use a dense Cauchy transform to optimize the approximation ratio.

Corollary 17.4.5. Given \(A \in \mathbb{R}^{n \times d}\), there exists an algorithm which takes \(n d \cdot \operatorname{poly}(k)+\) \((n+d) \cdot 2^{\widetilde{O}\left(k^{2}\right)}\) time and outputs two matrices \(U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}\) such that
\[
\|U V-A\|_{1} \leq O\left(\operatorname{poly}(k) \log n \log ^{2} d\right) \min _{\operatorname{rank}-k}\left\|A_{k}-A\right\|_{1}
\]
holds with probability 9/10.

The poly \((k)\) factor in the above corollary is much smaller than that in Theorem 17.4.3.

\subsection*{17.4.3 \(\operatorname{poly}(k, \log d)\)-approximation for an arbitrary matrix \(A\)}

In this section, we explain how to get an \(O(\log d) \cdot \operatorname{poly}(k)\) approximation.
```

Algorithm 17.3 poly $(k) \log d$-approximation Algorithm
procedure L1LowRankApproxPolykLogD $(A, n, d, k) \quad \triangleright$ Theorem 17.4.6
Set $s \leftarrow \widetilde{O}\left(k^{5}\right)$.
Choose sparse Cauchy matrices $S \in \mathbb{R}^{s \times n}$ and compute $S \cdot A$.
Implicitly obtain $B=U_{B} V_{B}$ by finding $V_{B}=S A \in \mathbb{R}^{s \times d}$ and $U_{B} \in \mathbb{R}^{n \times s}$ where
$\forall i \in[n]$, row vector $\left(U_{B}\right)^{i}$ gives an $O(1)$ approximation to $\min _{x \in \mathrm{R}^{1 \times s}}\left\|x S A-A^{i}\right\|_{1}$.
$U, V \leftarrow \mathrm{~L} 1 \mathrm{LowRAnkApproxB}\left(U_{B}, V_{B}, n, d, k, s\right) . \quad \triangleright$ Theorem 17.4.19
return $U, V$.
end procedure

```

Intuitively, our algorithm has two stages. In the first stage, we just want to find a low rank matrix \(B\) which is a good approximation to \(A\). Then, we can try to find a rank- \(k\) approximation to \(B\). Since now \(B\) is a low rank matrix, it is much easier to find a rank- \(k\) approximation to \(B\). The procedure L1LowRankApproxB \(\left(U_{B}, V_{B}, n, d, k, s\right)\) corresponds to Theorem 17.4.19.

Theorem 17.4.6. Given matrix \(A \in \mathbb{R}^{n \times d}\), for any \(k \geq 1\), there exists an algorithm which takes \(\operatorname{nnz}(A)+(n+d) \cdot \operatorname{poly}(k)\) time to output two matrices \(U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}\) such that
\[
\|U V-A\|_{1} \leq \operatorname{poly}(k) \log d \min _{\text {rank }-k}\left\|A_{k}-A\right\|_{1}
\]
holds with probability 9/10.

Proof. We define
\[
\mathrm{OPT}:=\min _{\operatorname{rank}-k A_{k}}\left\|A_{k}-A\right\|_{1} .
\]

The main idea is to replace the given \(n \times d\) matrix \(A\) with another low rank matrix \(B\) which also has size \(n \times d\). Choose \(S \in \mathbb{R}^{s \times n}\) to be a Cauchy matrix, where \(s \leq \operatorname{poly}(k)\) (note that, if \(S\) is a dense Cauchy transform matrix, computing \(S A\) takes \(O(s \mathrm{nnz}(A))\) time, while if \(S\) is a sparse Cauchy transform matrix, computing \(S A\) takes \(O(\mathrm{nnz}(A))\) time). Then \(B\) is obtained by taking each row of \(A\) and replacing it with its closest point (in \(\ell_{1}\)-distance) in the row span of \(S A\). By using Part II of Corollary 17.4.2, we have,
\[
\min _{U \in \mathbb{R}^{n \times k}, Z \in \mathbb{R}^{k \times s}}\|U Z S A-A\|_{1} \leq O(\sqrt{s} \text { poly }(k) \log d) \mathrm{OPT}
\]

We define \(B\) to be the product of two matrices \(U_{B} \in \mathbb{R}^{n \times s}\) and \(V_{B} \in \mathbb{R}^{s \times d}\). We define \(V_{B}\) to be \(S A\) and \(U_{B}\) to be such that for any \(i \in[n],\left(U_{B}\right)^{i}\) gives an \(O(1)\)-approximation to problem \(\min _{x \in \mathbb{R}^{1 \times s}}\left\|x S A-A^{i}\right\|_{1}\), i.e.,
\[
\left\|\left(U_{B}\right)^{i} S A-A^{i}\right\|_{1} \leq O(1) \min _{x \in \mathbb{R}^{1 \times s}}\left\|x S A-A^{i}\right\|_{1}, \forall i \in[n],
\]
which means
\[
\left\|U_{B} S A-A\right\|_{1} \leq O(1) \min _{X \in \mathbb{R}^{n \times s}}\|X S A-A\|_{1}
\]

For a fixed \(S A \in \mathbb{R}^{s \times d}\), we can compute \(D \in \mathbb{R}^{d \times d}\), which is a sampling and rescaling matrix corresponding to Lewis weights of \((S A)^{\top}\), and let \(m=O(s \log s)\) be the number of nonzero entries on the diagonal of \(D\).

Define \(\widehat{X}=\underset{X \in \mathbb{R}^{n \times s}}{\arg \min }\|X S A D-A D\|_{1}\), thus by Lemma 17.5.5 and Lemma 17.5.2, we have
\[
\|\widehat{X} S A-A\|_{1} \leq O(1) \min _{X \in \mathbb{R}^{n \times s}}\|X S A-A\|_{1} .
\]

Notice that computing Lewis weights takes \(d \operatorname{poly}(s)\) time. We can use \(\ell_{1}\)-regression solver and linear programming to find \(\widehat{X} \in \mathbb{R}^{n \times s}\) in \((n+d)\) poly \((s)\) time. Thus \(U_{B}\) can be found in \(O(\operatorname{nnz}(A))+(n+d)\) poly \((s)\) time.

By the definition of \(B\), it is an \(n \times d\) matrix. Naïvely we can write down \(B\) after finding \(U_{B}\) and \(V_{B}\). The time for writing down \(B\) is \(O(n d)\). To avoid this, we can just keep a factorization \(U_{B}\) and \(V_{B}\). We are still able to run algorithm L1LowRANKApproxB. Because \(s=\operatorname{poly}(k)\), the running time of algorithm L1LowRANKAPPROXB is still \(O(\mathrm{nnz}(A))+(n+\) d) \(\operatorname{poly}(k)\).

By the definition of \(B\), we have that \(B\) has rank at most \(s\). Suppose we then solve \(\ell_{1}\)-low rank approximation problem for rank- \(s\) matrix \(B\), finding a rank- \(k g\)-approximation matrix \(U V\). Due to Lemma 17.4.15, we have that if \(B\) is an \(f\)-approximation solution to \(A\), then \(U V\) is also an \(O(f g)\)-approximation solution to \(A\),
\[
\|U V-A\|_{1} \leq O(\log d) \cdot \operatorname{poly}(k) \cdot g \mathrm{OPT} .
\]

Using Theorem 17.4.19 we have that \(g=\operatorname{poly}(s)\), which completes the proof.

\subsection*{17.4.4 \(\widetilde{O}(k)\)-approximation for an arbitrary matrix \(A\)}
```

Algorithm 17.4 $\widetilde{O}(k)$-approximation Algorithm
procedure L1LowRANKApproxK $(A, n, d, k)$
$\triangleright$ Theorem 17.4.7
$r \leftarrow O(k \log k), m \leftarrow t_{1} \leftarrow O(r \log r), t_{2} \leftarrow O(m \log m)$.
Guess a diagonal matrix $R \in \mathbb{R}^{d \times d}$ with only $r$ s. $\quad$ s selects $r$ columns of
$A \in \mathbb{R}^{n \times d}$.
Compute a sampling and rescaling matrix $D \in \mathbb{R}^{n \times n}, T_{1} \in \mathbb{R}^{n \times n}$ corresponding to
the Lewis weights of $A R$, and let them have $m, t_{1}$ nonzero entries on the diagonals,
respectively.
Compute a sampling and rescaling matrix $T_{2}^{\top} \in \mathbb{R}^{d \times d}$ according to the Lewis weights
of $(D A)^{\top}$, and let it have $t_{2}$ nonzero entries on the diagonal.
Solve $\min _{X, Y}\left\|T_{1} A R X Y D A T_{2}-T_{1} A T_{2}\right\|_{1}$.
Take the best solution $X, Y$ over all guesses of $R$.
return $A R X, Y D A$.
end procedure

```

Theorem 17.4.7. Given matrix \(A \in \mathbb{R}^{n \times d}\), there exists an algorithm that takes poly \((n)\). \(d^{\widetilde{O}(k)} \cdot 2^{\widetilde{O}\left(k^{2}\right)}\) time and outputs two matrices \(U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}\) such that
\[
\|U V-A\|_{1} \leq \widetilde{O}(k) \min _{\mathrm{rank}-k A_{k}}\left\|A_{k}-A\right\|_{1}
\]
holds with probability 9/10.

Proof. We define
\[
\mathrm{OPT}:=\min _{\operatorname{rank}-k A_{k}}\left\|A_{k}-A\right\|_{1}
\]

Let \(U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times d}\) satisfy
\[
\left\|U^{*} V^{*}-A\right\|_{1}=\mathrm{OPT}
\]

Let \(S^{\top} \in \mathbb{R}^{d \times d}\) denote the sampling and rescaling matrix corresponding to the Lewis weights of \(\left(V^{*}\right)^{\top}\), where the number of nonzero entries on the diagonal of \(S\) is \(s=r=\) \(O(k \log k)\). Let \(R^{\top} \in \mathbb{R}^{d \times d}\) denote a diagonal matrix such that \(\forall i \in[d]\), if \(S_{i, i} \neq 0\), then \(R_{i, i}=1\), and if \(S_{i, i}=0\), then \(R_{i, i}=0\). Since rowspan \(\left(R^{\top} A^{\top}\right)=\operatorname{rowspan}\left(S^{\top} A^{\top}\right)\),
\[
\min _{Z \in \mathbb{R}^{m \times k}, V \in \mathbb{R}^{k \times d}}\|A S Z V-A\|_{1}=\min _{Z \in \mathbb{R}^{m \times k}, V \in \mathbb{R}^{k \times d}}\|A R Z V-A\|_{1} .
\]

Combining with Part III of Corollary 17.4.2, there exists a rank- \(k\) solution in the column span of \(A R\), which means,
\[
\begin{equation*}
\min _{Z \in \mathbb{R}^{m \times k}, V \in \mathbb{R}^{k \times d}}\|A R Z V-A\|_{1} \leq O(\sqrt{r}) \mathrm{OPT} \tag{17.10}
\end{equation*}
\]

Because the number of 1 s of \(R\) is \(r\), and the size of the matrix is \(d \times d\), there are \(\binom{d}{r}=d^{\widetilde{O}(k)}\) different choices for locations of 1 s on the diagonal of \(R\). We cannot compute \(R\) directly, but we can guess all the choices of locations of 1 s . Regarding \(R\) as selecting \(r\) columns of \(A\), then there are \(d^{\widetilde{O}(k)}\) choices. There must exist a "correct" way of selecting a subset of columns over all all choices. After trying all of them, we will have chosen the right one.

For a fixed guess \(R\), we can compute \(D \in \mathbb{R}^{n \times n}\), which is a sampling and rescaling matrix corresponding to the Lewis weights of \(A R\), and let \(m=O\left(k \log ^{2} k\right)\) be the number of nonzero entries on the diagonal of \(D\).

By Equation (17.10), there exists a \(W \in \mathbb{R}^{r \times k}\) such that,
\[
\begin{equation*}
\min _{V \in \mathbb{R}^{k \times d}}\|A R W V-A\|_{1} \leq O(\sqrt{r}) \mathrm{OPT} . \tag{17.11}
\end{equation*}
\]

We define \(\widehat{V}_{i}=\underset{V_{i} \in \mathbb{R}^{k \times d}}{\arg \min }\left\|D A R W V_{i}-D A_{i}\right\|_{2}, \forall i \in[d]\), which means \(\widehat{V}_{i}=(D A R W)^{\dagger} D A_{i} \in \mathbb{R}^{k}\). Then \(\widehat{V}=(D A R W)^{\dagger} D A \in \mathbb{R}^{k \times d}\). We define \(V^{*}=\underset{V \in \mathbb{R}^{k \times d}}{\arg \min }\|A R W V-A\|_{1}\). Then, by

Claim 17.3.5, it has
\[
\|D A R W \widehat{V}-D A\|_{1} \leq O(\sqrt{m}) \min _{V \in \mathbb{R}^{k \times d}}\|D A R W V-D A\|_{1}
\]

By applying Lemma 17.5.5, Lemma 17.5.2 and Equation (17.11), we can show
\[
\|A R W \widehat{V}-A\|_{1} \leq O(\sqrt{m})\left\|A R W V^{*}-A\right\|_{1} \leq O(\sqrt{m r}) \mathrm{OPT} \leq \widetilde{O}(k) \mathrm{OPT}
\]

Plugging \(\widehat{V}=(D A R W)^{\dagger} D A\) into \(\|A R W \widehat{V}-A\|_{1}\), we obtain that
\[
\left\|A R W(D A R W)^{\dagger} D A-A\right\|_{1} \leq \widetilde{O}(k) \mathrm{OPT}
\]
and it is clear that,
\[
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times m}}\|A R X Y D A-A\|_{1} \leq\left\|A R W(D A R W)^{\dagger} D A-A\right\|_{1} \leq \widetilde{O}(k) \mathrm{OPT}
\]

Recall that we guessed \(R\), so it is known. We can compute \(T_{1} \in \mathbb{R}^{n \times n}\), which is a sampling and rescaling diagonal matrix corresponding to the Lewis weights of \(A R\), and \(t_{1}=O(r \log r)\) is the number of nonzero entries on the diagonal of \(T_{1}\).

Also, \(D A\) is known, and the number of nonzero entries in \(D\) is \(m=O\left(k \log ^{2} k\right)\). We can compute \(T_{2}^{\top} \in \mathbb{R}^{d \times d}\), which is a sampling and rescaling matrix corresponding to the Lewis weights of \((D A)^{\top}\), and \(t_{2}=O(m \log m)\) is the number of nonzero entries on the diagonal of \(T_{2}\).
\[
\text { Define } X^{*}, Y^{*}=\underset{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times m}}{\arg \min }\left\|T_{1}(A R) X Y(D A) T_{2}-T_{1} A T_{2}\right\|_{1} \text {. Thus, using Lemma 17.5.4, }
\] we have
\[
\left\|(A R) X^{*} Y^{*}(D A)-A\right\|_{1} \leq \widetilde{O}(k) \mathrm{OPT}
\]

To find \(X^{*}, Y^{*}\), we need to solve this minimization problem
\[
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times m}}\left\|T_{1}(A R) X Y(D A) T_{2}-T_{1} A T_{2}\right\|_{1},
\]
which can be solved by a polynomial system verifier (see more discussion in Section 17.3.1 and 17.3.5).

In the next paragraphs, we explain how to solve the above problem by using a polynomial system verifier. Notice that \(T_{1}(A R)\) is known and \((D A) T_{2}\) is also known. First, we create \(r \times k\) variables for the matrix \(X\), i.e., one variable for each entry of \(X\). Second, we create \(k \times m\) variables for matrix \(Y\), i.e., one variable for each entry of \(Y\). Putting it all together and creating \(t_{1} \times t_{2}\) variables \(\sigma_{i, j}, \forall i \in\left[t_{1}\right], j \in\left[t_{2}\right]\) for handling the unknown signs, we write down the following optimization problem
\[
\begin{aligned}
\min _{X, Y} & \sum_{i=1}^{t_{1}} \sum_{j=1}^{t_{2}} \sigma_{i, j}\left(T_{1}(A R) X Y(D A) T_{2}\right)_{i, j} \\
\text { s.t } & \sigma_{i, j}^{2}=1, \forall i \in\left[t_{1}\right], j \in\left[t_{2}\right] \\
& \sigma_{i, j} \cdot\left(T_{1}(A R) X Y(D A) T_{2}\right)_{i, j} \geq 0, \forall i \in\left[t_{1}\right], j \in\left[t_{2}\right] .
\end{aligned}
\]

Notice that the number of constraints is \(O\left(t_{1} t_{2}\right)=\widetilde{O}\left(k^{2}\right)\), the maximum degree is \(O(1)\), and the number of variables \(O\left(t_{1} t_{2}+k m+r k\right)=\widetilde{O}\left(k^{2}\right)\). Thus the running time is,
\[
(\# \text { constraints } \cdot \# \text { degree })^{O(\# \text { variables })}=2^{\widetilde{O}\left(k^{2}\right)} .
\]

To use a polynomial system verifier, we need to discuss the bit complexity. Suppose that all entries are multiples of \(\delta\), and the maximum is \(\Delta\), i.e., each entry
\[
\in\{-\Delta, \cdots,-2 \delta,-\delta, 0, \delta, 2 \delta, \cdots, \Delta\}
\]
and \(\Delta / \delta=2^{\operatorname{poly}(n d)}\). Then the running time is \(O(\operatorname{poly}(\Delta / \delta)) \cdot 2^{\widetilde{O}\left(k^{2}\right)}=\operatorname{poly}(n d) 2^{\widetilde{O}\left(k^{2}\right)}\).
Also, a polynomial system verifier is able to tell us whether there exists a solution in a semi-algebraic set. In order to find the solution, we need to do a binary search over the cost \(C\). In each step of the binary search we use a polynomial system verifier to determine if there exists a solution in,
\[
\begin{aligned}
& \sum_{i=1}^{t_{1}} \sum_{j=1}^{t_{2}} \sigma_{i, j}\left(T_{1}(A R) X Y(D A) T_{2}\right)_{i, j} \leq C \\
& \sigma_{i, j}^{2}=1, \forall i \in\left[t_{1}\right], j \in\left[t_{2}\right] \\
& \sigma_{i, j} \cdot\left(T_{1}(A R) X Y(D A) T_{2}\right)_{i, j} \geq 0, \forall i \in\left[t_{1}\right], j \in\left[t_{2}\right]
\end{aligned}
\]

In order to do binary search over the cost, we need to know an upper bound on the cost and also a lower bound on the minimum nonzero cost. The upper bound on the cost is \(C_{\max }=O(n d \Delta)\), and the minimum nonzero cost is \(C_{\min }=2^{-\Omega(\operatorname{poly}(n d))}\). Thus, the total number of steps for binary search is \(O\left(\log \left(C_{\max } / C_{\min }\right)\right)\). Overall, the running time is
\[
n d^{\widetilde{O}(k)} \cdot 2^{\widetilde{O}\left(k^{2}\right)} \cdot \log (\Delta / \delta) \cdot \log \left(C_{\max } / C_{\min }\right)=\operatorname{poly}(n) d^{\widetilde{O}(k)} 2^{\widetilde{O}\left(k^{2}\right)}
\]

This completes the proof.

Instead of solving an \(\ell_{1}\) problem at the last step of L1LowRANkApproxK by using a polynomial system verifier, we can just solve a Frobenius norm minimization problem. This slightly improves the running time and pays an extra \(\operatorname{poly}(k)\) factor in the approximation ratio. Thus, we obtain the following corollary,

Corollary 17.4.8. Given matrix \(A \in \mathbb{R}^{n \times d}\), there exists an algorithm that takes poly \((n) \cdot d^{\widetilde{O}(k)}\) time which outputs two matrices \(U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}\) such that
\[
\mid U V-A\left\|_{1} \leq \operatorname{poly}(k) \min _{\text {rank }-k}\right\| A_{k}-A \|_{1}
\]
holds with probability 9/10.

\subsection*{17.4.5 Rank-3k and \(O(1)\)-approximation algorithm for an arbitrary matrix \(A\)}

In this section, we show how to output a rank- \(3 k\) solution that is able to achieve an \(O(1)\)-approximation.
```

Algorithm 17.5 Bicriteria $O(1)$-approximation Algorithm
procedure L1LowRankApproxBicriteria $(A, n, d, k) \quad \triangleright$ Theorem 17.4.9
$U_{B}, V_{B} \leftarrow \min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|U V-A\|_{F}$.
$r \leftarrow O(k \log k)$.
Guess a diagonal matrix $D \in \mathbb{R}^{n \times n}$ with $r$ nonzero entries.
Guess matrix $D U \in \mathbb{R}^{r \times 2 k}$.
Find $V_{A}$ by solving $\min _{V}\|D U V-D(A-B)\|_{1}$.
Find $U_{A}$ by solving $\min _{U}\left\|U V_{A}-(A-B)\right\|_{1}$.
Take the best solution $\left[\begin{array}{ll}U_{A} & U_{B}\end{array}\right],\left[\begin{array}{l}V_{A} \\ V_{B}\end{array}\right]$ over all guesses.
return $\left[\begin{array}{ll}U_{A} & U_{B}\end{array}\right],\left[\begin{array}{l}V_{A} \\ V_{B}\end{array}\right]$.
end procedure

```

Theorem 17.4.9. Given matrix \(A \in \mathbb{R}^{n \times d}\), for any \(k \geq 1\), there exists an algorithm which takes \((n d)^{\widetilde{O}\left(k^{2}\right)}\) time to output two matrices \(U \in \mathbb{R}^{n \times 3 k}, V \in \mathbb{R}^{3 k \times d}\),
\[
\|U V-A\|_{1} \lesssim \min _{\text {rank }-k}\left\|A_{k}-A\right\|_{1}
\]
holds with probability 9/10.

Proof. We define OPT to be
\[
\underset{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}{\arg \min }\|U V-A\|_{1}
\]

Solving the Frobenius norm problem, we can find a factorization of a rank- \(k\) matrix \(B=U_{B} V_{B}\) where \(U_{B} \in \mathbb{R}^{n \times k}, V_{B} \in \mathbb{R}^{k \times d}\), and \(B\) satisfies \(\|A-B\|_{1} \leq \alpha_{B}\) OPT for an \(\alpha_{B}=\sqrt{n d}\).

We define \(U^{*} \in \mathbb{R}^{n \times 2 k}, V^{*} \in \mathbb{R}^{2 k \times d}\) to be the optimal solution, i.e.,
\[
U^{*} V^{*}=\underset{U \in \mathbb{R}^{n \times 2 k}, V \in \mathbb{R}^{2 k \times d}}{\arg \min }\|U V-(A-B)\|_{1} .
\]

Let \(D\) be a sampling and rescaling diagonal matrix corresponding to the Lewis weights of \(U^{*}\), and let the number of nonzero entries on the diagonal be \(t=O(k \log k)\).

By Lemma 17.5.5 and Lemma 17.5.2, the solution of \(\min _{V \in \mathbb{R}^{k \times d}}\left\|D U^{*} V-D(A-B)\right\|_{1}\) together with \(U^{*}\) gives an \(O(1)\)-approximation to \(A-B\). In order to compute \(D\) we need to know \(U^{*}\). Although we do not know \(U^{*}\), there still exists a way to figure out the Lewis weights. The idea is the same as in the previous discussion Lemma 17.4.10 "Guessing Lewis weights". By Claim 17.4.11 and Claim 17.4.12, the total number of possible \(D\) is \(n^{O(t)}\).

Lemma 17.4.10 tries to find a rank- \(k\) solution when all the entries in \(A\) are integers at most poly \((n d)\). Here we focus on a bicriteria algorithm that outputs a rank- \(3 k\) matrix without such a strong bit complexity assumption. We can show a better claim 17.4.13, which is that the total number of possible \(D U^{*}\) is \(N^{\widetilde{O}\left(k^{2}\right)}\) where \(N=\operatorname{poly}(n)\) is the number of choices for a single entry in \(U^{*}\).

We explain how to obtain an upper bound on \(N\). Consider the optimum \(\| U^{*} V^{*}-\) \((A-B) \|_{1}\). We can always change the basis so assume \(V^{*}\) is an Auerbach basis (i.e., an \(\ell_{1}\)-well-conditioned basis discussed in Section 17.1), so \(e\|x\|_{1} \geq\left\|x V^{*}\right\|_{1} \geq\|x\|_{1} / f\), where \(e, f=\operatorname{poly}(k)\). Then no entry of \(U^{*}\) is larger than \(2 f\|A-B\|_{1}\), otherwise we could replace \(U^{*}\) with 0 and get a better solution. Also any entry smaller than \(\|A-B\|_{1} /\left(e n k \alpha_{B} 100\right)\) can be replaced with 0 as this will incur additive error at most OPT /100. So if we round to integer multiples of \(\|A-B\|_{1} /\left(e n k \alpha_{B} 100\right)\) we only have \(O\left(e n k \alpha_{B} f\right)\) possibilities for each
entry of \(U^{*}\) and still have an \(O(1)\)-approximation. We will just refer to this rounded \(U^{*}\) as \(U^{*}\), abusing notation.

Let \(\mathcal{U}\) denote the set of all the matrices \(U\) that we guess. From the above discussion, we conclude that, there exists a \(U \in \mathcal{U}\) such that \(\left\|U V^{*}-(A-B)\right\|_{1} \leq O(\mathrm{OPT})\).

For each guess of \(D U^{*}\) and \(D\), we find \(V_{A}, U_{A}\) in the following way. We find \(V_{A}\) by using a linear program to solve,
\[
\min _{V \in \mathbb{R}^{k \times d}}\left\|D U^{*} V-D(A-B)\right\|_{1}
\]

Given \(V_{A}\) and \(A\), we write down a linear program to solve this problem,
\[
\min _{U \in \mathbb{R}^{n \times k}}\left\|U V_{A}-(A-B)\right\|_{1},
\]
which takes \(\operatorname{poly}(n d k)\) time. Then we obtain \(U_{A}\).
Recall that \(V_{B}, U_{B}\) are the two factors of \(B\) and it is a rank- \(k, \alpha_{B}\)-approximation solution to \(\min _{U, V}\|U V-A\|_{1}\). Then we have
\[
\left\|\left[\begin{array}{ll}
U_{A} & U_{B}
\end{array}\right]\left[\begin{array}{l}
V_{A} \\
V_{B}
\end{array}\right]-A\right\|_{1}=\left\|U_{A} V_{A}-\left(A-U_{B} V_{B}\right)\right\|_{1}=\left\|U_{A} V_{A}-(A-B)\right\|_{1}
\]

Because there must exist a pair \(U_{A}, V_{A}\) satisfying \(\left\|U_{A} V_{A}-(A-B)\right\|_{1} \leq O(\mathrm{OPT})\), it follows that by taking the best solution \(\left[\begin{array}{ll}U_{A} & U_{B}\end{array}\right]\left[\begin{array}{l}V_{A} \\ V_{B}\end{array}\right]\) over all guesses, we obtain an \(O(1)\) approximation solution.

Overall, the running time is \((n d)^{\widetilde{O}\left(k^{2}\right)}\).
Lemma 17.4.10. Given an \(n \times d\) matrix \(A\) with integers bounded by \(\operatorname{poly}(n)\), for any \(k \geq 1\), there exists an algorithm which takes \((n d)^{\widetilde{O}\left(k^{3}\right)}\) time to output two matrices \(U \in \mathbb{R}^{n \times k}\),
\(V \in \mathbb{R}^{k \times d}\), such that
\[
\|U V-A\|_{1} \lesssim \min _{\text {rank }-k}\left\|A_{k}-A\right\|_{1}
\]

Proof. We define \(U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times d}\) to be the optimal solution, i.e.,
\[
U^{*}, V^{*}=\underset{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}{\arg \min }\|U V-A\|_{1},
\]
and define OPT \(=\left\|U^{*} V^{*}-A\right\|_{1}\).
Let \(D\) denote a sampling and rescaling diagonal matrix corresponding to the Lewis weights of \(U^{*}\), and let the number of nonzero entries on the diagonal be \(t=O(k \log k)\).

By Lemma 17.5.5 and Lemma 17.5.2, the solution to \(\min _{V \in \mathbb{R}^{2 k \times d}}\left\|D U^{*} V-D A\right\|_{1}\) together with \(U^{*}\) gives an \(O(1)\)-approximation to \(A\). In order to compute \(D\), we need to know \(U^{*}\). Although we do not know \(U^{*}\), there still exists a way to figure out the Lewis weights. We call the idea "Guessing Lewis weights". We will explain this idea in the next few paragraphs.

First, we can guess the nonzero entries on the diagonal, because the number of choices is small.

Claim 17.4.11. The number of possible choice of \(\operatorname{supp}(D)\) is at most \(n^{O(t)}\).

Proof. The matrix has dimension \(n \times n\) and the number of nonzero entries is \(t\). Thus the number of possible choices is at most \(\sum_{i=1}^{t}\binom{n}{t}=n^{O(t)}\).

Second, we can guess the value of each probability. For each probability, it is trivially at most 1. If the probability is less than \(1 /(\operatorname{poly}(n) k \log k)\), then we will never sample that row with high probability. It means that we can truncate the probability if it is below that
threshold. We can also round each probability to \(2^{-i}\) which only loses another constant factor in the approximation ratio. Thus, we have:

Claim 17.4.12. The total number of possible \(D\) is \(n^{O(t)}=n^{\widetilde{O}(k)}\).

Since \(\Delta / \delta \leq \operatorname{poly}(n)\) and the entries of \(A\) are in
\[
\{-\Delta,-\Delta+\delta, \ldots,-2 \delta,-\delta, 0, \delta, 2 \delta, \cdots, \Delta-\delta, \Delta\}
\]
we can lower bound the cost of \(\left\|U^{*} V-A\right\|_{1}\) given that it is non-zero by \((n d \Delta / \delta)^{-O(k)}\) (if it is zero then \(A\) has rank at most \(k\) and we output \(A\) ) using Lemma 4.1 in [CW09] and relating entrywise \(\ell_{1}\)-norm to Frobenius norm. We can assume \(V\) is an \(\ell_{1}\) well-conditioned basis, since we can replace \(U^{*}\) with \(U^{*} R^{-1}\) and \(V\) with \(R V\) for any invertible linear transformation \(R\). By properties of such basis, we can discretize the entries of \(U^{*}\) to integer multiples of \((n d \Delta / \delta)^{-O(k)}\) while preserving relative error. Hence we can correctly guess each entry of \(D U^{*}\) in \(\left(n^{O(k)}\right)\) time.

Claim 17.4.13. The total number of possible \(D U^{*}\) is \(n^{\widetilde{O}\left(k^{3}\right)}\).

In the following, let \(D U\) denote a guess of \(D U^{*}\). Now the problem remaining is to solve \(\min _{V \in \mathbb{R}^{k \times d}}\|D U V-D A\|_{1}\). Since we already know \(D A\) can be computed, and we know \(D U\), we can solve this multiple regression problem by running linear programming. Thus the running time of this step is in poly \((n d)\). After we get such a solution \(V\), we use a linear program to solve \(\min _{U \in \mathbb{R}^{n \times k}}\|D U V-D A\|_{1}\). Then we can get \(U\).

After we guess all the choices of \(D\) and \(D U^{*}\), we must find a solution \(U, V\) which gives an \(O(1)\) approximation. The total running time is \(n^{\widetilde{O}(k)} \cdot n^{\widetilde{O}\left(k^{3}\right)} \cdot \operatorname{poly}(n d)=(n d)^{\widetilde{O}\left(k^{3}\right)}\).

\subsection*{17.4.6 CUR decomposition for an arbitrary matrix \(A\)}
```

Algorithm 17.6 CUR Decomposition Algorithm
procedure L1LowRAnkApproxCUR $(A, n, d, k) \quad \triangleright$ Theorem 17.4.14
$U_{B}, V_{B} \leftarrow$ L1LowRankApproxPolykLogd $(A, n, d, k)$.
Let $D_{1} \in \mathbb{R}^{n \times n}$ be the sampling and rescaling diagonal matrix corresponding to the
Lewis weights of $B_{1}=U_{B} \in \mathbb{R}^{n \times k}$, and let $D_{1}$ have $d_{1}=O(k \log k)$ nonzero entries.
Let $D_{2}^{\top} \in \mathbb{R}^{d \times d}$ be the sampling and rescaling diagonal matrix corresponding to the
Lewis weights of $B_{2}^{\top}=\left(\left(D_{1} B_{1}\right)^{\dagger} D_{1} A\right)^{\top} \in \mathbb{R}^{d \times k}$, and let $D_{2}$ have $d_{2}=O(k \log k)$ nonzero
entries.
$C \leftarrow A D_{2}, U \leftarrow\left(B_{2} D_{2}\right)^{\dagger}\left(D_{1} B_{1}\right)^{\dagger}$, and $R \leftarrow D_{1} A$.
return $C, U, R$.
end procedure

```

Theorem 17.4.14. Given matrix \(A \in \mathbb{R}^{n \times d}\), for any \(k \geq 1\), there exists an algorithm which takes \(O(\operatorname{nnz}(A))+(n+d)\) poly \((k)\) time to output three matrices \(C \in \mathbb{R}^{n \times c}\) with columns from \(A, U \in \mathbb{R}^{c \times r}\), and \(R \in \mathbb{R}^{r \times d}\) with rows from \(A\), such that \(\operatorname{rank}(C U R)=k, c=O(k \log k)\), \(r=O(k \log k)\), and
\[
\|C U R-A\|_{1} \leq \operatorname{poly}(k) \log d \min _{\mathrm{rank}-k}\left\|A_{k}-A\right\|_{1},
\]
holds with probability 9/10.

Proof. We define
\[
\mathrm{OPT}:=\min _{\operatorname{rank}-k A_{k}}\left\|A_{k}-A\right\|_{1}
\]

Due to Theorem 17.4.6, we can output two matrices \(U_{B} \in \mathbb{R}^{n \times k}, V_{B} \in \mathbb{R}^{k \times d}\) such that \(U_{B} V_{B}\) gives a rank- \(k\), and \(\operatorname{poly}(k) \log d\)-approximation solution to \(A\), i.e.,
\[
\begin{equation*}
\left\|U_{B} V_{B}-A\right\|_{1} \leq \operatorname{poly}(k) \log d \mathrm{OPT} \tag{17.12}
\end{equation*}
\]

By Section 17.3.3, we can compute \(D_{1} \in \mathbb{R}^{n \times n}\) which is a sampling and rescaling matrix corresponding to the Lewis weights of \(B_{1}=U_{B}\) in \(O(n\) poly \((k))\) time, and there are \(d_{1}=\) \(O(k \log k)\) nonzero entries on the diagonal of \(D_{1}\).

Define \(V^{*} \in \mathbb{R}^{k \times d}\) to be the optimal solution of \(\min _{V \in \mathbb{R}^{k \times d}}\left\|B_{1} V-A\right\|_{1}, \widehat{V}=\left(D_{1} B_{1}\right)^{\dagger} D_{1} A \in\) \(\mathbb{R}^{k \times d}, U_{1} \in \mathbb{R}^{n \times k}\) to be the optimal solution of \(\min _{U \in \mathbb{R}^{n \times k}}\|U \widehat{V}-A\|_{1}\), and \(V^{\prime}\) to be the optimal solution of \(\min _{V \in \mathbb{R}^{k \times d}}\left\|D_{1} A-D_{1} B_{1} V\right\|_{1}\).

By Claim 17.3.5, we have
\[
\left\|D_{1} B_{1} \widehat{V}-D_{1} A\right\|_{1} \leq \sqrt{d_{1}}\left\|D_{1} B_{1} V^{\prime}-D_{1} A\right\|_{1}
\]

Due to Lemma 17.5.5 and Lemma 17.5.2, with constant probability, we have
\[
\left\|B_{1} \widehat{V}-A\right\|_{1} \leq \sqrt{d_{1}} \alpha_{D_{1}}\left\|B_{1} V^{*}-A\right\|_{1}
\]
where \(\alpha_{D_{1}}=O(1)\).
Now, we can show,
\[
\begin{array}{rlr}
\left\|U_{1} \widehat{V}-A\right\|_{1} & \leq\left\|B_{1} \widehat{V}-A\right\|_{1} & \text { by } U_{1}=\underset{U \in \mathbb{R}^{n \times k}}{\arg \min }\|U \widehat{V}-A\|_{1} \\
& \lesssim \sqrt{d_{1}}\left\|B_{1} V^{*}-A\right\|_{1} & \\
& \leq \sqrt{d_{1}}\left\|U_{B} V_{B}-A\right\|_{1} & \\
& \leq \operatorname{poly}(k) \log d \mathrm{OPT} . & \text { by Equation }(17.12) \tag{17.13}
\end{array}
\]

We define \(B_{2}=\widehat{V}\), then we replace \(\widehat{V}\) by \(B_{2} \in \mathbb{R}^{k \times d}\) and look at this objective function,
\[
\min _{U \in \mathbb{R}^{n \times k}}\left\|U B_{2}-A\right\|_{1},
\]
where \(U^{*}\) denotes the optimal solution. We use a sketching matrix to sketch the RHS of matrix \(U B_{2}-A \in \mathbb{R}^{n \times d}\). Let \(D_{2}^{\top} \in \mathbb{R}^{d \times d}\) denote a sampling and rescaling diagonal matrix corresponding to the Lewis weights of \(B_{2}^{\top} \in \mathbb{R}^{d \times k}\), and let the number of nonzero entries on the diagonal of \(D_{2}\) be \(d_{2}=O(k \log k)\). We define \(\widehat{U}=A D_{2}\left(B_{2} D_{2}\right)^{\dagger} \in \mathbb{R}^{n \times k}, U^{\prime} \in \mathbb{R}^{n \times k}\) to be the optimal solution of \(\min _{U \in \mathbb{R}^{n \times k}}\left\|\left(U B_{2}-A\right) D_{2}\right\|_{1}\). Recall that \(U_{1} \in \mathbb{R}^{n \times k}\) is the optimal of \(\min _{U \in \mathbb{R}^{n \times k}}\left\|U B_{2}-A\right\|_{1}\).

By Claim 17.3.5, we have
\[
\left\|\widehat{U} B_{2} D_{2}-A D_{2}\right\|_{1} \leq \sqrt{d_{2}}\left\|U^{\prime} B_{2} D_{2}-A D_{2}\right\|_{1}
\]

According to Lemma 17.5.2 and Lemma 17.5.5, with constant probability,
\[
\left\|\widehat{U} B_{2}-A\right\|_{1} \leq \alpha_{D_{2}} \sqrt{d_{2}}\left\|U_{1} B_{2}-A\right\|_{1},
\]
where \(\alpha_{B_{2}}=O(1)\).
We have
\[
\begin{aligned}
& \left\|\widehat{U} B_{2}-A\right\|_{1} \\
\leq & \sqrt{d_{2}} \alpha_{D_{2}}\left\|U_{1} B_{2}-A\right\|_{1} \\
= & \sqrt{d_{2}} \alpha_{D_{2}}\left\|U_{1} \widehat{V}-A\right\|_{1} \\
\leq & \operatorname{poly}(k) \log (d) \text { OPT } .
\end{aligned} \quad \text { by } B_{2}=\widehat{V}
\]

Notice that \(\widehat{U} B_{2}=A D_{2}\left(B_{2} D_{2}\right)^{\dagger}\left(D_{1} B_{1}\right)^{\dagger} D_{1} A\). Setting
\[
C=A D_{2} \in \mathbb{R}^{n \times d_{2}}, U=\left(B_{2} D_{2}\right)^{\dagger}\left(D_{1} B_{1}\right)^{\dagger} \in \mathbb{R}^{d_{2} \times d_{1}}, \text { and } R=D_{1} A \in \mathbb{R}^{d_{1} \times d}
\]
we get the desired CUR decomposition,
\[
\|\underbrace{A D_{2}}_{C} \cdot \underbrace{\left(B_{2} D_{2}\right)^{\dagger}\left(D_{1} B_{1}\right)^{\dagger}}_{U} \cdot \underbrace{D_{1} A}_{R}-A\|_{1} \leq \operatorname{poly}(k) \log (d) \mathrm{OPT} .
\]
with \(\operatorname{rank}(C U R)=k\). Overall, the running time is \(O(\mathrm{nnz}(A))+(n+d) \operatorname{poly}(k)\).

\subsection*{17.4.7 Rank- \(r\) matrix \(B\)}

\subsection*{17.4.7.1 Properties}

Lemma 17.4.15. Given matrix \(A \in \mathbb{R}^{n \times d}\), let \(\mathrm{OPT}=\min _{\operatorname{rank}-k}\left\|A-A_{k}\right\|_{1}\). For any \(r \geq k\), if rank-r matrix \(B \in \mathbb{R}^{n \times d}\) is an \(f\)-approximation to \(A\), i.e.,
\[
\|B-A\|_{1} \leq f \cdot \mathrm{OPT}
\]
and \(U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}\) is a \(g\)-approximation to \(B\), i.e.,
\[
\|U V-B\|_{1} \leq g \cdot \min _{\text {rank }-k}\left\|B_{k}-B\right\|_{1}
\]
then,
\[
\|U V-A\|_{1} \lesssim g f \cdot \mathrm{OPT}
\]

Proof. We define \(\widetilde{U} \in \mathbb{R}^{n \times k}, \widetilde{V} \in \mathbb{R}^{k \times d}\) to be two matrices, such that
\[
\|\widetilde{U} \widetilde{V}-B\|_{1} \leq g \min _{\operatorname{rank}-k B_{k}}\left\|B_{k}-B\right\|_{1}
\]
and also define,
\[
\widehat{U}, \widehat{V}=\underset{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}{\arg \min }\|U V-B\|_{1} \text { and } U^{*}, V^{*}=\underset{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}{\arg \min }\|U V-A\|_{1},
\]

Then,
\[
\begin{array}{rlr}
\|\tilde{U} \tilde{V}-A\|_{1} & \leq\|\widetilde{U} \widetilde{V}-B\|_{1}+\|B-A\|_{1} & \text { by triangle inequality } \\
& \leq g\|\widehat{U} \widehat{V}-B\|_{1}+\|B-A\|_{1} & \text { by definition } \\
& \leq g\left\|U^{*} V^{*}-B\right\|_{1}+\|B-A\|_{1} & \text { by }\|\widehat{U} \widehat{V}-B\|_{1} \leq\left\|U^{*} V^{*}-B\right\|_{1} \\
& \leq g\left\|U^{*} V^{*}-A\right\|_{1}+g\|B-A\|_{1}+\|B-A\|_{1} & \\
& =g \mathrm{OPT}+(g+1)\|B-A\|_{1} & \\
& & \text { by triangle inequality } \\
& \leq g \mathrm{OPT}+(g+1) f \cdot \mathrm{OPT} & \text { by } B \text { is } f \text {-approximation of OPT } \\
& \lesssim g f \mathrm{OPT} . &
\end{array}
\]

This completes the proof.

Lemma 17.4.16. Given a matrix \(B \in \mathbb{R}^{n \times d}\) with rank \(r\), for any \(1 \leq k<r\), for any fixed \(U^{*} \in \mathbb{R}^{n \times k}\), choose a Cauchy matrix \(S\) with \(m=O(r \log r)\) rows and rescaled by \(\Theta(1 / m)\). With probability . 999 for all \(V \in \mathbb{R}^{k \times d}\), we have
\[
\left\|S U^{*} V-S B\right\|_{1} \geq\left\|U^{*} V-B\right\|_{1}
\]

Proof. This follows by definitions in Section 17.5 and Lemma 17.5.15.

Lemma 17.4.17. Given a matrix \(B \in \mathbb{R}^{n \times d}\) with rank \(r\), for any \(1 \leq k<r\), for any fixed \(U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times d}\), choose a Cauchy matrix \(S\) with \(m\) rows and rescaled by \(\Theta(1 / m)\). We have
\[
\left\|S U^{*} V^{*}-S B\right\|_{1} \leq O(r \log r)\left\|U^{*} V^{*}-B\right\|_{1}
\]
with probability . 999 .

Proof. Let \(U\) denote a well-conditioned basis of \(\left[U^{*} V^{*}, B\right]\), then \(U\) has \(\widetilde{d}=O(r)\) columns. We have
\[
\begin{array}{rlr}
\|S U\|_{1} & =\sum_{i=1}^{\widetilde{d}} \sum_{j=1}^{m}\left|\left(S U_{i}\right)_{j}\right| & \\
& =\sum_{i=1}^{\tilde{d}} \sum_{j=1}^{m}\left|\frac{1}{m} \sum_{l=1}^{n} S_{j, l} U_{l, i}\right| & \\
& =\frac{1}{m} \sum_{i=1}^{\tilde{d}} \sum_{j=1}^{m}\left|c_{i, j}\right| & \text { by } S_{j, l} \sim C(0,1) \\
& =\frac{1}{m} \sum_{i=1}^{\widetilde{d}} \sum_{j=1}^{m}\left\|U_{i}\right\|_{1} \cdot w_{i+(j-1) d}, &
\end{array}
\]
where the last step follows since each \(w_{i}\) can be thought of as a clipped half-Cauchy random variable. Define \(d^{\prime}=m \widetilde{d}\). Define event \(\xi_{i}\) to be the situation when \(w_{i}<D\) (we will choose \(D\) later), and define event \(\xi=\xi_{1} \cap \xi_{2} \cap \cdots \cap \xi_{d^{\prime}}\). Using a similar proof as Lemma 17.5.6, which is also similar to previous work [Ind06, SW11, \(\mathrm{CDMI}^{+}\)13], we obtain that
\[
\operatorname{Pr}\left[\sum_{i=1}^{m}\left\|S U_{i}\right\|_{1} \geq \sum_{i=1}^{m}\left\|U_{i}\right\|_{1} t\right] \lesssim \frac{\log d^{\prime}}{t}+\frac{d^{\prime}}{D}
\]

Choosing \(t=\Theta\left(\log d^{\prime}\right)\) and \(D=\Theta\left(d^{\prime}\right)\), we have
\[
\operatorname{Pr}\left[\sum_{i=1}^{m}\left\|S U_{i}\right\|_{1} \geq \sum_{i=1}^{m}\left\|U_{i}\right\|_{1} O\left(\log d^{\prime}\right)\right] \leq \frac{1}{C}
\]
for a constant \(C\). Condition on the above event. Let \(y=U x\), for some \(x \in \mathbb{R}^{d}\). Then for
any \(y\),
\[
\begin{array}{rlr}
\|S y\|_{1} & =\|S U x\|_{1} & \\
& \leq \sum_{j=1}^{d^{\prime}}\left\|S U_{j} x_{j}\right\|_{1} & \text { by triangle inequality } \\
& =\sum_{j=1}^{d^{\prime}}\left|x_{j}\right| \cdot\left\|S U_{j}\right\|_{1} & \\
& \lesssim\|x\|_{\infty} \log \left(d^{\prime}\right) \sum_{j=1}^{d^{\prime}}\left\|U_{j}\right\|_{1} & \\
& \lesssim r \log r\|y\|_{1} &
\end{array}
\]
where the last step follows by \(\sum_{j=1}^{d^{\prime}}\left\|U_{j}\right\|_{1} \leq d^{\prime}\) and \(\|x\|_{\infty} \leq\|U x\|_{1}=\|y\|_{1}\). Choosing \(C=1000\) completes the proof.

Lemma 17.4.18. Given a matrix \(M \in \mathbb{R}^{n \times d}\) with rank \(O(r)\), choose a random matrix \(S \in \mathbb{R}^{m \times n}\) with each entry drawn from a standard Cauchy distribution and scaled by \(\Theta(1 / m)\). We have that
\[
\|S M\|_{1} \leq O(r \log r)\|M\|_{1},
\]
holds with probability . 999 .

Proof. Let \(U \in \mathbb{R}^{O(r)}\) be the well-conditioned basis of \(M\). Then each column of \(M\) can be expressed by \(U x\) for some \(x\). We then follow the same proof as that of Lemma 17.4.17.

\subsection*{17.4.7.2 \(\operatorname{poly}(k, r)\)-approximation for rank- \(r\) matrix \(B\)}

Theorem 17.4.19. Given a factorization of a rank-r matrix \(B=U_{B} V_{B} \in \mathbb{R}^{n \times d}\), where \(U_{B} \in\) \(\mathbb{R}^{n \times r}, V_{B} \in \mathbb{R}^{r \times d}\), for any \(1 \leq k \leq r\) there exists an algorithm which takes \((n+d) \cdot \operatorname{poly}(k)\)
```

Algorithm 17.7 poly $(r, k)$-approximation Algorithm for Rank- $r$ matrix $B$
procedure L1LowRANkApproxB $\left(U_{B}, V_{B}, n, d, k, r\right)$
$\triangleright$ Theorem 17.4.19
Set $s \leftarrow \widetilde{O}(r), r^{\prime} \leftarrow \widetilde{O}(r), t_{1} \leftarrow \widetilde{O}(r), t_{2} \leftarrow \widetilde{O}(r)$.
Choose dense Cauchy matrices $S \in \mathbb{R}^{s \times n}, R \in \mathbb{R}^{d \times r^{\prime}}, T_{1} \in \mathbb{R}^{t_{1} \times n}, T_{2} \in \mathbb{R}^{d \times t_{2}}$.
Compute $S \cdot U_{B} \cdot V_{B}, U_{B} \cdot V_{B} \cdot R$ and $T_{1} \cdot U_{B} \cdot V_{B} \cdot T_{2}$.
Compute $X Y=\arg \min _{X, Y}\left\|T_{1} U_{B} V_{B} R X Y S U_{B} V_{B} T_{2}-T_{1} U_{B} V_{B} T_{2}\right\|_{F}$.
return $U_{B} V_{B} R X, Y S U_{B} V_{B}$.
end procedure

```
time to output two matrices \(U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}\) such that
\[
\|U V-B\|_{1} \leq \operatorname{poly}(r) \min _{\text {rank }-k}\left\|B_{k}-B\right\|_{1},
\]
holds with probability 9/10.

Proof. We define
\[
\mathrm{OPT}=\min _{\text {rank }-k}\left\|B_{k}-B\right\|_{1}
\]

Choose \(S \in \mathbb{R}^{s \times n}\) to be a dense Cauchy transform matrix with \(s=O(r \log r)\). Using Lemma 17.4.18, Lemma 17.5.15, and combining with Equation (17.8), we have
\[
\min _{U \in \mathbb{R}^{n \times k}, Z \in \mathbb{R}^{k \times s}}\|U Z S B-B\|_{1} \leq \sqrt{s} O(r \log r) \mathrm{OPT}=O\left(r^{1.5} \log ^{1.5} r\right) \mathrm{OPT}
\]

Let \(\alpha_{s}=O\left(r^{1.5} \log ^{1.5} r\right)\).
We define \(U^{*}, Z^{*}=\underset{U \in \mathbb{R}^{n \times k}, Z \in \mathbb{R}^{k \times d}}{\arg \min }\|U Z S B-B\|_{1}\). For the fixed \(Z^{*} \in \mathbb{R}^{k \times s}\), choose a dense Cauchy transform matrix \(R \in \mathbb{R}^{d \times r^{\prime}}\) with \(r^{\prime}=O(r \log r)\) and sketch on the right of \((U Z S B-B)\). We obtain the minimization problem, \(\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z^{*} S B R-B R\right\|_{1}\).

Define \(\widehat{U}^{j}=B^{j} R\left(\left(Z^{*} S B\right) R\right)^{\dagger} \in \mathbb{R}^{k}, \forall j \in[n]\). Then \(\widehat{U}=B R\left(\left(Z^{*} S B\right) R\right)^{\dagger} \in \mathbb{R}^{n \times k}\). Due to Claim 17.3.5,
\[
\sum_{j=1}^{n}\left\|B^{j} R\left(\left(Z^{*} S B\right) R\right)^{\dagger} Z^{*} S B R-B^{j} R\right\|_{1} \leq O\left(\sqrt{r^{\prime}}\right) \sum_{j=1}^{n} \min _{U^{j} \in \mathbb{R}^{k}}\left\|U^{j} Z^{*} S B R-B^{j} R\right\|_{1}
\]
which is equivalent to
\[
\left\|B R\left(\left(Z^{*} S B\right) R\right)^{\dagger} Z^{*} S B R-B R\right\|_{1} \leq O\left(\sqrt{r^{\prime}}\right) \min _{U \in \mathbb{R}^{n \times k}}\left\|U Z^{*} S B R-B R\right\|_{1}
\]
where \(B R\) is an \(n \times r^{\prime}\) matrix and \(S B\) is an \(s \times d\) matrix. Both of them can be computed in \((n+d) \operatorname{poly}(r)\) time.

Using Lemma 17.5.2, Lemma 17.5.15, Lemma 17.4.18, we obtain,
\[
\left\|B R\left(\left(Z^{*} S B\right) R\right)^{\dagger} Z^{*} S B-B\right\|_{1} \leq O\left(\sqrt{r^{\prime}} \alpha_{r^{\prime}}\right) \min _{U \in \mathbb{R}^{n \times k}}\left\|U Z^{*} S B-B\right\|_{1}
\]
where \(\alpha_{r^{\prime}}=\widetilde{O}(r)\).
We define \(X^{*} \in \mathbb{R}^{r \times k}, Y^{*} \in \mathbb{R}^{k \times s}\),
\[
X^{*}, Y^{*}=\underset{X \in \mathbb{R}^{r^{\prime} \times k}, Y \in \mathbb{R}^{k \times s}}{\arg \min }\|B R X Y S B-B\|_{1} .
\]

Then,
\[
\begin{aligned}
\left\|B R X^{*} Y^{*} S B-B\right\|_{1} & \leq\left\|B R\left(\left(Z^{*} S B\right) R\right)^{\dagger} Z^{*} S B-B\right\|_{1} \\
& \leq O\left(\sqrt{r^{\prime}} \alpha_{r^{\prime}}\right) \min _{U \in \mathbb{R}^{n \times k}}\left\|U Z^{*} S B-B\right\|_{1} \\
& =O\left(\sqrt{r^{\prime}} \alpha_{r^{\prime}}\right) \min _{U \in \mathbb{R}^{n \times k}, Z \in \mathbb{R}^{k \times s}}\|U Z S B-B\|_{1} \\
& \leq O\left(\sqrt{r^{\prime}} \alpha_{r^{\prime}} \alpha_{s}\right) \mathrm{OPT} .
\end{aligned}
\]

It means that \(B R X, Y S B\) gives an \(O\left(\alpha_{r^{\prime}} \alpha_{s} \sqrt{r^{\prime}}\right)\)-approximation to the original problem.

Thus it suffices to use Lemma 17.4.20 to solve
\[
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\|B R X Y S B-B\|_{1}
\]
by losing an extra \(\operatorname{poly}(r)\) approximation ratio. Therefore, we finish the proof.
Lemma 17.4.20. Suppose we are given \(S \in \mathbb{R}^{s \times n}, R \in \mathbb{R}^{d \times r^{\prime}}\), and a factorization of a rank-r matrix \(B=U_{B} V_{B} \in \mathbb{R}^{n \times d}\), where \(U_{B} \in \mathbb{R}^{n \times r}, V_{B} \in \mathbb{R}^{r \times d}\). Then for any \(1 \leq k \leq r\), there exists an algorithm which takes \((n+d) \operatorname{poly}\left(r, r^{\prime}, s\right)\) time to output two matrices \(X^{\prime} \in\) \(\mathbb{R}^{r^{\prime \times k}}, Y^{\prime} \in \mathbb{R}^{k \times s}\) such that
\[
\left\|B R X^{\prime} \cdot Y^{\prime} S B-B\right\|_{1} \leq \operatorname{poly}(r) \min _{X \in \mathbb{R}^{r^{\prime} \times k}, Y \in \mathbb{R}^{k \times s}}\|B R X Y S B-B\|_{1}
\]
holds with probability at least . 999 .

Proof. Choosing dense Cauchy matrices \(T_{1} \in \mathbb{R}^{t_{1} \times n}, T_{2}^{\top} \in \mathbb{R}^{t_{2} \times d}\) to sketch on both sides, we get the problem
\[
\begin{equation*}
\min _{X \in \mathbb{R}^{r} \times k, Y \in \mathbb{R}^{k \times s}}\left\|T_{1} B R X Y S B T_{2}-T_{1} B T_{2}\right\|_{1} \tag{17.14}
\end{equation*}
\]
where \(t_{1}=\widetilde{O}(r)\) and \(t_{2}=\widetilde{O}(r)\).
Define \(X^{\prime}, Y^{\prime}\) to be the optimal solution of
\[
\min _{X \in \mathbb{R}^{r^{\prime} \times k}, Y \in \mathbb{R}^{k \times s}}\left\|T_{1} B R X Y S B T_{2}-T_{1} B T_{2}\right\|_{F}
\]

Define \(\widetilde{X}, \widetilde{Y}\) to the be the optimal solution of
\[
\min _{X \in \mathbb{R}^{r^{\prime} \times k, Y \in \mathbb{R}^{k \times s}}}\|B R X Y S B-B\|_{1}
\]

By Claim 17.3.6,
\(\left\|T_{1} B R X^{\prime} Y^{\prime} S B T_{2}-T_{1} B T_{2}\right\|_{1} \leq \sqrt{t_{1} t_{2}} \min _{X \in \mathbb{R}^{r \prime} \times k, Y \in \mathbb{R}^{k \times s}}\left\|T_{1} B R X Y S B T_{2}-T_{1} B T_{2}\right\|_{1}\).

By Lemma 17.5.4, Lemma 17.4.18 and Lemma 17.5.15, we have
\[
\left\|B R X^{\prime} \cdot Y^{\prime} S B-B\right\|_{1} \leq \sqrt{t_{1} t_{2}} \cdot \widetilde{O}\left(r^{2}\right)\|B R \tilde{X} \cdot \widetilde{Y} S B-B\|_{1}
\]

This completes the proof.

\subsection*{17.5 Contraction and Dilation Bound for \(\ell_{1}\)}

This section presents the essential lemmas for \(\ell_{1}\)-low rank approximation. Section 17.5.1 gives some basic definitions. Section 17.5 .2 shows some properties implied by contraction and dilation bounds. Section 17.5.3 presents the no dilation lemma for a dense Cauchy transform. Section 17.5 .4 and 17.5 .5 presents the no contraction lemma for dense Cauchy transforms. Section 17.5 .6 and 17.5 .7 contains the results for sparse Cauchy transforms and Lewis weights.

\subsection*{17.5.1 Definitions}

Definition 17.5.1. Given a matrix \(M \in \mathbb{R}^{n \times d}\), if matrix \(S \in \mathbb{R}^{m \times n}\) satisfies
\[
\|S M\|_{1} \leq c_{1}\|M\|_{1}
\]
then \(S\) has at most \(c_{1}\)-dilation on \(M\).

Definition 17.5.2. Given a matrix \(U \in \mathbb{R}^{n \times k}\), if matrix \(S \in \mathbb{R}^{m \times n}\) satisfies
\[
\forall x \in \mathbb{R}^{k},\|S U x\|_{1} \geq \frac{1}{c_{2}}\|U x\|_{1},
\]
then \(S\) has at most \(c_{2}\)-contraction on \(U\).

Definition 17.5.3. Given matrices \(U \in \mathbb{R}^{n \times k}, A \in \mathbb{R}^{n \times d}\), let \(V^{*}=\arg \min _{V \in \mathbb{R}^{k \times d}}\|U V-A\|_{1}\).
If matrix \(S \in \mathbb{R}^{m \times n}\) satisfies
\[
\forall V \in \mathbb{R}^{k \times d},\|S U V-S A\|_{1} \geq \frac{1}{c_{3}}\|U V-A\|_{1}-c_{4}\left\|U V^{*}-A\right\|_{1}
\]
then \(S\) has at most \(\left(c_{3}, c_{4}\right)\)-contraction on \((U, A)\).

Definition 17.5.4. A \(\left(c_{5}, c_{6}\right) \ell_{1}\)-subspace embedding for the column space of an \(n \times k\) matrix \(U\) is a matrix \(S \in \mathbb{R}^{m \times n}\) for which all \(x \in \mathbb{R}^{k}\)
\[
\frac{1}{c_{5}}\|U x\|_{1} \leq\|S U x\|_{1} \leq c_{6}\|U x\|_{1} .
\]

Definition 17.5.5. Given matrices \(U \in \mathbb{R}^{n \times k}, A \in \mathbb{R}^{n \times d}\), let \(V^{*}=\arg \min _{V \in \mathbb{R}^{k \times d}}\|U V-A\|_{1}\). Let \(S \in \mathbb{R}^{m \times n}\). If for all \(c \geq 1\), and if for any \(\widehat{V} \in \mathbb{R}^{k \times d}\) which satisfies
\[
\|S U \widehat{V}-S A\|_{1} \leq c \cdot \min _{V \in \mathbb{R}^{k \times d}}\|S U V-S A\|_{1}
\]
it holds that
\[
\|U \widehat{V}-A\|_{1} \leq c \cdot c_{7} \cdot\left\|U V^{*}-A\right\|_{1},
\]
then \(S\) provides a \(c_{7}\)-multiple-regression-cost preserving sketch of \((U, A)\).
Definition 17.5.6. Given matrices \(L \in \mathbb{R}^{n \times m_{1}}, N \in \mathbb{R}^{m_{2} \times d}, A \in \mathbb{R}^{n \times d}, k \geq 1\), let
\[
X^{*}=\arg \min _{\text {rank }-k X}\|L X N-A\|_{1} .
\]

Let \(S \in \mathbb{R}^{m \times n}\). If for all \(c \geq 1\), and if for any rank \(-k \widehat{X} \in \mathbb{R}^{m_{1} \times m_{2}}\) which satisfies
\[
\|S L \widehat{X} N-S A\|_{1} \leq c \cdot \min _{\text {rank }-k X}\|S L X N-S A\|_{1}
\]
it holds that
\[
\|L \widehat{X} N-A\|_{1} \leq c \cdot c_{8} \cdot\left\|L X^{*} N-A\right\|_{1}
\]
then \(S\) provides a \(c_{8}\)-restricted-multiple-regression-cost preserving sketch of \((L, N, A, k)\).

\subsection*{17.5.2 Properties}

Lemma 17.5.1. Given matrices \(A \in \mathbb{R}^{n \times d}, U \in \mathbb{R}^{n \times k}\), let \(V^{*}=\arg \min _{V \in \mathbb{R}^{k \times d}}\|U V-A\|_{1}\). If \(S \in \mathbb{R}^{m \times n}\) has at most \(c_{1}\)-dilation on \(U V^{*}-A\), i.e.,
\[
\left\|S\left(U V^{*}-A\right)\right\|_{1} \leq c_{1}\left\|U V^{*}-A\right\|_{1}
\]
and it has at most \(c_{2}\)-contraction on \(U\), i.e.,
\[
\forall x \in \mathbb{R}^{k},\|S U x\|_{1} \geq \frac{1}{c_{2}}\|U x\|_{1},
\]
then \(S\) has at most \(\left(c_{2}, c_{1}+\frac{1}{c_{2}}\right)\)-contraction on \((U, A)\), i.e.,
\[
\forall V \in \mathbb{R}^{k \times d},\|S U V-S A\|_{1} \geq \frac{1}{c_{2}}\|U V-A\|_{1}-\left(c_{1}+\frac{1}{c_{2}}\right)\left\|U V^{*}-A\right\|_{1},
\]

Proof. Let \(A \in \mathbb{R}^{n \times d}, U \in \mathbb{R}^{n \times k}, S \in \mathbb{R}^{m \times n}\) be the same as that described in the lemma. Then \(\forall V \in \mathbb{R}^{k \times d}\)
\[
\begin{aligned}
\|S U V-S A\|_{1} & \geq\left\|S U V-S U V^{*}\right\|_{1}-\left\|S U V^{*}-S A\right\|_{1} \\
& \geq\left\|S U V-S U V^{*}\right\|_{1}-c_{1}\left\|U V^{*}-A\right\|_{1} \\
& =\left\|S U\left(V-V^{*}\right)\right\|_{1}-c_{1}\left\|U V^{*}-A\right\|_{1} \\
& =\sum_{j=1}^{d}\left\|S U\left(V-V^{*}\right)_{j}\right\|_{1}-c_{1}\left\|U V^{*}-A\right\|_{1} \\
& \geq \sum_{j=1}^{d} \frac{1}{c_{2}}\left\|U\left(V-V^{*}\right)_{j}\right\|_{1}-c_{1}\left\|U V^{*}-A\right\|_{1} \\
& =\frac{1}{c_{2}}\left\|U V-U V^{*}\right\|_{1}-c_{1}\left\|U V^{*}-A\right\|_{1} \\
& \geq \frac{1}{c_{2}}\|U V-A\|_{1}-\frac{1}{c_{2}}\left\|U V^{*}-A\right\|_{1}-c_{1}\left\|U V^{*}-A\right\|_{1} \\
& =\frac{1}{c_{2}}\|U V-A\|_{1}-\left(\left(\frac{1}{c_{2}}+c_{1}\right)\left\|U V^{*}-A\right\|_{1}\right) .
\end{aligned}
\]

The first inequality follows by the triangle inequality. The second inequality follows since \(S\) has at most \(c_{1}\) dilation on \(U V^{*}-A\). The third inequality follows since \(S\) has at most \(c_{2}\) contraction on \(U\). The fourth inequality follows by the triangle inequality.

Lemma 17.5.2. Given matrices \(A \in \mathbb{R}^{n \times d}, U \in \mathbb{R}^{n \times k}\), let \(V^{*}=\arg \min _{V \in \mathbb{R}^{k \times d}}\|U V-A\|_{1}\). If \(S \in \mathbb{R}^{m \times n}\) has at most \(c_{1}\)-dilation on \(U V^{*}-A\), i.e.,
\[
\left\|S\left(U V^{*}-A\right)\right\|_{1} \leq c_{1}\left\|U V^{*}-A\right\|_{1}
\]
and has at most \(c_{2}\)-contraction on \(U\), i.e.,
\[
\forall x \in \mathbb{R}^{k},\|S U x\|_{1} \geq \frac{1}{c_{2}}\|U x\|_{1},
\]
then \(S\) provides a \(\left(2 c_{1} c_{2}+1\right)\)-multiple-regression-cost preserving sketch of \((U, A)\), i.e., for all \(c \geq 1\), for any \(\widehat{V} \in \mathbb{R}^{k \times d}\) which satisfies
\[
\|S U \widehat{V}-S A\|_{1} \leq c \cdot \min _{V \in \mathbb{R}^{k \times d}}\|S U V-S A\|_{1}
\]
it has
\[
\|U \widehat{V}-A\|_{1} \leq c \cdot\left(2 c_{1} c_{2}+1\right) \cdot\left\|U V^{*}-A\right\|_{1}
\]

Proof. Let \(S \in \mathbb{R}^{m \times n}, A \in \mathbb{R}^{n \times d}, U \in \mathbb{R}^{n \times k}, V^{*}, \widehat{V} \in \mathbb{R}^{k \times d}\), and \(c\) be the same as stated in the lemma.
\[
\begin{aligned}
\|U \widehat{V}-A\|_{1} & \leq c_{2}\|S U \widehat{V}-S A\|_{1}+\left(1+c_{1} c_{2}\right)\left\|U V^{*}-A\right\|_{1} \\
& \leq c_{2} c \min _{V \in \mathbb{R}^{k \times d}}\|S U V-S A\|_{1}+\left(1+c_{1} c_{2}\right)\left\|U V^{*}-A\right\|_{1} \\
& \leq c_{2} c\left\|S U V^{*}-S A\right\|_{1}+\left(1+c_{1} c_{2}\right)\left\|U V^{*}-A\right\|_{1} \\
& \leq c_{1} c_{2} c\left\|U V^{*}-A\right\|_{1}+\left(1+c_{1} c_{2}\right)\left\|U V^{*}-A\right\|_{1} \\
& \leq c \cdot\left(1+2 c_{1} c_{2}\right)\left\|U V^{*}-A\right\|_{1} .
\end{aligned}
\]

The first inequality follows by Lemma 17.5.1. The second inequality follows by the guarantee of \(\widehat{V}\). The fourth inequality follows since \(S\) has at most \(c_{1}\)-dilation on \(U V^{*}-A\). The fifth inequality follows since \(c \geq 1\).

Lemma 17.5.3. Given matrices \(L \in \mathbb{R}^{n \times m_{1}}, N \in \mathbb{R}^{m_{2} \times d}, A \in \mathbb{R}^{n \times d}, k \geq 1\), let
\[
X^{*}=\arg \min _{\operatorname{rank}-k X}\|L X N-A\|_{1}
\]

If \(S \in \mathbb{R}^{m \times n}\) has at most \(c_{1}\)-dilation on \(L X^{*} N-A\), i.e.,
\[
\left\|S\left(L X^{*} N-A\right)\right\|_{1} \leq c_{1}\left\|L X^{*} N-A\right\|_{1}
\]
and has at most \(c_{2}\)-contraction on L, i.e.,
\[
\forall x \in \mathbb{R}^{m_{1}}\|S L x\|_{1} \geq\|L x\|_{1}
\]
then \(S\) provides a \(\left(2 c_{1} c_{2}+1\right)\)-restricted-multiple-regression-cost preserving sketch of ( \(\left.L, N, A, k\right)\), i.e., for all \(c \geq 1\), for any rank \(-k \widehat{X} \in \mathbb{R}^{m_{1} \times m_{2}}\) which satisfies
\[
\|S L \widehat{X} N-S A\|_{1} \leq c \cdot \min _{\text {rank }-k X}\|S L X N-S A\|_{1}
\]
it holds that
\[
\|L \widehat{X} N-A\|_{1} \leq c \cdot\left(2 c_{1} c_{2}+1\right) \cdot\left\|L X^{*} N-A\right\|_{1}
\]

Proof. Let \(S \in \mathbb{R}^{m \times n}, L \in \mathbb{R}^{n \times m_{1}}, \widehat{X} \in \mathbb{R}^{m_{1} \times m_{2}}, X^{*} \in \mathbb{R}^{m_{1} \times m_{2}}, N \in \mathbb{R}^{m_{2} \times d}, A \in \mathbb{R}^{n \times d}\), and
\(c \geq 1\) be the same as stated in the lemma.
\[
\begin{aligned}
\|S L \widehat{X} N-S A\|_{1} & \geq\left\|S L \widehat{X} N-S L X^{*} N\right\|_{1}-\left\|S L X^{*} N-S A\right\|_{1} \\
& \geq \frac{1}{c_{2}}\left\|L\left(\widehat{X} N-X^{*} N\right)\right\|_{1}-c_{1}\left\|L X^{*} N-A\right\|_{1} \\
& \geq \frac{1}{c_{2}}\|L \widehat{X} N-A\|_{1}-\frac{1}{c_{2}}\left\|L X^{*} N-A\right\|_{1}-c_{1}\left\|L X^{*} N-A\right\|_{1} \\
& =\frac{1}{c_{2}}\|L \widehat{X} N-A\|_{1}-\left(\frac{1}{c_{2}}+c_{1}\right)\left\|L X^{*} N-A\right\|_{1}
\end{aligned}
\]

The inequality follows by the triangle inequality. The second inequality follows since \(S\) has at most \(c_{2}\)-contraction on \(L\), and it has at most \(c_{1}\)-dilation on \(L X^{*} N-A\). The third inequality follows by the triangle inequality.

It follows that
\[
\begin{aligned}
\|L \widehat{X} N-A\|_{1} & \leq c_{2}\|S L \widehat{X} N-S A\|_{1}+\left(1+c_{1} c_{2}\right)\left\|L X^{*} N-A\right\|_{1} \\
& \leq c_{2} c \cdot \min _{\text {rank }-k X}\|S L X N-S A\|_{1}+\left(1+c_{1} c_{2}\right)\left\|L X^{*} N-A\right\|_{1} \\
& \leq c_{2} c \cdot\left\|S L X^{*} N-S A\right\|_{1}+\left(1+c_{1} c_{2}\right)\left\|L X^{*} N-A\right\|_{1} \\
& \leq c c_{1} c_{2} \cdot\left\|L X^{*} N-A\right\|_{1}+\left(1+c_{1} c_{2}\right)\left\|L X^{*} N-A\right\|_{1} \\
& \leq c \cdot\left(1+2 c_{1} c_{2}\right)\left\|L X^{*} N-A\right\|_{1}
\end{aligned}
\]

The first inequality directly follows from the previous one. The second inequality follows from the guarantee of \(\widehat{X}\). The fourth inequality follows since \(S\) has at most \(c_{1}\) dilation on \(L X^{*} N-A\). The fifth inequality follows since \(c \geq 1\).

Lemma 17.5.4. Given matrices \(L \in \mathbb{R}^{n \times m_{1}}, N \in \mathbb{R}^{m_{2} \times d}, A \in \mathbb{R}^{n \times d}, k \geq 1\), let
\[
X^{*}=\arg \min _{\text {rank }-k X}\|L X N-A\|_{1}
\]

Let \(T_{1} \in \mathbb{R}^{t_{1} \times n}\) have at most \(c_{1}\)-dilation on \(L X^{*} N-A\), and at most \(c_{2}\)-contraction on \(L\). Let
\[
\tilde{X}=\arg \min _{\operatorname{rank}-k X}\left\|T_{1} L X N-T_{1} A\right\|_{1}
\]

Let \(T_{2}^{\top} \in \mathbb{R}^{t_{2} \times d}\) have at most \(c_{1}^{\prime}\)-dilation on \(\left(T_{1} L \widetilde{X} N-T_{1} A\right)^{\top}\), and at most \(c_{2}^{\prime}\)-contraction on \(N^{\top}\). Then, for all \(c \geq 1\), for any \(\operatorname{rank}-k \widehat{X} \in \mathbb{R}^{m_{1} \times m_{2}}\) which satisfies
\[
\left\|T_{1}(L \widehat{X} N-S A) T_{2}\right\|_{1} \leq c \cdot \min _{\text {rank }-k X}\left\|T_{1}(L X N-A) T_{2}\right\|_{1},
\]
it has
\[
\|L \widehat{X} N-A\|_{1} \leq c \cdot\left(2 c_{1} c_{2}+1\right)\left(2 c_{1}^{\prime} c_{2}^{\prime}+1\right) \cdot\left\|L X^{*} N-A\right\|_{1} .
\]

Proof. Apply Lemma 17.5.3 for sketch matrix \(T_{2}\). Then for any \(c \geq 1\), any rank \(-k \widehat{X} \in\) \(\mathbb{R}^{m_{1} \times m_{2}}\) which satisfies
\[
\left\|T_{1}(L \widehat{X} N-A) T_{2}\right\|_{1} \leq c \cdot \min _{\text {rank }-k X}\left\|T_{1}(L X N-A) T_{2}\right\|_{1}
\]
has
\[
\left\|T_{1}(L \widehat{X} N-A)\right\|_{1} \leq c \cdot\left(2 c_{1}^{\prime} c_{2}^{\prime}+1\right) \cdot\left\|T_{1}(L \widetilde{X} N-A)\right\|_{1} .
\]

Apply Lemma 17.5 .3 for sketch matrix \(T_{1}\). Then for any \(c \geq 1\), any rank \(-k \widehat{X} \in\) \(\mathbb{R}^{m_{1} \times m_{2}}\) which satisfies
\[
\left\|T_{1}(L \widehat{X} N-A)\right\|_{1} \leq c\left(2 c_{1}^{\prime} c_{2}^{\prime}+1\right) \cdot \min _{\text {rank }-k}\left\|T_{1}(L \widetilde{X} N-A)\right\|_{1},
\]
has
\[
\|L \widehat{X} N-A\|_{1} \leq c \cdot\left(2 c_{1} c_{2}+1\right)\left(2 c_{1}^{\prime} c_{2}^{\prime}+1\right) \cdot\left\|L X^{*} N-A\right\|_{1} .
\]

Lemma 17.5.5. Given matrices \(M \in \mathbb{R}^{n \times d}, U \in \mathbb{R}^{n \times t}, d \geq t=\operatorname{rank}(U), n \geq d \geq r=\) \(\operatorname{rank}(M)\), if sketching matrix \(S \in \mathbb{R}^{m \times n}\) is drawn from any of the following probability distributions of matrices, with .99 probability \(S\) has at most \(c_{1}\)-dilation on \(M\), i.e.,
\[
\|S M\|_{1} \leq c_{1}\|M\|_{1}
\]
and \(S\) has at most \(c_{2}\)-contraction on \(U\), i.e.,
\[
\forall x \in \mathbb{R}^{t},\|S U x\|_{1} \geq \frac{1}{c_{2}}\|U x\|_{1}
\]
where \(c_{1}, c_{2}\) are parameters depend on the distribution over \(S\).
(I) \(S \in \mathbb{R}^{m \times n}\) is a dense Cauchy matrix: a matrix with i.i.d. entries from the standard Cauchy distribution. If \(m=O(t \log t)\), then \(c_{1} c_{2}=O(\log d)\). If \(m=O((t+r) \log (t+\) \(r))\), then \(c_{1} c_{2}=O(\min (\log d, r \log r))\).
(II) \(S \in \mathbb{R}^{m \times n}\) is a sparse Cauchy matrix: \(S=T D\), where \(T \in \mathbb{R}^{m \times n}\) has each column i.i.d. from the uniform distribution on standard basis vectors of \(\mathbb{R}^{m}\), and \(D \in \mathbb{R}^{n \times n}\) is a diagonal matrix with i.i.d. diagonal entries following a standard Cauchy distribution. If \(m=O\left(t^{5} \log ^{5} t\right)\), then \(c_{1} c_{2}=O\left(t^{2} \log ^{2} t \log d\right)\). If \(m=O\left((t+r)^{5} \log ^{5}(t+r)\right)\), then \(c_{1} c_{2}=O\left(\min \left(t^{2} \log ^{2} t \log d, r^{3} \log ^{3} r\right)\right)\).
(III) \(S \in \mathbb{R}^{m \times n}\) is a sampling and rescaling matrix (notation \(S \in \mathbb{R}^{n \times n}\) denotes a diagonal sampling and rescaling matrix with \(m\) non-zero entries): If \(S\) samples and reweights \(m=O(t \log t)\) rows of \(U\), selecting each with probability proportional to the \(i^{\text {th }}\) row's \(\ell_{1}\) Lewis weight and reweighting by the inverse probability, then \(c_{1} c_{2}=O(1)\).

In the above, if we replace \(S\) with \(\sigma \cdot S\) where \(\sigma \in \mathbb{R} \backslash\{0\}\) is any scalar, then the relation between \(m\) and \(c_{1} c_{2}\) can be preserved.

For (I), if \(m=O(t \log t)\), then \(c_{1} c_{2}=O(\log d)\) is implied by Lemma 17.5.15 and Lemma 17.5.7. If \(m=O((t+r) \log (t+r)), c_{1} c_{2}=O(r \log r)\) is implied by [SW11].

For (II), if \(m=O\left(t^{5} \log ^{5} t\right)\), then \(c_{1} c_{2}=O\left(t^{2} \log ^{2} t \log d\right)\) is implied by Corollary 17.5.19 and Lemma 17.5.17. If \(m=O\left((t+r)^{5} \log ^{5}(t+r)\right), c_{1} c_{2}=O\left(r^{3} \log ^{3} r\right)\) is implied by [MM13].

For (III), it is implied by [CP15] and Lemma 17.5.21.

\subsection*{17.5.3 Cauchy embeddings, no dilation}

Lemma 17.5.6. Define \(U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times d}\) to be the optimal solution of \(\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}} \| U V\) \(A \|_{1}\). Choose a Cauchy matrix \(S\) with \(m\) rows and rescaled by \(\Theta(1 / m)\). We have that
\[
\left\|S U^{*} V^{*}-S A\right\|_{1} \leq O(\log d)\left\|U^{*} V^{*}-A\right\|_{1}
\]
holds with probability at least 99/100.

Proof. The proof technique has been used in [Ind06] and [CDMI \(\left.{ }^{+} 13\right]\). Fix the optimal \(U^{*}\) and \(V^{*}\), then
\[
\begin{array}{rlr}
\left\|S U^{*} V^{*}-S A\right\|_{1} & =\sum_{i=1}^{d}\left\|S\left(U^{*} V_{i}^{*}-A_{i}\right)\right\|_{1} \\
& =\sum_{i=1}^{d} \sum_{j=1}^{m}\left|\sum_{l=1}^{n} \frac{1}{m} S_{j, l} \cdot\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right| & \quad \text { where } S_{j, l} \sim C(0,1) \\
& =\sum_{i=1}^{d} \sum_{j=1}^{m} \frac{1}{m}\left|c_{i, j}\right| \quad \text { where } c_{i, j} \sim C\left(0,\left\|U^{*} V_{i}^{*}-A_{i}\right\|_{1}\right) \\
& =\frac{1}{m} \sum_{i=1}^{d} \sum_{j=1}^{m}\left\|U^{*} V_{i}^{*}-A_{i}\right\|_{1} \cdot w_{i+d(j-1)} . \quad \text { where } w_{i+d(j-1)} \sim|C(0,1)| \tag{17.15}
\end{array}
\]
where the last step follows since each \(w_{i}\) can be thought of as a clipped half-Cauchy random variable. Define \(d^{\prime}=m d\). Define event \(\xi_{i}\) to be the situation in which \(w_{i}<D\) (we will decide upon \(D\) later \()\), and define event \(\xi=\xi_{1} \cap \xi_{2} \cap \cdots \cap \xi_{d^{\prime}}\). Then it is clear that \(\xi \cap \xi_{i}=\xi, \forall i \in\left[d^{\prime}\right]\).

Using the probability density function (pdf) of a Cauchy and because \(\tan ^{-1} x \leq x\), we can lower bound \(\operatorname{Pr}\left[\right.\) event \(\xi_{i}\) holds \(]\) in the following sense,
\[
\operatorname{Pr}\left[\xi_{i}\right]=\frac{2}{\pi} \tan ^{-1}(D)=1-\frac{2}{\pi} \tan ^{-1}(1 / D) \geq 1-\frac{2}{\pi D} .
\]

By a union bound over all \(i \in\left[d^{\prime}\right]\), we can lower bound \(\operatorname{Pr}[\) event \(\xi\) holds \(]\),
\[
\begin{equation*}
\operatorname{Pr}[\xi] \geq 1-\sum_{i=1}^{d^{\prime}} \operatorname{Pr}\left[\bar{\xi}_{i}\right] \geq 1-\frac{2 d^{\prime}}{\pi D} \tag{17.16}
\end{equation*}
\]

By Bayes rule and \(\xi=\xi \cap \xi_{i}, \operatorname{Pr}\left[\xi \mid \xi_{i}\right] \operatorname{Pr}\left[\xi_{i}\right]=\operatorname{Pr}\left[\xi \cap \xi_{i}\right]=\operatorname{Pr}[\xi]\), which implies that \(\operatorname{Pr}\left[\xi \mid \xi_{i}\right]=\) \(\operatorname{Pr}[\xi] / \operatorname{Pr}\left[\xi_{i}\right]\). First, we can lower bound \(\mathbb{E}\left[w_{i} \mid \xi_{i}\right]\),
\[
\begin{array}{rlr}
\mathbb{E}\left[w_{i} \mid \xi_{i}\right] & =\mathbb{E}\left[w_{i} \mid \xi_{i} \cap \xi\right] \operatorname{Pr}\left[\xi \mid \xi_{i}\right]+\mathbb{E}\left[w_{i} \mid \xi \cap \bar{\xi}\right] \operatorname{Pr}\left[\bar{\xi} \mid \xi_{i}\right] & \\
& \geq \mathbb{E}\left[w_{i} \mid \xi_{i} \cap \xi\right] \operatorname{Pr}\left[\xi \mid \xi_{i}\right] & \text { by } w_{i} \geq 0 \text { and } \operatorname{Pr}[] \geq 0 \\
& =\mathbb{E}\left[w_{i} \mid \xi\right] \operatorname{Pr}\left[\xi \mid \xi_{i}\right] & \text { by } \xi=\xi \cap \xi_{i} .
\end{array}
\]

The above equation implies that
\[
\begin{array}{rlr}
\mathbb{E}\left[w_{i} \mid \xi\right] & \leq \frac{\mathbb{E}\left[w_{i} \mid \xi_{i}\right]}{\operatorname{Pr}\left[\xi \mid \xi_{i}\right]} \\
& =\frac{\mathbb{E}\left[w_{i} \mid \xi_{i}\right] \operatorname{Pr}\left[\xi_{i}\right]}{\operatorname{Pr}\left[\xi \cap \xi_{i}\right]} & \text { by Bayes rule } \operatorname{Pr}\left[\xi \mid \xi_{i}\right] \operatorname{Pr}\left[\xi_{i}\right]=\operatorname{Pr}\left[\xi \cap \xi_{i}\right] \\
& =\frac{\mathbb{E}\left[w_{i} \mid \xi_{i}\right] \operatorname{Pr}\left[\xi_{i}\right]}{\operatorname{Pr}[\xi]} & \text { by } \xi=\xi \cap \xi_{i} .
\end{array}
\]

Using the pdf of a Cauchy, \(\mathbb{E}\left[w_{i} \mid \xi_{i}\right]=\frac{1}{\pi} \log \left(1+D^{2}\right) / \operatorname{Pr}\left[\xi_{i}\right]\) and plugging it into the lower bound of \(\mathbb{E}\left[w_{i} \mid \xi\right]\),
\[
\mathbb{E}\left[w_{i} \mid \xi\right] \leq \frac{\mathbb{E}\left[w_{i} \mid \xi_{i}\right] \operatorname{Pr}\left[\xi_{i}\right]}{\operatorname{Pr}[\xi]}=\frac{\frac{1}{\pi} \log \left(1+D^{2}\right)}{\operatorname{Pr}[\xi]} \leq \frac{\frac{1}{\pi} \log \left(1+D^{2}\right)}{1-\frac{2 d}{\pi D}} \lesssim \log (D),
\]
where the third step follows since \(\operatorname{Pr}[\xi] \geq 1-\frac{2 d^{\prime}}{\pi D}\) and the last step follows by choosing \(D=\Theta\left(d^{\prime}\right)\).

We can conclude
\[
\begin{align*}
\mathbb{E}\left[\left\|S U^{*} V^{*}-S A\right\|_{1} \mid \xi\right] & =\frac{1}{m} \sum_{i=1}^{d} \sum_{j=1}^{m}\left\|U^{*} V_{i}^{*}-A_{i}\right\|_{1} \cdot \mathbb{E}\left[w_{i+d(j-1)} \mid \xi\right] \\
& \lesssim\left(\log d^{\prime}\right) \cdot\left\|U^{*} V^{*}-A\right\|_{1} . \tag{17.17}
\end{align*}
\]

For simplicity, define \(X=\left\|S U^{*} V^{*}-S A\right\|_{1}\) and \(\gamma=\left\|U^{*} V^{*}-A\right\|_{1}\). By Markov's inequality and because \(\operatorname{Pr}[X \geq \gamma t \mid \bar{\xi}] \leq 1\), we have
\[
\begin{aligned}
& \operatorname{Pr}[X \geq \gamma t] \\
& =\operatorname{Pr}[X \geq \gamma t \mid \xi] \operatorname{Pr}[\xi]+\operatorname{Pr}[X \geq \gamma t \mid \bar{\xi}] \operatorname{Pr}[\bar{\xi}] \\
& \leq \operatorname{Pr}[X \geq \gamma t \mid \xi]+\operatorname{Pr}[\bar{\xi}] \\
& \leq \frac{\mathbb{E}[X \mid \xi]}{\gamma t}+\operatorname{Pr}[\bar{\xi}] \quad \text { by Markov's inequality } \\
& \leq \frac{\mathbb{E}[X \mid \xi]}{\gamma t}+\frac{2 d^{\prime}}{\pi D} \quad \text { by Equation (17.16) } \\
& \lesssim \frac{\log d^{\prime}}{t}+\frac{2 d^{\prime}}{\pi D} \quad \quad \text { by Equation (17.17) } \\
& \leq .01,
\end{aligned}
\]
where choosing \(t=\Theta\left(\log d^{\prime}\right)\) and \(D=\Theta\left(d^{\prime}\right)\). Since \(k \leq d\) and \(m=\operatorname{poly}(k)\), we have \(t=\Theta(\log d)\), which completes the proof.

Lemma 17.5.7. Given any matrix \(M \in \mathbb{R}^{n \times d}\), if matrix \(S \in \mathbb{R}^{m \times n}\) has each entry drawn from an i.i.d. standard Cauchy distribution and is rescaled by \(\Theta(1 / m)\), then
\[
\|S M\|_{1} \leq O(\log d)\|M\|_{1}
\]
holds with probability at least 99/100.

Proof. Just replace the matrix \(U^{*} V^{*}-A\) in the proof of Lemma 17.5.6 with \(M\). Then we can get the result directly.

\subsection*{17.5.4 Cauchy embeddings, no contraction}

We prove that if we choose a Cauchy matrix \(S\), then for a fixed optimal solution \(U^{*}\) of \(\min _{U, V}\|U V-A\|_{1}\), and for all \(V\), we have that with high probability \(\left\|S U^{*} V-S A\right\|_{1}\) is lower bounded by \(\left\|U^{*} V-A\right\|_{1}\) up to some constant.

Lemma 17.5.8. Define \(U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times d}\) to be the optimal solution of \(\min _{U \in \mathbb{R}^{n \times k} V \in \mathbb{R}^{k \times d}} \| U V-\) \(A \|_{1}\), where \(A \in \mathbb{R}^{n \times d}\). Let \(m=O(k \log k), S \in \mathbb{R}^{m \times n}\) be a random matrix with each entry an i.i.d. standard Cauchy random variable, scaled by \(\Theta(1 / m)\). Then with probability at least 0.95,
\[
\forall V \in \mathbb{R}^{k \times d},\left\|U^{*} V-A\right\|_{1} \lesssim\left\|S U^{*} V-S A\right\|_{1}+O(\log (d))\left\|U^{*} V^{*}-A\right\|_{1}
\]

Proof. This follows by Lemmas 17.5.6, 17.5.15, 17.5.1.

\subsection*{17.5.5 Cauchy embeddings, \(k\)-dimensional subspace}

The goal of this section is to prove Lemma 17.5.15.
Before getting into the details, we first give the formal definition of an \((\alpha, \beta) \ell_{1}\) well-conditioned basis and \(\epsilon\)-net.

Definition 17.5.7 ( \(\ell_{1}\) Well-conditioned basis). [ \(\left.\mathrm{DDH}^{+} 09\right]\) A basis \(U\) for the range of \(A\) is \((\alpha, \beta)\)-conditioned if \(\|U\|_{1} \leq \alpha\) and for all \(x \in \mathbb{R}^{k},\|x\|_{\infty} \leq \beta\|U x\|_{1}\). We will say \(U\) is well-conditioned if \(\alpha\) and \(\beta\) are low-degree polynomials in \(k\), independent of \(n\).

Note that a well-conditioned basis implies the following result.

Fact 17.5.9. There exist \(\alpha, \beta \geq 1\) such that,
\[
\forall x \in \mathbb{R}^{k}, \frac{1}{k \beta}\|x\|_{1} \leq\|U x\|_{1} \leq \alpha\|x\|_{1} .
\]

Proof. The lower bound can be proved in the following sense,
\[
\|U x\|_{1} \geq \frac{1}{\beta}\|x\|_{\infty} \geq \frac{1}{\beta} \frac{1}{k}\|x\|_{1},
\]
where the first step follows by the properties of a well-conditioned basis, and the second step follows since \(k\|x\|_{\infty} \geq\|x\|_{1}\). Then we can show an upper bound,
\[
\|U x\|_{1} \leq\|U\|_{1} \cdot\|x\|_{1} \leq \alpha\|x\|_{1}
\]
where the first step follows by \(\|U x\|_{1} \leq\|U\|_{1}\|x\|_{1}\), and the second step follows using \(\|U\|_{1} \leq\) \(\alpha\).

Definition 17.5.8 ( \(\epsilon\)-net). Define \(\mathcal{N}\) to an \(\epsilon\)-net where, for all the \(x \in \mathbb{R}^{k}\) that \(\|x\|_{1}=1\), for any vectors two \(x, x^{\prime} \in \mathcal{N},\left\|x-x^{\prime}\right\|_{1} \geq \epsilon\), for any vector \(y \notin \mathcal{N}\), there exists an \(x \in \mathcal{N}\) such that \(\|x-y\|_{1} \leq \epsilon\). Then the size of \(\mathcal{N}\) is \(\left(\frac{1}{\epsilon}\right)^{O(k)}=2^{O(k \log (1 / \epsilon))}\).

Lemma 17.5.10 (Lemma 6 in [SW11]). There is a constant \(c_{0}>0\) such that for any \(t \geq 1\) and any constant \(c>c_{0}\), if \(S\) is a \(t \times n\) matrix whose entries are i.i.d. standard Cauchy random variables scaled by \(c / t\), then, for any fixed \(y \in \mathbb{R}^{n}\),
\[
\operatorname{Pr}\left[\|S y\|_{1}<\|y\|_{1}\right] \leq 1 / 2^{t}
\]

Lemma 17.5.11. Suppose we are given a well-conditioned basis \(U \in \mathbb{R}^{n \times k}\). If \(S\) is a \(t \times n\) matrix whose entries are i.i.d. standard Cauchy random variables scaled by \(\Theta(1 / t)\), then with probability \(1-2^{-\Omega(t)}\), for all vectors \(x \in \mathcal{N}\) we have that \(\|S U x\|_{1} \geq\|U x\|_{1}\).

Proof. First, using Lemma 17.5.10, we have for any fixed vector \(y \in \mathbb{R}^{n}, \operatorname{Pr}\left[\|S y\|<\|y\|_{1}\right] \leq\) \(1 / 2^{t}\). Second, we can rewrite \(y=U x\). Then for any fixed \(x \in \mathcal{N}, \operatorname{Pr}\left[\|S U x\|<\|U x\|_{1}\right] \leq 1 / 2^{t}\). Third, choosing \(t \gtrsim k \log (1 / \epsilon)\) and taking a union bound over all the vectors in the \(\epsilon\)-net \(\mathcal{N}\) completes the proof.

Lemma 17.5.12 (Lemma 7 in [SW11]). Let \(S\) be a \(t \times n\) matrix whose entries are i.i.d standard Cauchy random variables, scaled by \(c / t\) for a constant \(c\), and \(t \geq 1\). Then there is a constant \(c^{\prime}=c^{\prime}(c)>0\) such that for any fixed set of \(\left\{y_{1}, y_{2}, \cdots, y_{k}\right\}\) of \(d\) vectors in \(\left\{y \in \mathbb{R}^{n}: x \in \mathbb{R}^{k}, U x=y\right\}\),
\[
\operatorname{Pr}\left[\sum_{i=1}^{k}\left\|S y_{i}\right\|_{1} \geq c^{\prime} \log (t k) \sum_{i=1}^{k}\left\|y_{i}\right\|_{1}\right] \leq \frac{1}{1000} .
\]

Using Lemma 17.5.12 and the definition of an \((\alpha, \beta)\) well-conditioned basis, we can show the following corollary.

Corollary 17.5.13. Suppose we are given an \((\alpha, \beta) \ell_{1}\) well-conditioned basis \(U \in \mathbb{R}^{n \times k}\). Choose \(S\) to be an i.i.d. Cauchy matrix with \(t=O(k \log k)\) rows, and rescale each entry by \(\Theta(1 / t)\). Then with probability \(99 / 100\), for all vectors \(x \in \mathbb{R}^{k}\),
\[
\|S U x\|_{1} \leq O(\alpha \beta \log (k)) \cdot\|U x\|_{1}
\]

Proof. Define event \(E\) to be the situation when
\[
\sum_{i=1}^{k}\left\|S U_{i}\right\|_{1}<c^{\prime} \log (t k) \sum_{i=1}^{k}\left\|U_{i}\right\|_{1}
\]
holds, and \(U\) is a well-conditioned basis. Using Lemma 17.5.12, we can show that event \(E\) holds with probability \(999 / 1000\). We condition on Event \(E\) holding. Then for any \(y=U x\) for an \(x \in \mathbb{R}^{k}\), we have
\[
\begin{array}{rlr}
\|S y\|_{1} & =\|S U x\|_{1} & \\
& \leq \sum_{j=1}^{k}\left\|S U_{j} x_{j}\right\|_{1} & \text { by triangle inequality } \\
& =\sum_{j=1}^{k}\left|x_{j}\right| \cdot\left\|S U_{j}\right\|_{1} & \\
& \leq\|x\|_{\infty} c^{\prime} \log (t k) \sum_{j=1}^{k}\left\|U_{j}\right\|_{1} & \text { by Lemma } 17.5 .12 \\
& =\|x\|_{\infty} c^{\prime} \log (t k) \alpha & \text { by }\|U\|_{1}=\sum_{j=1}^{k}\left\|U_{j}\right\|_{1} \leq \alpha \\
& \leq \beta\|U x\|_{1} c^{\prime} \log (t k) \alpha & \text { by }\|x\|_{\infty} \leq \beta\|U x\|_{1} \\
& \leq \beta\|y\|_{1} c^{\prime} \log (t k) \alpha & \text { by } U x=y .
\end{array}
\]

This completes the proof.

Lemma 17.5.14. Given an \((\alpha, \beta) \ell_{1}\) well-conditioned basis, condition on the following two events,
1. For all \(x \in \mathcal{N},\|S U x\|_{1} \geq\|U x\|_{1}\). (Lemma 17.5.10)
2. For all \(x \in \mathbb{R}^{k},\|S U x\|_{1} \leq O(\alpha \beta \log k)\|U x\|_{1}\). (Corollary 17.5.13)

Then, for all \(w \in \mathbb{R}^{k},\|S U w\|_{1} \gtrsim\|U w\|_{1}\).

Proof. For any \(w \in \mathbb{R}^{k}\) we can write it as \(w=\ell \cdot z\) where \(\ell\) is some scalar and \(z\) has \(\|z\|_{1}=1\). Define \(y=\underset{y^{\prime} \in \mathcal{N}}{\arg \min }\left\|y^{\prime}-z\right\|_{1}\).

We first show that if \(U\) is an \((\alpha, \beta)\) well-conditioned basis for \(\ell_{1}\), then \(\|U(y-z)\|_{1} \leq\) \(\alpha \beta k \epsilon\|U y\|_{1}\),
\[
\begin{array}{rr} 
& \|U(y-z)\|_{1} \\
\leq & \\
\leq\|y-z\|_{1} & \text { by }\|U(y-z)\|_{1} \leq \alpha\|y-z\|_{1} \\
\leq & \text { by }\|y-z\|_{1} \leq \epsilon \\
\leq \alpha \epsilon\|y\|_{1} & \text { by }\|y\|_{1}=1 \\
\leq \alpha \epsilon\|U y\|_{1} \beta k & \text { by }\|U y\|_{1} \geq \frac{1}{\beta k}\|y\|_{1} .
\end{array}
\]

Because \(\epsilon<1 /\left(\alpha \beta k^{c+1}\right)\), we have
\[
\begin{equation*}
\|U(y-z)\|_{1} \leq \frac{1}{k^{c}}\|U y\|_{1} . \tag{17.18}
\end{equation*}
\]

Using the triangle inequality, we can lower bound \(\|U y\|_{1}\) by \(\|U z\|_{1}\) up to some constant,
\[
\begin{equation*}
\|U y\|_{1} \geq\|U z\|_{1}-\|U(y-z)\|_{1} \geq\|U z\|_{1}-\frac{1}{k^{c}}\|U y\|_{1} \tag{17.19}
\end{equation*}
\]
which implies
\[
\begin{equation*}
\|U y\|_{1} \geq .99\|U z\|_{1} . \tag{17.20}
\end{equation*}
\]

Thus,
\[
\begin{array}{rr} 
& \|S U z\|_{1} \\
\geq & \|S U y\|_{1}-\|S U(z-y)\|_{1}
\end{array} \text { by triangle inequality }
\]
by rescaling \(z\) to \(w\), we complete the proof.
Lemma 17.5.15. Given matrix \(U \in \mathbb{R}^{n \times k}\), let \(t=O(k \log k)\), and let \(S \in \mathbb{R}^{t \times n}\) be a random matrix with entries drawn i.i.d. from a standard Cauchy distribution, where each entry is rescaled by \(\Theta(1 / t)\). With probability .99,
\[
\forall x \in \mathbb{R}^{k},\|S U x\|_{1} \gtrsim\|U x\|_{1} .
\]

Proof. We can compute a well-conditioned basis for \(U\), and denote it \(U^{\prime}\). Then, \(\forall x \in \mathbb{R}^{k}\), there exists \(y \in \mathbb{R}^{k}\) such that \(U x=U^{\prime} y\). Due to Lemma 17.5.14, with probability .99, we have
\[
\forall y \in \mathbb{R}^{k},\left\|S U^{\prime} y\right\|_{1} \gtrsim\left\|U^{\prime} y\right\|_{1} .
\]

\subsection*{17.5.6 Sparse Cauchy transform}

This section presents the proof of two lemmas related to the sparse Cauchy transform. We first prove the no dilation result in Lemma 17.5.16. Then we show how to get the no contraction result in Lemma 17.5.18.

Lemma 17.5.16. Given matrix \(A \in \mathbb{R}^{n \times d}\), let \(U^{*}, V^{*}\) be the optimal solutions of \(\min _{U, V} \| U V-\) \(A \|_{1}\). Let \(\Pi=\sigma \cdot S C \in \mathbb{R}^{m \times n}\), where \(S \in \mathbb{R}^{m \times n}\) has each column vector chosen independently and uniformly from the \(m\) standard basis vectors of \(\mathbb{R}^{m}\), where \(C\) is a diagonal matrix with diagonals chosen independently from the standard Cauchy distribution, and \(\sigma\) is a scalar. Then
\[
\left\|\Pi U^{*} V^{*}-\Pi A\right\|_{1} \lesssim \sigma \cdot \log (m d) \cdot\left\|U^{*} V^{*}-A\right\|_{1}
\]
holds with probability at least . 999 .

Proof. We define \(\Pi=\sigma \cdot S C \in \mathbb{R}^{m \times n}\), as in the statement of the lemma. Then,
\[
\begin{aligned}
& \left\|\Pi\left(U^{*} V^{*}-A\right)\right\|_{1} \\
= & \sum_{i=1}^{d}\left\|S D\left(U^{*} V_{i}^{*}-A_{i}\right)\right\|_{1} \\
= & \sum_{i=1}^{d}\left\|\left[\begin{array}{cccc}
S_{11} & S_{12} & \cdots & S_{1 n} \\
S_{21} & S_{22} & \cdots & S_{2 n} \\
\cdots & \cdots & \cdots & \cdots \\
S_{m 1} & S_{m 2} & \cdots & S_{m n}
\end{array}\right] \cdot\left[\begin{array}{cccc}
c_{1} & 0 & 0 & 0 \\
0 & c_{2} & 0 & 0 \\
0 & 0 & \cdots & 0 \\
0 & 0 & 0 & c_{n}
\end{array}\right] \cdot\left(U^{*} V_{i}^{*}-A_{i}\right)\right\|_{1} \\
= & \sum_{i=1}^{d}\left\|\left[\begin{array}{cccc}
c_{1} S_{11} & c_{2} S_{12} & \cdots & c_{n} S_{1 n} \\
c_{1} S_{21} & c_{2} S_{22} & \cdots & c_{n} S_{2 n} \\
\cdots & \cdots & \cdots & \cdots \\
c_{1} S_{m 1} & c_{2} S_{m 2} & \cdots & c_{n} S_{m n}
\end{array}\right] \cdot\left(U^{*} V_{i}^{*}-A_{i}\right)\right\|_{1} \\
= & \sum_{i=1}^{d}\left\|\sum_{l=1}^{n} c_{l} S_{1 l} \cdot\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}, \sum_{l=1}^{n} c_{l} S_{2 l} \cdot\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}, \cdots, \sum_{l=1}^{n} c_{l} S_{m l} \cdot\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right\|_{1} \\
= & \sum_{i=1}^{d} \sum_{j=1}^{m}\left|\sum_{l=1}^{n} c_{l} S_{j l} \cdot\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right| \\
= & \sum_{i=1}^{d} \sum_{j=1}^{m}\left|\widetilde{w}_{i j} \cdot \sum_{l=1}^{n}\right| S_{j l}\left(U^{*} V_{i}^{*}-A_{i}\right)_{l} \| \text { where } \widetilde{w}_{i j} \sim C(0,1) \\
= & \sum_{i=1}^{d} \sum_{j=1}^{m} \sum_{l=1}^{n}\left|S_{j l}\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right| \cdot\left|\widetilde{w}_{i j}\right| \\
= & \sum_{i=1}^{d} \sum_{j=1}^{m} \sum_{l=1}^{n}\left|S_{j l}\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right| \cdot w_{i+(j-1) d} \text { where } w_{i+(j-1) d} \sim|C(0,1)|,
\end{aligned}
\]
where the last step follows since each \(w_{i}\) can be thought of as a clipped half-Cauchy random variable. Define \(d^{\prime}=m d\). Define event \(\xi_{i}\) to be the situation when \(w_{i}<D\) (we will decide upon \(D\) later \()\). Define event \(\xi=\xi_{1} \cap \xi_{2} \cap \cdots \cap \xi_{d^{\prime}}\). By choosing \(D=\Theta\left(d^{\prime}\right)\), we can conclude
that,
\[
\begin{aligned}
\mathbb{E}\left[\left\|\Pi U^{*} V^{*}-\Pi A\right\|_{1} \mid \xi\right] & =\sum_{i=1}^{d} \sum_{j=1}^{m} \sum_{l=1}^{n}\left|S_{j l}\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right| \cdot \mathbb{E}\left[w_{i+(j-1) d} \mid \xi\right] \\
& \lesssim \sum_{i=1}^{d} \sum_{j=1}^{m} \sum_{l=1}^{n}\left|S_{j l}\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right| \cdot \log \left(d^{\prime}\right) \\
& =\sum_{i=1}^{d} \sum_{l=1}^{n}\left|\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right|_{1} \cdot \log \left(d^{\prime}\right) \\
& =\log \left(d^{\prime}\right) \cdot\left\|U^{*} V^{*}-A\right\|_{1}
\end{aligned}
\]
where the third step follows from \(\sum_{j=1}^{m}\left|S_{j l}\right|=1, \forall l \in[n]\).
Thus, we can show that
\[
\operatorname{Pr}\left[\left\|\Pi U^{*} V^{*}-\Pi A\right\|_{1} \lesssim \log \left(d^{\prime}\right)\left\|U^{*} V^{*}-A\right\|_{1}\right] \geq 0.999
\]

Lemma 17.5.17. Given any matrix \(M \in \mathbb{R}^{n \times d}\). Let \(\Pi=\sigma \cdot S C \in \mathbb{R}^{m \times n}\), where \(S \in \mathbb{R}^{m \times n}\) has each column vector chosen independently and uniformly from the \(m\) standard basis vectors of \(\mathbb{R}^{m}\), where \(C\) is a diagonal matrix with diagonals chosen independently from the standard Cauchy distribution, and \(\sigma\) is a scalar. Then
\[
\|\Pi M\|_{1} \lesssim \sigma \cdot \log (m d) \cdot\|M\|_{1}
\]
holds with probability at least . 999 .

Proof. Just replace the matrix \(U^{*} V^{*}-A\) in the proof of Lemma 17.5.16 with \(M\). Then we can get the result directly.

We already provided the proof of Lemma 17.5.16. It remains to prove Lemma 17.5.18.

Lemma 17.5.18. Given matrix \(A \in \mathbb{R}^{n \times d}\), let \(U^{*}, V^{*}\) be the optimal solution of \(\min _{U, V} \| U V-\) \(A \|_{1}\). Let \(\Pi=\sigma \cdot S C \in \mathbb{R}^{m \times n}\), where \(S \in \mathbb{R}^{m \times n}\) has each column vector chosen independently and uniformly from the \(m\) standard basis vectors of \(\mathbb{R}^{m}\), and where \(C\) is a diagonal matrix with diagonals chosen independently from the standard Cauchy distribution. Then with probability at least .999 , for all \(V \in \mathbb{R}^{k \times d}\),
\[
\left\|\Pi U^{*} V-\Pi A\right\|_{1} \geq\left\|U^{*} V-A\right\|_{1}-O(\sigma \cdot \log (m d))\left\|U^{*} V^{*}-A\right\|_{1}
\]

Notice that \(m=O\left(k^{5} \log ^{5} k\right)\) and \(\sigma=O\left(k^{2} \log ^{2} k\right)\) according to Theorem 2 in [MM13].

We start by using Theorem 2 in [MM13] to generate the following Corollary.

Corollary 17.5.19. Given \(U \in \mathbb{R}^{n \times k}\) with full column rank, let \(\Pi=\sigma \cdot S C \in \mathbb{R}^{m \times n}\) where \(S \in \mathbb{R}^{m \times n}\) has each column chosen independently and uniformly from the \(m\) standard basis vectors of \(\mathbb{R}^{m}\), and where \(C \in \mathbb{R}^{n \times n}\) is a diagonal matrix with diagonals chosen independently from the standard Cauchy distribution. Then with probability .999, for all \(x \in \mathbb{R}^{k}\), we have
\[
\|\Pi U x\|_{1} \geq\|U x\|_{1} .
\]

Notice that \(m=O\left(k^{5} \log ^{5} k\right)\) and \(\sigma=O\left(k^{2} \log ^{2} k\right)\) according to Theorem 2 in [MM13].

We give the proof of Lemma 17.5.18.

Proof. Follows by Corollary 17.5.19, Lemma 17.5.16, Lemma 17.5.1.

\subsection*{17.5.7 \(\ell_{1}\)-Lewis weights}

In this section we show how to use Lewis weights to get a better dilation bound. The goal of this section is to prove Lemma 17.5.24 and Lemma 17.5.23. Notice that our algorithms in Section 17.4 use Lewis weights in several different ways. The first way is using Lewis weights to show an existence result, which means we only need an existential result for Lewis weights. The second way is only guessing the nonzero locations on the diagonal of a sampling and rescaling matrix according to the Lewis weights. The third way is guessing the values on the diagonal of a sampling and rescaling matrix according to the Lewis weights. The fourth way is computing the Lewis weights for a known low dimensional \(\operatorname{matrix}(n \times \operatorname{poly}(k)\) or \(\operatorname{poly}(k) \times d)\). We usually do not need to optimize the running time of computing Lewis weights for a low-rank matrix to have input-sparsity time.

Claim 17.5.20. Given matrix \(A \in \mathbb{R}^{n \times d}\), let \(B=U^{*} V^{*}-A\). For any distribution \(p=\) \(\left(p_{1}, p_{2}, \ldots, p_{n}\right)\) define random variable \(X\) such that \(X=\left\|B_{i}\right\|_{1} / p_{i}\) with probability \(p_{i}\). Then take any \(m\) independent samples \(X^{1}, X^{2}, \ldots, X^{m}\), let \(Y=\frac{1}{m} \sum_{j=1}^{m} X^{j}\). We have
\[
\operatorname{Pr}\left[Y \leq 1000\|B\|_{1}\right] \geq .999
\]

Proof. We can compute the expectation of \(X^{j}\), for any \(j \in[m]\)
\[
\mathbb{E}\left[X^{j}\right]=\sum_{i=1}^{n} \frac{\left\|B_{i}\right\|_{1}}{p_{i}} \cdot p_{i}=\|B\|_{1}
\]

Then, \(\mathbb{E}[Y]=\frac{1}{m} \sum_{j=1}^{m} \mathbb{E}\left[X^{j}\right]=\|B\|_{1}\). Using Markov's inequality, we have
\[
\operatorname{Pr}\left[Y \geq 1000\|B\|_{1}\right] \leq .001
\]

Lemma 17.5.21. Given matrix \(M \in \mathbb{R}^{n \times d}\), let \(S \in \mathbb{R}^{n \times n}\) be any sampling and rescaling diagonal matrix. Then with probability at least .999,
\[
\|S M\|_{1} \lesssim\|M\|_{1} .
\]

Proof. Just replace the matrix \(B\) in the proof of Claim 17.5.20 with \(M\). Then we can get the result directly.

Using Theorem 1.1 of [CP15], we have the following result,

Claim 17.5.22. Given matrix \(A \in \mathbb{R}^{n \times d}\), for any fixed \(U^{*} \in \mathbb{R}^{n \times k}\) and \(V^{*} \in \mathbb{R}^{k \times d}\), choose \(D \in \mathbb{R}^{n \times n}\) to be the sampling and rescaling diagonal matrix with \(m=O(k \log k)\) nonzeros according to the Lewis weights of \(U^{*}\). Then with probability .999, for all \(V\),
\[
\left\|U^{*} V^{*}-U^{*} V\right\|_{1} \leq\left\|D U^{*} V^{*}-D U^{*} V\right\|_{1} \lesssim\left\|U^{*} V^{*}-U^{*} V\right\|_{1}
\]

Lemma 17.5.23. Given matrix \(A \in \mathbb{R}^{n \times d}, U^{*} \in \mathbb{R}^{n \times k}\), define \(V^{*} \in \mathbb{R}^{k \times d}\) to be the optimal solution of \(\min _{V \in \mathbb{R}^{k \times d}}\left\|U^{*} V-A\right\|_{1}\). Choose a sampling and rescaling diagonal matrix \(D \in \mathbb{R}^{n \times n}\) with \(m=O(k \log k)\) non-zero entries according to the Lewis weights of \(U^{*}\). Then with probability at least .99, we have: for all \(V \in \mathbb{R}^{k \times d}\),
\[
\left\|D U^{*} V-D A\right\|_{1} \lesssim\left\|U^{*} V^{*}-U^{*} V\right\|_{1}+O(1)\left\|U^{*} V^{*}-A\right\|_{1} \lesssim\left\|U^{*} V-A\right\|_{1}
\]
holds with probability at least . 99 .

Proof. Using the above two claims, we have with probability at least .99 , for all \(V \in \mathbb{R}^{k \times d}\),
\[
\begin{aligned}
\left\|D U^{*} V-D A\right\|_{1} & \leq\left\|D U^{*} V-D U^{*} V^{*}\right\|_{1}+\left\|D U^{*} V^{*}-D A\right\|_{1} \\
& \lesssim\left\|D U^{*} V-D U^{*} V^{*}\right\|_{1}+O(1)\left\|U^{*} V^{*}-A\right\|_{1} \\
& \lesssim\left\|U^{*} V-U^{*} V^{*}\right\|_{1}+O(1)\left\|U^{*} V^{*}-A\right\|_{1} \\
& \leq\left\|U^{*} V-A\right\|_{1}+\left\|U^{*} V^{*}-A\right\|_{1}+O(1)\left\|U^{*} V^{*}-A\right\|_{1} \\
& \lesssim\left\|U^{*} V-A\right\|_{1}
\end{aligned}
\]
where the first step follows from triangle inequality, the second step follows from Claim 17.5.20, the third step follows from Claim 17.5.22, the forth step follows from triangle inequality.

Lemma 17.5.24. Given matrix \(A \in \mathbb{R}^{n \times d}\), define \(U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times d}\) to be the optimal solution of \(\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|U V-A\|_{1}\). Choose a sampling and rescaling diagonal matrix \(D \in\) \(\mathbb{R}^{n \times n}\) with \(m=O(k \log k)\) non-zero entries according to the Lewis weights of \(U^{*}\). For all \(V \in \mathbb{R}^{k \times d}\) we have
\[
\left\|U^{*} V-A\right\|_{1} \lesssim\left\|D U^{*} V-D A\right\|_{1}+O(1)\left\|U^{*} V^{*}-A\right\|_{1},
\]
holds with probability at least . 99 .

Proof. Follows by Claim 17.5.20, Lemma 17.5.23, Lemma 17.5.1.

\subsection*{17.6 Hardness Results for Cauchy Matrices, Row Subset Selection, OSE}

Section 17.6 .1 presents some inapproximability results by using random Cauchy matrices. Section 17.6.2 is a warmup for inapproximability results for row subset selection problems. Section 17.6 .3 shows the inapproximability results by using any linear oblivious subspace embedding(OSE), and also shows inapproximability results for row subset selection.

\subsection*{17.6.1 Hard instance for Cauchy matrices}

The goal of this section is to prove Theorem 17.6.2. Before stating the result, we first introduce some useful tools in our analysis.

Lemma 17.6.1 (Cauchy Upper Tail Inequality, Lemma 3 of \(\left.\left[\mathrm{CDMI}^{+} 13\right]\right)\). For \(i \in[m]\), let \(C_{i}\) be \(m\) random Cauchy variables from \(C(0,1)\) (not necessarily independent), and \(\gamma_{i}>0\) with \(\gamma=\sum_{i \in[m]} \gamma_{i}\). Let \(X=\sum_{i \in[m]} \gamma_{i}\left|C_{i}\right|\). Then, for any \(t \geq 1\),
\[
\operatorname{Pr}[X>\gamma t] \leq O(\log (m t) / t)
\]

Theorem 17.6.2. Let \(k \geq 1\). There exist matrices \(A \in \mathbb{R}^{d \times d}\) such that for any \(o(\log d) \geq\) \(t \geq 1\), where \(c\) can be any constant smaller than \(1 / 3\), for random Cauchy matrices \(S \in \mathbb{R}^{t \times d}\) where each entry is sampled from an i.i.d. Cauchy distribution \(C(0, \gamma)\) where \(\gamma\) is an arbitrary real number, with probability .99 we have
\[
\min _{U \in \mathbb{R}^{d \times t}}\|U S A-A\|_{1} \geq \Omega(\log d /(t \log t))_{\text {rank }-k} \min _{A^{\prime}}\left\|A^{\prime}-A\right\|_{1} .
\]

Proof. We define matrix \(A \in \mathbb{R}^{d \times d}\)
\[
A=B+I=\alpha \cdot\left[\begin{array}{ccccc}
1 & 1 & 1 & \cdots & 1 \\
0 & 0 & 0 & \cdots & 0 \\
0 & 0 & 0 & \cdots & 0 \\
\cdots & \cdots & \cdots & \cdots & \cdots \\
0 & 0 & 0 & \cdots & 0
\end{array}\right]+\left[\begin{array}{ccccc}
1 & 0 & 0 & \cdots & 0 \\
0 & 1 & 0 & \cdots & 0 \\
0 & 0 & 1 & \cdots & 0 \\
\cdots & \cdots & \cdots & \cdots & \cdots \\
0 & 0 & 0 & \cdots & 1
\end{array}\right]
\]
where \(\alpha=\Theta(\log d)\).So, if we only fit the first row of \(A\), we can get approximation cost at most \(O(d)\).

For \(t>0\), let \(S \in \mathbb{R}^{t \times d}\) denote a random Cauchy matrix where \(S_{i, j}\) denotes the entry
in the \(i\) th row and \(j\) th column. Then \(S A\) is
\[
S A=S B+S I=\alpha \cdot\left[\begin{array}{ccccc}
S_{1,1} & S_{1,1} & S_{1,1} & \cdots & S_{1,1} \\
S_{2,1} & S_{2,1} & S_{2,1} & \cdots & S_{2,1} \\
S_{3,1} & S_{3,1} & S_{3,1} & \cdots & S_{3,1} \\
\cdots & \cdots & \cdots & \cdots & \cdots \\
S_{t, 1} & S_{t, 1} & S_{t, 1} & \cdots & S_{t, 1}
\end{array}\right]+\left[\begin{array}{ccccc}
S_{1,1} & S_{1,2} & S_{1,3} & \cdots & S_{1, d} \\
S_{2,1} & S_{2,2} & S_{2,3} & \cdots & S_{2, d} \\
S_{3,1} & S_{3,2} & S_{3,3} & \cdots & S_{3, d} \\
\cdots & \cdots & \cdots & \cdots & \cdots \\
S_{t, 1} & S_{t, 2} & S_{t, 3} & \cdots & S_{t, d}
\end{array}\right] .
\]

Since \(\forall \gamma\),
\[
\min _{U \in \mathbb{R}^{d \times t}}\|U S A-A\|_{1}=\min _{U \in \mathbb{R}^{d \times t}}\|\gamma U S A-A\|_{1},
\]
without loss of generality, we can let \(S_{i, j} \sim C(0,1)\). Then we want to argue that, if we want to use \(S A\) to fit the first row of \(A\), with high probability, the cost will be \(\Omega(d \log d)\).

Let \(b \in \mathbb{R}^{d}\) denote the first row of \(A\). Then, we want to use \(S A\) to fit the first row of \(A\), which is a \(d\)-dimensional vector that has entry \(\Theta(\log d)\) on each position. The problem is equivalent to
\[
\min _{x \in \mathbb{R}^{t}}\left\|(S A)^{\top} x-b\right\|_{1} .
\]

First, we want to show that for any \(x\) in \(\mathbb{R}^{t}\), if \(x^{\top} S A\) fits the first row of \(A\) very well, then the \(\ell_{1}\) and \(\ell_{2}\) norm of vector \(x\) must have reasonable size.

Claim 17.6.3. Define \(A^{1}\) to be the first row of matrix \(A\). With probability .999, for any column vector \(x \in \mathbb{R}^{t \times 1}\), if \(\left\|x^{\top} S A-A^{1}\right\|_{1} \leq o(d \log d)\), then

Property (I) : \(\|x\|_{2} \geq \Omega(1 / t \log t) ;\)
Property (II) : \(\|x\|_{1} \leq O(\log d)\).

Proof. Consider the absolute value of the \(i\)-th coordinate of \(x^{\top} S A\). We can rewrite \(\left|\left\langle(S A)_{i}, x\right\rangle\right|\) in the following sense,
\[
\begin{align*}
\left|\left\langle(S A)_{i}, x\right\rangle\right| & =\left|\left\langle(S B)_{i}, x\right\rangle+\left\langle(S I)_{i}, x\right\rangle\right| \\
& =\left|\left\langle(S B)_{1}, x\right\rangle+\left\langle(S I)_{i}, x\right\rangle\right| \\
& =\left|\left\langle(S B)_{1}, x\right\rangle+\left\langle S_{i}, x\right\rangle\right| \tag{17.21}
\end{align*}
\]
where the second step follows because \((S B)_{i}=(S B)_{1}, \forall i \in[n]\), and the last step follows because \((S I)_{i}=S_{i}, \forall i \in[n]\).

We start by proving Property I. Using the triangle inequality and Equation (17.21),
\[
\begin{aligned}
&\left|\left\langle(S A)_{i}, x\right\rangle\right| \\
& \leq\left|\left\langle(S B)_{1}, x\right\rangle\right|+\left|\left\langle S_{i}, x\right\rangle\right| \\
& \leq\left\|(S B)_{1}\right\|_{2}\|x\|_{2}+\left\|S_{i}\right\|_{2}\|x\|_{2} \\
& \leq\left\|(S B)_{1}\right\|_{1}\|x\|_{2}+\left\|S_{i}\right\|_{1}\|x\|_{2} . \\
& \text { by Cauchy-Schwarz inequality }
\end{aligned}
\]

Then, according to Lemma 17.6.1, with probability .99999, we have \(\left\|(S B)_{1}\right\|_{1} \leq\) \(O(t \log t \log d)\) and for a fixed \(i \in[d],\left\|S_{i}\right\|_{1} \leq O(t \log t)\). Applying the Chernoff bound, with probability \(1-2^{-\Omega(d)}\), there are a constant fraction of \(i\) such that \(\left\|S_{i}\right\|_{1}=O(t \log t)\). Taking the union bound, with probability .9999, there exists a constant fraction of \(i\) such that \(\left|\left\langle(S A)_{i}, x\right\rangle\right| \leq O(t \log t \log d)\|x\|_{2}\). Because \(A_{1, i} \geq \alpha, \forall i \in[d]\) where \(\alpha=\Theta(\log d)\), we need \(\|x\|_{2} \geq \Omega(1 / t \log t)\). Otherwise, the total cost on this constant fraction of coordinates will be at least \(\Omega(d \log d)\).

For Property II, for a fixed \(x \in \mathbb{R}^{t}\), we have
\[
\begin{aligned}
& \left|\left\langle(S A)_{i}, x\right\rangle\right| \\
= & \left|\left\langle(S B)_{1}, x\right\rangle+\left\langle S_{i}, x\right\rangle\right| \quad \text { by Equation }(17.21) \\
= & \left|\Theta(\log d)\|x\|_{1} w_{1}^{\prime}(x)+\|x\|_{1} w_{i}^{\prime}(x)\right|
\end{aligned}
\]
where \(w_{i}^{\prime}(x) \sim C(0,1)\), and for different \(x, w_{i}^{\prime}(x)\) are different. Then for a fixed \(x \in \mathbb{R}^{t}\) with probability at least \(0.9,\left|w_{i}^{\prime}(x)\right|=\Omega(1)\), and with probability \(0.5, w_{1}^{\prime}(x)\) and \(w_{i}^{\prime}(x)\) have the same sign. Since these two events are independent, with probability at least 0.45 , we have
\[
\left|\left\langle(S A)_{i}, x\right\rangle\right| \geq\|x\|_{1} \cdot \Omega(1)
\]

Applying the Chernoff bound, with probability at least \(1-2^{\Theta(d)}\), there exists a \(3 / 10\) fraction of \(i\) such that \(\left|\left\langle(S A)_{i}, x\right\rangle\right| \geq\|x\|_{1} \Omega(1)\).

We build an \(\epsilon\)-net \(\mathcal{N}\) for \(x \in \mathbb{R}^{t}\) on an \(\ell_{1}\)-norm unit ball, where \(\epsilon=1 /\left(t^{2} \log ^{2} d\right)\). Thus the size of the net is \(|\mathcal{N}|=2^{\widetilde{\Theta}(t)}\). Consider \(y\) to be an arbitrary vector, let \(y /\|y\|_{1}=x+\delta\), where \(x\) is the closest point to \(y /\|y\|_{1}\) and \(x \in \mathcal{N}\). For any \(\delta \in \mathbb{R}^{t}\),
\[
\begin{array}{rlr}
\left|\left\langle(S A)_{i}, \delta\right\rangle\right| & =\left|\left\langle(S B)_{1}, \delta\right\rangle+\left\langle(S I)_{i}, \delta\right\rangle\right| & \\
& \leq\left|\left\langle(S B)_{1}, \delta\right\rangle\right|+\left|\left\langle(S I)_{i}, \delta\right\rangle\right| & \text { by triangle inequality } \\
& \leq\left\|(S B)_{1}\right\|_{2}\|\delta\|_{2}+\left\|S_{i}\right\|_{2}\|\delta\|_{2} & \text { by Cauchy-Schwarz inequality } \\
& \leq\left\|(S B)_{1}\right\|_{2}\|\delta\|_{1}+\left\|S_{i}\right\|_{2}\|\delta\|_{1} . & \text { by }\|\cdot\|_{2} \leq\|\cdot\|_{1}
\end{array}
\]

As we argued before, With probability .99999, we have
\[
\begin{equation*}
\left\|(S B)_{1}\right\|_{2} \leq\left\|(S B)_{1}\right\|_{1} \leq O(t \log t \log d) \tag{17.22}
\end{equation*}
\]
and with probability \(1-2^{-\Theta(d)}\), there is a \(9 / 10\) fraction of \(i \in[d],\left\|S_{i}\right\|_{2} \leq\left\|S_{i}\right\|_{1}=O(t \log t)\). Therefore, with probability .999 , for any \(\delta \in \mathbb{R}^{t}\), there exists a \(9 / 10\) fraction of \(i\) such that \(\left|\left\langle(S A)_{i}, \delta\right\rangle\right| \leq \Theta(t \log t \log d)\|\delta\|_{1}\).

Therefore, with probability . \(99, \forall y \in \mathbb{R}^{t}\), due to the pigeonhole principle, there is a \(3 / 10+9 / 10-1=1 / 5\) fraction of \(i\) such that
\[
\begin{array}{rlr} 
& \left|\left\langle(S A)_{i}, y\right\rangle\right| & \\
= & \|y\|_{1} \cdot\left|\left\langle(S A)_{i}, y /\|y\|_{1}\right\rangle\right| & \\
= & \|y\|_{1} \cdot\left|\left\langle(S A)_{i}, x+\delta\right\rangle\right| & \text { by } y /\|y\|_{1}=x+\delta \\
\geq & \|y\|_{1} \cdot\left(\left|\left\langle(S B)_{i}, x\right\rangle+\left\langle(S I)_{i}, x\right\rangle\right|-\left|\left\langle(S B)_{1}, \delta\right\rangle+\left\langle(S I)_{i}, \delta\right\rangle\right|\right) & \text { by triangle inequality } \\
\geq & \|y\|_{1}(\Omega(1)-\epsilon O(t \log t \log d)) & \\
\geq & \|y\|_{1} \Omega(1) &
\end{array}
\]

So \(\|y\|_{1}\) should be \(O(\log d)\). Otherwise, the total cost on this \(1 / 5\) fraction of coordinates is at least \(\Omega(d \log d)\).

Combining Property (I) and (II) completes the proof.

Next, we need to show the following claim is true,
Claim 17.6.4. For any d independent Cauchy random variables \(x_{1}, x_{2}, \cdots, x_{d}\) from \(C(0,1)\), with probability \(1-1 / \operatorname{poly}(t \log d)\), for any \(j \in[1,2, \cdots, \log d-\Theta(\log \log (t \log d))]\), there are \(\Omega\left(d / 2^{j}\right)\) variables belonging to \(\left(2^{j}, 2^{j+1}\right]\).

Proof. For each Cauchy random variable \(x_{i}\), we have for any \(j \in[1,2, \cdots, \Theta(\log d)]\),
\[
\operatorname{Pr}\left[\left|x_{i}\right| \in\left(2^{j}, 2^{j+1}\right]\right]=\Theta\left(1 / 2^{j}\right) .
\]

We define the indicator random variable \(z_{i, j}\)
\[
z_{i, j}= \begin{cases}1 & \text { if }\left|x_{i}\right| \in\left(2^{j}, 2^{j+1}\right] \\ 0 & \text { otherwise }\end{cases}
\]

We define \(z_{j}=\sum_{i=1}^{d} z_{i, j}\). It is clear that \(\mathbb{E}\left[z_{i, j}\right]=\Theta\left(1 / 2^{j}\right)\). We use a Chernoff bound,
\[
\operatorname{Pr}[X<(1-\delta) \mu]<\left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right)^{\mu}
\]
and set \(X=z_{j}, \delta=1 / 2, \mu=\mathbb{E}\left[z_{j}\right]\). Then, this probability is at most \(2^{-\Omega(\mu)}=2^{-\Omega\left(\mathbb{E}\left[z_{j}\right]\right)}\). For any \(j \in[1, \log d-\Theta(\log \log (t \log d))]\), we have \(\mathbb{E}\left[z_{j}\right]=d \mathbb{E}\left[z_{i j}\right]=\Omega\left(d / 2^{j}\right)=\Omega(\log (t \log d))\). Thus, this probability is at most \(1 / \operatorname{poly}(t \log d)\). Overall, we have
\[
\operatorname{Pr}\left[z_{j} \gtrsim d / 2^{j}\right] \geq 1-1 / \operatorname{poly}(t \log d)
\]

Taking a union bound over \(\Theta(\log d)\) such \(j\), we complete the proof.

Claim 17.6.5. Let \(1>c_{1}>c_{2}>1 / 3>0\) be three arbitrary constants. We fix the first column of \((S I)^{\top} \in \mathbb{R}^{d \times t}\). All the rows are grouped together according to the value in the first column. Let \(R_{j}\) be the set of rows for which the entry in the first column is \(\in\left(2^{j}, 2^{j+1}\right]\). With probability at least \(1-O(1 / \operatorname{poly}(t \log (d)))\), for any \(j \in\left[c_{2} \log d, c_{1} \log d\right]\), the following event holds. There exist \(\Omega\left(d / 2^{j}\right)\) rows such that the first coordinate \(\in\left(2^{j}, 2^{j+1}\right]\) and all the other coordinates are at most \(O\left(d^{1 / 3}\right)\).

Proof. Let \(R_{j}\) be a subset of rows of \(S^{\top}\), such that for any row in \(R_{j}\), the first coordinate is \(\in\left(2^{j}, 2^{j+1}\right]\). By Claim 17.6.4, we have that with probability \(1-1 / \operatorname{poly}(t \log d)\), for any \(j \in\left[c_{2} \log d, c_{1} \log d\right],\left|R_{j}\right| \geq \Omega\left(d / 2^{j}\right)\). We want to show that for any \(j \in\left[c_{2} \log d, c_{1} \log d\right]\),
there exists a constant fraction of rows in \(R_{j}\) such that, the remaining coordinates are at most \(O\left(d^{1 / 3}\right)\).

For a fixed row in \(R_{j}\) and a fixed coordinate in that row, the probability that the absolute value of this entry is at least \(\Omega\left(d^{1 / 3}\right)\) is at most \(O\left(1 / d^{1 / 3}\right)\). By taking a union bound, with probability at least \(1-O\left(t / d^{1 / 3}\right)\), the row has every coordinate of absolute value at \(\operatorname{most} O\left(d^{1 / 3}\right)\).

By applying a Chernoff bound, for a fixed subset \(R_{j}\) of rows, the probability that there is a constant fraction of these rows for which every coordinate except the first one in absolute value is at most \(O\left(d^{1 / 3}\right)\), is at least \(1-2^{-\Theta\left(\left|R_{j}\right|\right)} \geq 1-2^{-\Theta\left(d^{1-c_{1}}\right)} \geq 1-O(1 / \operatorname{poly}(t \log (d)))\).

After taking a union over all the \(j\), we complete the proof.

Claim 17.6.6. With probability at least \(1-O(1 / t)\), the absolute value of any coordinate in column vector \((S B)_{1} \in \mathbb{R}^{t \times 1}\) is at most \(O\left(t^{2} \log d\right)\).

Proof. Because each entry is sampled from a Cauchy distribution \(C(0,1)\) scaled by \(O(\log d)\), then with probability at most \(O\left(1 / t^{2}\right)\), the absolute value of one coordinate is at least \(\Omega\left(t^{2} \log d\right)\). Taking a union over all the coordinates, we complete the proof.

Because \(\|x\|_{1} \leq O(\log d),\|x\|_{2} \geq \Omega(1 / t \log t)\), there exists one coordinate \(i\) such that the absolute value of it is at least \(\Omega(1 / t \log t)\) and all the other coordinates are most \(O(\log d)\). We can assume that \(i=1\) for now. (To remove this assumption, we can take a union bound over all the possibilities for \(i \in[t]\).) According to Claim 17.6.5 and 17.6.6, let \(\widehat{R}_{j}\) denote a subset of \(R_{j}\), which is a "good" constant fraction of \(R_{j}\). Considering the \(\ell_{1}\)-norm of all
coordinates \(l \in \widehat{R}_{j} \subset R_{j} \subset[d]\), we have
\[
\begin{aligned}
& \sum_{l \in \widehat{R}_{j}}\left|\left((S A)^{\top} x\right)_{l}-O(\log d)\right| \\
\geq & \sum_{l \in \widehat{R}_{j}}\left|\left\langle(S A)_{l}, x\right\rangle-O(\log d)\right| \\
\geq & \sum_{l \in \widehat{R}_{j}}\left(\left|(S I)_{1, l} \cdot x_{1}\right|-\sum_{j=2}^{t}\left|(S I)_{j, l} \cdot x_{j}\right|-\left|\left\langle(S B)_{l}, x\right\rangle\right|-O(\log d)\right) \\
\geq & \sum_{l \in \widehat{R}_{j}}\left(\Omega\left(\frac{2^{j}}{t \log t}\right)-O\left(t d^{1 / 3} \log d\right)-\left|\left\langle(S B)_{1}, x\right\rangle\right|-O(\log d)\right) \\
\geq & \sum_{l \in \widehat{R}_{j}}\left(\Omega\left(\frac{2^{j}}{t \log t}\right)-O\left(t d^{1 / 3} \log d\right)-\left\|(S B)_{1}\right\|_{1}\|x\|_{1}-O(\log d)\right) \\
\geq & \sum_{l \in \widehat{R}_{j}}\left(\Omega\left(\frac{2^{j}}{t \log t}\right)-O\left(t d^{1 / 3} \log d\right)-O\left(t \log t \log ^{2} d\right)-O(\log d)\right) \\
\gtrsim & \sum_{l \in \widehat{R}_{j}} 2^{j} /(t \log t) \\
\gtrsim & d / 2^{j} \cdot 2^{j} /(t \log t) \\
\gtrsim & d /(t \log t) .
\end{aligned}
\]

The second inequality follows by the triangle inequality. The third inequality follows by \((S B)_{1}=(S B)_{l},\left|x_{1}\right|=\Omega(1 / t \log t),(S I)_{1, l} \in\left[2^{j}, 2^{j+1}\right)\), and \(\forall j \neq 1,\left|x_{j}\right|<O(\log d),(S I)_{j, l} \leq\) \(O\left(d^{1 / 3}\right)\). The fourth inequality follows by Cauchy-Schwarz and \(\|\cdot\|_{2} \leq\|\cdot\|_{1}\). The fifth inequality follows by Equation (17.22) and Claim 17.6.3. The sixth inequality follows by \(t=o(\log d)\) where \(c\) is a constant smaller than \(1 / 3\) and \(2^{j} \geq d^{c_{2}}>\operatorname{poly}(t)\). The seventh inequality follows from \(\left|\widehat{R}_{j}\right| \geq \Omega\left(d / 2^{j}\right)\).

Since there are \(c 1-c 2\) different \(j\), the total cost is \(\Omega(d \log d /(t \log t))\). The gap then is \(\Omega(\log d /(t \log t))\). This completes the proof of Theorem 17.6.2.

\subsection*{17.6.2 Hard instance for row subset selection}

Theorem 17.6.7. Let \(\epsilon \in(0,1)\). There exists a value \(k \geq 1\) and matrix \(A \in \mathbb{R}^{(d-1) \times d}\) such that, for any subset \(R\) of rows of \(A\), letting \(B\) be the \(\ell_{1}\)-projection of each row of \(A\) onto \(R\), we have
\[
\|A-B\|_{1}>(2-\epsilon)_{\text {rank }-k} \min _{A^{\prime}}\left\|A^{\prime}-A\right\|_{1},
\]
unless \(|R| \gtrsim \epsilon d\).

Proof. We construct the \((d-1) \times d\) matrix \(A\) in the following way. The first column is all 1 s , and then the remaining \((d-1) \times(d-1)\) matrix is the identity matrix.
\[
A=\left[\begin{array}{cccccc}
1 & 1 & 0 & 0 & \cdots & 0 \\
1 & 0 & 1 & 0 & \cdots & 0 \\
1 & 0 & 0 & 1 & \cdots & 0 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\
1 & 0 & 0 & 0 & \cdots & 1
\end{array}\right]
\]

Note that \(\min _{\operatorname{rank}-1}{ }_{A^{\prime}}\left\|A-A^{\prime}\right\|_{1}=\mathrm{OPT} \leq d-1\), since one could choose \(A^{\prime}\) to have the first column all 1 s and remaining entries 0 . On the other hand, consider any subset of \(r\) rows. We can permute the columns and preserve the entrywise \(\ell_{1}\)-norm so w.l.o.g., we can take the first \(r\) rows for \(R\).

Because we are taking the first \(r\) rows, we do not pay any cost on the first rows. To minimize the cost on the \(i\)-th row, where \(i \in\{r+1, r+2, \ldots, d\}\), let \(v_{i}\) denote the row vector we use for the \(i\)-th row. Then \(v_{i}\) can be written as a linear combination of the first \(r\) rows \(\left\{A^{1}, A^{2}, \ldots, A^{r}\right\}\),
\[
v_{i}=\sum_{j=1}^{r} \alpha_{i, j} A^{j}
\]

Then the cost of using \(v_{i}\) to approximate the \(i\)-th row of A is:
\[
\begin{aligned}
\left\|v_{i}-A^{i}\right\|_{1} & =(\text { cost in } 1 \text { st col })+(\operatorname{cost} \text { in } 2 \mathrm{nd}, 3 \mathrm{rd}, \ldots, r+1 \text { th cols })+(\operatorname{cost} \text { in } i+1 \text { th col }) \\
& =\left|\sum_{j=1}^{r} \alpha_{i, j}-1\right|+\sum_{j=1}^{r}\left|\alpha_{i, j}\right|+1 \\
& \geq\left|\sum_{j=1}^{r} \alpha_{i, j}-1-\sum_{j=1}^{r} \alpha_{i, j}\right|+1 \text { by triangle inequality } \\
& =2 .
\end{aligned}
\]

Hence, the cost of using \(v_{i}\) to approximate the \(i\)-th row of \(A\) is at least 2. So in total, across these \((d-1-r)\) rows, the algorithm pays at least \(2(d-1-r)\) cost, which needs to be at most \(C(d-1)\), and therefore \(r \geq(d-1)(1-C / 2)=\Omega(\epsilon d)\). Choosing \(C=2-\epsilon\) completes the proof.

Theorem 17.6.8. There exists a value \(k \geq 1\) and matrix \(A \in \mathbb{R}^{(d-1) \times d}\) such that, there is no algorithm that is able to output a rank - \(k\) matrix \(B\) in the row span of \(A\) satisfying
\[
\|A-B\|_{1}<2\left(1-\Theta\left(\frac{1}{d}\right)\right) \min _{\text {rank }-k}\left\|A^{\prime}-A\right\|_{1} .
\]

Proof. We use the same matrix \(A\) as in the previous theorem.
\[
A=\left[\begin{array}{cccccc}
1 & 1 & 0 & 0 & \cdots & 0 \\
1 & 0 & 1 & 0 & \cdots & 0 \\
1 & 0 & 0 & 1 & \cdots & 0 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\
1 & 0 & 0 & 0 & \cdots & 1
\end{array}\right]
\]

Let vector \(v\) be \(\left(\sum_{i=1}^{d-1} \beta_{i}, \beta_{1}, \ldots, \beta_{d-1}\right)\). For the \(i\)-th row of \(A\), we use \(\alpha_{j} \cdot v\) to cancel the cost of that row, where \(\alpha_{j}\) is a scalar. Then for any \(\beta_{1}, \ldots, \beta_{d-1}, \alpha_{1}, \ldots, \alpha_{d-1}\), we can
compute the entire residual cost,
\[
\begin{aligned}
f(\alpha, \beta) & =\sum_{j=1}^{d-1}\left\|A^{j}-v \alpha_{j}\right\|_{1} \\
& =\sum_{j=1}^{d-1}\left(\left|1-\alpha_{j}\left(\sum_{i=1}^{d-1} \beta_{i}\right)\right|+\left|1-\alpha_{j} \beta_{j}\right|+\sum_{i \neq j}\left|\alpha_{j} \beta_{i}\right|\right) .
\end{aligned}
\]

In the next few paragraphs, we will show that optimizing \(f(\alpha, \beta)\) over some extra constraints for \(\alpha, \beta\) does not change the optimality. Without loss of generality, we can assume that \(\sum_{i=1}^{d-1} \beta_{i}=1\). If not we can just rescale all the \(\alpha_{i}\). Consider a fixed index \(j\). All the terms related to \(\alpha_{j}\) and \(\beta_{j}\) are,
\[
\left|1-\alpha_{j}\left(\sum_{i=1}^{d-1} \beta_{i}\right)\right|+\left|1-\alpha_{j} \beta_{j}\right|+\sum_{i \neq j}\left|\alpha_{j} \beta_{i}\right|+\left|\alpha_{i} \beta_{j}\right| .
\]

We first show a simple proof by assuming \(\beta_{j} \geq 0\). Later, we will prove the general version which does not have any assumptions.

Handling a special case We can optimize \(f(\alpha, \beta)\) in the following sense,
\[
\begin{array}{ll}
\min & f(\alpha, \beta) \\
\text { s.t. } & \sum_{j=1}^{d-1} \beta_{j}=1 \\
& \alpha_{j} \geq 0, \beta_{j} \geq 0, \forall j \in[d-1] .
\end{array}
\]

For each \(j\), we consider three cases. Case I, if \(\alpha_{j} \geq \frac{1}{\beta_{j}}\), then
\[
\begin{aligned}
& \left|1-\alpha_{j}\right|+\left|1-\alpha_{j} \beta_{j}\right|+\sum_{i \neq j} \alpha_{j} \beta_{i}= \\
= & 2 \alpha_{j}-2 \geq 2\left(1 / \beta_{j}-1\right) \geq 2\left(1-\beta_{j}\right),
\end{aligned}
\]
where the last step follows by \(\beta_{j}+1 / \beta_{j} \geq 2\). Case II, if \(1 \leq \alpha_{j}<1 / \beta_{j}\), then
\[
\begin{aligned}
& \left|1-\alpha_{j}\right|+\left|1-\alpha_{j} \beta_{j}\right|+\sum_{i \neq j} \alpha_{j} \beta_{i}=\quad \alpha_{j}-1+1-\alpha_{j} \beta_{j}+\sum_{i \neq j} \alpha_{j} \beta_{i} \\
= & 2 \alpha_{j}\left(1-\beta_{j}\right) \geq 2\left(1-\beta_{j}\right) .
\end{aligned}
\]

Case III, if \(\alpha_{j}<1\), then
\[
\begin{aligned}
& \left|1-\alpha_{j}\right|+\left|1-\alpha_{j} \beta_{j}\right|+\sum_{i \neq j} \alpha_{j} \beta_{i}=\quad 1-\alpha_{j}+1-\alpha_{j} \beta_{j}+\sum_{i \neq j} \alpha_{j} \beta_{i} \\
= & 2\left(1-\alpha_{j} \beta_{j}\right) \geq 2\left(1-\beta_{j}\right) .
\end{aligned}
\]

Putting it all together, we have
\[
f(\alpha, \beta) \geq \sum_{j=1}^{d-1} 2\left(1-\beta_{j}\right)=2(d-1)-2 \sum_{j=1}^{d-1} \beta_{j}=2(d-2) .
\]

To handle the case where \(\beta_{i}\) can be negative Without loss of generality, we can assume that \(\sum_{i=1}^{d-1} \beta_{i}=1\). Notice that we can also assume \(\beta_{i} \neq 0\), otherwise it means we do not choose that row. We split all \(\beta_{i}\) into two disjoint sets \(S\) and \(\bar{S}\). For any \(i \in S, \beta_{i}>0\) and for any \(i \in \bar{S}, \beta_{i}<0\).

As a first step, we discuss the case when all the \(j\) are in set \(S\). Case I, if \(1-\) \(\alpha_{j} \sum_{i=1}^{d-1} \beta_{i}<0\) and \(1-\alpha_{j} \beta_{j}<0\), then it means \(\alpha_{j} \geq \max \left(\frac{1}{\sum_{i=1}^{d-1} \beta_{i}}, \frac{1}{\beta_{j}}\right)\). The cost of that row is,
\[
\begin{aligned}
& =\alpha_{j} \sum_{i=1}^{d-1} \beta_{i}-1+\alpha_{j} \beta_{j}-1+\sum_{i \neq j}\left|\alpha_{j} \beta_{i}\right| \\
& =2 \alpha_{j} \beta_{j}+2 \alpha_{j} \sum_{i \in S \backslash j} \beta_{i}-2 .
\end{aligned}
\]

If \(\beta_{j} \geq 1\), then \(\alpha_{j} \geq 1\). The cost is at least \(2 \sum_{i \in S \backslash j} \beta_{i}\). If \(\beta_{j}<1\), then \(\alpha_{j} \geq 1 / \beta_{j}\), and the cost is at least \(2 \frac{1}{\beta_{j}} \sum_{i \in S \backslash j} \beta_{i}\). If there are \(C\) such \(j\) in Case I, then the total cost of Case I is at least 0 if \(C=1\), and \(2(C-1)\) if \(C \geq 2\).

Case II, if \(1-\alpha_{j} \sum_{i=1}^{d-1} \beta_{i}<0\) and \(1-\alpha_{j} \beta_{j}>0\), then it means \(1<\alpha_{j}<1 / \beta_{j}\). The cost of that row is,
\[
\begin{aligned}
& =\alpha_{j} \sum_{i=1}^{d-1} \beta_{i}-1+1-\alpha_{j} \beta_{j}+\left|\alpha_{j}\right| \sum_{i \neq j}\left|\beta_{j}\right| \\
& =2 \alpha_{j} \sum_{i \in S \backslash j} \beta_{i} \\
& \geq 2 \sum_{i \in S \backslash j} \beta_{i} .
\end{aligned}
\]

Similarly to before, if there are \(C\) such \(j\) in Case II, then the total cost of Case II is at least 0 if \(C=1\), and \(2(C-1)\) if \(C \geq 2\).
\[
\text { Case III, if } 1-\alpha_{j} \sum_{i=1}^{d-1} \beta_{i}>0 \text { and } 1-\alpha_{j} \beta_{j}<0 \text {, then it means } 1 / \beta_{j}<\alpha_{j}<\frac{1}{\sum_{i=1}^{d-1} \beta_{i}} \text {. }
\] The cost of the row is,
\[
\begin{aligned}
& =1-\alpha_{j} \sum_{i=1}^{d-1} \beta_{i}+\alpha_{j} \beta_{j}-1+\left|\alpha_{j}\right| \sum_{i \neq j}\left|\beta_{j}\right| \\
& =2\left|\alpha_{j}\right| \sum_{i \in \bar{S}}\left|\beta_{i}\right| \\
& =2\left|\alpha_{j}\right|\left(\sum_{i \in S}\left|\beta_{i}\right|-1\right) \\
& \geq 2 \frac{1}{\beta_{j}} \sum_{i \in S \backslash j} \beta_{i} .
\end{aligned}
\]

If there are \(C\) such \(j\) in Case III, then the total cost of Case III is at least 0 if \(C=1\), and \(2(C \cdot(C-1))\) if \(C \geq 2\).

Case IV, if \(1-\alpha_{j} \sum_{i=1}^{d-1} \beta_{i}>0\) and \(1-\alpha_{j} \beta_{j}>0\), then it means \(\alpha_{j} \leq \min \left(1 / \beta_{j}, \frac{1}{\sum_{i=1}^{d-1} \beta_{i}}\right)\). The cost for the case \(\alpha_{j}<0\) is larger than the cost for case \(\alpha>0\). Thus we can ignore the case \(\alpha_{j}<0\). The cost of the row is,
\[
\begin{aligned}
& =1-\alpha_{j} \sum_{i=1}^{d-1} \beta_{i}+1-\alpha_{j} \beta_{j}+\left|\alpha_{j}\right| \sum_{i \neq j}\left|\beta_{j}\right| \\
& =2-2 \alpha_{j} \beta_{j}+2 \alpha_{j} \sum_{i \in \bar{S}}\left|\beta_{i}\right| .
\end{aligned}
\]

If \(\beta_{j}<\sum_{i \in \bar{S}}\left|\beta_{i}\right|\), we know that the cost is at least 2 . Otherwise, using \(\alpha_{j} \leq 1\), we have the cost is at least \(2-2 \beta_{j}+\sum_{i \in \bar{S}}\left|\beta_{i}\right|\). Let \(T\) denote the set of those \(j\) with \(\beta_{j} \geq \sum_{i \in \bar{S}}\left|\beta_{i}\right|\). If \(|T|=1\), we know that cost is at least 0 . If \(|T| \geq 2\), then the cost is at least
\[
\begin{aligned}
& =\sum_{j \in T}\left(2-2 \beta_{j}+2 \sum_{i \in \bar{S}}\left|\beta_{i}\right|\right) \\
& \geq 2|T|-2 \sum_{j \in S} \beta_{j}+2|T| \sum_{i \in \bar{S}}\left|\beta_{i}\right| \\
& \geq 2|T|-2+2(|T|-1) \sum_{i \in \bar{S}}\left|\beta_{i}\right| \\
& \geq 2(C-1) .
\end{aligned}
\]

Now we discuss the case where \(j \in \bar{S}\), which means \(\beta_{j}<0\).
Case V if \(1-\alpha_{j} \sum_{i=1}^{d-1} \beta_{i}<0\) and \(1-\alpha_{j} \beta_{j}<0\), then it means \(\alpha_{j}>1 / \sum_{i=1}^{d-1} \beta_{i}\) and \(\alpha_{j}<1 / \beta_{j}\). Notice that this case will never happen, because \(\sum_{i=1}^{d-1} \beta_{i}=1\) and \(\alpha_{j}<0\).

Case VI if \(1-\alpha_{j} \sum_{i=1}^{d-1} \beta_{i}<0\) and \(1-\alpha_{j} \beta_{j}<0\), then it means \(\alpha_{j} \geq \max \left(1 / \sum_{i=1}^{d-1} \beta_{i}, 1 / \beta_{j}\right)\).

Because of \(\beta_{j}<0\), then \(\alpha_{j} \geq 1\). The cost of that is,
\[
\begin{aligned}
& =\alpha_{j} \sum_{i=1}^{d-1} \beta_{i}-1+1-\alpha_{j} \beta_{j}+\left|\alpha_{j}\right| \sum_{i \neq j}\left|\beta_{i}\right| \\
& =2 \alpha_{j} \sum_{i \in S}\left|\beta_{i}\right| \\
& \geq 2 .
\end{aligned}
\]
\[
\text { by } \alpha_{j} \geq 1 \text { and } \sum_{i \in S}\left|\beta_{i}\right| \geq 1
\]

Case VII if \(1-\alpha_{j} \sum_{i=1}^{d-1} \beta_{i}>0\) and \(1-\alpha_{j} \beta_{j}<0\), then it means \(\alpha_{j}<\min \left(1 / \beta_{j}, 1 / \sum_{i=1}^{d-1} \beta_{i}\right)\). Because \(\beta_{j}<0\) and \(\sum_{i=1}^{d-1} \beta_{i}=1\), thus \(\alpha_{j}<1 / \beta_{j}\). The cost of that row is,
\[
\begin{aligned}
& =1-\alpha_{j} \sum_{i=1}^{d-1} \beta_{i}+\alpha_{j} \beta_{j}-1+\left|\alpha_{j}\right| \sum_{i \neq j}\left|\beta_{i}\right| \\
& =2\left|\alpha_{j}\right| \sum_{i \in \bar{S} \backslash j}\left|\beta_{i}\right| \\
& \geq 2 \frac{1}{\left|\beta_{j}\right|} \sum_{i \in \bar{S} \backslash j}\left|\beta_{i}\right| .
\end{aligned}
\]

If there are \(C\) such \(j\) in Case VII, then the total cost of Case VII is at least 0 if \(C=1\), and \(2(C \cdot(C-1))\) if \(C \geq 2\).

Case VIII if \(1-\alpha_{j} \sum_{i=1}^{d-1} \beta_{i}>0\) and \(1-\alpha_{j} \beta_{j}>0\), then it means \(1 / \beta_{j}<\alpha_{j}<\) \(1 / \sum_{i=1}^{d-1} \beta_{i}\). The cost of that row,
\[
\begin{aligned}
& =1-\alpha_{j} \sum_{i=1}^{d-1} \beta_{i}+1-\alpha_{j} \beta_{j}+\left|\alpha_{j}\right| \sum_{i \neq j}\left|\beta_{i}\right| \\
& =2-2 \alpha_{j} \beta_{j}+2\left|\alpha_{j}\right| \sum_{i \in \bar{S} \backslash j}\left|\beta_{i}\right| .
\end{aligned}
\]

If \(\alpha_{j}>0\), then the cost is always at least 2 . If \(\alpha_{j}<0\), the cost is at least,
\[
2-2\left|\alpha_{j} \beta_{j}\right|+2\left|\alpha_{j}\right| \sum_{i \in \bar{S} \backslash j}\left|\beta_{i}\right| .
\]

If \(\sum_{i \in \bar{S} \backslash j}\left|\beta_{i}\right|>\left|\beta_{j}\right|\), we also have cost at least 2. Otherwise, we have
\[
2-2\left|\alpha_{j}\right|\left(\left|\beta_{j}\right|-\sum_{i \in \bar{S} \backslash j}\left|\beta_{i}\right|\right)= \begin{cases}2-2\left|\beta_{j}\right|+2 \sum_{i \in \bar{S} \backslash j}\left|\beta_{i}\right| & \text { if } 1 /\left|\beta_{j}\right| \leq 1 \\ 2 \frac{1}{\left|\beta_{j}\right|} \sum_{i \in \bar{S} \backslash j}\left|\beta_{i}\right| & \text { if } 1 /\left|\beta_{j}\right|>1 .\end{cases}
\]

Let \(C\) denote the number of such \(j\) with \(1 /\left|\beta_{j}\right| \leq 1\). If \(C=1\), the cost is at least 0 . If \(C \geq 2\), the cost is at least \(2 C\). Let \(C^{\prime}\) denote the number of such \(j\) with \(1 /\left|\beta_{j}\right|>1\). If \(C^{\prime}=1\), the cost is at least 0 , if \(C^{\prime} \geq 2\), the cost is at least \(2\left(C^{\prime}\left(C^{\prime}-1\right)\right)\). Overall, putting all the eight cases together, we complete the proof.

\subsection*{17.6.3 Hard instance for oblivious subspace embedding and more row subset selection}

The goal in this section is to prove Theorem 17.6.28. By applying Yao's minmax principle, it suffices to prove Theorem 17.6.25.

\subsection*{17.6.3.1 Definitions}

We first give the definition of total variation distance and Kullback-Leibler divergence.

Definition 17.6.1. [LPW09, Ver14] The total variation distance between two probability measures \(P\) and \(Q\) on the measurable space \((\mathcal{X}, \mathcal{F})\) is defined as,
\[
D_{\mathrm{TV}}(P, Q)=\sup _{\mathcal{A} \in \mathcal{F}}|P(\mathcal{A})-Q(\mathcal{A})|
\]

The Kullback-Leibler( KL) divergence of \(P\) and \(Q\) is defined as,
\[
D_{\mathrm{KL}}(P \| Q)=\underset{P}{\mathbb{E}}\left(\log \frac{\mathrm{~d} P}{\mathrm{~d} Q}\right)=\int_{x}\left(\log \frac{\mathrm{~d} P}{\mathrm{~d} Q}\right) \mathrm{d} P
\]

Lemma 17.6.9. [Pin60, Tsy09, CK11] Pinsker's inequality states that, if \(P\) and \(Q\) are two probability distributions on a measurable \((\mathcal{X}, \mathcal{F})\), then
\[
D_{\mathrm{TV}}(P, Q) \leq \sqrt{\frac{1}{2} D_{\mathrm{KL}}(P \| Q)}
\]

Lemma 17.6.10. For any Gaussian random variable \(x \sim N\left(0, \sigma^{2}\right)\), we have, for any \(a>0\)
\[
\operatorname{Pr}[|x|<a] \lesssim \frac{a}{\sigma}
\]

Proof.
\[
\begin{array}{rlr}
\operatorname{Pr}[|x|<a] & =\operatorname{erf}\left(\frac{a}{\sigma \sqrt{2}}\right) & \\
& \leq\left(1-e^{-\frac{-4 a^{2}}{2 \sigma^{2} \pi}}\right)^{\frac{1}{2}} & \text { by } 1-\exp \left(-4 x^{2} / \pi\right) \geq \operatorname{erf}(x)^{2} \\
& \leq\left(\frac{4 a^{2}}{2 \sigma^{2} \pi}\right)^{\frac{1}{2}} & \text { by } 1-e^{-x} \leq x \\
& \lesssim \frac{a}{\sigma} . &
\end{array}
\]

Lemma 17.6.11. Let \(V \in \mathbb{R}^{n \times m}\) be a matrix with orthonormal columns. Let \(\mathcal{H}^{\prime}\) be a distribution over \(V^{\top} A\), where \(A \in \mathbb{R}^{n \times k}\) is a random matrix with each entry i.i.d. Gaussian \(N(0,1)\). Denote \(\mathcal{H}\) as a distribution over \(H \in \mathbb{R}^{m \times k}\), where each entry of \(H\) is drawn from i.i.d. Gaussian \(N(0,1)\). Then, \(\mathcal{H}\) and \(\mathcal{H}^{\prime}\) are the same distribution.

Proof. It is clear that each entry of \(V^{\top} A\) is a random Gaussian variable from \(N(0,1)\), so our goal is to prove the entries of \(V^{\top} A\) are fully independent. Since the the \(j^{\text {th }}\) column of \(V^{\top} A\) only depends on \(A_{j}\), the variables from different columns are fully independent. Now, we look at one column of \(x=V^{\top} A_{j}\). The density function is
\[
f(x)=\frac{1}{\sqrt{(2 \pi)^{m}\left|V^{\top} V\right|}} \exp \left(-\frac{1}{2} x^{\top} V^{\top} V x\right)=\frac{1}{\sqrt{(2 \pi)^{m}}} \exp \left(-\frac{1}{2} x^{\top} x\right),
\]
which is exactly the density function of \(N\left(0, I_{m}\right)\). Thus \(x \in N\left(0, I_{m}\right)\). Therefore, all the entries of \(V^{\top} A\) are fully independent.

Lemma 17.6.12 (Matrix Determinant Lemma). Suppose \(A \in \mathbb{R}^{n \times n}\) is an invertible matrix, and \(u, v \in \mathbb{R}^{n}\). Then,
\[
\left|A+u v^{\top}\right|=\left(1+v^{\top} A^{-1} u\right)|A| .
\]

Lemma 17.6.13 (KL divergence between two multivariate Gaussians \(\left[\mathrm{PP}^{+} 08\right]^{1}\) ). Given two d-dimensional multivariate Gaussian distribution \(N\left(\mu_{1}, \Sigma_{1}\right)\) and \(N\left(\mu_{2}, \Sigma_{2}\right)\), then
\[
D_{\mathrm{KL}}\left(N\left(\mu_{1}| | \Sigma_{1}\right), N\left(\mu_{2}, \Sigma_{2}\right)\right)=\frac{1}{2}\left(\log \frac{\left|\Sigma_{2}\right|}{\left|\Sigma_{1}\right|}-d+\operatorname{tr}\left(\Sigma_{2}^{-1} \Sigma_{1}\right)+\left(\mu_{2}-\mu_{1}\right)^{\top} \Sigma_{2}^{-1}\left(\mu_{2}-\mu_{1}\right)\right) .
\]

Lemma 17.6.14 (Sherman-Morrison formula). Suppose \(A \in \mathbb{R}^{n \times n}\) is an invertible matrix and \(u, v \in \mathbb{R}^{n}\). Suppose furthermore that \(1+v^{\top} A^{-1} u \neq 0\). Then the Sherman-Morrison formula states that
\[
\begin{equation*}
\left(A+u v^{\top}\right)^{-1}=A^{-1}-\frac{A^{-1} u v^{\top} A^{-1}}{1+v^{\top} A^{-1} u} . \tag{17.23}
\end{equation*}
\]

\subsection*{17.6.3.2 Main results}

Lemma 17.6.15. Let \(V \in \mathbb{R}^{n \times m}\) be a matrix with orthonormal columns, and let \(A \in \mathbb{R}^{n \times k}\) be a random matrix with each entry drawn from i.i.d. Gaussian \(N(0,1)\). We denote the distribution \(\mathcal{D}_{i}\) over \(D_{i} \in \mathbb{R}^{(m+1) \times k}\) where
\[
D_{i}=\left[\begin{array}{c}
V^{\top} A \\
A^{i}
\end{array}\right]
\]

If \(\left\|\left(V^{\top}\right)_{i}\right\|_{2}^{2}<\frac{1}{2}\), then
\[
D_{\mathrm{TV}}\left(\mathcal{D}_{i}, \mathcal{G}\right) \leq O\left(k\left\|V^{i}\right\|_{2}\right)+2^{-\Theta(k)}
\]
where \(\mathcal{G}\) is a distribution over \(G \in \mathbb{R}^{(m+1) \times k}\), where each entry of \(G\) is drawn from the i.i.d. Gaussian \(N(0,1)\).

\footnotetext{
\({ }^{1}\) http://stats.stackexchange.com/questions/60680/kl-divergence-between-two-multivariate-gaussians
}

Proof. Let \(\mathcal{H}\) be a distribution over \(H \in \mathbb{R}^{m \times k}\), where each entry of \(H\) is an i.i.d. Gaussian \(N(0,1)\). Let \(\mathcal{H}^{\prime}\) be a distribution over \(V^{\top} A\). According to Lemma 17.6.11, \(\mathcal{H}\) and \(\mathcal{H}^{\prime}\) are the same distribution.

We define matrix \(V^{\neg i} \in \mathbb{R}^{n \times m}\) as
\[
V^{\neg i}=\left[\begin{array}{lllllll}
\left(V^{1}\right)^{\top} & \left(V^{2}\right)^{\top} & \cdots & \left(V^{i-1}\right)^{\top} & 0 & \left(V^{i+1}\right)^{\top} & \cdots
\end{array}\left(V^{n}\right)^{\top}\right]^{\top},
\]
where \(V^{i}\) denotes the \(i^{\text {th }}\) row of \(V \in \mathbb{R}^{n \times m}, \forall i \in[n]\).
Let \(\mathcal{G}\) be a distribution over \(G \in \mathbb{R}^{(m+1) \times k}\), where each entry of \(G\) is an i.i.d. Gaussian \(N(0,1)\). Let \(\mathcal{P}_{i}\) be a distribution over \(P_{i} \in \mathbb{R}^{(m+1) \times k}\), where
\[
P_{i}=\left[\begin{array}{c}
\left(V^{\neg i}\right)^{\top} A \\
A^{i}
\end{array}\right]
\]

Let \(\widehat{\mathcal{P}}_{i}\) be a distribution over \(\widehat{P}_{i} \in \mathbb{R}^{m \times k}\), where \(\widehat{P}_{i}=\left(V^{\neg i}\right)^{\top} A\). Then we have:
\[
\begin{equation*}
D_{\mathrm{TV}}\left(\mathcal{P}_{i}, \mathcal{G}\right)=D_{\mathrm{TV}}\left(\widehat{\mathcal{P}}_{i}, \mathcal{H}^{\prime}\right)=D_{\mathrm{TV}}\left(\widehat{\mathcal{P}}_{i}, \mathcal{H}\right) \tag{17.24}
\end{equation*}
\]

The first equality is because \(\left(V^{\neg i}\right)^{\top} A\) is independent from \(A^{i}\). The second equality follows the Lemma 17.6.11.

Claim 17.6.16. If \(\left\|V^{i}\right\|_{2}^{2} \leq \frac{1}{2}\),
\[
D_{\mathrm{KL}}\left(\widehat{\mathcal{P}}_{i} \| \mathcal{H}\right)=-\frac{k}{2}\left(\log \left(1-\left\|V^{i}\right\|_{2}^{2}\right)+\left\|V^{i}\right\|_{2}^{2}\right) \leq k\left\|V^{i}\right\|_{2}^{2}
\]

Proof. Let \(\widehat{P}_{i} \sim \widehat{\mathcal{P}}_{i}, H \sim \mathcal{H}\). Notice that different columns of \(\widehat{P}_{i}\) are i.i.d, and all entries of \(H\) are fully independent. We can look at one column of \(\widehat{P}_{i}\) and \(H\). Since \(\widehat{P}_{i}=\left(V^{\neg i}\right)^{\top} A\), it is easy to see that its column is drawn from \(N\left(0, \Sigma_{1}\right)\) where \(\Sigma_{1}=I_{m}-\left(V^{i}\right)^{\top} V^{i}\). Since \(H\) is fully independent, each column of \(H\) is drawn from \(N\left(0, \Sigma_{2}\right)\) where \(\Sigma_{2}=I_{m}\). Let \(p(x)\)
be the pdf of the column of \(\widehat{P}_{i}\), and let \(q(x)\) be the pdf of the column of \(H\). We have the following calculation \(\left[\mathrm{PP}^{+} 08\right]^{2}\),
\[
\begin{aligned}
D_{\mathrm{KL}}\left(\widehat{\mathcal{P}}_{i} \| \mathcal{H}\right) & =\frac{k}{2}\left(\log \frac{\left|\Sigma_{2}\right|}{\left|\Sigma_{1}\right|}-m+\operatorname{tr}\left(\Sigma_{2}^{-1} \Sigma_{1}\right)\right) \\
& =\frac{k}{2}\left(\log \frac{\left|I_{m}\right|}{\left|I_{m}-\left(V^{i}\right)^{\top} V^{i}\right|}-m+\operatorname{tr}\left(I_{m}-\left(V^{i}\right)^{\top} V^{i}\right)\right) \\
& =\frac{k}{2}\left(-\log \left|I_{m}-\left(V^{i}\right)^{\top} V^{i}\right|-m+m-\left\|V^{i}\right\|_{2}^{2}\right) \\
& =\frac{k}{2}\left(-\log \left(1-\left\|V^{i}\right\|_{2}^{2}\right)-m+m-\left\|V^{i}\right\|_{2}^{2}\right) \\
& =-\frac{k}{2}\left(\log \left(1-\left\|V^{i}\right\|_{2}^{2}\right)+\left\|V^{i}\right\|_{2}^{2}\right) \\
& \leq-\frac{k}{2} \cdot 2\left\|V^{i}\right\|_{2}^{2} \\
& =k\left\|V^{i}\right\|_{2}^{2} .
\end{aligned}
\]

The first equality is due to Lemma 17.6.13. The sixth equality follows by \(\Sigma_{2}=I_{m}\) and \(\Sigma_{1}=I_{m}-\left(V^{i}\right)^{\top} V^{i}\). The eighth equality follows by Lemma 17.6.12. The first inequality follows by \(\log (1-x)+x \geq-2 x\), when \(0<x<1 / 2\).

According to Lemma 17.6.9, we have
\[
\begin{equation*}
D_{\mathrm{TV}}\left(\widehat{\mathcal{P}}_{i}, \mathcal{H}\right) \leq \sqrt{\frac{1}{2} D_{\mathrm{KL}}\left(\widehat{\mathcal{P}}_{i} \| \mathcal{H}\right)} \leq \sqrt{k}\left\|V^{i}\right\|_{2} \tag{17.25}
\end{equation*}
\]

Now, we want to argue that \(D_{\mathrm{TV}}\left(\mathcal{D}_{i}, \mathcal{P}_{i}\right)\) is small, where \(\mathcal{D}_{i}\) is a distribution over \(D_{i} \in\) \(\mathbb{R}^{(m+1) \times k}\) that
\[
D_{i}=\left[\begin{array}{c}
V^{\top} A \\
A^{i}
\end{array}\right]
\]

\footnotetext{
\({ }^{2}\) http://stats.stackexchange.com/questions/60680/kl-divergence-between-two-multivariate-gaussians
}

For a fixed \(x \in \mathbb{R}^{k}\), let \(\widehat{\mathcal{D}}_{i}(x)\) be a distribution over \(\widehat{D}_{i}(x) \in \mathbb{R}^{m \times k}\), where \(\widehat{D}_{i}(x)=\left(V^{\neg i}\right)^{\top} A+\) \(\left(V^{i}\right)^{\top} x\). Let \(p(x)\) be the pdf of \(\left(A^{i}\right)^{\top}\), then
\[
\begin{equation*}
D_{\mathrm{TV}}\left(\mathcal{D}_{i}, \mathcal{P}_{i}\right)=\int D_{\mathrm{TV}}\left(\widehat{\mathcal{D}}_{i}(x), \widehat{\mathcal{P}}_{i}\right) p(x) \mathrm{d} x \tag{17.26}
\end{equation*}
\]

Now we look at the \(j^{\text {th }}\) column of \(\widehat{D}_{i}(x)\) and the \(j^{\text {th }}\) column of \(\widehat{P}_{i}\). The distribution of the previous one is over \(N\left(\left(V^{i}\right)^{\top} x_{j}, \Sigma_{2}\right)\), and the latter distribution as we said before is \(N\left(0, \Sigma_{2}\right)\), where \(\Sigma_{2}=I_{m}-\left(V^{i}\right)^{\top} V^{i}\). Now we can argue that the KL divergence between them is bounded:
\[
\begin{aligned}
& D_{\mathrm{KL}}\left(N\left(\left(V^{i}\right)^{\top} x_{j}, \Sigma_{2}\right) \| N\left(0, \Sigma_{2}\right)\right) \\
= & \frac{1}{2}\left(\log \frac{\left|\Sigma_{2}\right|}{\left|\Sigma_{2}\right|}-m+\operatorname{tr}\left(\Sigma_{2}^{-1} \Sigma_{2}\right)+x_{j}^{2} V^{i} \Sigma_{2}^{-1}\left(V^{i}\right)^{\top}\right) \\
= & \frac{1}{2}\left(-m+\operatorname{tr}\left(I_{m}\right)+x_{j}^{2} V^{i} \Sigma_{2}^{-1}\left(V^{i}\right)^{\top}\right) \\
= & \frac{1}{2} x_{j}^{2} V^{i} \Sigma_{2}^{-1}\left(V^{i}\right)^{\top} \\
= & \frac{1}{2} x_{j}^{2} V^{i}\left(I_{m}-\left(V^{i}\right)^{\top} V^{i}\right)^{-1}\left(V^{i}\right)^{\top} \\
= & \frac{1}{2} x_{j}^{2} V^{i}\left(I_{m}+\frac{\left(V^{i}\right)^{\top} V^{i}}{1-V^{i}\left(V^{i}\right)^{\top}}\right)\left(V^{i}\right)^{\top} \\
= & \frac{1}{2} x_{j}^{2}\left(\left\|V^{i}\right\|_{2}^{2}+\frac{\left\|V^{i}\right\|_{2}^{2}}{1-\left\|V^{i}\right\|_{2}^{2}}\right) \\
= & \frac{1}{2} x_{j}^{2} \frac{\left\|V^{i}\right\|_{2}^{2}}{1-\left\|V^{i}\right\|_{2}^{2}} .
\end{aligned}
\]

The first equality is due to Lemma 17.6.13. The fourth equality follows by \(\Sigma_{2}=I_{m}-\left(V^{i}\right)^{\top} V^{i}\). The fifth equality follows by Lemma 17.6.14.

By summing the KL divergence on all the columns up,
\[
D_{\mathrm{KL}}\left(\widehat{\mathcal{D}}_{i}(x) \| \widehat{\mathcal{P}}_{i}\right)=\frac{1}{2}\|x\|_{2}^{2} \frac{\left\|V^{i}\right\|_{2}^{2}}{1-\left\|V^{i}\right\|_{2}^{2}}
\]

Applying Lemma 17.6.9 again, we get
\[
D_{\mathrm{TV}}\left(\widehat{\mathcal{D}}_{i}(x), \widehat{\mathcal{P}}_{i}\right) \leq \sqrt{\frac{1}{2} D_{\mathrm{KL}}\left(\widehat{\mathcal{D}}_{i}(x) \| \widehat{\mathcal{P}}_{i}\right)}=\frac{\|x\|_{2}\left\|V^{i}\right\|_{2}}{2 \sqrt{1-\left\|V^{i}\right\|_{2}^{2}}}
\]

Plugging it into Equation (17.26), we get
\[
\begin{aligned}
D_{\mathrm{TV}}\left(\mathcal{D}_{i}, \mathcal{P}_{i}\right) & =\int D_{\mathrm{TV}}\left(\widehat{\mathcal{D}}_{i}(x), \widehat{\mathcal{P}}_{i}\right) p(x) \mathrm{d} x \\
& =\int_{\|x\|_{2} \leq 10 k} D_{\mathrm{TV}}\left(\widehat{\mathcal{D}}_{i}(x), \widehat{\mathcal{P}}_{i}\right) p(x) \mathrm{d} x+\int_{\|x\|_{2}>10 k} D_{\mathrm{TV}}\left(\widehat{\mathcal{D}}_{i}(x), \widehat{\mathcal{P}}_{i}\right) p(x) \mathrm{d} x \\
& \leq \int_{\|x\|_{2} \leq 10 k} D_{\mathrm{TV}}\left(\widehat{\mathcal{D}}_{i}(x), \widehat{\mathcal{P}}_{i}\right) p(x) \mathrm{d} x+\int_{\|x\|_{2}>10 k} p(x) \mathrm{d} x \\
& \leq \int_{\|x\|_{2} \leq 10 k} \frac{\|x\|_{2}\left\|V^{i}\right\|_{2}}{2 \sqrt{1-\left\|V^{i}\right\|_{2}^{2}}} p(x) \mathrm{d} x+\int_{\|x\|_{2}>10 k} p(x) \mathrm{d} x \\
& \leq \frac{10 k\left\|V^{i}\right\|_{2}}{2 \sqrt{1-\left\|V^{i}\right\|_{2}^{2}}}+\int_{\|x\|_{2}>10 k} p(x) \mathrm{d} x \\
& \leq \frac{10 k\left\|V^{i}\right\|_{2}}{2 \sqrt{1-\left\|V^{i}\right\|_{2}^{2}}}+2^{-\Theta(k)}
\end{aligned}
\]

The first inequality just follows from the fact that total variation distance is smaller than 1. The second inequality is what we plugged in. The last inequality follows from the fact that \(x \sim N\left(0, I_{k}\right)\) and from the tail bounds of a Gaussian.

Together with Equation (17.24) and Equation (17.25), we can get
\[
\begin{aligned}
D_{\mathrm{TV}}\left(\mathcal{D}_{i}, \mathcal{G}\right) & \leq D_{\mathrm{TV}}\left(\mathcal{D}_{i}, \mathcal{P}_{i}\right)+D_{\mathrm{TV}}\left(\mathcal{G}, \mathcal{P}_{i}\right) \\
& =D_{\mathrm{TV}}\left(\mathcal{D}_{i}, \mathcal{P}_{i}\right)+D_{\mathrm{TV}}\left(\widehat{\mathcal{P}}_{i}, \mathcal{H}\right) \\
& \leq \frac{10 k\left\|V^{i}\right\|_{2}}{2 \sqrt{1-\left\|V^{i}\right\|_{2}^{2}}}+2^{-\Theta(k)}+\sqrt{k}\left\|V^{i}\right\|_{2} \\
& \leq 10 k\left\|V^{i}\right\|_{2}+2^{-\Theta(k)}+\sqrt{k}\left\|V^{i}\right\|_{2} .
\end{aligned}
\]

The last inequality follows by \(\left\|V^{i}\right\|_{2}^{2} \leq \frac{1}{2}\). Then, we have completed the proof.

Lemma 17.6.17. \(A \in \mathbb{R}^{r \times k}(r \geq k)\) is a random matrix for which each entry is i.i.d. \(N(0,1)\). With probability at least \(1-e^{-\Theta(r)}\), the maximum singular value \(\|A\|_{2}\) is at most \(O(\sqrt{r})\).

Proof. Since \(A \in \mathbb{R}^{k \times r}\) is a random matrix with each entry i.i.d. \(N(0,1)\), this follows by standard arguments (Proposition 2.4 in [RV10]). Since \(r \geq k\), with probability at least \(1-e^{-\Theta(r)},\|A\|_{2}\) is at most \(O(\sqrt{r})\).

Definition 17.6.2. Let \(V \in \mathbb{R}^{n \times r}\) be a matrix with orthonormal columns, and let each entry of \(A \in \mathbb{R}^{k \times r}\) be a random variable drawn from an i.i.d. Gaussian \(N(0,1)\). Define event \(\widehat{\mathcal{E}}(A, V, \beta, \gamma)\) to be: \(\forall y \in \mathbb{R}^{n},\|y\|_{1} \leq O\left(k^{\gamma}\right)\) and each coordinate of \(y\) has absolute value at most \(1 / k^{\beta}\), and also \(A V^{\top} y\) has at most \(O(k / \log k)\) coordinates with absolute value at least \(\Omega(1 / \log k)\), and \(\|A\|_{2} \leq O(\sqrt{r})\).

Lemma 17.6.18. For any \(k \geq 1\), and any constants \(c_{2} \geq c_{1} \geq 1\), let \(k \leq r=O\left(k^{c_{1}}\right), r \leq\) \(n=O\left(k^{c_{2}}\right)\), let \(V \in \mathbb{R}^{n \times r}\) be a matrix with orthonormal columns, and let each entry of \(A \in \mathbb{R}^{k \times r}\) be a random variable drawn from i.i.d. Gaussian \(N(0,1)\). Furthermore, if \(\beta\) and \(\gamma\) are two constants which satisfy \(\beta>\gamma>0\) and \(\beta+\gamma<1\),
\[
\operatorname{Pr}[\widehat{\mathcal{E}}(A, V, \beta, \gamma)] \geq 1-2^{-\Theta(k)}
\]

Proof. Due to Lemma 17.6 .17 , with probability at least \(1-2^{-\Theta(r)},\|A\|_{2} \leq O(\sqrt{r})\), so we can restrict attention to \(\|A\|_{2} \leq O(\sqrt{r})\) in the following proof.

Take any \(y \in \mathbb{R}^{n}\) which has each non-zero coordinate with absolute value at most \(1 / k^{\beta}\), and write it as \(y=\sum_{j=j_{0}}^{+\infty} y^{j}\), where the coordinates in \(y^{j}\) have absolute value in the
range \(\left[2^{-j-1}, 2^{-j}\right.\) ), and the supports of different \(y^{j}\) are disjoint. Since each coordinate of \(y\) is at most \(1 / k^{\beta}, 2^{-j_{0}}=1 / k^{\beta}\). Since \(\|y\|_{1} \leq O\left(k^{\gamma}\right)\), the support size of \(y^{j} \in \mathbb{R}^{n}\) is at most \(s_{j} \leq O\left(2^{j+1} k^{\gamma}\right)\), so it follows that
\[
\left\|y^{j}\right\|_{2} \leq \sqrt{s_{j} \cdot 2^{-2 j}} \leq \sqrt{2^{-j+1} k^{\gamma}} .
\]

We also know \(\left\|y^{j}\right\|_{2} \geq \sqrt{s_{j} \cdot 2^{-2 j-2}}\). Then we can conclude that \(y^{j}\) has 2-norm \(\Theta\left(2^{-j} \sqrt{s_{j}}\right)\). Now we state an \(\epsilon\)-net for all the possible \(y^{j}\) : let \(\epsilon=O\left(1 /\left(n r k^{3}\right)\right)=O\left(1 /\left(k^{c_{1}+c_{2}+3}\right)\right)\). Let \(\mathcal{N}_{j} \subset \mathbb{R}^{n}\) be the following:
\(\mathcal{N}_{j}=\left\{p \in \mathbb{R}^{n} \mid \exists q \in \mathbb{Z}^{n}\right.\), s.t. \(p=\epsilon q,\|p\|_{1} \leq O\left(k^{\gamma}\right), \forall i \in[n]\), either \(\left|p_{i}\right| \in\left[2^{-j-1}, 2^{-j}\right)\) or \(\left.p_{i}=0\right\}\).
Obviously, for any \(y^{j}\), there exists \(p \in \mathcal{N}_{j}\) such that \(\left\|y^{j}-p\right\|_{\infty} \leq \epsilon=O\left(1 /\left(k^{c_{1}+c_{2}+3}\right)\right)\), since \(n \leq O\left(k^{c_{2}}\right),\left\|y^{j}-p\right\|_{2} \leq\left\|y^{j}-p\right\|_{1} \leq n\left\|y^{j}-p\right\|_{\infty} \leq O\left(1 / k^{c_{1}+3}\right)\). Now let us consider the size of \(\mathcal{N}_{j}\). If \(p \in \mathcal{N}_{j}\), the choice of one coordinate of \(p\) is at most \(2 / \epsilon+1\). And since \(\|p\|_{1} \leq O\left(k^{\gamma}\right)\) and each coordinate of \(p\) has absolute value at least \(2^{-j-1}\), the number of supports of \(p\) is at most \(O\left(2^{j+1} k^{\gamma}\right)\). Therefore,
\[
\begin{aligned}
\left|\mathcal{N}_{j}\right| & \leq(n+1)^{O\left(2^{j+1} k^{\gamma}\right)} \cdot(2 / \epsilon+1)^{O\left(2^{j+1} k^{\gamma}\right)} \\
& \leq 2^{O\left(2^{j+1} k^{\gamma}(\log n+\log (2 / \epsilon+1))\right.} \\
& \leq 2^{O\left(2^{j+1} k^{\gamma} \log k\right)}
\end{aligned}
\]

The last inequality follows from \(n \leq O\left(k^{c_{2}}\right), r \leq O\left(k^{c_{1}}\right), \epsilon=O\left(1 /\left(r n k^{3}\right)\right)\).
We define an event \(\mathcal{E}\left(y^{j}\right)\) to be: \(A V^{\top} y^{j}\) has \(k / \log ^{2} k\) coordinates which are each at least \(2 / \log ^{2} k\) in absolute value. Now, we want to show,

Claim 17.6.19. For any \(j \geq j_{0}\), for a fixed \(y^{j} \in \mathbb{R}^{n}\),
\[
\operatorname{Pr}\left[\mathcal{E}\left(y^{j}\right) \text { happens }\right] \leq 2^{k} e^{-\Theta\left(k /\left(\left\|y^{j}\right\|_{2}^{2} \log ^{6} k\right)\right)} \leq e^{-\Theta\left(2^{j-1} k^{1-\gamma} / \log ^{6} k\right)}
\]

Proof. We let \(p\) be the probability that the absolute value of a single coordinate of \(A V^{\top} y^{j}\) is at least \(1 / \log ^{2} k\). Notice that each coordinate of \(A V^{\top} y^{j}\) is i.i.d. Gaussian \(N\left(0,\left\|V^{\top} y^{j}\right\|_{2}^{2}\right)\) and because for any Gaussian random variable \(g, \operatorname{Pr}[|g| \geq t] \leq \exp \left(-\Theta\left(t^{2} / \sigma^{2}\right)\right)\), then \(p \leq\) \(\exp \left(-1 /\left(\left\|V^{\top} y^{j}\right\|_{2}^{2} \log ^{4} k\right)\right)\), by plugging \(\sigma^{2}=\left\|V^{\top} y^{j}\right\|_{2}^{2}\) and \(t=1 / \log ^{2} k\). So the probability \(A V^{\top} y^{j}\) has \(k / \log ^{2} k\) coordinates which are each at least \(1 / \log ^{2} k\) in absolute value is,
\[
\begin{aligned}
\sum_{i=k / \log ^{2} k}^{k} p^{i}(1-p)^{k-i}\binom{k}{i} & \leq \sum_{i=k / \log ^{2} k}^{k} p^{i}\binom{k}{k / 2} \leq 2 p^{k / \log ^{2} k} 2^{k} \\
& \leq e^{-\Theta\left(k t^{2} /\left(\sigma^{2} \log ^{2} k\right)\right)} 2^{k} \\
& \leq e^{-\Theta\left(k /\left(\left\|V^{\top} y^{j}\right\|_{2}^{2} \log ^{6} k\right)\right)} 2^{k} \\
& \leq e^{-\Theta\left(k /\left(\left\|y^{j}\right\|_{2}^{2} \log ^{6} k\right)\right)} 2^{k} \\
& \leq e^{-\Theta\left(k /\left(2^{-2 j} s_{j} \log ^{6} k\right)\right)} 2^{k} \\
& \leq e^{-\Theta\left(k /\left(2^{-j+1} k^{\gamma} \log ^{6} k\right)\right)} 2^{k} \\
& =e^{-\Theta\left(k /\left(2^{-j+1} k^{\gamma} \log ^{6} k\right)\right)} \\
& \leq e^{-\Theta\left(2^{j-1} k^{\left.1-\gamma / \log ^{6} k\right)}\right.}
\end{aligned}
\]
where the second step follows from \(p \leq \exp \left(-\Theta\left(t^{2} / \sigma^{2}\right)\right)\), the forth step follows from \(\left\|V^{\top} y^{j}\right\|_{2}^{2} \leq\) \(\|V\|_{2}^{2}\left\|y^{j}\right\|_{2}^{2} \leq\left\|y^{j}\right\|_{2}^{2}\), the fifth step follows from \(\left\|y^{j}\right\|_{2}^{2} \leq 2^{-2 j} s_{j}\), the sixth step follows from \(s_{j} \leq O\left(2^{j+1} k^{\gamma}\right)\), the seventh step follows from \(2^{-j+1} k^{\gamma} \log ^{6} k \leq 2^{-j_{0}+1} k^{\gamma} \log ^{6} k \leq\) \(2 k^{\gamma-\beta} \log ^{6} k=o(1)\).

For \(j_{1}=\left\lceil 100\left(c_{1}+c_{2}+1\right) \log k\right\rceil=\Theta(\log k)\), consider \(j \in\left[j_{1}, \infty\right)\). We have \(\left\|\sum_{j=j_{1}}^{\infty} y^{j}\right\|_{2}\) is at most \(\Theta\left(2^{-j_{1}} \sqrt{n}\right) \leq 1 / k^{100\left(1+c_{1}\right)}\), and so
\[
\begin{aligned}
\left\|A V^{\top} \sum_{j=j_{1}}^{\infty} y^{j}\right\|_{1} & \leq \sqrt{k}\left\|A V^{\top} \sum_{j=j_{1}}^{\infty} y^{j}\right\|_{2} \\
& \leq \sqrt{k}\|A\|_{2}\|V\|_{2}\left\|\sum_{j=j_{1}}^{\infty} y^{j}\right\|_{2} \\
& \leq \sqrt{k} \cdot \sqrt{r} \cdot 1 / k^{100\left(1+c_{1}\right)} \leq 1 / k^{50} .
\end{aligned}
\]

The last inequality follows from \(r=O\left(k^{c_{1}}\right)\). So the contribution of \(y^{j}\) to \(\left\|A V^{\top} y^{j}\right\|_{1}\) for all \(j \geq j_{1}\) is at most \(1 / k^{50}\). Thus, if we only consider those \(j\) which contribute, i.e., \(j_{0} \leq j \leq j_{1}\), we have \(O(\log k)\) values of \(j\). Then we can only construct \(O(\log k)\) nets \(\mathcal{N}_{j_{0}}, \mathcal{N}_{j_{0}+1}, \cdots, \mathcal{N}_{j_{1}}\). Since the size of net \(\mathcal{N}_{j}\) is \(2^{\Theta\left(2^{j+1} k^{\gamma} \log k\right)}\), by combining Claim 17.6.19 and taking a union bound, we have
\[
\begin{aligned}
\operatorname{Pr}\left[\exists y^{j} \in \bigcup_{j=j_{0}}^{j_{1}} \mathcal{N}_{j}, \mathcal{E}\left(y^{j}\right) \text { happens }\right] & \leq \sum_{j=j_{0}}^{j_{1}} 2^{\Theta\left(2^{j+1} k^{\gamma} \log k\right)} \cdot e^{-\Theta\left(2^{j-1} k^{1-\gamma} / \log ^{6} k\right)} \\
& \leq e^{-\Theta\left(2^{j 0-1} k^{1-\gamma} / \log ^{6} k\right)} \\
& \leq e^{-\Theta\left(k^{1+\beta-\gamma} / \log ^{6} k\right)} \\
& \leq 2^{-\Theta(k)}
\end{aligned}
\]

The second inequality follows since \(k^{\gamma}=o\left(k^{1-\gamma}\right)\). The third inequality follows since \(2^{j_{0}} \geq k^{\beta}\). The fourth inequality follows since \(1+\beta-\gamma>1\).

Then, \(\forall j_{0} \leq j \leq j_{1}, \forall \widetilde{y}^{j} \notin \mathcal{N}_{j}\), there exists a vector \(\widehat{y}^{j}\) in \(\mathcal{N}_{j}\), such that \(\left\|\widehat{y}^{j}-\widetilde{y}^{j}\right\|_{2} \leq\)
\(1 / k^{3+c_{1}}\). We can upper bound the \(\ell_{\infty}\) norm of \(A V^{\top} \widehat{y}^{j}-A V^{\top} \widetilde{y}^{j}\) in the following sense,
\[
\begin{array}{rlr}
\left\|A V^{\top} \widehat{y}^{j}-A V^{\top} \widetilde{y}^{j}\right\|_{\infty} & \leq\left\|A V^{\top} \widehat{y}^{j}-A V^{\top} \widetilde{y}^{j}\right\|_{2} & \text { by }\|\cdot\|_{\infty} \leq\|\cdot\|_{2} \\
& \leq\|A\|_{2} \cdot\left\|\widehat{y}^{j}-\widetilde{y}^{j}\right\|_{2} & \\
& \leq \sqrt{r} / k^{3+c_{1}} & \text { by Claim 17.6.17 and }\left\|\widehat{y}^{j}-\widetilde{y}^{j}\right\|_{2} \leq 1 / k^{3+c_{1}} \\
& =1 / k^{2} . & \text { by } r \leq O\left(k^{c_{1}}\right)
\end{array}
\]

We let \(Y=\left\{y \in \mathbb{R}^{n} \quad \mid \quad\|y\|_{1} \leq O\left(k^{\gamma}\right)\right.\) and each coordinate of \(\left.y \leq 1 / k^{\beta}\right\}\). Since \(1 / k^{2}<\) \(1 / \log ^{2} k\) we can conclude that,
\[
\begin{equation*}
\operatorname{Pr}\left[\exists y \in Y, j \geq j_{0}, \mathcal{E}\left(y^{j}\right) \text { happens }\right] \leq 2^{-\Theta(k)} \tag{17.27}
\end{equation*}
\]

Recalling \(y=\sum_{j} y^{j}\), by Equation (17.27), for any \(y\), with probability at most \(2^{-\Theta(k)}\), there are at most \(O(\log k) \cdot k / \log ^{2} k \leq O(k / \log k)\) coordinates for which \(A V^{\top} \sum_{j=j_{0}}^{j_{1}} y^{j}\) (the same statement also holds for \(A V^{\top} \sum_{j=j_{0}}^{\infty} y^{j}=A V^{\top} y\) since we argued that there is negligible "contribution" for those \(\left.j>j_{1}\right)\) is at least \(O(\log k) / \log ^{2} k=O(1 / \log k)\) on that coordinate. Summarizing,
\[
\operatorname{Pr}[\widehat{\mathcal{E}}(A, V, \beta, \gamma)] \geq 1-2^{-\Theta(k)}
\]

Lemma 17.6.20. For any \(t, k \geq 1\), and any constants \(c_{2} \geq c_{1} \geq 1\), let \(k \leq r=O\left(k^{c_{1}}\right), r \leq\) \(n=O\left(k^{c_{2}}\right)\), let \(V \in \mathbb{R}^{n \times r}\) be a matrix with orthonormal columns, and let each entry of \(A \in \mathbb{R}^{k \times r}, v_{1}, v_{2}, \cdots, v_{t} \in \mathbb{R}^{k}\) be an i.i.d. Gaussian \(N(0,1)\) random variable. For a constant \(\alpha \in(0,0.5)\) which can be arbitrarily small, if \(\widehat{\mathcal{E}}(A, V, 0.5+\alpha / 2,0.5-\alpha)\) happens, then with probability at least \(1-2^{-\Theta(t k)}\), there are at least \(\lceil t / 10\rceil\) such \(j \in[t]\) that \(\forall x \in \mathbb{R}^{r}\) either \(\left\|A x-v_{j}\right\|_{1} \geq \Omega\left(k^{0.5-\alpha}\right)\) or \(\|V x\|_{1} \geq \Omega\left(k^{0.5-\alpha}\right)\) holds.

Proof. For convenience, we define \(\gamma=0.5-\alpha\) which can be an arbitrary constant in \((0,0.5)\). We let constant \(\beta=0.5+\alpha / 2\). Then we have \(\beta+\gamma<1\) and \(\beta>\gamma\). Let \(v \in \mathbb{R}^{k}\) be a random vector with each entry drawn from i.i.d. Gaussian \(N(0,1)\). Suppose \(\widehat{\mathcal{E}}(A, V, \beta, \gamma)\) happens.

For any \(x \in \mathbb{R}^{r}\), we can find a \(y \in \mathbb{R}^{n}\) such that \(y=V x\). Since \(V \in \mathbb{R}^{n \times r}\) has orthonormal columns, \(x=V^{\top} V x=V^{\top} y\). Then our goal is to argue that with high probability if \(\left\|A V^{\top} y-v\right\|_{1} \leq O\left(k^{\gamma}\right)\), then \(\|y\|_{1}>\Omega\left(k^{\gamma}\right)\). Take any \(y \in \mathbb{R}^{n}\) which can be expressed as \(V x\), and decompose it as \(y=y_{0}+y_{1}\), where \(y_{0} \in \mathbb{R}^{n}\) and \(y_{1} \in \mathbb{R}^{n}\) have disjoint supports, \(y_{0}\) has each coordinate with absolute value greater than \(1 / k^{\beta}\), and \(y_{1}\) has each coordinate with absolute value at most \(1 / k^{\beta}\).

Now we create an \(\epsilon\)-net for \(y_{0}\) : let \(\epsilon=O\left(1 /\left(r n k^{3}\right)\right)=O\left(1 / k^{c_{1}+c_{2}+3}\right)\), we denote \(\mathcal{N} \subset \mathbb{R}^{n}\) as follows:
\[
\mathcal{N}=\left\{p \in \mathbb{R}^{n} \mid \exists q \in \mathbb{Z}^{n} \text {, s.t. } p=\epsilon q,\|p\|_{1} \leq O\left(k^{\gamma}\right), \forall i \in[n], \text { either }\left|p_{i}\right|>1 / k^{\beta} \text { or } p_{i}=0\right\}
\]

Obviously, for any \(y_{0}\), there exists \(p \in \mathcal{N}\) such that \(\left\|y_{0}-p\right\|_{\infty} \leq \epsilon=O\left(1 /\left(k^{c_{1}+c_{2}+3}\right)\right)\), since \(n \leq O\left(k^{c_{2}}\right),\left\|y_{0}-p\right\|_{2} \leq\left\|y_{0}-p\right\|_{1} \leq n\left\|y_{0}-p\right\|_{\infty} \leq O\left(1 / k^{c_{1}+3}\right)\). Now let us consider the size of \(\mathcal{N}\). If \(p \in \mathcal{N}\), the number of choices of one coordinate of \(p\) is at most \(O\left(k^{\gamma} / \epsilon\right)\). And since \(\|p\|_{1} \leq O\left(k^{\gamma}\right)\) and each coordinate of \(p\) has absolute value at least \(1 / k^{\beta}\), the number of supports of \(p\) is at most \(O\left(k^{\gamma+\beta}\right)\). Therefore,
\[
\begin{aligned}
|\mathcal{N}| & \leq(n+1)^{O\left(k^{\gamma+\beta}\right)} \cdot O\left(k^{\gamma} / \epsilon\right)^{O\left(k^{\gamma+\beta}\right)} \\
& \leq 2^{O\left(k^{\gamma+\beta} \log k\right)} .
\end{aligned}
\]

The last inequality follows from \(n \leq O\left(k^{c_{2}}\right), r \leq O\left(k^{c_{1}}\right), \epsilon=O\left(1 /\left(r n k^{3}\right)\right)\).
For \(y_{0} \in \mathbb{R}^{n}\), we define event \(\mathcal{E}_{1}\left(y_{0}\right)\) as: \(\exists\) valid \(y_{1} \in \mathbb{R}^{n},\left\|A V^{\top} y-v\right\|_{1} \leq O\left(k^{\gamma}\right)\), where \(y=y_{0}+y_{1}\) is the decomposition of a possible \(y \in \mathbb{R}^{n}\). Here \(y_{1}\) is valid means that there
exists \(y \in \mathbb{R}^{n}\) such that \(\|y\|_{1} \leq O\left(k^{\gamma}\right)\) and the decomposition of \(y\) is \(y=y_{0}+y_{1}\). We define event \(\mathcal{E}_{2}\left(y_{0}\right)\) as: \(\exists\) valid \(y_{1}\), the absolute value of \(A V^{\top} y-v\) is at most \(O\left(1 / k^{\gamma}\right)\) for at least \(k-O\left(k^{2 \gamma}\right)\) coordinates \(i\) in \([k]\). We define event \(\mathcal{E}_{3}\left(y_{0}\right)\) as: at least \(k-O\left(k^{2 \gamma}\right)-O(k / \log k)\) coordinates of \(A y_{0}-v\) have absolute value at most \(O(1 / \log k)\).

Claim 17.6.21. For \(y_{0} \in \mathbb{R}^{n}\),
\[
\operatorname{Pr}\left[\mathcal{E}_{3}\left(y_{0}\right) \text { happens }\right] \geq \operatorname{Pr}\left[\mathcal{E}_{2}\left(y_{0}\right) \text { happens }\right] \geq \operatorname{Pr}\left[\mathcal{E}_{1}\left(y_{0}\right) \text { happens }\right] .
\]

Proof. If \(\mathcal{E}_{1}\left(y_{0}\right)\) happens, then there exists a valid \(y_{1} \in \mathbb{R}^{n}\) such that \(y=y_{0}+y_{1}\) and \(\left\|A V^{\top} y-v\right\|_{1} \leq O\left(k^{\gamma}\right)\). For this \(y\), there are at least \(k-O\left(1 / k^{2 \gamma}\right)\) coordinates of \(A V^{\top} y-v\) with absolute value at most \(O\left(1 / k^{\gamma}\right)\). Otherwise, \(\left\|A V^{\top} y-v\right\|_{1}>\Omega\left(k^{\gamma}\right)\). Thus, \(\mathcal{E}_{1}\left(y_{0}\right)\) implies \(\mathcal{E}_{2}\left(y_{0}\right)\).

Now we want to show \(\mathcal{E}_{2}\left(y_{0}\right)\) implies \(\mathcal{E}_{3}\left(y_{0}\right)\). We suppose \(\mathcal{E}_{2}\left(y_{0}\right)\) happens. Then there is a valid \(y_{1}\) such that there are at least \(k-O\left(1 / k^{2 \gamma}\right)\) coordinates of \(A V^{\top} y-v\) with absolute value at most \(O\left(1 / k^{\gamma}\right)\), where \(y=y_{0}+y_{1}\). Recall that the event \(\widehat{\mathcal{E}}(A, V, \beta, \gamma)\) happens, for any valid \(y_{1}\) there are at most \(O(k / \log k)\) coordinates of \(A V^{\top} y_{1}\) are at least \(\Omega(1 / \log k)\). Therefore,
\[
A V^{\top} y_{0}-v=\underbrace{A V^{\top} y-v}_{\substack{\geq k-O\left(k^{2 \gamma}\right) \text { coordinates } \\
\text { each } \leq O\left(1 / k^{\gamma}\right)}}-\underbrace{A V^{\top} y_{1}}_{\begin{array}{c}
\leq O(k / \log k) \text { coordinates } \\
\text { each } \geq \Omega(1 / \log k)
\end{array}} .
\]

Therefore, at least \(k-O\left(k^{2 \gamma}\right)-O(k / \log k)\) coordinates of \(A y_{0}-v\) in absolute value is at \(\operatorname{most} O(1 / \log k)+O\left(1 / k^{\gamma}\right)=O(1 / \log k)\)

Claim 17.6.22. Define event \(\mathcal{F}_{1}\) to be the situation for which there exists \(1 / 2\) of the coordinates of \(v \in \mathbb{R}^{k}\) that are at least \(1 / 100\). The probability this event \(\mathcal{F}_{1}\) holds is at least \(1-2^{-\Theta(k)}\).

Proof. Note that \(\overline{\mathcal{F}}_{1}\) means there exist \(k / 2\) of the coordinates of \(v \in \mathbb{R}^{k}\) which are at most \(1 / 100\). Using Lemma 17.6.10, for each single coordinate the probability it is smaller than \(1 / 100\) is \(1 / 200\). Then the probability more than half the coordinates are no more than \(1 / 100\) is
\[
\sum_{i=k / 2}^{k} p^{i}(1-p)^{k-i}\binom{k}{i} \leq \sum_{i=k / 2}^{k} p^{i}\binom{k}{k / 2} \leq 2 p^{k / 2} 2^{k} \leq(1 / 200)^{\Theta(k)}=2^{-\Theta(k)}
\]

Conditioned on \(\mathcal{F}_{1}\) happening, if \(\mathcal{E}_{3}\left(y_{0}\right)\) happens, then due to the pigeonhole principle, there are at least \(\frac{k}{2}-O\left(k^{2 \gamma}\right)-O(k / \log k)\) coordinates of \(A V^{\top} y_{0}-v\) that are at most \(O(1 / \log k)\) and the corresponding coordinate of \(v\) is larger than \(1 / 100\). Now, let us look at the probability on a single coordinate of \(A V^{\top} y_{0}-v\).

Claim 17.6.23. If the \(i^{\text {th }}\) coordinate of \(v \in \mathbb{R}^{k}\) is at least \(1 / 100\). Then \(\operatorname{Pr}\left[\left|\left(A V^{\top} y_{0}\right)_{i}-(v)_{i}\right| \leq\right.\) \(O(1 / \log k)] \leq O(1 / \log k)\).

Proof. Since \(\left|(v)_{i}\right|>1 / 100\) and \(v\) is independent from \(A \in \mathbb{R}^{k \times r}\), for \(0<\eta<1 / 100\), \(\operatorname{Pr}\left[\left|\left(A V^{\top} y_{0}\right)_{i}-(v)_{i}\right| \leq \eta\right]\) is always upper bounded by \(\operatorname{Pr}\left[\left(A x_{0}\right)_{i} \in[1 / 100-\eta, 1 / 100+\eta]\right.\). Thus, it suffices to prove an upper bound for \(\operatorname{Pr}\left[\left(A V^{\top} x_{0}\right)_{i} \in[1 / 100-\eta, 1 / 100+\eta]\right.\). Let \(f(x)\)
be the pdf of \(N\left(0, \sigma^{2}\right)\), where \(\sigma^{2}=\left\|V^{\top} y_{0}\right\|_{2}^{2}, V \in \mathbb{R}^{n \times r}, y_{0} \in \mathbb{R}^{n}\). Then
\[
\begin{aligned}
& \operatorname{Pr}\left[\left(A V^{\top} y_{0}\right)_{i} \in[1 / 100-\eta, 1 / 100+\eta]\right] \\
= & \int_{1 / 100-\eta}^{1 / 100+\eta} f(x) \mathrm{d} x \\
\leq & f(1 / 200) \int_{1 / 100-\eta}^{1 / 100+\eta} \mathrm{d} x \\
\leq & O(\eta) .
\end{aligned}
\]
where the last step follows since \(f(1 / 200) \leq 200\). We set \(\eta=O(1 / \log k)\), then we get the statement.

Claim 17.6.24. Conditioned on \(\mathcal{F}_{1}\), for a fixed \(y_{0} \in \mathbb{R}^{n}\), with probability at most \(2^{-\Theta(k)}\), there are at least \(\frac{k}{10}\) coordinates of \(A V^{\top} y_{0}-v \in \mathbb{R}^{k}\) which are at most \(O(1 / \log k)\) and the corresponding coordinate of \(v\) is larger than \(1 / 100\).

Proof. We look at the coordinate \(i \in[k]\) which has \(\left|(v)_{i}\right|>1 / 100\). The probability \(\left\|\left(A V^{\top} y_{0}\right)_{i}-(v)_{i}\right\| \leq O(1 / \log k)\) is at most \(O(1 / \log k)\). Due to the independence between different coordinates of \(v\), since there are at least \(k / 2\) coordinates of \(v\) satisfying that they have absolute value greater than \(1 / 100\), with probability at most \(2^{-\Theta(k)}\), there are at least \(\frac{1}{5} \cdot \frac{k}{2}=k / 10\) coordinates of \(A V^{\top} y_{0}-v\) which are at most \(O(1 / \log k)\).

Because \(\mathcal{E}_{3}\left(y_{0}\right)\) implies the event described in the above claim when conditioning on \(\mathcal{F}_{1}\), for a fixed \(y_{0}\), the probability that \(\mathcal{E}_{3}\left(y_{0}\right)\) holds is at most \(2^{-\Theta(k)}\). Since \(\gamma+\beta<1\), the \(|\mathcal{N}| \leq 2^{k^{o(1)}}\), we can take a union bound over the \(\mathcal{N}\) :
\[
\operatorname{Pr}\left[\exists y_{0} \in \mathcal{N}, \mathcal{E}_{1}\left(y_{0}\right) \text { happens }\right] \leq 2^{-\Theta(k)}
\]

It means that with probability at least \(1-2^{-\Theta(k)}\), for any \(y=y_{0}+y_{1}\) with \(y_{0} \in \mathcal{N}\), \(\left\|A V^{\top} y-v\right\|_{1}>\Omega\left(k^{\gamma}\right)\). Let \(\widetilde{y}=\widetilde{y}_{0}+\sum \widetilde{y}^{j}\), where \(\widetilde{y}_{0} \notin \mathcal{N}\). We can find \(\widehat{y}_{0} \in \mathcal{N}\) which is the closest to \(\widetilde{y}_{0}\). Denote \(\widehat{y}=\widehat{y}+\sum \widetilde{y}^{j}\). Then,
\[
\begin{array}{rlr}
\left\|A V^{\top} \widehat{y}-v\right\|_{1} & \leq\left\|A V^{\top} \widetilde{y}-v\right\|_{1}+\left\|A V^{\top} \widetilde{y}-A V^{\top} \widehat{y}\right\|_{1} & \text { by triangle inequality } \\
& \leq\left\|A V^{\top} \widetilde{y}-v\right\|_{1}+\sqrt{k}\left\|A V^{\top} \widetilde{y}-A V^{\top} \widehat{y}\right\|_{2} & \text { by }\|\cdot\|_{1} \leq \sqrt{\operatorname{dim}}\|\cdot\|_{2} \\
& \leq\left\|A V^{\top} \widetilde{y}-v\right\|_{1}+\sqrt{k}\|A\|_{2}\|\widetilde{y}-\widehat{y}\|_{2} & \\
& \leq\left\|A V^{\top} \widetilde{y}-v\right\|_{1}+\sqrt{k} \cdot \sqrt{r} \cdot\left\|\widetilde{y}_{0}-\widehat{y}_{0}\right\|_{2} & \\
& \leq\left\|A V^{\top} \widetilde{y}-v\right\|_{1}+\sqrt{k} \cdot \sqrt{r} \cdot \epsilon & \text { by }\|A\|_{2} \leq \sqrt{r} \\
& =\left\|A V^{\top} \widetilde{y}-v\right\|_{1}+\sqrt{k} \cdot \sqrt{r} \cdot 1 / k^{c_{1}+3} & \text { by }\left\|\widetilde{y}_{0}-\widehat{y}_{0}\right\|_{2} \leq \epsilon \\
& =\left\|A V^{\top} \widetilde{y}-v\right\|_{1}+1 / k . & \text { by } \epsilon^{\prime}=1 / k^{c_{1}+3} \\
& & \text { by } r=O\left(k^{c_{1}}\right)
\end{array}
\]
and so if \(\left\|A V^{\top} \widetilde{y}-v\right\|_{1}\) is at most \(O\left(k^{\gamma}\right)\) then \(\|A \widehat{y}-v\|_{1}\) is at most \(O\left(k^{\gamma}\right)\).
For \(j \in[t]\), we now use notation \(\mathcal{E}_{4}\left(v_{j}\right)\) to denote the event: \(\exists y_{0} \in \mathbb{R}^{n}\) with \(\left\|y_{0}\right\|_{1} \leq\) \(O\left(k^{\gamma}\right)\) and each non-zero coordinate of \(y_{0}\) is greater than \(1 / k^{\beta}\), at least \(k-O\left(k^{2 \gamma}\right)-O(k / \log k)\) coordinates of \(A V^{\top} y_{0}-v_{j}\) in absolute value are at most \(O(1 / \log k)\). Based on the previous argument, conditioned on \(\widehat{\mathcal{E}}(A, V, \beta, \gamma)\),
\[
\operatorname{Pr}\left[\mathcal{E}_{4}\left(v_{j}\right)\right] \leq 2^{-\Theta(k)}
\]

Also notice that, conditioned on \(\widehat{\mathcal{E}}(A, V, \beta, \gamma), \forall j \in[t], \mathcal{E}_{4}\left(v_{j}\right)\) are independent. Thus, conditioned on \(\widehat{\mathcal{E}}(A, V, \beta, \gamma)\), due to the Chernoff bound, the probability that there are \(\lceil t / 10\rceil\) such \(j\) that \(\mathcal{E}_{4}\left(v_{j}\right)\) happens is at most \(2^{-\Theta(k t)}\). We define \(\mathcal{E}_{5}\left(v_{j}\right)\) to be the event: \(\exists y \in\) \(\mathbb{R}^{n},\left\|A V^{\top} y-v_{j}\right\| \leq O\left(k^{\gamma}\right)\) with \(\|y\|_{1} \leq O\left(k^{\gamma}\right)\). Similar to the proof of Claim 17.6.21, conditioned on \(\widehat{\mathcal{E}}(A, V, \beta, \gamma), \mathcal{E}_{5}\left(v_{j}\right)\) implies \(\mathcal{E}_{4}\left(v_{j}\right)\). Thus, conditioned on \(\widehat{\mathcal{E}}(A, V, \beta, \gamma)\), the
probability that there are \(\lceil t / 10\rceil\) such \(j\) that \(\mathcal{E}_{5}\left(v_{j}\right)\) happens is at most \(2^{-\Theta(t k)}\). Then, we complete the proof.

Theorem 17.6.25. For any \(k \geq 1\), and any constants \(c_{1}, c_{2}\) which satisfy \(c_{2}-2>c_{1}>1\), let \(r=\Theta\left(k^{c_{1}}\right), n=\Theta\left(k^{c_{2}}\right)\), and let \(\mathcal{A}(k, n)\) denote a distribution over \(n \times(k+n)\) matrices where each entry of the first \(n \times k\) matrix is i.i.d. Gaussian \(N(0,1)\) and the next \(n \times n\) matrix is an identity matrix. For any fixed \(r \times n\) matrix \(S\) and a random matrix \(\widehat{A} \sim \mathcal{A}(k, n)\), with probability at least \(1-O\left(k^{1+\frac{c_{1}-c_{2}}{2}}\right)-2^{-\Theta(k)}\), there is no algorithm that is able to output a matrix \(B \in \mathbb{R}^{n \times r}\) such that
\[
\|B S \widehat{A}-\widehat{A}\|_{1} \leq O\left(k^{0.5-\epsilon}\right) \min _{\text {rank }-k}\left\|A^{\prime}-\widehat{A}\right\|_{1},
\]
where \(\epsilon>0\) is a constant which can be arbitrarily small.

Proof. For convenience, we let \(\gamma=0.5-\epsilon\) be a constant which can be arbitrarily close to 0.5. Since the last \(n\) columns of \(\widehat{A}\) is an identity matrix, we can fit the first \(k\) columns of \(\widehat{A}\), so we have
\[
\min _{\operatorname{rank}-k}\left\|A^{\prime}-\widehat{A}\right\|_{1} \leq n
\]

Now, we want to argue that, for a fixed \(S\), with high probability, for any rank- \(k n \times r\) matrix \(B\), the cost
\[
\|B S \widehat{A}-\widehat{A}\|_{1} \geq \Omega\left(n \cdot k^{\gamma}\right)
\]

Thus, the approximation gap will be at least \(\Omega\left(k^{\gamma}\right)\).

We denote the SVD of \(S=U_{S} \Sigma_{S} V_{S}^{\top}\) where \(U_{S} \in \mathbb{R}^{r \times r}, \Sigma_{S} \in \mathbb{R}^{r \times r}, V_{S} \in \mathbb{R}^{n \times r}\). Then, we can rewrite
\[
\begin{align*}
\|B S \widehat{A}-\widehat{A}\|_{1} & =\left\|B U_{S} \Sigma_{S} V_{S}^{\top} \widehat{A}-\widehat{A}\right\|_{1} \\
& =\sum_{l=1}^{n}\left\|\left(B U_{S} \Sigma_{S}\right)^{l}\left(V_{S}^{\top} \widehat{A}\right)-\widehat{A}^{l}\right\|_{1} \\
& \geq \sum_{l:\left\|V_{S}^{l}\right\|_{2}^{2} \leq 2 r / n}\left\|\left(B U_{S} \Sigma_{S}\right)^{l}\left(V_{S}^{\top} \widehat{A}\right)-\widehat{A}^{l}\right\|_{1} . \tag{17.28}
\end{align*}
\]

The first equality follows from the SVD of \(S\). The second equality follows from the fact that the \(\ell_{1}\)-norm of a matrix is the sum of \(\ell_{1}\)-norms of rows. The third inequality follows since we just look at the cost on a part of the rows.

We use \(\beta_{l}\) to denote \(\left(B U_{S} \Sigma_{S}\right)^{l}\). We look at a fixed row \(l\), then the cost on this row is:
\[
\begin{align*}
& \left\|\beta_{l}\left(V_{S}^{\top} \widehat{A}\right)-\widehat{A}^{l}\right\|_{1} \\
= & \left\|\beta_{l}\left(V_{S}^{\top} \widehat{A}\right)_{[1: k]}-\left(\widehat{A}^{l}\right)_{[1: k]}\right\|_{1}+\left\|\beta_{l}\left(V_{S}^{\top} \widehat{A}\right)_{[k+1: k+n]}-\left(\widehat{A}^{l}\right)_{[k+1: n+k]}\right\|_{1} \\
\geq & \left\|\beta_{l}\left(V_{S}^{\top} \widehat{A}\right)_{[1: k]}-\left(\widehat{A}^{l}\right)_{[1: k]}\right\|_{1}+\left\|\beta_{l}\left(V_{S}^{\top} \widehat{A}\right)_{[k+1: k+n]}\right\|_{1}-\left\|\left(\widehat{A}^{l}\right)_{[k+1: n+k]}\right\|_{1} \\
\geq & \left\|\beta_{l}\left(V_{S}^{\top} \widehat{A}\right)_{[1: k]}-\left(\widehat{A}^{l}\right)_{[1: k]}\right\|_{1}+\left\|\beta_{l}\left(V_{S}^{\top} \widehat{A}\right)_{[k+1: k+n]}\right\|_{1}-1 \\
\geq & \left\|\beta_{l}\left(V_{S}^{\top} \widehat{A}\right)_{[1: k]}-\left(\widehat{A}^{l}\right)_{[1: k]}\right\|_{1}+\left\|\beta_{l} V_{S}^{\top}\right\|_{1}-1 . \tag{17.29}
\end{align*}
\]
where \(\left(V_{S}^{\top} \widehat{A}\right)_{[1: k]}\) denotes the first \(k\) columns of \(\left(V_{S}^{\top} \widehat{A}\right)\), and similarly, \(\left(V_{S}^{\top} \widehat{A}\right)_{[k+1: k+n]}\) denotes the last \(n\) columns of \(\left(V_{S}^{\top} \widehat{A}\right)\). The first equality is because we can compute the sum of \(\ell_{1}\) norms on the first \(k\) coordinates and \(\ell_{1}\) norm on the last \(n\) coordinates. The first inequality follows from the triangle inequality. The second inequality follows since the last \(n\) columns of \(\widehat{A}\) form an identity, so there is exactly one 1 on the last \(n\) columns in each row. The third
inequality follows since the last \(n\) columns of \(\widehat{A}\) form an identity. Let \(\mathcal{D}_{l}\) be a distribution over \(D_{l} \in \mathbb{R}^{(r+1) \times k}\), where
\[
D_{l}=\left[\begin{array}{c}
\left(V_{S}^{\top} \widehat{A}\right)_{[1: k]} \\
\left(\widehat{A}^{l}\right)_{[1: k]}
\end{array}\right] .
\]

Let \(\mathcal{G}\) be a distribution over \(G \in \mathbb{R}^{(r+1) \times k}\) where each entry of \(G\) is drawn from i.i.d. \(N(0,1)\). According to Lemma 17.6.15, we have
\[
\begin{equation*}
D_{\mathrm{TV}}\left(\mathcal{D}_{l}, \mathcal{G}\right) \leq O\left(k\left\|\left(V_{S}\right)^{l}\right\|_{2}\right)+2^{-\Theta(k)} \tag{17.30}
\end{equation*}
\]

Let \(A=G^{[1: r]}, v=G^{r+1}\). Due to Lemma 17.6.18, with probability at least \(1-2^{-\Theta(k)}\), \(\widehat{\mathcal{E}}\left(A^{\top}, V_{S}, 0.75-\gamma / 2, \gamma\right)\) happens. Then conditioned on \(\widehat{\mathcal{E}}\left(A^{\top}, V_{S}, 0.75-\gamma / 2, \gamma\right)\), due to Lemma 17.6.20, with probability at most \(2^{-\Theta(k)}\), there exists \(\beta_{l}\) such that
\[
\left\|\beta_{l} A-v\right\|_{1}+\left\|\beta_{l} V_{S}^{\top}\right\|_{1}=o\left(k^{\gamma}\right)
\]

Combined with Equation (17.30), we can get that for a fixed \(l\), with probability at most \(O\left(k\left\|\left(V_{S}\right)^{l}\right\|_{2}\right)+2^{-\Theta(k)}\), there exists \(\beta_{l}\) such that
\[
\left\|\beta_{l}\left(V_{S}^{\top} \widehat{A}\right)_{[1: k]}-\left(\widehat{A}^{l}\right)_{[1: k]}\right\|_{1}+\left\|\beta_{l} V_{S}^{\top}\right\|_{1}=o\left(k^{\gamma}\right)
\]

When \(\left\|\left(V_{S}\right)^{l}\right\|_{2}^{2} \leq 2 r / n=\Theta\left(k^{c_{1}-c_{2}}\right)\), this probability is at most \(\Theta\left(k^{1+\left(c_{1}-c_{2}\right) / 2}\right)+2^{-\Theta(k)}\). Since \(\sum_{l=1}^{n}\left\|\left(V_{S}\right)^{l}\right\|_{2}^{2}=r\), there are at most \(n / 2\) such \(l\) that \(\left\|\left(V_{S}\right)^{l}\right\|_{2}^{2}>2 r / n\) which means that there are at least \(n / 2\) such \(l\) that \(\left\|\left(V_{S}\right)^{l}\right\|_{2}^{2} \leq 2 r / n\). Let \(s\) be the number of \(l\) such that \(\left\|\left(V_{S}\right)^{l}\right\|_{2}^{2} \leq 2 r / n\), then \(s>n / 2\). Let \(t\) be a random variable which denotes that the number of \(l\) which satisfies \(\left\|\left(V_{S}\right)^{l}\right\|_{2}^{2} \leq 2 r / n\) and achieve
\[
\left\|\beta_{l}\left(V_{S}^{\top} \widehat{A}\right)_{[1: k]}-\left(\widehat{A}^{l}\right)_{[1: k]}\right\|_{1}+\left\|\beta_{l} V_{S}^{\top}\right\|_{1}=o\left(k^{\gamma}\right)
\]
at the same time. Then,
\[
\mathbb{E}[t] \leq\left(O\left(k\left\|\left(V_{S}\right)^{l}\right\|_{2}\right)+2^{-\Theta(k)}\right) s=\left(O(k \sqrt{2 r / n})+2^{-\Theta(k)}\right) s .
\]

Due to a Markov inequality,
\[
\operatorname{Pr}[t>s / 2>n / 4] \leq O(k \sqrt{2 r / n})+2^{-\Theta(k)}=O\left(k^{1+\left(c_{1}-c_{2}\right) / 2}\right)+2^{-\Theta(k)} .
\]

The equality follows since \(r=\Theta\left(k^{c_{1}}\right), n=\Theta\left(k^{c_{2}}\right)\). Plugging it into Equation (17.28), now we can conclude, with probability at least \(1-O\left(k^{1+\left(c_{1}-c_{2}\right) / 2}\right)-2^{-\Theta(k)}, \forall B \in \mathbb{R}^{n \times r}\)
\[
\|B S \widehat{A}-\widehat{A}\|_{1} \geq \sum_{l:\left\|V_{S}^{l}\right\|_{2}^{2} \leq 2 r / n}\left\|\left(B U_{S} \Sigma_{S}\right)^{l}\left(V_{S}^{\top} \widehat{A}\right)-\widehat{A}^{l}\right\|_{1} \geq n / 4 \cdot \Omega\left(k^{\gamma}\right)=\Omega\left(n \cdot k^{\gamma}\right)
\]

Theorem 17.6.26 (Hardness for row subset selection). For any \(k \geq 1\), any constant \(c \geq 1\), let \(n=O\left(k^{c}\right)\), and let \(\mathcal{A}(k, n)\) denote the same distribution stated in Theorem 17.6.25. For matrix \(\widehat{A} \sim \mathcal{A}(k, n)\), with positive probability, there is no algorithm that is able to output \(B \in \mathbb{R}^{n \times(n+k)}\) in the row span of any \(r=n / 2\) rows of \(\widehat{A}\) such that
\[
\|\widehat{A}-B\|_{1} \leq O\left(k^{0.5-\alpha}\right) \min _{\text {rank }-k}\left\|A^{\prime}-\widehat{A}\right\|_{1},
\]
where \(\alpha \in(0,0.5)\) is a constant which can be arbitrarily small.

Proof. For convenience, we define \(\gamma=0.5-\alpha\) which can be an arbitrary constant in \((0,0.5)\). Since the last \(n\) columns of \(\widehat{A}\) is an identity matrix, we can fit the first \(k\) columns of \(\widehat{A}\), so we have
\[
\min _{\operatorname{rank}-k A^{\prime}}\left\|A^{\prime}-\widehat{A}\right\|_{1} \leq n
\]

We want to argue that \(\forall B \in \mathbb{R}^{n \times(k+n)}\) in the row span of any \(r=n / 2\) rows of \(\widehat{A}\),
\[
\|\widehat{A}-B\|_{1} \geq \Omega\left(n \cdot k^{\gamma}\right)
\]

Let \(A^{\top} \in \mathbb{R}^{n \times k}\) be the first \(k\) columns of \(\widehat{A}\), and let \(S \subset[n]\) be a set of indices of chosen rows of \(\widehat{A}\) with \(k \leq|S| \leq r\). Let \(M^{S} \in \mathbb{R}^{k \times n}\) with the \(i^{\text {th }}\) column \(M_{i}^{S}=A_{i}\) if \(i \in S\) and \(M_{i}^{S}=0\) otherwise. We use \(\widehat{M}^{S} \in \mathbb{R}^{k \times r}\) to be \(M^{S}\) without those columns of zeros, so it is a random matrix with each entry i.i.d. Gaussian \(N(0,1)\). Then the minimum cost of using a matrix in the span of rows of \(\widehat{A}\) with index in \(S\) to fit \(\widehat{A}\) is at least:
\[
\sum_{l \notin S} \min _{x_{l} \in \mathbb{R}^{n}}\left(\left\|M^{S} x_{l}-A_{l}\right\|_{1}+\left\|x_{l}\right\|_{1}-1\right) .
\]

The part of \(\left\|M^{S} x_{l}-A_{l}\right\|_{1}\) is just the cost on the \(l^{\text {th }}\) row of the first \(k\) columns of \(\widehat{A}\), and the part of \(\left\|x_{l}\right\|_{1}-1\) is just the lower bound of the cost on the \(l^{\text {th }}\) row of the last \(n\) columns of \(\widehat{A}\).

Claim 17.6.27. \(A, \widehat{M}^{S} \in \mathbb{R}^{k \times n}, \gamma \in(0,0.5)\),
\[
\operatorname{Pr}\left[\widehat{\mathcal{E}}\left(\widehat{M}^{S}, I_{r}, 0.75-\gamma / 2, \gamma\right) \mid \widehat{\mathcal{E}}\left(A, I_{n}, 0.75-\gamma / 2, \gamma\right)\right]=1 .
\]

Proof. Suppose \(\widehat{\mathcal{E}}\left(A, I_{n}, 0.75-\gamma / 2, \gamma\right)\) happens. Since \(M^{S}\) has just a subset of columns of \(A,\left\|M^{S}\right\|_{2} \leq\|A\|_{2} \leq \sqrt{n}\). Notice that \(\forall x \in \mathbb{R}^{n}\) with \(\|x\|_{1} \leq O\left(k^{\gamma}\right)\) and each non-zero coordinate of \(x\) is at most \(O\left(1 / k^{0.75-\gamma / 2}\right), M^{S} x \equiv M^{S} x^{S} \equiv A x^{S}\), where \(x^{S} \in \mathbb{R}^{n}\) has \(x^{S}{ }_{i}=x_{i}\) if \(i \in S\) and \(x^{S}{ }_{i}=0\) otherwise. Because \(\left\|x^{S}\right\|_{1} \leq O\left(k^{\gamma}\right)\) and \(x^{S}\) has each coordinate in absolute value at most \(O\left(1 / k^{0.75-\gamma / 2}\right), A x^{S}\) has at most \(O(k / \log k)\) coordinates in absolute value at least \(\Omega(1 / \log k)\). So, \(M^{S} x=A x^{S}\) has at most \(O(k / \log k)\) coordinates in absolute value at least \(\Omega(1 / \log k)\).

We denote \(\operatorname{cost}(S, l)=\min _{x_{l} \in \mathbb{R}^{n}}\left(\left\|M^{S} x_{l}-A_{l}\right\|_{1}+\left\|x_{l}\right\|_{1}-1\right)\). Since \(\forall l \notin S, A_{l}\) are independent, and they are independent from \(M^{S}\), due to Lemma 17.6.20,
\[
\begin{equation*}
\operatorname{Pr}\left[\sum_{l \notin S} \operatorname{cost}(S, l) \leq O\left(n \cdot k^{\gamma}\right) \mid \widehat{\mathcal{E}}\left(\widehat{M}^{S}, I_{r}, 0.75-\gamma / 2, \gamma\right)\right] \leq 2^{-\Theta(r k)} \tag{17.31}
\end{equation*}
\]

Now we just want to upper bound the following:
\[
\begin{aligned}
& \operatorname{Pr}\left[\exists S \subset[n],|S| \leq r, \sum_{l \notin S} \operatorname{cost}(S, l) \leq O\left(n \cdot k^{\gamma}\right)\right] \\
\leq & \operatorname{Pr}\left[\exists S \subset[n],|S| \leq r, \sum_{l \notin S} \operatorname{cost}(S, l) \leq O\left(n \cdot k^{\gamma}\right) \mid \widehat{\mathcal{E}}\left(A, I_{n}, 0.75-\gamma / 2, \gamma\right)\right] \\
+ & \operatorname{Pr}\left[\neg \widehat{\mathcal{E}}\left(A, I_{n}, 0.75-\gamma / 2, \gamma\right)\right] \\
\leq & \sum_{S \subset[n],|S| \leq r} \operatorname{Pr}\left[\sum_{l \notin S} \operatorname{cost}(S, l) \leq O\left(n \cdot k^{\gamma}\right) \mid \widehat{\mathcal{E}}\left(A, I_{n}, 0.75-\gamma / 2, \gamma\right)\right] \\
+ & \operatorname{Pr}\left[\neg \widehat{\mathcal{E}}\left(A, I_{n}, 0.75-\gamma / 2, \gamma\right)\right] \\
\leq & \sum_{S \subset[n],|S| \leq r} \operatorname{Pr}\left[\sum_{l \notin S} \operatorname{cost}(S, l) \leq O\left(n \cdot k^{\gamma}\right) \mid \widehat{\mathcal{E}}\left(A, I_{n}, 0.75-\gamma / 2, \gamma\right)\right]+2^{-\Theta(k)} \\
\leq & \sum_{S \subset[n],|S| \leq r} \operatorname{Pr}\left[\sum_{l \notin S} \operatorname{cost}(S, l) \leq O\left(n \cdot k^{\gamma}\right) \mid \widehat{\mathcal{E}}\left(\widehat{M^{S}}, I_{r}, 0.75-\gamma / 2, \gamma\right)\right]+2^{-\Theta(k)} \\
\leq & (n+1)^{r} 2^{-\Theta(r k)}+2^{-\Theta(k)} \\
\leq & 2^{-\Theta(r k)}+2^{-\Theta(k)} \\
\leq & 2^{-\Theta(k)} .
\end{aligned}
\]

The second inequality follows by a union bound. The third inequality follows by Lemma 17.6.18. The fourth inequality follows by Claim 17.6.27. The fifth inequality is due to Equation (17.31). The sixth inequality follows by \(n \leq O\left(k^{c}\right), r=n / 2\). Thus, with probability at least \(1-2^{-\Theta(k)}\),
\(\forall B \in \mathbb{R}^{n \times(n+k)}\) which is in the span of any \(r \leq n / 2\) rows of \(\widehat{A}\),
\[
\|B-\widehat{A}\|_{1} \geq \Omega\left(n \cdot k^{\gamma}\right)
\]

Then, we have completed the proof.

Definition 17.6.3. Given a matrix \(A \in \mathbb{R}^{n \times d}\), a matrix \(S \in \mathbb{R}^{r \times n}, k \geq 1\) and \(\gamma \in\left(0, \frac{1}{2}\right)\), we say that an algorithm \(\mathcal{M}(A, S, k, \gamma)\) which outputs a matrix \(B \in \mathbb{R}^{n \times r}\) "succeeds", if
\[
\|B S A-A\|_{1} \leq k^{\gamma} \cdot \min _{\operatorname{rank}-k}\left\|A^{\prime}-A\right\|_{1},
\]
holds.

Theorem 17.6.28 (Hardness for oblivious embedding). Let \(\Pi\) denote a distribution over matrices \(S \in \mathbb{R}^{r \times n}\). For any \(k \geq 1\), any constant \(\gamma \in\left(0, \frac{1}{2}\right)\), arbitrary constants \(c_{1}, c_{2}>0\) and \(\min (n, d) \geq \Omega\left(k^{c_{2}}\right)\), if for all \(A \in \mathbb{R}^{n \times d}\), it holds that
\[
\operatorname{Pr}_{S \sim \Pi}[\mathcal{M}(A, S, k, \gamma) \text { succeeds }] \geq \Omega\left(1 / k^{c_{1}}\right)
\]

Then \(r\) must be at least \(\Omega\left(k^{c_{2}-2 c_{1}-2}\right)\).

Proof. We borrow the idea from [NN14, PSW17]. We use Yao's minimax principle [Yao77] here. Let \(\mathcal{D}\) be an arbitrary distribution over \(\mathbb{R}^{n \times d}\), then
\[
\operatorname{Pr}_{A \sim \mathcal{D}, S \sim \Pi}[\mathcal{M}(A, S, k, \gamma)] \geq 1-\delta
\]

It means that there is a fixed \(S_{0}\) such that
\[
\operatorname{Pr}_{A \sim D}\left[\mathcal{M}\left(A, S_{0}, k, \gamma\right)\right] \geq 1-\delta
\]

Therefore, we want to find a hard distribution \(\mathcal{D}_{\text {hard }}\) that if
\[
\operatorname{Pr}_{A \sim \mathcal{D}_{\text {hard }}}\left[\mathcal{M}\left(A, S_{0}, k, \gamma\right)\right] \geq 1-\delta,
\]
\(S_{0}\) must have at least some larger poly \((k)\) rows.
Here, we just use the distribution \(\mathcal{A}\left(k, \Omega\left(k^{c_{2}}\right)\right)\) described in Theorem 17.6.25 as our hard distribution. We can just fill zeros to expand the size of matrix to \(n \times d\). We can complete the proof by using Theorem 17.6.25.

Remark 17.6.1. Actually, in Lemma 17.6.18 and Lemma 17.6.20, the reason we need \(\beta>\gamma>\) 0 is that we want \(k^{\beta-\gamma}=\omega(\operatorname{poly}(\log k))\), and the reason we need \(\beta+\gamma<1\) is that we want \(k^{\beta+\gamma} \operatorname{poly}(\log k)=o(k)\). Thus we can replace all the \(k^{\gamma}\) by \(\sqrt{k} / \operatorname{poly}(\log k)\), e.g., \(\sqrt{k} / \log ^{20} k\), and replace all the \(k^{\beta}\) by \(\sqrt{k}\) poly \((\log k)\) with a smaller poly \((\log k)\), e.g., \(\sqrt{k} \log ^{10} k\). Our proofs still work. Therefore, if we replace the approximation ratio in Theorem 17.6.25, Theorem 17.6.26, and Theorem 17.6 .28 to be \(\sqrt{k} / \log ^{c} k\) where \(c\) is a sufficiently large constant, the statements are still correct.

\section*{17.7 \(\ell_{p}\)-Low Rank Approximation}

This section presents some fundamental lemmas for \(\ell_{p}\)-low rank approximation problems. Using these lemmas, all the algorithms described for \(\ell_{1}\)-low rank approximation problems can be extended to \(\ell_{p}\)-low rank approximation directly. We only state the important Lemmas in this section, due to most of the proofs in this section being identical to the proofs in Section 17.5.

\subsection*{17.7.1 Definitions}

This section is just a generalization of Section 17.5 . 1 to the \(\ell_{p}\) setting when \(1<p<2\).
Definition 17.7.1. Given a matrix \(M \in \mathbb{R}^{n \times d}\), if matrix \(S \in \mathbb{R}^{m \times n}\) satisfies
\[
\|S M\|_{p}^{p} \leq c_{1}\|M\|_{p}^{p}
\]
then \(S\) has at most \(c_{1}\)-dilation on \(M\).
Definition 17.7.2. Given a matrix \(U \in \mathbb{R}^{n \times k}\), if matrix \(S \in \mathbb{R}^{m \times n}\) satisfies
\[
\forall x \in \mathbb{R}^{k},\|S U x\|_{p}^{p} \geq \frac{1}{c_{2}}\|U x\|_{p}^{p}
\]
then \(S\) has at most \(c_{2}\)-contraction on \(U\).
Definition 17.7.3. Given matrices \(U \in \mathbb{R}^{n \times k}, A \in \mathbb{R}^{n \times d}\), denote \(V^{*}=\arg \min _{V \in \mathbb{R}^{k \times d}} \| U V-\) \(A \|_{p}^{p}\). If matrix \(S \in \mathbb{R}^{m \times n}\) satisfies
\[
\forall V \in \mathbb{R}^{k \times d},\|S U V-S A\|_{p}^{p} \geq \frac{1}{c_{3}}\|U V-A\|_{p}^{p}-c_{4}\left\|U V^{*}-A\right\|_{p}^{p}
\]
then \(S\) has at most \(\left(c_{3}, c_{4}\right)\)-contraction on \((U, A)\).
Definition 17.7.4. A \(\left(c_{5}, c_{6}\right) \ell_{p}\)-subspace embedding for the column space of an \(n \times k\) matrix \(U\) is a matrix \(S \in \mathbb{R}^{m \times n}\) for which all \(x \in \mathbb{R}^{k}\)
\[
\frac{1}{c_{5}}\|U x\|_{p}^{p} \leq\|S U x\|_{p}^{p} \leq c_{6}\|U x\|_{p}^{p}
\]

Definition 17.7.5. Given matrices \(U \in \mathbb{R}^{n \times k}, A \in \mathbb{R}^{n \times d}\), denote \(V^{*}=\arg \min _{V \in \mathbb{R}^{k \times d}} \| U V-\) \(A \|_{p}^{p}\). Let \(S \in \mathbb{R}^{m \times n}\). If for all \(c \geq 1\), and if for any \(\widehat{V} \in \mathbb{R}^{k \times d}\) which satisfies
\[
\|S U \widehat{V}-S A\|_{p}^{p} \leq c \cdot \min _{V \in \mathbb{R}^{k \times d}}\|S U V-S A\|_{p}^{p}
\]
it holds that
\[
\|U \widehat{V}-A\|_{p}^{p} \leq c \cdot c_{7} \cdot\left\|U V^{*}-A\right\|_{p}^{p}
\]
then \(S\) provides a \(c_{7}\)-multiple-regression-cost preserving sketch of \((U, A)\).

Definition 17.7.6. Given matrices \(L \in \mathbb{R}^{n \times m_{1}}, N \in \mathbb{R}^{m_{2} \times d}, A \in \mathbb{R}^{n \times d}, k \geq 1\), let
\[
X^{*}=\arg \min _{\text {rank }-k X}\|L X N-A\|_{p}^{p}
\]

Let \(S \in \mathbb{R}^{m \times n}\). If for all \(c \geq 1\), and if for any rank \(-k \widehat{X} \in \mathbb{R}^{m_{1} \times m_{2}}\) which satisfies
\[
\|S L \widehat{X} N-S A\|_{p}^{p} \leq c \cdot \min _{\operatorname{rank}-k X}\|S L X N-S A\|_{p}^{p}
\]
it holds that
\[
\|L \widehat{X} N-A\|_{p}^{p} \leq c \cdot c_{8} \cdot\left\|L X^{*} N-A\right\|_{p}^{p}
\]
then \(S\) provides a \(c_{8}\)-restricted-multiple-regression-cost preserving sketch of \((L, N, A, k)\).

\subsection*{17.7.2 Properties}

Lemma 17.7.1. Given matrices \(A \in \mathbb{R}^{n \times d}, U \in \mathbb{R}^{n \times k}\), let \(V^{*}=\arg \min _{V \in \mathbb{R}^{k \times d}}\|U V-A\|_{p}^{p}\). If \(S \in \mathbb{R}^{m \times n}\) has at most \(c_{1}\)-dilation on \(U V^{*}-A\), i.e.,
\[
\left\|S\left(U V^{*}-A\right)\right\|_{p}^{p} \leq c_{1}\left\|U V^{*}-A\right\|_{p}^{p}
\]
and it has at most \(c_{2}\)-contraction on \(U\), i.e.,
\[
\forall x \in \mathbb{R}^{k},\|S U x\|_{p}^{p} \geq \frac{1}{c_{2}}\|U x\|_{p}^{p}
\]
then \(S\) has at most \(\left(2^{2 p-2} c_{2}, c_{1}+2^{1-p} \frac{1}{c_{2}}\right)\)-contraction on \((U, A)\), i.e.,
\[
\forall V \in \mathbb{R}^{k \times d},\|S U V-S A\|_{p}^{p} \geq \frac{1}{c_{2}}\|U V-A\|_{p}^{p}-\left(c_{1}+\frac{1}{c_{2}}\right)\left\|U V^{*}-A\right\|_{p}^{p}
\]

Proof. Let \(A \in \mathbb{R}^{n \times d}, U \in \mathbb{R}^{n \times k}\), and \(S \in \mathbb{R}^{m \times n}\) be the same as that described in the lemma. Then \(\forall V \in \mathbb{R}^{k \times d}\)
\[
\begin{aligned}
\|S U V-S A\|_{p}^{p} & \geq 2^{1-p}\left\|S U V-S U V^{*}\right\|_{p}^{p}-\left\|S U V^{*}-S A\right\|_{p}^{p} \\
& \geq 2^{1-p}\left\|S U V-S U V^{*}\right\|_{p}^{p}-c_{1}\left\|U V^{*}-A\right\|_{p}^{p} \\
& =2^{1-p}\left\|S U\left(V-V^{*}\right)\right\|_{p}^{p}-c_{1}\left\|U V^{*}-A\right\|_{p}^{p} \\
& =2^{1-p} \sum_{j=1}^{d}\left\|S U\left(V-V^{*}\right)_{j}\right\|_{p}^{p}-c_{1}\left\|U V^{*}-A\right\|_{p}^{p} \\
& \geq 2^{1-p} \sum_{j=1}^{d} \frac{1}{c_{2}}\left\|U\left(V-V^{*}\right)_{j}\right\|_{p}^{p}-c_{1}\left\|U V^{*}-A\right\|_{p}^{p} \\
& =2^{1-p} \frac{1}{c_{2}}\left\|U V-U V^{*}\right\|_{p}^{p}-c_{1}\left\|U V^{*}-A\right\|_{p}^{p} \\
& \geq 2^{2-2 p} \frac{1}{c_{2}}\|U V-A\|_{p}^{p}-2^{1-p} \frac{1}{c_{2}}\left\|U V^{*}-A\right\|_{p}^{p}-c_{1}\left\|U V^{*}-A\right\|_{p}^{p} \\
& =2^{2-2 p} \frac{1}{c_{2}}\|U V-A\|_{p}^{p}-\left(\left(2^{1-p} \frac{1}{c_{2}}+c_{1}\right)\left\|U V^{*}-A\right\|_{p}^{p}\right) .
\end{aligned}
\]

The first inequality follows by Fact 17.7.8. The second inequality follows since \(S\) has at most \(c_{1}\) dilation on \(U V^{*}-A\). The third inequality follows since \(S\) has at most \(c_{2}\) contraction on \(U\). The fourth inequality follows by Fact 17.7.8.

Lemma 17.7.2. Given matrices \(A \in \mathbb{R}^{n \times d}, U \in \mathbb{R}^{n \times k}\), let \(V^{*}=\arg \min _{V \in \mathbb{R}^{k \times d}}\|U V-A\|_{p}^{p}\). If \(S \in \mathbb{R}^{m \times n}\) has at most \(c_{1}\)-dilation on \(U V^{*}-A\), i.e.,
\[
\left\|S\left(U V^{*}-A\right)\right\|_{p}^{p} \leq c_{1}\left\|U V^{*}-A\right\|_{p}^{p}
\]
and has at most \(c_{2}\)-contraction on \(U\), i.e.,
\[
\forall x \in \mathbb{R}^{k},\|S U x\|_{p}^{p} \geq \frac{1}{c_{2}}\|U x\|_{p}^{p}
\]
then \(S\) provides a \(2^{p-1}\left(2 c_{1} c_{2}+1\right)\)-multiple-regression-cost preserving sketch of \((U\), A), i.e., for all \(c \geq 1\), for any \(\widehat{V} \in \mathbb{R}^{k \times d}\) which satisfies
\[
\|S U \widehat{V}-S A\|_{p}^{p} \leq c \cdot \min _{V \in \mathbb{R}^{k \times d}}\|S U V-S A\|_{p}^{p}
\]
it has
\[
\|U \widehat{V}-A\|_{p}^{p} \leq c \cdot 2^{p-1}\left(2 c_{1} c_{2}+1\right) \cdot\left\|U V^{*}-A\right\|_{p}^{p}
\]

Proof. Let \(S \in \mathbb{R}^{m \times n}, A \in \mathbb{R}^{n \times d}, U \in \mathbb{R}^{n \times k}, V^{*}, \widehat{V} \in \mathbb{R}^{k \times d}\), and \(c\) be the same as stated in the lemma.
\[
\begin{aligned}
\|U \widehat{V}-A\|_{p}^{p} & \leq 2^{2 p-2} c_{2}\|S U \widehat{V}-S A\|_{p}^{p}+\left(2^{p-1}+2^{2 p-2} c_{1} c_{2}\right)\left\|U V^{*}-A\right\|_{p}^{p} \\
& \leq 2^{2 p-2} c_{2} c \min _{V \in \mathbb{R}^{k \times d}}\|S U V-S A\|_{p}^{p}+\left(2^{p-1}+2^{2 p-2} c_{1} c_{2}\right)\left\|U V^{*}-A\right\|_{p}^{p} \\
& \leq 2^{2 p-2} c_{2} c\left\|S U V^{*}-S A\right\|_{p}^{p}+\left(2^{p-1}+2^{2 p-2} c_{1} c_{2}\right)\left\|U V^{*}-A\right\|_{p}^{p} \\
& \leq 2^{2 p-2} c_{1} c_{2} c\left\|U V^{*}-A\right\|_{p}^{p}+\left(2^{p-1}+2^{2 p-2} c_{1} c_{2}\right)\left\|U V^{*}-A\right\|_{p}^{p} \\
& \leq c \cdot 2^{p-1}\left(1+2 c_{1} c_{2}\right)\left\|U V^{*}-A\right\|_{p}^{p} .
\end{aligned}
\]

The first inequality follows by Lemma 17.7.1. The second inequality follows by the guarantee of \(\widehat{V}\). The fourth inequality follows since \(S\) has at most \(c_{1}\)-dilation on \(U V^{*}-A\). The fifth inequality follows since \(c \geq 1\).

Lemma 17.7.3. Given matrices \(L \in \mathbb{R}^{n \times m_{1}}, N \in \mathbb{R}^{m_{2} \times d}, A \in \mathbb{R}^{n \times d}, k \geq 1\), let
\[
X^{*}=\arg \min _{\operatorname{rank}-k X}\|L X N-A\|_{p}^{p}
\]

If \(S \in \mathbb{R}^{m \times n}\) has at most \(c_{1}\)-dilation on \(L X^{*} N-A\), i.e.,
\[
\left\|S\left(L X^{*} N-A\right)\right\|_{p}^{p} \leq c_{1}\left\|L X^{*} N-A\right\|_{p}^{p}
\]
and has at most \(c_{2}\)-contraction on \(L\), i.e.,
\[
\forall x \in \mathbb{R}^{m_{1}}\|S L x\|_{p}^{p} \geq\|L x\|_{p}^{p}
\]
then \(S\) provides a \(2^{p-1}\left(2 c_{1} c_{2}+1\right)\)-restricted-multiple-regression-cost preserving sketch of \((L, N, A, k)\), i.e., for all \(c \geq 1\), for any rank \(-k \widehat{X} \in \mathbb{R}^{m_{1} \times m_{2}}\) which satisfies
\[
\|S L \widehat{X} N-S A\|_{p}^{p} \leq c \cdot \min _{\text {rank }-k X}\|S L X N-S A\|_{p}^{p}
\]
it has
\[
\|L \widehat{X} N-A\|_{p}^{p} \leq c \cdot 2^{p-1}\left(2 c_{1} c_{2}+1\right) \cdot\left\|L X^{*} N-A\right\|_{p}^{p}
\]

Proof. Let \(S \in \mathbb{R}^{m \times n}, L \in \mathbb{R}^{n \times m_{1}}, \widehat{X} \in \mathbb{R}^{m_{1} \times m_{2}}, X^{*} \in \mathbb{R}^{m_{1} \times m_{2}}, N \in \mathbb{R}^{m_{2} \times d}, A \in \mathbb{R}^{n \times d}\) and
\(c \geq 1\) be the same as stated in the lemma.
\[
\begin{aligned}
\|S L \widehat{X} N-S A\|_{p}^{p} & \geq 2^{1-p}\left\|S L \widehat{X} N-S L X^{*} N\right\|_{p}^{p}-\left\|S L X^{*} N-S A\right\|_{p}^{p} \\
& \geq 2^{1-p} \frac{1}{c_{2}}\left\|L\left(\widehat{X} N-X^{*} N\right)\right\|_{p}^{p}-c_{1}\left\|L X^{*} N-A\right\|_{p}^{p} \\
& \geq 2^{2-2 p} \frac{1}{c_{2}}\|L \widehat{X} N-A\|_{p}^{p}-2^{1-p} \frac{1}{c_{2}}\left\|L X^{*} N-A\right\|_{1}-c_{1}\left\|L X^{*} N-A\right\|_{p}^{p} \\
& =2^{2-2 p} \frac{1}{c_{2}}\|L \widehat{X} N-A\|_{p}^{p}-\left(2^{1-p} \frac{1}{c_{2}}+c_{1}\right)\left\|L X^{*} N-A\right\|_{p}^{p}
\end{aligned}
\]

The inequality follows from the Fact 17.7.8. The second inequality follows since \(S\) has at most \(c_{2}\)-contraction on \(L\), and it has at most \(c_{1}\)-dilation on \(L X^{*} N-A\). The third inequality follows by Fact 17.7.8.

It follows that
\[
\begin{aligned}
\|L \widehat{X} N-A\|_{p}^{p} & \leq 2^{2 p-2} c_{2}\|S L \widehat{X} N-S A\|_{p}^{p}+\left(2^{p-1}+2^{2 p-2} c_{1} c_{2}\right)\left\|L X^{*} N-A\right\|_{p}^{p} \\
& \leq 2^{2 p-2} c_{2} c \cdot \min _{\text {rank }-k}\|S L X N-S A\|_{p}^{p}+\left(2^{p-1}+2^{2 p-2} c_{1} c_{2}\right)\left\|L X^{*} N-A\right\|_{p}^{p} \\
& \leq 2^{2 p-2} c_{2} c \cdot\left\|S L X^{*} N-S A\right\|_{p}^{p}+\left(2^{p-1}+2^{2 p-2} c_{1} c_{2}\right)\left\|L X^{*} N-A\right\|_{p}^{p} \\
& \leq 2^{2 p-2} c c_{1} c_{2} \cdot\left\|L X^{*} N-A\right\|_{p}^{p}+\left(2^{p-1}+2^{2 p-2} c_{1} c_{2}\right)\left\|L X^{*} N-A\right\|_{p}^{p} \\
& \leq c \cdot 2^{p-1}\left(1+2 c_{1} c_{2}\right)\left\|L X^{*} N-A\right\|_{p}^{p}
\end{aligned}
\]

The first inequality directly follows from the previous one. The second inequality follows from the guarantee of \(\widehat{X}\). The fourth inequality follows since \(S\) has at most \(c_{1}\) dilation on \(L X^{*} N-A\). The fifth inequality follows since \(c \geq 1\).

Lemma 17.7.4. Given matrices \(L \in \mathbb{R}^{n \times m_{1}}, N \in \mathbb{R}^{m_{2} \times d}, A \in \mathbb{R}^{n \times d}, k \geq 1\), let
\[
X^{*}=\arg \min _{\operatorname{rank}-k X}\|L X N-A\|_{p}^{p}
\]

Let \(T_{1} \in \mathbb{R}^{t_{1} \times n}\) have at most \(c_{1}\)-dilation on \(L X^{*} N-A\), and have at most \(c_{2}\)-contraction on L. Let
\[
\tilde{X}=\arg \min _{\operatorname{rank}-k X}\left\|T_{1} L X N-T_{1} A\right\|_{p}^{p}
\]

Let \(T_{2}^{\top} \in \mathbb{R}^{t_{2} \times d}\) have at most \(c_{1}^{\prime}\)-dilation on \(\left(T_{1} L \widetilde{X} N-T_{1} A\right)^{\top}\), and at most \(c_{2}^{\prime}\)-contraction on \(N^{\top}\). Then, for all \(c \geq 1\), for any rank \(-k \widehat{X} \in \mathbb{R}^{m_{1} \times m_{2}}\) which satisfies
\[
\left\|T_{1}(L \widehat{X} N-S A) T_{2}\right\|_{p}^{p} \leq c \cdot \min _{\text {rank }-k X}\left\|T_{1}(L X N-A) T_{2}\right\|_{p}^{p},
\]
it holds that
\[
\|L \widehat{X} N-A\|_{p}^{p} \leq c \cdot 2^{2 p-2}\left(2 c_{1} c_{2}+1\right)\left(2 c_{1}^{\prime} c_{2}^{\prime}+1\right) \cdot\left\|L X^{*} N-A\right\|_{p}^{p}
\]

Proof. Apply Lemma 17.5 .3 for sketching matrix \(T_{2}\). Then for any \(c \geq 1\), any rank \(-k \widehat{X} \in\) \(\mathbb{R}^{m_{1} \times m_{2}}\) which satisfies
\[
\left\|T_{1}(L \widehat{X} N-A) T_{2}\right\|_{p}^{p} \leq c \cdot \min _{\operatorname{rank}-k X}\left\|T_{1}(L X N-A) T_{2}\right\|_{p}^{p}
\]
it has
\[
\left\|T_{1}(L \widehat{X} N-A)\right\|_{p}^{p} \leq c \cdot 2^{p-1}\left(2 c_{1}^{\prime} c_{2}^{\prime}+1\right) \cdot\left\|T_{1}(L \widetilde{X} N-A)\right\|_{p}^{p}
\]

Apply Lemma 17.5.3 for sketch matrix \(T_{1}\). Then for any \(c \geq 1\), any rank \(-k \widehat{X} \in\) \(\mathbb{R}^{m_{1} \times m_{2}}\) which satisfies
\[
\left\|T_{1}(L \widehat{X} N-A)\right\|_{p}^{p} \leq c 2^{p-1}\left(2 c_{1}^{\prime} c_{2}^{\prime}+1\right) \cdot \min _{\text {rank }-k X}\left\|T_{1}(L \widetilde{X} N-A)\right\|_{p}^{p}
\]
it has
\[
\|L \widehat{X} N-A\|_{p}^{p} \leq c \cdot 2^{2 p-2}\left(2 c_{1} c_{2}+1\right)\left(2 c_{1}^{\prime} c_{2}^{\prime}+1\right) \cdot\left\|L X^{*} N-A\right\|_{p}^{p} .
\]

Lemma 17.7.5. Given matrices \(M \in \mathbb{R}^{n \times d}, U \in \mathbb{R}^{n \times t}, d \geq t=\operatorname{rank}(U), n \geq d \geq r=\) \(\operatorname{rank}(M)\). If sketching matrix \(S \in \mathbb{R}^{m \times n}\) is drawn from any of the following probability distributions on matrices, with .99 probability, \(S\) has at most \(c_{1}\)-dilation on \(M\), i.e.,
\[
\|S M\|_{p}^{p} \leq c_{1}\|M\|_{p}^{p}
\]
and \(S\) has at most \(c_{2}\)-contraction on \(U\), i.e.,
\[
\forall x \in \mathbb{R}^{t},\|S U x\|_{p}^{p} \geq \frac{1}{c_{2}}\|U x\|_{p}^{p}
\]
where \(c_{1}, c_{2}\) are parameters depend on the distribution over \(S\).
(I) \(S \in \mathbb{R}^{m \times n}\) is a dense matrix with entries drawn from a p-stable distribution: a matrix with i.i.d. standard \(p\)-stable random variables. If \(m=O(t \log t)\), then \(c_{1} c_{2}=O(\log d)\).
(II) \(S \in \mathbb{R}^{m \times n}\) is a sparse matrix with some entries drawn from a \(p\)-stable distribution: \(S=\) \(T D\), where \(T \in \mathbb{R}^{m \times n}\) has each column drawn i.i.d. from the uniform distribution over standard basis vectors of \(\mathbb{R}^{m}\), and \(D \in \mathbb{R}^{n \times n}\) is a diagonal matrix with each diagonal entry drawn from i.i.d. from the standard p-stable distribution. If \(m=O\left(t^{5} \log ^{5} t\right)\), then \(c_{1} c_{2}=O\left(t^{2 / p} \log ^{2 / p} t \log d\right)\). If \(m=O\left((t+r)^{5} \log ^{5}(t+r)\right)\), then \(c_{1} c_{2}=O\left(\min \left(t^{2 / p} \log ^{2 / p} t \log d, r^{3 / p} \log ^{3 / p} r\right)\right)\).
(III) \(S \in \mathbb{R}^{m \times n}\) is a sampling and rescaling matrix (notation \(S \in \mathbb{R}^{n \times n}\) denotes a diagonal sampling and rescaling matrix with \(m\) non-zero entries): If \(S\) samples and reweights \(m=O(t \log t \log \log t)\) rows of \(U\), selecting each with probability proportional to the \(i^{\text {th }}\) row's \(\ell_{p}\) Lewis weight and reweighting by the inverse probability, then \(c_{1} c_{2}=O(1)\).

In the above, if we replace \(S\) with \(\sigma \cdot S\) where \(\sigma \in \mathbb{R} \backslash\{0\}\) is any scalar, then the relation between \(m\) and \(c_{1} c_{2}\) can be preserved.

For (I), it is implied by Lemma E.17, Lemma E.19. Also see from [SW11] \({ }^{3}\).
For (II), if \(m=O\left(t^{5} \log ^{5} t\right)\), then \(c_{1} c_{2}=O\left(t^{2 / p} \log ^{2 / p} t \log d\right)\) is implied by Corollary 17.5.19 and Lemma 17.7.13 and Theorem 4 in [MM13]. If \(m=O\left((t+r)^{5} \log ^{5}(t+r)\right)\), \(c_{1} c_{2}=O\left(r^{3 / p} \log ^{3 / p} r\right)\) is implied by [MM13].

For (III), it is implied by [CP15] and Lemma 17.5.21.

\footnotetext{
\({ }^{3}\) Full version.
}

\subsection*{17.7.3 Tools and inequalities}

Lemma 17.7.6 (Lemma 9 in [MM13], Upper Tail Inequality for \(p\)-stable Distributions). Let \(p \in(1,2)\) and \(m \geq 3 . \forall i \in[m]\), let \(X_{i}\) be \(m\) (not necessarily independent) random variables sampled from \(D_{p}\), and let \(\gamma_{i}>0\) with \(\gamma=\sum_{i=1}^{m} \gamma_{i}\). Let \(X=\sum_{i=1}^{m} \gamma_{i}\left|X_{i}\right|^{p}\). Then for any \(t \geq 1\),
\[
\operatorname{Pr}\left[X \geq t \alpha_{p} \gamma\right] \leq \frac{2 \log (m t)}{t}
\]

We first review some facts about the \(p\)-norm and \(q\)-norm,

Fact 17.7.7. For any \(p \geq q>0\) and any \(x \in \mathbb{R}^{k}\),
\[
\|x\|_{p} \leq\|x\|_{q} \leq k^{\frac{1}{q}-\frac{1}{p}}\|x\|_{p}
\]

We provide the triangle inequality for the \(p\)-norm,
Fact 17.7.8. For any \(p \in(1,2)\), for any \(x, y \in \mathbb{R}^{k}\),
\[
\|x+y\|_{p} \leq\|x\|_{p}+\|y\|_{p}, \text { and }\|x+y\|_{p}^{p} \leq 2^{p-1}\left(\|x\|_{p}^{p}+\|y\|_{p}^{p}\right) .
\]

Fact 17.7.9 (Hölder's inequality). For any \(x, y \in \mathbb{R}^{k}\), if \(\frac{1}{p}+\frac{1}{q}=1\), then \(\left|x^{\top} y\right| \leq\|x\|_{p}\|y\|_{q}\).

We give the definition of a well-conditioned basis for \(\ell_{p}\),
Definition 17.7.7. Let \(p \in(1,2)\). A basis \(U\) for the range of \(A\) is \((\alpha, \beta, p)\)-conditioned if \(\|U\|_{p} \leq \alpha\) and for all \(x \in \mathbb{R}^{k},\|x\|_{q} \leq \beta\|U x\|_{p}\). We will say \(U\) is well-conditioned if \(\alpha\) and \(\beta\) are low-degree polynomials in \(k\), independent of \(n\).

Proof. We first show an upper bound,
\[
\|U x\|_{p} \leq\|U\|_{p} \cdot\|x\|_{p} \leq \alpha\|x\|_{p}
\]

Then we show a lower bound,
\[
\|U x\|_{p} \geq \frac{1}{\beta}\|x\|_{q}
\]

For any \(p\) and \(q\) with \(1 / p+1 / q=1\), by Hölder's inequality we have
\[
\left|x^{\top} y\right| \leq\|x\|_{p} \cdot\|y\|_{q}
\]
choosing \(y\) to be the vector that has 1 everywhere, \(\|x\|_{1} \leq\|x\|_{p} k^{1 / q}\)

\subsection*{17.7.4 Dense \(p\)-stable transform}

This section states the main tools for the dense \(p\)-stable transform. The proof is identical to that for the dense Cauchy transform.

Lemma 17.7.10. Given matrix \(A \in \mathbb{R}^{n \times d}\) and \(p \in(1,2)\), define \(U^{*} \in \mathbb{R}^{n \times k}\), \(V^{*} \in \mathbb{R}^{k \times d}\) to be an optimal solution of \(\min _{U \in \mathbb{R}^{\times k}, V \in \mathbb{R}^{k \times d}}\|U V-A\|_{p}\). Choose a p-stable distribution matrix \(S \in \mathbb{R}^{m \times n}\), rescaled by \(\Theta\left(1 / m^{1 / p}\right)\). Then we have
\[
\left\|S U^{*} V^{*}-S A\right\|_{p}^{p} \lesssim \log (m d)\left\|U^{*} V^{*}-A\right\|_{p}^{p}
\]
with probability at least 99/100.

Proof. Let \(P(0,1)\) denote the \(p\)-stable distribution. Then,
\[
\begin{array}{rlr}
\left\|S U^{*} V^{*}-S A\right\|_{p}^{p} & \leq \sum_{i=1}^{d}\left\|S\left(U^{*} V_{i}^{*}-A_{i}\right)\right\|_{p}^{p} & \\
& =\sum_{i=1}^{d} \sum_{j=1}^{m}\left|\sum_{l=1}^{n} \frac{1}{m} S_{j, l}\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right|^{p} & \text { where } S_{j, l} \sim P(0,1) \\
& =\frac{1}{m} \sum_{i=1}^{d} \sum_{j=1}^{m}\left|\widetilde{w}_{i j}\left(\sum_{l=1}^{n}\left|\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right|^{p}\right)^{1 / p}\right|^{p} & \text { where } \widetilde{w}_{i j} \sim P(0,1) \\
& =\frac{1}{m} \sum_{i=1}^{d} \sum_{j=1}^{m} \sum_{l=1}^{n}\left|\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right|^{p} \cdot\left|\widetilde{w}_{i j}\right|^{p} & \\
& =\frac{1}{m} \sum_{i=1}^{d} \sum_{j=1}^{m}\left\|U^{*} V_{i}^{*}-A_{i}\right\|_{p}^{p} \cdot w_{i+(j-1) d}^{p} \quad & \text { where } w_{i+(j-1) d} \sim|P(0,1)|
\end{array}
\]
where the last step follows since each \(w_{i}\) can be thought of as a clipped half- \(p\)-stable random variable. Define \(X\) to be \(\sum_{i=1}^{d} \sum_{j=1}^{m}\left\|U^{*} V_{i}^{*}-A_{i}\right\|_{p}^{p} \cdot w_{i+(j-1) d}^{p}\) and \(\gamma\) to be \(\sum_{i=1}^{d} \sum_{j=1}^{m} \| U^{*} V_{i}^{*}-\)
\(A_{i} \|_{p}^{p}\). Then applying Lemma 17.7.6,
\[
\operatorname{Pr}\left[X \geq t \alpha_{p} \gamma\right] \leq \frac{2 \log (m d t)}{t}
\]

Choosing \(t=\Theta(\log (m d))\), we have with probability .999,
\[
X \lesssim \log (m d) \alpha_{p} \gamma=\log (m d) \alpha_{p} \sum_{i=1}^{d}\left\|U^{*} V_{i}^{*}-A_{i}\right\|_{p}^{p}
\]
where the last step follows by definition of \(\gamma\). Thus, we can conclude that with probability .999, \(\left\|\Pi\left(U^{*} V^{*}-A\right)\right\|_{p}^{p} \lesssim \log (m d)\left\|U^{*} V^{*}-A\right\|_{p}^{p}\).

Lemma 17.7.11. Given matrix \(A \in \mathbb{R}^{n \times d}\) and \(p \in(1,2)\), define \(U^{*} \in \mathbb{R}^{n \times k}\) to be the optimal solution of \(\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|U V-A\|_{p}\). Choose a matrix of i.i.d. p-stable random variables \(S \in \mathbb{R}^{m \times n}\). Then for all \(V \in \mathbb{R}^{k \times n}\), we have
\[
\left\|S U^{*} V-S A\right\|_{p}^{p} \gtrsim\left\|U^{*} V-A\right\|_{p}^{p}-O(\log (m d))\left\|U^{*} V^{*}-A\right\|_{p}^{p}
\]

Lemma 17.7.12. Let \(p \in(1,2)\). Given an \((\alpha, \beta) \ell_{p}\) well-conditioned basis, condition on the following two events,
1. For all \(x \in \mathcal{N},\|S U x\|_{p} \gtrsim\|U x\|_{p}\).
2. For all \(x \in \mathbb{R}^{k},\|S U x\|_{p} \leq \operatorname{poly}(k)\|U x\|_{p}\).

Then for all \(x \in \mathbb{R}^{k},\|S U x\|_{p} \gtrsim\|U x\|_{p}\).

Proof. The proof is identical to Lemma 17.5.14 in Section 17.5.

\subsection*{17.7.5 Sparse \(p\)-stable transform}

This section states the main tools for the sparse \(p\)-stable transform. The proof is identical to that of the sparse Cauchy transform.

Lemma 17.7.13. Let \(p \in(1,2)\). Given matrix \(A \in \mathbb{R}^{n \times d}\) with \(U^{*}, V^{*}\) an optimal solution of \(\min _{U, V}\|U V-A\|_{p}\), let \(\Pi=\sigma \cdot S C \in \mathbb{R}^{m \times n}\), where \(S \in \mathbb{R}^{m \times n}\) has each column vector chosen independently and uniformly from the \(m\) standard basis vectors of \(\mathbb{R}^{m}\), where \(C\) is a diagonal matrix with diagonals chosen independently from the standard \(p\)-stable distribution, and \(\sigma\) is a scalar. Then
\[
\left\|\Pi U^{*} V^{*}-\Pi A\right\|_{p}^{p} \lesssim \sigma \cdot \log (m d) \cdot\left\|U^{*} V^{*}-A\right\|_{p}^{p}
\]
holds with probability at least . 999.

Proof. We define \(\Pi=\sigma \cdot S C \in \mathbb{R}^{m \times n}\) as in the statement of the lemma. Then by the
definition of \(\Pi\), we have,
\[
\begin{aligned}
& \left\|\Pi\left(U^{*} V^{*}-A\right)\right\|_{p}^{p} \\
= & \sum_{i=1}^{d}\left\|S C\left(U^{*} V_{i}^{*}-A_{i}\right)\right\|_{p}^{p} \\
= & \sum_{i=1}^{d}\left\|\left[\begin{array}{cccc}
S_{11} & S_{12} & \cdots & S_{1 n} \\
S_{21} & S_{22} & \cdots & S_{2 n} \\
\cdots & \cdots & \cdots & \cdots \\
S_{m 1} & S_{m 2} & \cdots & S_{m n}
\end{array}\right] \cdot\left[\begin{array}{cccc}
c_{1} & 0 & 0 & 0 \\
0 & c_{2} & 0 & 0 \\
0 & 0 & \cdots & 0 \\
0 & 0 & 0 & c_{n}
\end{array}\right] \cdot\left(U^{*} V_{i}^{*}-A_{i}\right)\right\|_{p}^{p} \\
= & \sum_{i=1}^{d}\left\|\left[\begin{array}{cccc}
c_{1} S_{11} & c_{2} S_{12} & \cdots & c_{n} S_{1 n} \\
c_{1} S_{21} & c_{2} S_{22} & \cdots & c_{n} S_{2 n} \\
\cdots & \cdots & \cdots & \cdots \\
c_{1} S_{m 1} & c_{2} S_{m 2} & \cdots & c_{n} S_{m n}
\end{array}\right] \cdot\left(U^{*} V_{i}^{*}-A_{i}\right)\right\|_{p}^{p} \\
= & \sum_{i=1}^{d}\left\|\sum_{l=1}^{n} c_{l} S_{1 l} \cdot\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}, \sum_{l=1}^{n} c_{l} S_{2 l} \cdot\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}, \cdots, \sum_{l=1}^{n} c_{l} S_{m l} \cdot\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right\|_{p}^{p} \\
= & \sum_{i=1}^{d} \sum_{j=1}^{m}\left|\sum_{l=1}^{n} c_{l} S_{j l} \cdot\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right|^{p} \text { by }^{p} a X+b Y \text { and }\left(|a|^{p}+|b|^{p}\right)^{1 / p} Z \text { are identically distributed } \\
= & \sum_{i=1}^{d} \sum_{j=1}^{m}\left|\widetilde{w}_{i j} \cdot\left(\sum_{l=1}^{n}\left|S_{j l}\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right|^{p}\right)^{1 / p}\right|^{p} \text { where } \widetilde{w}_{i j} \sim P(0,1) \\
= & \sum_{i=1}^{d} \sum_{j=1}^{m} \sum_{l=1}^{n}\left|S_{j l}\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right|^{p} \cdot\left|\widetilde{w}_{i j}\right|^{p} \\
= & \sum_{i=1}^{d} \sum_{j=1}^{m} \sum_{l=1}^{n}\left|S_{j l}\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right|^{p} \cdot w_{i+(j-1) d}^{p}, \text { where } w_{i+(j-1) d} \sim|P(0,1)|
\end{aligned}
\]
where the last step follows since each \(w_{i}\) can be thought of as a clipped half- \(p\)-stable random variable. Define \(X\) to be \(\sum_{i=1}^{d} \sum_{j=1}^{m} \sum_{l=1}^{n}\left|S_{j l}\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right|^{p} \cdot w_{i+(j-1) d}^{p}\) and \(\gamma\) to be \(\sum_{i=1}^{d} \sum_{j=1}^{m} \sum_{l=1}^{n}\left|S_{j l}\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right|^{p}\). Then applying Lemma 17.7.6,
\[
\operatorname{Pr}\left[X \geq t \alpha_{p} \gamma\right] \leq \frac{2 \log (m d t)}{t}
\]

Choosing \(t=\Theta(\log (m d))\), we have with probability .999,
\[
X \lesssim \log (m d) \alpha_{p} \gamma=\log (m d) \alpha_{p} \sum_{i=1}^{d} \sum_{l=1}^{n}\left|\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right|^{p},
\]
where the last steps follows by
\[
\begin{aligned}
\gamma & =\sum_{i=1}^{d} \sum_{j=1}^{m} \sum_{l=1}^{n}\left|S_{j l}\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right|^{p} \\
& =\sum_{i=1}^{d} \sum_{j=1}^{m} \sum_{l=1}^{n}\left|S_{j l}\right|^{p}\left|\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right|^{p} \\
& =\sum_{i=1}^{d} \sum_{l=1}^{n}\left|\left(U^{*} V_{i}^{*}-A_{i}\right)_{l}\right|^{p} .
\end{aligned}
\]

Thus, we can conclude that with probability .999,
\[
\left\|\Pi\left(U^{*} V^{*}-A\right)\right\|_{p}^{p} \lesssim \log (m d)\left\|U^{*} V^{*}-A\right\|_{p}^{p}
\]

Lemma 17.7.14. Given matrix \(A \in \mathbb{R}^{n \times d}\) with \(U^{*}, V^{*}\) an optimal solution of \(\min _{U, V} \| U V-\) \(A \|_{p}\), let \(\Pi=\sigma \cdot S C \in \mathbb{R}^{m \times n}\), where \(S \in \mathbb{R}^{m \times n}\) has each column vector chosen independently and uniformly from the \(m\) standard basis vectors of \(\mathbb{R}^{m}\), and where \(C\) is a diagonal matrix with diagonals chosen independently from the standard p-stable distribution. Then with probability at least .999, for all \(V \in \mathbb{R}^{k \times d}\),
\[
\left\|\Pi U^{*} V-\Pi A\right\|_{p}^{p} \geq\left\|U^{*} V-A\right\|_{p}^{p}-O(\sigma \log (m d))\left\|U^{*} V^{*}-A\right\|_{p}^{p}
\]

Notice that \(m=O\left(k^{5} \log ^{5} k\right)\) and \(\sigma=O\left((k \log k)^{2 / p}\right)\) according to Theorem 4 in [MM13].

\subsection*{17.7.6 \(\quad \ell_{p}\)-Lewis weights}

This section states the main tools for \(\ell_{p}\)-Lewis weights. The proof is identical to \(\ell_{1}\)-Lewis weights.

Lemma 17.7.15. For any \(p \in(1,2)\). Given matrix \(A \in \mathbb{R}^{n \times d}\), define \(U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times d}\) to be an optimal solution of \(\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|U V-A\|_{p}\). Choose a diagonal matrix \(D \in \mathbb{R}^{n \times n}\) according to the Lewis weights of \(U^{*}\). We have that
\[
\left\|D U^{*} V^{*}-D A\right\|_{p}^{p} \lesssim\left\|U^{*} V^{*}-A\right\|_{p}^{p}
\]
holds with probability at least . 99 .

Lemma 17.7.16. Let \(p \in(1,2)\). Given matrix \(A \in \mathbb{R}^{n \times d}\), define \(U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times d}\) to be an optimal solution of \(\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|U V-A\|_{p}\). Choose a sampling and rescaling matrix \(D \in \mathbb{R}^{n \times n}\) according to the Lewis weights of \(U^{*}\). For all \(V \in \mathbb{R}^{k \times d}\) we have
\[
\left\|D U^{*} V-D A\right\|_{p}^{p} \gtrsim\left\|U^{*} V-A\right\|_{p}^{p}-O(1)\left\|U^{*} V^{*}-A\right\|_{p}^{p},
\]
holds with probability at least .99.

\subsection*{17.8 EMD-Low Rank Approximation}

In this section we explain how to embed EMD to \(\ell_{1}\). For more detailed background on the Earth-Mover Distance(EMD) problem, we refer the reader to [IT03, AIK08, ABIW09, IP11, BIRW16] and [SL09, LOG16]. Section 17.8.1 introduces some necessary notation and definitions for Earth-Mover Distance. Section 17.8.2 presents the main result for the EarthMover distance low rank approximation problem.

\subsection*{17.8.1 Definitions}

Consider any two non-negative vectors \(x, y \in \mathbb{R}_{+}^{[\Delta]^{2}}\) such that \(\|x\|_{1}=\|y\|_{1}\). Let \(\Gamma(x, y)\) be the set of functions \(\gamma:[\Delta]^{2} \times[\Delta]^{2} \rightarrow \mathbb{R}_{+}\), such that for any \(i, j \in[\Delta]^{2}\) we have \(\sum_{l} \gamma(i, l)=x_{i}\) and \(\sum_{l} \gamma(l, j)=y_{j}\); that is, \(\Gamma\) is the set of possible "flows" from \(x\) to \(y\). Then we define
\[
\operatorname{EMD}(x, y)=\min _{\gamma \in \Gamma} \sum_{i, j \in[\Delta]^{2}} \gamma(i, j)\|i-j\|_{1}
\]
to be the min cost flow from \(x\) to \(y\), where the cost of an edge is its \(\ell_{1}\) distance.
Using the \(\operatorname{EMD}(\cdot, \cdot)\) metric, for general vectors \(w\), we define \(\|\cdot\|_{\text {EEMD }}\) distance (which is the same as [SL09]),
\[
\|w\|_{\text {EEMD }}=\min _{\substack{x-y+z=w \\\|x\|_{1}=\|y\|_{1} \\ x, y \geq 0}} \operatorname{EMD}(x, y)+2 \Delta\|z\|_{1} .
\]

Using \(\|\cdot\|_{\text {EEMD }}\) distance, for general matrices \(X \in \mathbb{R}^{n \times d}\), we define the \(\|\cdot\|_{1, \text { EEMD }}\) distance,
\[
\|X\|_{1, \mathrm{EEMD}}=\sum_{i=1}^{d}\left\|X_{i}\right\|_{\mathrm{EEMD}}
\]
where \(X_{i}\) denotes the \(j\)-th column of matrix \(X\).

\subsection*{17.8.2 Analysis of no contraction and no dilation bound}

Lemma 17.8.1. Given matrix \(A \in \mathbb{R}^{n \times d}\) and \(U^{*}, V^{*}=\underset{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}{\arg \min }\|U V-A\|_{1, \mathrm{EEMD}}\), there exist sketching matrices \(S \in \mathbb{R}^{m \times n}\) such that, with probability .999 , for all \(V \in \mathbb{R}^{k \times d}\),
\[
\left\|S\left(U^{*} V-A\right)\right\|_{1} \geq\left\|U^{*} V-A\right\|_{1, \text { EEMD }}
\]
holds.

Proof. Using Lemma 1 in [IT03], there exists a constant \(C>0\) such that for all \(i \in[d]\),
\[
\begin{equation*}
C\left\|S U^{*} V_{i}-S A_{i}\right\|_{1} \geq\left\|U^{*} V_{i}-A_{i}\right\|_{\text {EEMD }} \tag{17.32}
\end{equation*}
\]

Then taking a summation over all \(d\) terms and rescaling the matrix \(S\), we obtain,
\[
\sum_{i=1}^{d}\left\|S\left(U^{*} V_{i}-A_{i}\right)\right\|_{1} \geq\left\|U^{*} V_{i}-A_{i}\right\|_{\text {EEMD }}
\]
which completes the proof.

Lemma 17.8.2. Given matrix \(A \in \mathbb{R}^{n \times d}\) and \(U^{*}, V^{*}=\underset{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}{\arg \min }\|U V-A\|_{1, \mathrm{EEMD}}\), there exist sketching matrices \(S \in \mathbb{R}^{m \times n}\) such that
\[
\left\|S\left(U^{*} V^{*}-A\right)\right\|_{1} \leq O(\log n)\left\|U^{*} V^{*}-A\right\|_{1, \text { EEMD }}
\]
holds with probability at least . 999.

Proof. Using Lemma 2 in [IT03], we have for any \(i \in[d]\),
\[
\begin{equation*}
\mathbb{E}\left[\left\|S U^{*} V_{i}^{*}-S A_{i}\right\|_{1}\right] \leq O(\log n)\left\|U^{*} V_{i}^{*}-A_{i}\right\|_{\text {EEMD }} \tag{17.33}
\end{equation*}
\]

Then using that the expectation is linear, we have
\[
\begin{aligned}
\mathbb{E}\left[\left\|S U^{*} V^{*}-S A\right\|_{1}\right] & =\mathbb{E}\left[\sum_{i=1}^{d}\left\|S U^{*} V_{i}^{*}-S A_{i}\right\|_{1}\right] \\
& =\sum_{i=1}^{d} \mathbb{E}\left[\left\|S U^{*} V_{i}^{*}-S A_{i}\right\|_{1}\right] \\
& \leq \sum_{i=1}^{d} O(\log n)\left\|U^{*} V_{i}^{*}-A_{i}\right\|_{\text {EEMD }} \quad \text { by Equation }(17.33) \\
& =O(\log n)\left\|U^{*} V^{*}-A\right\|_{1, \text { EEMD }} .
\end{aligned}
\]

Using Markov's inequality, we can complete the proof.

Theorem 17.8.3. Given a matrix \(A \in \mathbb{R}^{n \times d}\), there exists an algorithm running in \(\operatorname{poly}(k, n, d)\) time that is able to output \(U \in \mathbb{R}^{n \times k}\) and \(V \in \mathbb{R}^{k \times d}\) such that
\[
\|U V-A\|_{1, \mathrm{EEMD}} \leq \operatorname{poly}(k) \cdot \log d \cdot \log n \min _{\text {rank }-k A_{k}}\left\|A_{k}-A\right\|_{1, \mathrm{EEMD}}
\]
holds with probability . 99 .

Proof. First using Lemma 17.8.1 and Lemma 17.8.2, we can reduce the original problem into an \(\ell_{1}\)-low rank approximation problem by choosing \(m=\operatorname{poly}(n)\). Second, we can use our \(\ell_{1}\)-low rank approximation algorithm to solve it. Notice that all of our \(\ell_{1}\)-low rank approximation algorithms can be applied here. If we apply Theorem 17.4.6, we complete the proof.

Our current \(\|\cdot\|_{1, \text { EEMD }}\) is column-based. We can also define it to be row-based. Then we get a slightly better result by applying the \(\ell_{1}\)-low rank algorithm.

Corollary 17.8.4. Given a matrix \(A \in \mathbb{R}^{n \times d}\), there exists an algorithm running in \(\operatorname{poly}(k, n, d)\) time that is able to output \(U \in \mathbb{R}^{n \times k}\) and \(V \in \mathbb{R}^{k \times d}\) such that
\[
\|U V-A\|_{1, \mathrm{EEMD}} \leq \operatorname{poly}(k) \cdot \log ^{2} d \min _{\text {rank }-k}\left\|A_{k}-A\right\|_{1, \mathrm{EEMD}}
\]
holds with probability . 99 .

\subsection*{17.9 Hardness}

This section presents our hardness results. Section 17.9.1 states several useful tools from literature. Section 17.9 .2 shows that, it is NP-hard to get some multiplicative error. Assuming ETH is true, we provide a stronger hardness result in Section 17.9.3. Section 17.9.4 extends the result from the rank- 1 case to the rank- \(k\) case.

\subsection*{17.9.1 Previous results}

Definition 17.9.1 \(\left(\|A\|_{\infty \rightarrow 1},[G V 15]\right)\). Given matrix \(A \in \mathbb{R}^{n \times d}\),
\[
\|A\|_{\infty \rightarrow 1}=\min _{x \in\{-1,+1\}^{n}, y \in\{-1,+1\}^{d}} x^{\top} A y
\]

The following lemma says that computing \(\|A\|_{\infty \rightarrow 1}\) for matrix \(A\) with entries in \(\{-1,+1\}\) is equivalent to computing a best \(\{-1,+1\}\) matrix which is an \(\ell_{1}\) norm rank- 1 approximation to \(A\).

Lemma 17.9.1 (Lemma 3 of [GV15]). Given matrix \(A \in\{-1,+1\}^{n \times d}\),
\[
\|A\|_{\infty \rightarrow 1}+\min _{x \in\{-1,+1\}^{n}, y \in\{-1,+1\}^{d}}\left\|A-x y^{\top}\right\|_{1}=n d .
\]

Lemma 17.9.2 (Theorem 2 of [GV15]). Given \(A \in\{-1,+1\}^{n \times d}\), we have
\[
\min _{x \in\{-1,+1\}^{n}, y \in\{-1,+1\}^{d}}\left\|A-x y^{\top}\right\|_{1}=\min _{x \in \mathbb{R}^{n}, y \in \mathbb{R}^{d}}\left\|A-x y^{\top}\right\|_{1} .
\]

Combining with Lemma 17.9.1 and Lemma 17.9.2, it implies that computing \(\|A\|_{\infty \rightarrow 1}\) for \(A \in\{-1,+1\}^{n \times d}\) is equivalent to computing the best \(\ell_{1}\) norm rank-1 approximation to the matrix \(A\) :
\[
\|A\|_{\infty \rightarrow 1}+\min _{x \in \mathbb{R}^{n}, y \in \mathbb{R}^{d}}\left\|A-x y^{\top}\right\|_{1}=n d
\]

Theorem 17.9.3 (NP-hard result, Theorem 1 of [GV15]). Computing \(\|A\|_{\infty \rightarrow 1}\) for matrix \(A \in\{-1,+1\}^{n \times d}\) is NP-hard.

The proof of the above theorem in [GV15] is based on the reduction from MAXCUT problem. The above theorem implies that computing \(\min _{x \in \mathbb{R}^{n}, y \in \mathbb{R}^{d}}\left\|A-x y^{\top}\right\|_{1}\) is also NP-hard.

\subsection*{17.9.2 Extension to multiplicative error \(\ell_{1}\)-low rank approximation}

The previous result only shows that solving the exact problem is NP-hard. This section presents a stronger hardness result, which says that, it is still NP-hard even if the goal is to find a solution that is able to achieve some multiplicative error. The proof in this section and the next section are based on the reduction from the MAX-CUT problem. For recent progress on MAX-CUT problem, we refer the readers to [GW95, BGS98, TSSW00, Hås01, KKMO07, FLP15].

Theorem 17.9.4. Given \(A \in\{-1,+1\}^{n \times d}\), computing an \(\widehat{x} \in \mathbb{R}^{n}, \widehat{y} \in \mathbb{R}^{d}\) s.t.
\[
\left\|A-\widehat{x}^{\top} \widehat{y}\right\|_{1} \leq\left(1+\frac{1}{n d}\right) \min _{x \in \mathbb{R}^{n}, y \in \mathbb{R}^{d}}\left\|A-x^{\top} y\right\|_{1}
\]
is NP-hard.

MAX-CUT decision problem: Given a positive integer \(c^{*}\) and an unweighted graph \(G=(V, E)\) where \(V\) is the set of vertices of \(G\) and \(E\) is the set of edges of \(G\), the goal is to determine whether there is a cut of \(G\) has at least \(c^{*}\) edges.

Lemma 17.9.5. MAX-CUT decision problem is NP-hard.

We give the definition for the Hadamard matrix,
Definition 17.9.2. The Hadamard matrix \(H_{p}\) of size \(p \times p\) is defined recursively : \(\left[\begin{array}{cc}H_{p / 2} & H_{p / 2} \\ H_{p / 2} & -H_{p / 2}\end{array}\right]\) with \(H_{2}=\left[\begin{array}{ll}+1 & +1 \\ +1 & -1\end{array}\right]\).

For simplicity, we use \(H\) to denote \(H_{p}\) in the rest of the proof.

Recall the reduction shown in [GV15] which is from MAX-CUT to computing \(\|\cdot\|_{\infty \rightarrow 1}\) for \(\{-1,+1\}\) matrices. We do the same thing: for a given graph \(G=(V, E)\), we construct a matrix \(A \in\{-1,+1\}^{n \times d}\) where \(n=p|E|\) and \(d=p|V|\). Notice that \(p=\operatorname{poly}(|E|,|V|)\) is a parameter which will be determined later, and also \(p\) is a power of 2 .

We divide the matrix \(A\) into \(|E| \times|V|\) blocks, and each block has size \(p \times p\). For \(e \in[|E|]\), if the \(e^{t h}\) edge has endpoints \(i \in[|V|], j \in[|V|]\) and \(i<j\), let all the \(p \times p\) elements of \((e, i)\) block of \(A\) be 1 , all the \(p \times p\) elements of \((e, j)\) block of \(A\) be -1 , and all the \((e, l)\) block of \(A\) be \(p \times p\) Hadamard matrix \(H\) for \(l \neq i, j\).

Claim 17.9.6 (Lower bound of \(\|A\|_{\infty \rightarrow 1}\), proof of Theorem 1 of [GV15]). If there is a cut of \(G\) with cut size at least \(c\),
\[
\|A\|_{\infty \rightarrow 1} \geq 2 p^{2} c-|E \| V| p^{3 / 2}
\]

Claim 17.9.7 (Upper bound of \(\|A\|_{\infty \rightarrow 1}\), proof of Theorem 1 of [GV15]). If the max cut of \(G\) has fewer than \(c\) edges,
\[
\|A\|_{\infty \rightarrow 1} \leq 2 p^{2}(c-1)+|E \| V| p^{3 / 2}
\]

Remark 17.9.1. In [GV15], they set \(p\) as a power of 2 and \(p>|E|^{2}|V|^{2}\). This implies
\[
\forall c \in[|E|], 2 p^{2}(c-1)+|E||V| p^{3 / 2}<2 p^{2} c-|E||V| p^{3 / 2}
\]

Therefore, according to Claim 17.9.6 and Claim 17.9.7, if we can know the precise value of \(\|A\|_{\infty \rightarrow 1}\), we can decide whether \(G\) has a cut with cut size at least \(c^{*}\).

For convenience, we use \(T^{*}\) to denote \(\|A\|_{\infty \rightarrow 1}\) and use \(L^{*}\) to denote
\[
\min _{x \in \mathbb{R}^{n}, y \in \mathbb{R}^{d}}\left\|A-x^{\top} y\right\|_{1} .
\]

Also, we use \(L\) to denote a \(\left(1+\frac{1}{n d}\right)\) relative error approximation to \(L^{*}\), which means:
\[
L^{*} \leq L \leq\left(1+\frac{1}{n d}\right) L^{*}
\]

We denote \(T\) as \(n d-L\).

Proof of Theorem 17.9.4. Because \(L^{*} \leq L \leq\left(1+\frac{1}{n d}\right) L^{*}\), we have:
\[
n d-L^{*} \geq n d-L \geq n d-\left(1+\frac{1}{n d}\right) L^{*}
\]

Due to Lemma 17.9.1 and the definition of \(T\), it has:
\[
T^{*} \geq T \geq T^{*}-\frac{1}{n d} L^{*}
\]

Notice that \(A\) is a \(\{-1,+1\}\) matrix, we have
\[
L^{*} \leq\|A\|_{1} \leq 2 n d
\]

Thus,
\[
T^{*} \geq T \geq T^{*}-2
\]

It means
\[
T+2 \geq T^{*} \geq T
\]

According to Claim 17.9.11, if \(G\) has a cut with cut size at least \(c\), we have:
\[
T+2 \geq T^{*} \geq 2 p^{2} c-|E||V| p^{3 / 2}
\]

That is
\[
T \geq 2 p^{2} c-|E||V| p^{3 / 2}-2
\]

According to Claim 17.9.12, if the max cut of \(G\) has fewer than \(c\) edges,
\[
T \leq T^{*} \leq 2 p^{2}(c-1)+|E||V| p^{3 / 2}
\]

Let \(p\) be a power of 2 and \(p>|E|^{3}|V|^{3}\), we have
\[
2 p^{2}(c-1)+|E||V| p^{3 / 2}<2 p^{2} c-|E||V| p^{3 / 2}-2
\]

Therefore, we can decide whether \(G\) has a cut with size at least \(c\) based on the value of \(T\).
Thus, if we can compute \(\widehat{x} \in \mathbb{R}^{n}, \widehat{y} \in \mathbb{R}^{d}\) s.t.
\[
\left\|A-\widehat{x}^{\top} \widehat{y}\right\|_{1} \leq\left(1+\frac{1}{n d}\right) \min _{x \in \mathbb{R}^{n}, y \in \mathbb{R}^{d}}\left\|A-x^{\top} y\right\|_{1}
\]
in polynomial time, it means we can compute \(L\) and \(T\) in polynomial time, and we can solve MAX-CUT decision problem via the value of \(T\), which leads to a contradiction.

\subsection*{17.9.3 Using the ETH assumption}

The goal of this section is to prove Theorem 17.9.9. We first introduce the definition of 3SAT and Exponential Time Hypothesis(ETH). For the details and background of 3SAT problem, we refer the readers to [AB09].

Definition 17.9.3 (3SAT problem). Given an \(r\) variables and \(m\) clauses conjunctive normal form CNF formula with size of each clause at most 3, the goal is to decide whether there exists an assignment for the \(r\) boolean variables to make the CNF formula be satisfied.

Hypothesis 17.9.8 (Exponential Time Hypothesis (ETH) [IPZ98]). There is a \(\delta>0\) such that 3SAT problem defined in Definition 17.9.3 cannot be solved in \(O\left(2^{\delta r}\right)\) running time.

The main lower bound is stated as follows:

Theorem 17.9.9. Unless ETH (see Hypothesis 17.9.8) fails, for arbitrarily small constant \(\gamma>0\), given some matrix \(A \in\{-1,+1\}^{n \times d}\), there is no algorithm can compute \(\widehat{x} \in \mathbb{R}^{n}, \widehat{y} \in\) \(\mathbb{R}^{d}\) s.t.
\[
\left\|A-\widehat{x}^{\top} \widehat{y}\right\|_{1} \leq\left(1+\frac{1}{\log ^{1+\gamma} n d}\right) \min _{x \in \mathbb{R}^{n}, y \in \mathbb{R}^{d}}\left\|A-x^{\top} y\right\|_{1}
\]
in \((n d)^{O(1)}\) running time.

Before we prove our lower bound, we introduce the following theorem which is used in our proof.

Definition 17.9.4 (MAX-CUT decision problem). Given a positive integer \(c^{*}\) and an unweighted graph \(G=(V, E)\) where \(V\) is the set of vertices of \(G\) and \(E\) is the set of edges of \(G\), the goal is to determine whether there is a cut of \(G\) has at least \(c^{*}\) edges.

Theorem 17.9.10 (Theorem 6.1 in [FLP15]). There exist constants \(a, b \in(0,1)\) and \(a>b\), such that, for a given MAX-CUT (see Definition 17.9.4) instance graph \(G=(E, V)\) which is an n-vertices 5 -regular graph, if there is an algorithm in time \(2^{o(n)}\) which can distinguish the following two cases:
1. At least one cut of the instance has at least a|E| edges,
2. All cuts of the instance have at most \(b|E|\) edges,
then ETH (see Hypothesis 17.9.8) fails.

Proof of Theorem 17.9.9. We prove it by contradiction. We assume, for any given \(A \in\) \(\{-1,+1\}^{n \times d}\), there is an algorithm can compute \(\widehat{x} \in \mathbb{R}^{n}, \widehat{y} \in \mathbb{R}^{d}\) s.t.
\[
\left\|A-\widehat{x}^{\top} \widehat{y}\right\|_{1} \leq\left(1+\frac{1}{W}\right) \min _{x \in \mathbb{R}^{n}, y \in \mathbb{R}^{d}}\left\|A-x^{\top} y\right\|_{1}
\]
in time poly \((n d)\), where \(W=\log ^{1+\gamma} d\) for arbitrarily small constant \(\gamma>0\). Then, we show the following. There exist constants \(a, b \in[0,1], a>b\), for a given MAX-CUT instance \(G=(V, E)\) with \(|E|=O(|V|)\), such that we can distinguish whether \(G\) has a cut with size at least \(a|E|\) or all the cuts of \(G\) have size at most \(b|E|\) in \(2^{o(|V|)}\) time, which leads to a contradiction to Theorem 17.9.10.

Recall the reduction shown in [GV15] which is from MAX-CUT to computing \(\|\cdot\|_{\infty \rightarrow 1}\) for \(\{-1,+1\}\) matrices. We do similar things here: for a given graph \(G=(V, E)\) where \(|E|=O(|V|)\), we construct a matrix \(A \in\{-1,+1\}^{n \times d}\) where \(n=p|E|\) and \(d=p|V|\). Notice that \(p\) is a parameter which will be determined later, and also \(p\) is a power of 2 .

We divide the matrix \(A\) into \(|E| \times|V|\) blocks, and each block has size \(p \times p\). For \(e \in[|E|]\), if the \(e^{\text {th }}\) edge has endpoints \(i \in[|V|], j \in[|V|]\) and \(i<j\), let all the \(p \times p\) elements of \((e, i)\) block of \(A\) be 1 , all the \(p \times p\) elements of \((e, j)\) block of \(A\) be -1 , and all the \((e, l)\) block of \(A\) be \(p \times p\) Hadamard matrix \(H\) for \(l \neq i, j\).

We can construct the matrix in \(n d\) time, which is \(p^{2}|E||V|\). We choose \(p\) to be the smallest number of power of 2 which is larger than \(2^{\frac{2}{a-b}|V|^{1-\frac{\gamma}{10}}}\) for some \(\gamma>0\). Thus, the


We will show, if we can compute a \((1+1 / W)\)-approximation to \(A\), we can decide whether \(G\) has a cut with size at least \(a|E|\) or has no cut with size larger than \(b|E|\). For convenience, we use \(T^{*}\) to denote \(\|A\|_{\infty \rightarrow 1}\) and use \(L^{*}\) to denote
\[
\min _{x \in \mathbb{R}^{n}, y \in \mathbb{R}^{d}}\left\|A-x^{\top} y\right\|_{1} .
\]

Also, we use \(L\) to denote a \(\left(1+\frac{1}{W}\right)\) relative error approximation to \(L^{*}\), which means:
\[
L^{*} \leq L \leq\left(1+\frac{1}{W}\right) L^{*}
\]

We denote \(T\) as \(n d-L\).
Because \(L^{*} \leq L \leq\left(1+\frac{1}{W}\right) L^{*}\), we have:
\[
n d-L^{*} \geq n d-L \geq n d-\left(1+\frac{1}{W}\right) L^{*} .
\]

Due to Lemma 17.9.1 and the definition of \(T\), it has:
\[
T^{*} \geq T \geq T^{*}-\frac{1}{W} L^{*}
\]

Notice that \(A\) is a \(\{-1,+1\}\) matrix, we have
\[
L^{*} \leq\|A\|_{1} \leq 2 n d
\]

Thus,
\[
T^{*} \geq T \geq T^{*}-2 n d / W
\]

It means
\[
T+2 n d / W \geq T^{*} \geq T
\]

Claim 17.9.11 (Lower bound of \(\|A\|_{\infty \rightarrow 1}\), proof of Theorem 1 of [GV15]). If there is a cut of \(G\) with cut size at least \(c\),
\[
\|A\|_{\infty \rightarrow 1} \geq 2 p^{2} c-|E||V| p^{3 / 2}
\]

Claim 17.9.12 (Upper bound of \(\|A\|_{\infty \rightarrow 1}\), proof of Theorem 1 of [GV15]). If the max cut of \(G\) has fewer than \(c\) edges,
\[
\|A\|_{\infty \rightarrow 1} \leq 2 p^{2}(c-1)+|E \| V| p^{3 / 2} .
\]

According to Claim 17.9.11, if \(G\) has a cut with cut size at least \(a|E|\), we have:
\[
T+2 n d / W \geq T^{*} \geq 2 p^{2} a|E|-|E||V| p^{3 / 2}
\]

That is
\[
\begin{equation*}
T \geq 2 p^{2} a|E|-|E||V| p^{3 / 2}-2 n d / W \tag{17.34}
\end{equation*}
\]

According to Claim 17.9.12, if the max cut of \(G\) has fewer than \(b|E|\) edges,
\[
\begin{equation*}
T \leq T^{*} \leq 2 p^{2} b|E|+|E \| V| p^{3 / 2} \tag{17.35}
\end{equation*}
\]

Using these conditions \(p \geq 2^{\frac{2}{a-b}|V|^{1-\frac{\gamma}{10}}}, d=p|V|, W \geq \log ^{1+\gamma} d\), we can lower bound \(|W|\) by \(|V|\) up to some constant,
\[
\begin{array}{rlrl}
W & \geq \log ^{1+\gamma} d & \text { by } W \geq \log ^{1+\gamma} d \\
& =\log ^{1+\gamma}(p|V|) & \text { by } d=p|V| \\
& =(\log |V|+\log p)^{1+\gamma} & \\
& \geq\left(\log |V|+\frac{2}{a-b}|V|^{1-\frac{\gamma}{10}}\right)^{1+\gamma} & & \\
& \geq \frac{2}{a-b}|V| . & \text { by } p \geq 2^{\frac{2}{a-b}|V|^{1-\frac{\gamma}{10}}}  \tag{17.36}\\
& &
\end{array}
\]

Thus, we can upper bound \(1 / W\) in the following sense,
\[
\begin{equation*}
\frac{1}{W} \leq \frac{a-b}{2|V|} \leq \frac{a-b}{|V|}-p^{-\frac{1}{2}} \tag{17.37}
\end{equation*}
\]
where the first inequality follows by Equation (17.36) and the second inequality follows by \(p \geq 2^{\frac{2}{a-b}|V|^{1-\frac{\gamma}{10}}}, \gamma\) is sufficient small, and \(|V|\) is large enough.

Now, we can conclude,
\[
\begin{aligned}
& \frac{1}{W} \leq \frac{a-b}{|V|}-p^{-\frac{1}{2}} \\
\Longleftrightarrow & p^{2}|E||V| / W \leq(a-b) p^{2}|E|-|E||V| p^{\frac{3}{2}}
\end{aligned}
\]
by multiplying \(p^{2}|E||V|\) on both sides
\[
\Longleftrightarrow 2 n d / W \leq 2(a-b) p^{2}|E|-2|E||V| p^{\frac{3}{2}}
\]
by multiplying 2 on both sides and \(p^{2}|E||V|=n d\)
\[
\begin{equation*}
\Longleftrightarrow 2 p^{2} b|E|+|E||V| p^{3 / 2}<2 p^{2} a|E|-|E||V| p^{3 / 2}-2 n d / W \tag{17.38}
\end{equation*}
\]
by adding \(2 p^{2} b|E|+|E||V| p^{3 / 2}-2 n d / W\) on both sides
which implies that Equation (17.38) is equivalent to Equation (17.37). Notice that the LHS of (17.38) is exactly the RHS of (17.35) and the RHS of (17.38) is exactly the RHS of (17.34). Therefore, we can decide whether \(G\) has a cut with size larger than \(a|E|\) or has no cut with size larger than \(b|E|\).

Thus, if we can compute \(\widehat{x} \in \mathbb{R}^{n}, \widehat{y} \in \mathbb{R}^{d}\) s.t.
\[
\left\|A-\widehat{x}^{\top} \widehat{y}\right\|_{1} \leq\left(1+\frac{1}{\log ^{1+\gamma} d}\right) \min _{x \in \mathbb{R}^{n}, y \in \mathbb{R}^{d}}\left\|A-x^{\top} y\right\|_{1}
\]
in poly \((n d)\) time, which means we can compute \(L\) and \(T\) in poly \((n d)\) time. Notice that \(n d=p^{2}|E||V|,|E|=O(|V|), p \geq 2^{\frac{2}{a-b}|V|^{1-\frac{\gamma}{10}}}\), it means poly \((n d)=2^{O\left(|V|^{1-\frac{\gamma}{10}}\right)}\). Because we decide whether \(G\) has a cut with size larger than \(a|E|\) or has no cut with size larger than \(b|E|\) via the value of \(T\), we can solve it in \(2^{O\left(|V|^{\left.1-\frac{\gamma}{10}\right)}\right.}\) time which leads to a contradiction to Theorem 17.9.10.

\subsection*{17.9.4 Extension to the rank- \(k\) case}

This section presents a way of reducing the rank- \(k\) case to the rank- 1 case. Thus, we can obtain a lower bound for general \(k \geq 1\) under ETH.

Theorem 17.9.13. For any constants \(c_{1}>0, c_{2}>0\) and \(c_{3}>0\), and any constant \(c_{4} \geq\) \(10\left(c_{1}+c_{2}+c_{3}+1\right)\), given any matrix \(A \in \mathbb{R}^{n \times n}\) with absolute value of each entry bounded by \(n^{c_{1}}\), we define a block diagonal matrix \(\widetilde{A} \in \mathbb{R}^{(n+k-1) \times(n+k-1)}\) as
\[
\widetilde{A}=\left[\begin{array}{ccccc}
A & 0 & 0 & \cdots & 0 \\
0 & B & 0 & \cdots & 0 \\
0 & 0 & B & \cdots & 0 \\
\cdots & \cdots & \cdots & \cdots & \cdots \\
0 & 0 & 0 & \cdots & B
\end{array}\right]
\]
where \(B=n^{c_{4}}\). If \(\widehat{A}\) is an \(\ell_{1}\)-norm rank-k \(C\)-approximation solution to \(\widetilde{A}\), i.e.,
\[
\|\widehat{A}-\widetilde{A}\|_{1} \leq C \cdot \min _{\operatorname{rank}-k}\left\|\widehat{A}^{\prime} \quad \widehat{A^{\prime}}-\widetilde{A}\right\|_{1}
\]
where \(C \in\left[1, n^{c_{3}}\right]\), then there must exist \(j^{*} \in[n]\) such that
\[
\min _{v \in \mathbb{R}^{n}}\left\|\widehat{A}_{j^{*}}^{[1: n]} v^{\top}-A\right\|_{1} \leq C \cdot \min _{u, v \in \mathbb{R}^{n}}\left\|u v^{\top}-A\right\|_{1}+1 / n^{c_{2}},
\]
i.e., the first \(n\) coordinates of the column \(j^{*}\) of \(\widehat{A}\) can give an \(\ell_{1}\)-norm rank- \(1 C\)-approximation to \(A\).

Proof. The first observation is that because we can use a rank-1 matrix to fit \(A\) and use a rank- \((k-1)\) matrix to fit other \(B \mathrm{~s}\), we have
\[
\begin{equation*}
\min _{\operatorname{rank}-k \widehat{A}^{\prime}}\left\|\widehat{A}^{\prime}-\widetilde{A}\right\|_{1} \leq \min _{u, v \in \mathbb{R}^{n}}\left\|u v^{\top}-A\right\|_{1} \leq\|A\|_{1} . \tag{17.39}
\end{equation*}
\]

Claim 17.9.14. Let \(\widehat{A}\) denote the rank- \(k C\)-approximate solution to \(\widetilde{A}\). Let \(Z \in \mathbb{R}^{(n+k-1) \times(k-1)}\) denote the rightmost \(k-1\) columns of \(\widehat{A}\), then, \(\operatorname{rank}(Z)=k-1\).

Proof. Consider the \((k-1) \times(k-1)\) submatrix \(Z^{[n+1: n+k-1]}\) of \(Z\). Each element on the diagonal of this submatrix should be at least \(B-C\|A\|_{1}\), and each element not on the diagonal of the submatrix should be at most \(C\|A\|_{1}\). Otherwise \(\|\widehat{A}-\widetilde{A}\|_{1}>C\|A\|_{1}\) which will lead to a contradiction. Since \(B=n^{c_{4}}\) is sufficiently large, \(Z^{[n+1: n+k-1]}\) is diagonally dominant. Thus \(\operatorname{rank}(Z) \geq \operatorname{rank}\left(Z^{[n+1: n+k-1]}\right)=k-1\). Because \(Z\) only has \(k-1\) columns, \(\operatorname{rank}(Z)=k-1\).

Claim 17.9.15. \(\forall x \in \mathbb{R}^{k-1}, i \in[n], \exists j \in\{n+1, n+2, \cdots, n+k-1\}\) such that,
\[
\frac{\left|(Z x)_{j}\right|}{\left|(Z x)_{i}\right|} \geq \frac{B}{2(k-1) C\|A\|_{1}} .
\]

Proof. Without loss of generality, we can let \(\|x\|_{1}=1\). Thus, there exists \(j\) such that \(\left|x_{j}\right| \geq \frac{1}{k-1}\). So we have
\[
\begin{aligned}
\left|(Z x)_{n+j}\right| & =\left|\sum_{i=1}^{k-1} Z_{n+j, i} x_{i}\right| \\
& \geq\left|Z_{n+j, j} x_{j}\right|-\sum_{i \neq j}\left|Z_{n+j, i} x_{i}\right| \\
& \geq\left(B-C\|A\|_{1}\right)\left|x_{j}\right|-\sum_{i \neq j}\left|x_{i}\right| C\|A\|_{1} \\
& \geq\left(B-C\|A\|_{1}\right) /(k-1)-C\|A\|_{1} \\
& \geq \frac{B}{2(k-1)} .
\end{aligned}
\]

The second inequality follows because \(\left|Z_{n+j, j}\right| \geq B-C\|A\|_{1}\) and \(\forall i \neq j,\left|Z_{n+j, i}\right| \leq C\|A\|_{1}\) (otherwise \(\|\widehat{A}-\widetilde{A}\|_{1}>C\|A\|_{1}\) which leads to a contradiction.) The third inequality follows
from \(\left|x_{j}\right|>1 /(k-1)\) and \(\|x\|_{1}=1\). The fourth inequality follows since \(B\) is large enough such that \(\frac{B}{2(k-1)} \geq \frac{C\|A\|_{1}}{k-1}+C\|A\|_{1}\).

Now we consider any \(q \in[n]\). We have
\[
\left|(Z x)_{q}\right|=\sum_{i=1}^{k-1}\left|Z_{q, i} x_{i}\right| \leq \max _{i \in[k-1]}\left|Z_{q, i}\right| \cdot \sum_{i=1}^{k-1}\left|x_{i}\right| \leq C\|A\|_{1} .
\]

The last inequality follows that \(\|x\|_{1}=1\) and \(\forall i \in[k-1], Z_{q, i} \leq C\|A\|_{1}\). Otherwise, \(\|\widehat{A}-\widetilde{A}\|_{1}>C\|A\|_{1}\) which will lead to a contradiction.

Look at \(\left|(Z x)_{n+j}\right| /\left|(Z x)_{q}\right|\), it is greater than \(\frac{B}{2(k-1) C\|A\|_{1}}\).

We now look at the submatrix \(\widehat{A}_{[1: n]}^{[1: n]}\), we choose \(i^{*}, j^{*} \in[n]\) such that
\[
\left|\widehat{A}_{i^{*}, j^{*}}\right| \geq 1 / n^{c_{2}-2}
\]

If there is no such \(\left(i^{*}, j^{*}\right)\), it means that we already found a good rank-1 approximation to A
\[
\begin{aligned}
\|\mathbf{0}-A\|_{1} & \leq\left\|\mathbf{0}-\widehat{A}_{[1: n]}^{[1: n]}\right\|_{1}+\left\|\widehat{A}_{[1: n]}^{[1: n]}-A\right\|_{1} \\
& \leq\left\|\mathbf{0}-\widehat{A}_{[1: n]}^{[1: n]}\right\|_{1}+\|\widehat{A}-\widetilde{A}\|_{1} \\
& \leq 1 / n^{c_{2}}+\|\widehat{A}-\widetilde{A}\|_{1} \\
& \leq 1 / n^{c_{2}}+C \min _{\operatorname{rank}-k \widehat{A^{\prime}}}\left\|\widehat{A}^{\prime}-\widetilde{A}\right\|_{1} \\
& \leq 1 / n^{c_{2}}+C \min _{u, v \in \mathbb{R}^{n}}\left\|u v^{\top}-A\right\|_{1},
\end{aligned}
\]
where \(\mathbf{0}\) is an \(n \times n\) all zeros matrix. The second inequality follows since \(\widehat{A}_{[1: n]}^{[1: n]}-A\) is a submatrix of \(\widehat{A}-\widetilde{A}\). The third inequality follows since each entry of \(\widehat{A}_{[1: n]}^{[1: n]}\) should be
no greater than \(1 / n^{c_{2}-2}\) (otherwise, we can find \(\left(i^{*}, j^{*}\right)\) ). The last inequality follows from equation 17.39.

Claim 17.9.16. \(\widehat{A}_{j^{*}}\) is not in the column span of \(Z\), i.e.,
\[
\forall x \in \mathbb{R}^{k-1}, \widehat{A}_{j^{*}} \neq Z x
\]

Proof. If there is an \(x\) such that \(\widehat{A}_{j^{*}}=Z x\), it means \((Z x)_{i^{*}}=\widehat{A}_{i^{*}, j^{*}} \geq 1 / n^{c_{2}-2}\). Due to Claim 17.9.15, there must exist \(i^{\prime} \in\{n+1, n+2, \cdots, n+k-1\}\) such that \((Z x)_{i^{\prime}} \geq\) \(\frac{B}{2(k-1) C\|A\|_{1} n^{c_{2}-2}}\). Since \(B\) is sufficiently large, \(\widehat{A}_{i^{\prime}, j^{*}}=(Z x)_{i^{\prime}}>C\|A\|_{1}\) which implies that \(\|\widehat{A}-\widetilde{A}\|_{1}>C\|A\|_{1}\), and so leads to a contradiction.

Due to Claim 17.9.14 and Claim 17.9.16, the dimension of the subspace spanned by \(\widehat{A}_{j^{*}}\) and the column space of \(Z\) is \(k\). Since \(\widehat{A}\) has rank at most \(k\), it means that each column of \(\widehat{A}\) can be written as a linear combination of \(\widehat{A}_{j^{*}}\) and the columns of \(Z\).

Now consider the \(j^{\text {th }}\) column \(\widehat{A}_{j}\) of \(\widehat{A}\) for \(j \in[n]\). We write it as
\[
\widehat{A}_{j}=\alpha_{j} \cdot \widehat{A}_{j^{*}}+Z x^{j}
\]

Claim 17.9.17. \(\forall j \in[n], \alpha_{j} \leq 2 C\|A\|_{1} n^{c_{2}+2}\).

Proof. Otherwise, suppose \(\alpha_{j}>2 C\|A\|_{1} n^{c_{2}+2}\). We have
\[
\begin{aligned}
\left|\left(Z x^{j}\right)_{i^{*}}\right| & \geq \alpha_{j} \cdot\left|\widehat{A}_{i^{*}, j^{*}}\right|-\left|\widehat{A}_{i^{*}, j}\right| \\
& \geq \alpha_{j} \cdot \frac{1}{n^{c_{2}-2}}-\left|\widehat{A}_{i^{*}, j}\right| \\
& \geq \frac{1}{2} \alpha_{j} \cdot \frac{1}{n^{c_{2}-2}} .
\end{aligned}
\]

The second inequality follows from \(\left|\widehat{A}_{i^{*}, j^{*}}\right| \geq 1 / n^{c 2-2}\). The third inequality follows from \(\left|\widehat{A}_{i^{*}, j}\right| \leq\|A\|_{1}\) and \(\frac{1}{2} \alpha_{j} \cdot \frac{1}{n^{c_{2}-2}} \geq C\|A\|_{1} \geq\|A\|_{1}\).

Due to Claim 17.9.15, there exists \(i \in\{n+1, n+2, \cdots, n+k-1\}\) such that \(\left|\left(Z x^{j}\right)_{i}\right| \geq\) \(\frac{B}{2(k-1) C\|A\|_{1}} \cdot \frac{1}{2} \alpha_{j} \cdot \frac{1}{n^{c_{2}-2}}\). For sufficiently large \(B\), we can have \(\left|\left(Z x^{j}\right)_{i}\right| \geq \alpha_{j} B^{1 / 2}\). Then we look at
\[
\begin{aligned}
\left|\widehat{A}_{i, j}\right| & \geq\left|\left(Z x^{j}\right)_{i}\right|-\alpha_{j}\left|\widehat{A}_{i, j^{*}}\right| \\
& \geq \alpha_{j}\left(B^{1 / 2}-C\|A\|_{1}\right) \\
& \geq \alpha_{j} \frac{1}{2} B^{1 / 2} .
\end{aligned}
\]

The second inequality follows by \(\left|\widehat{A}_{i, j^{*}}\right| \leq C\|A\|_{1}\), otherwise \(\|\widehat{A}-\widetilde{A}\|_{1}>C\|A\|_{1}\) which will lead to a contradiction. The third inequality follows that \(B\) is sufficient large that \(\frac{1}{2} B^{1 / 2}>C\|A\|_{1}\).

Since \(\left|\widehat{A}_{i, j}\right| \geq \alpha_{j} \frac{1}{2} B^{1 / 2}>C\|A\|_{1}\), it contradicts to the fact \(\|\widehat{A}-\widetilde{A}\|_{1} \leq C\|A\|_{1}\).
Therefore, \(\forall j \in[n], \alpha_{j} \leq 2 C\|A\|_{1} n^{c_{2}+2}\).
Claim 17.9.18. \(\forall j \in[n], i \in\{n+1, n+2, \cdots, n+k-1\},\left|\left(Z x^{j}\right)_{i}\right| \leq 4 C^{2}\|A\|_{1}^{2} n^{c_{2}+2}\)

Proof. Consider \(j \in[n], i \in\{n+1, n+2, \cdots, n+k-1\}\), we have
\[
\begin{aligned}
\left|\left(Z x^{j}\right)_{i}\right| & \leq\left|\widehat{A}_{i, j}\right|+\alpha_{j}\left|\widehat{A}_{i, j^{*}}\right| \\
& \leq C\|A\|_{1}+\alpha_{j} \cdot C\|A\|_{1} \\
& \leq 4 C^{2}\|A\|_{1}^{2} n^{c_{2}+2}
\end{aligned}
\]

The second inequality follows by \(\left|\widehat{A}_{i, j}\right| \leq C\|A\|_{1}\) and \(\left|\widehat{A}_{i, j^{*}}\right| \leq C\|A\|_{1}\), otherwise the \(\|\widehat{A}-\widetilde{A}\|_{1}\) will be too large and leads to a contradiction. The third inequality is due to \(\alpha_{j}+1 \leq\) \(4 C\|A\|_{1} n^{c_{2}+2}\) via Claim 17.9.17.

Claim 17.9.19. \(\forall j \in[n],\left\|\widehat{A}_{j}^{[1: n]}-\alpha_{j} \cdot \widehat{A}_{j^{*}}^{1: n}\right\|_{\infty} \leq 1 / n^{c_{2}-2}\)

Proof. Due to Claim 17.9.18 and Claim 17.9.15, \(\forall i, j \in[n]\), we have
\[
\left|\left(Z x^{j}\right)_{i}\right| \leq \frac{4 C^{2}\|A\|_{1}^{2} n^{c_{2}+2}}{B /\left(2(k-1) C\|A\|_{1}\right)} \leq 1 / B^{1 / 2}
\]

The second inequality follows for a large enough \(B\).
Therefore, \(\forall i, j \in[n]\),
\[
\left|\widehat{A}_{i, j}-\alpha_{j} \cdot \widehat{A}_{i, j^{*}}\right| \leq\left|\left(Z x^{j}\right)_{i}\right| \leq 1 / B^{1 / 2} \leq 1 / n^{c_{2}-2} .
\]

The last inequality follows since \(B\) is large enough.

Now, let us show that \(\widehat{A}_{j^{*}}^{1: n}\) can provide a good rank-1 approximation to \(A\) :
\[
\begin{aligned}
\left\|\widehat{A}_{j^{*}}^{1: n} \alpha^{\top}-A\right\|_{1} & \leq\left\|\widehat{A}_{j^{*}}^{1: n} \alpha^{\top}-\widehat{A}_{[1: n]}^{[1: n]}\right\|_{1}+\left\|\widehat{A}_{[1: n]}^{[1: n]}-A\right\|_{1} \\
& =\sum_{j=1}^{n}\left\|\alpha_{j} \widehat{A}_{j^{*}}-\widehat{A}_{j}^{[1: n]}\right\|_{1}+\left\|\widehat{A}_{[1: n]}^{[1: n]}-A\right\|_{1} \\
& \leq n^{2} \cdot 1 / n^{c_{2}-2}+\left\|\widehat{A}_{[1: n]}^{[1: n]}-A\right\|_{1} \\
& \leq 1 / n^{c_{2}}+\|\widehat{A}-\widetilde{A}\|_{1} \\
& \leq 1 / n^{c_{2}}+C \min _{u, v \in \mathbb{R}^{n}}\left\|u v^{\top}-A\right\|_{1} .
\end{aligned}
\]

The first inequality follows by triangle inequality. The first equality is due to the linearity of \(\ell_{1}\) norm. The second inequality is due to Claim 17.9.19. The third inequality follows since \(\widehat{A}_{[1: n]}^{[1: n]}-A\) is a submatrix of \(\widehat{A}-\widetilde{A}\). The fourth inequality is due to the equation 17.39.

\subsection*{17.10 Experiments and Discussions}

In this section, we provide some counterexamples for the other heuristic algorithms such that, for those examples, the heuristic algorithms can output a solution with a very "bad" approximation ratio, i.e., \(n^{c}\), where \(c>0\) and the input matrix has size \(n \times n\). We not only observe that heuristic algorithms sometimes have very bad performance in practice, but also give a proof in theory.

\subsection*{17.10.1 Setup}

We provide some details of our experimental setup. We obtained the R package of [KK05, Kwa08, BDB13] from https://cran.r-project.org/web/packages/pcaL1/index. html. We also implemented our algorithm and the r1-pca algorithm [DZHZ06] using the R language. The version of the R language is 3.0 .2 . We ran experiments on a machine with Intel X5550@2.67GHz CPU and 24G memory. The operating system of that machine is Linux Ubuntu 14.04.5 LTS. All the experiments were done in single-threaded mode.


Figure 17.1: The \(x\)-axis is \(n\) where \(A \in \mathbb{R}^{n \times n}\), and the \(y\)-axis is \(\left\|A^{\prime}-A\right\|_{1}\) where \(\operatorname{rank}\left(A^{\prime}\right)=k\). This figure shows the performance of both our algorithm and [DZHZ06] on input matrix \(A\) defined as Equation (17.40).

\subsection*{17.10.2 Counterexample for [DZHZ06]}

The goal is to find a rank \(k=1\) approximation for matrix \(A\). For any \(\epsilon \in[0,0.5)\), we define \(A \in \mathbb{R}^{n \times n}\) as
\[
A=\left[\begin{array}{cc}
n^{1.5+\epsilon} & 0  \tag{17.40}\\
0 & 0
\end{array}\right]+\left[\begin{array}{ll}
0 & 0 \\
0 & B
\end{array}\right]
\]
where \(B \in \mathbb{R}^{(n-1) \times(n-1)}\) is all 1 s matrix. It is immediate that the optimal cost is at most \(n^{1.5+\epsilon}\). However, using the algorithm in [DZHZ06], the cost is at least \(\Omega\left(n^{2}\right)\). Thus, we
can conclude, using algorithm [DZHZ06] to solve \(\ell_{1}\) low rank approximation problem on \(A\) cannot achieve an approximation ratio better than \(n^{0.5-\epsilon}\).


Figure 17.2: The \(x\)-axis is \(n\) where \(A \in \mathbb{R}^{n \times n}\), and the \(y\)-axis is \(\left\|A^{\prime}-A\right\|_{1}\) where \(\operatorname{rank}\left(A^{\prime}\right)=k\). This figure shows the performance of both our algorithm and [BDB13] on input matrix \(A\) defined as Equation (17.41).

\subsection*{17.10.3 Counterexample for [BDB13]}

The goal is to find a rank \(k=1\) approximation for matrix \(A\). The input matrix \(A \in \mathbb{R}^{d \times d}\) for algorithm [BDB13] is defined to be,
\[
A=\left[\begin{array}{cc}
n^{1.5} & 0  \tag{17.41}\\
0 & 0
\end{array}\right]+\left[\begin{array}{ll}
0 & 0 \\
0 & B
\end{array}\right]
\]
where \(B \in \mathbb{R}^{(n-1) \times(n-1)}\) is an all 1s matrix. It is immediate that the optimal cost is at most \(n^{1.5}\). Now, let us look at the procedure of [BDB13]. Basically, the algorithm described
in [BDB13] is that they first find a rank \(n-1\) approximation via a best \(\ell_{1}\)-fit hyperplane algorithm, then they rotate it based on the right singular vectors of the rank \(n-1\) approximation matrix, and next they recursively do the same thing for the rotated matrix which has only \(n-1\) columns.

When running their algorithm on \(A\), they will fit an arbitrary column except the first column of \(A\). Without loss of generality, it just fits the last column of \(A\). After the rotation, the matrix will be an \(n \times(n-1)\) matrix:
\[
\left[\begin{array}{ccccc}
n^{1.5} & 0 & 0 & \cdots & 0 \\
0 & \sqrt{n-2} & 0 & \cdots & 0 \\
0 & \sqrt{n-2} & 0 & \cdots & 0 \\
\cdots & \cdots & \cdots & \cdots & \cdots \\
0 & \sqrt{n-2} & 0 & \cdots & 0
\end{array}\right] .
\]

Then, after the \(t^{\text {th }}\) iteration for \(t<(n-1)\), they will get an \(n \times(n-t)\) matrix:
\[
\left[\begin{array}{ccccc}
n^{1.5} & 0 & 0 & \cdots & 0 \\
0 & \sqrt{n-2} & 0 & \cdots & 0 \\
0 & \sqrt{n-2} & 0 & \cdots & 0 \\
\cdots & \cdots & \cdots & \cdots & \cdots \\
0 & \sqrt{n-2} & 0 & \cdots & 0
\end{array}\right]
\]

This means that their algorithm will run on an \(n \times 2\) matrix:
\[
\left[\begin{array}{cc}
n^{1.5} & 0 \\
0 & \sqrt{n-2} \\
0 & \sqrt{n-2} \\
\cdots & \cdots \\
0 & \sqrt{n-2}
\end{array}\right],
\]
in the last iteration. Notice that \(n \times \sqrt{n-2}<n^{1.5}\). This means that their algorithm will fit the first column which implies that their algorithm will output a rank-1 approximation to \(A\) by just fitting the first column of \(A\). But the cost of this rank- 1 solution is \((n-1)^{2}\). Since the optimal cost is at most \(n^{1.5}\), their algorithm cannot achieve an approximation ratio better than \(n^{0.5}\).


Figure 17.3: The \(x\)-axis is \(n\) where \(A \in \mathbb{R}^{n \times n}\), and the \(y\)-axis is \(\left\|A^{\prime}-A\right\|_{1}\) where \(\operatorname{rank}\left(A^{\prime}\right)=k\). This figure shows the performance of both our algorithm and [Kwa08] on input matrix \(A\) defined as Equation (17.42). Algorithm [Kwa08] has two ways of initialization, which have similar performance on matrix \(A\).

\subsection*{17.10.4 Counterexample for [Kwa08]}

We show that the algorithm [Kwa08] cannot achieve an approximation ratio better than \(\Theta(n)\) on the matrix \(A \in \mathbb{R}^{(2 n+1) \times(2 n+1)}\) defined as,
\[
A=\left[\begin{array}{ccc}
n^{1.5} & 0 & 0  \tag{17.42}\\
0 & B & 0 \\
0 & 0 & B
\end{array}\right]
\]
where \(B\) is an \(n \times n\) matrix that contains all 1 s. We consider the rank- 2 approximation problem for this input matrix \(A\). The optimal cost is at most \(n^{1.5}\).

We run their algorithm. Let \(x_{0}\) denote the initial random unit vector. Consider the sign vector \(s \in\{ \pm 1\}^{2 n+1}\) where each entry is the sign of the inner product between \(x_{0}\) and each column of \(A\). There are only three possibilities,
\[
s=\left\{\begin{array}{l}
\left(1,\{1\}^{n},\{1\}^{n}\right) \\
\left(1,\{1\}^{n},\{-1\}^{n}\right) \\
\left(1,\{-1\}^{n},\{1\}^{n}\right)
\end{array}\right.
\]

Case I, \(s=\left(1,\{1\}^{n},\{1\}^{n}\right)\). Define \(\widehat{u}=\sum_{i=1}^{2 n+1} s_{i} \cdot A_{i}=\left(n^{1.5}, n, \cdots, n\right)\). Let \(u=\widehat{u} /\|u\|_{2}=\) \(\widehat{u} /\left(\sqrt{3} n^{1.5}\right)=(1 / \sqrt{3}, 1 / \sqrt{3 n}, \cdots, 1 / \sqrt{3 n})\). Define matrix \(D\) to be \(A-u u^{\top} A\). Then, we can compute \(D\),
\[
\begin{aligned}
D & =A-u u^{\top} A \\
& =A-\left[\begin{array}{ccc}
1 / 3 & \frac{1}{3 \sqrt{n}} \mathbf{1}^{\top} & \frac{1}{3 \sqrt{n}} \mathbf{1}^{\top} \\
\frac{1}{3 \sqrt{n}} \mathbf{1} & \frac{1}{3 n} B & \frac{1}{3 n} B \\
\frac{1}{3 \sqrt{n}} \mathbf{1} & \frac{1}{3 n} B & \frac{1}{3 n} B
\end{array}\right] A \\
& =\left[\begin{array}{ccc}
n^{1.5} & 0 & 0 \\
0 & B & 0 \\
0 & 0 & B
\end{array}\right]-\left[\begin{array}{ccc}
\frac{n^{1.5}}{3} & \frac{\sqrt{n}}{3} \mathbf{1}^{\top} & \frac{\sqrt{n}}{3} \mathbf{1}^{\top} \\
\frac{n}{3} \mathbf{1} & \frac{1}{3} B & \frac{1}{3} B \\
\frac{n}{3} \mathbf{1} & \frac{1}{3} B & \frac{1}{3} B
\end{array}\right] \\
& =\left[\begin{array}{ccc}
2 n^{1.5} / 3 & -\frac{\sqrt{n}}{3} \mathbf{1}^{\top} & -\frac{\sqrt{n}}{3} \mathbf{1}^{\top} \\
-\frac{n}{3} \mathbf{1} & \frac{2}{3} B & -\frac{1}{3} B \\
-\frac{n}{3} \mathbf{1} & -\frac{1}{3} B & \frac{2}{3} B
\end{array}\right] .
\end{aligned}
\]

Now we need to take the linear combination of columns \(D=A-u u^{\top} A\). Let \(w\) denote another sign vector \(\{-1,+1\}^{2 n+1}\). Then let \(v\) denote the basis vector, \(v=\sum_{i=1}^{2 n+1} D_{i}\). There are three possibilities, Case \(\mathrm{I}(\mathrm{a})\), if \(w=\left(1,\{1\}^{n},\{1\}^{n}\right)\), then \(v\) is the all 0 vector. Using vector \(v\) to interpolate each column of \(D\), the cost we obtain is at least \(2 n^{2}\). Case \(\mathrm{I}(\mathrm{b})\), if \(w=\left(1,\{1\}^{n},\{-1\}^{n}\right)\), then \(v=\left(2 n^{1.5} / 3,\{2 / 3\}^{n},\{-4 / 3\}^{n}\right)\). We also obtain at least \(2 n^{2}\) cost if we use that \(v\) to interpolate each column of \(D\). Case \(\mathrm{I}(\mathrm{c})\), if \(w=\left(1,\{-1\}^{n},\{-1\}^{n}\right)\), then \(v=\left(0,\{-2 / 3\}^{n},\{-2 / 3\}^{n}\right)\). The cost is also at least \(2 n^{2}\).

Case II, \(s=\left(1,\{1\}^{n},\{-1\}^{n}\right)\). Define \(\mathbf{1}\) to be a length \(n\) all 1 s column vector. Define \(\widehat{u}=\sum_{i=1}^{2 n+1} s_{i} \cdot A_{i}=\left(n^{1.5},\{n\}^{n},\{-n\}^{n}\right)\). Let
\[
u=\widehat{u} /\|u\|_{2}=\widehat{u} /\left(\sqrt{3} n^{1.5}\right)=\left(1 / \sqrt{3},\{1 / \sqrt{3 n}\}^{n},\{-1 / \sqrt{3 n}\}^{n}\right)
\]

Define \((2 n+1) \times(2 n+1)\) matrix \(D\) to be \(A-u u^{\top} A\). Then, we can compute \(D\),
\[
\begin{aligned}
D & =A-u u^{\top} A \\
& =A-\left[\begin{array}{ccc}
1 / 3 & \frac{1}{3 \sqrt{n}} \mathbf{1}^{\top} & -\frac{1}{3 \sqrt{n}} \mathbf{1}^{\top} \\
\frac{1}{3 \sqrt{n}} \mathbf{1} & \frac{1}{3 n} B & -\frac{1}{3 n} B \\
-\frac{1}{3 \sqrt{n}} \mathbf{1} & -\frac{1}{3 n} B & \frac{1}{3 n} B
\end{array}\right] A \\
& =\left[\begin{array}{ccc}
n^{1.5} & 0 & 0 \\
0 & B & 0 \\
0 & 0 & B
\end{array}\right]-\left[\begin{array}{ccc}
n^{1.5} / 3 & \frac{\sqrt{n}}{3} \mathbf{1}^{\top} & -\frac{\sqrt{n}}{3} \mathbf{1}^{\top} \\
\frac{n}{3} \mathbf{1} & \frac{1}{3} B & -\frac{1}{3} B \\
-\frac{n}{3} \mathbf{1} & -\frac{1}{3} B & \frac{1}{3} B
\end{array}\right] \\
& =\left[\begin{array}{ccc}
2 n^{1.5} / 3 & -\frac{\sqrt{n}}{3} \mathbf{1}^{\top} & \frac{\sqrt{n}}{3} \mathbf{1}^{\top} \\
-\frac{n}{3} \mathbf{1} & \frac{2}{3} B & \frac{1}{3} B \\
\frac{n}{3} \mathbf{1} & \frac{1}{3} B & \frac{2}{3} B
\end{array}\right] .
\end{aligned}
\]

Similarly to the previous case, we can also discuss three cases.
Case III, \(s=\left(1,\{1\}^{n},\{-1\}^{n}\right)\). Define 1 to be a length \(n\) all 1 s column vector. Define \(\widehat{u}=\sum_{i=1}^{2 n+1} s_{i} \cdot A_{i}=\left(n^{1.5},\{-n\}^{n},\{-n\}^{n}\right)\). Let \(u=\widehat{u} /\|u\|_{2}=\widehat{u} /\left(\sqrt{3} n^{1.5}\right)=\) \(\left(1 / \sqrt{3},\{-1 / \sqrt{3 n}\}^{n},\{-1 / \sqrt{3 n}\}^{n}\right)\). Define matrix \(D\) to be \(A-u u^{\top} A\). Then, we can compute D,
\[
\begin{aligned}
D & =A-u u^{\top} A \\
& =A-\left[\begin{array}{ccc}
1 / 3 & \left.-\frac{1}{3 \sqrt{n}}\right) \mathbf{1}^{\top} & -\frac{1}{3 \sqrt{n}} \mathbf{1}^{\top} \\
-\frac{1}{3 \sqrt{n}} \mathbf{1} & \frac{1}{3 n} B & \frac{1}{3 n} B \\
-\frac{1}{3 \sqrt{n}} \mathbf{1} & \frac{1}{3 n} B & \frac{1}{3 n} B
\end{array}\right] A \\
& =\left[\begin{array}{ccc}
n^{1.5} & 0 & 0 \\
0 & B & 0 \\
0 & 0 & B
\end{array}\right]-\left[\begin{array}{ccc}
n^{1.5} / 3 & -\frac{\sqrt{n}}{3} \mathbf{1}^{\top} & -\frac{\sqrt{n}}{3} \mathbf{1}^{\top} \\
-\frac{n}{3} \mathbf{1} & \frac{1}{3} B & \frac{1}{3} B \\
-\frac{n}{3} \mathbf{1} & \frac{1}{3} B & \frac{1}{3} B
\end{array}\right] \\
& =\left[\begin{array}{ccc}
2 n^{1.5} / 3 & \frac{\sqrt{n}}{3} \mathbf{1}^{\top} & \frac{\sqrt{n}}{3} \mathbf{1}^{\top} \\
\frac{n}{3} \mathbf{1} & \frac{2}{3} B & \frac{1}{3} B \\
\frac{n}{3} \mathbf{1} & \frac{1}{3} B & \frac{2}{3} B
\end{array}\right] .
\end{aligned}
\]

Similarly to the previous case, we can also discuss three cases.

\subsection*{17.10.5 Counterexample for [KK05]}

We show that there exist matrices such that the algorithm of [KK05] cannot achieve an approximation ratio better than \(\Theta(n)\). Their algorithm has two different ways of initialization. We provide counterexamples for each of the initialization separately.

Random vector initialization We provide a counterexample matrix \(A \in \mathbb{R}^{n \times n}\) defined as,
\[
A=\left[\begin{array}{cccc}
n^{c} & 0 & \cdots & 0  \tag{17.43}\\
0 & 0 & \cdots & 0 \\
\cdots & \cdots & \cdots & \cdots \\
0 & 0 & \cdots & 0
\end{array}\right]+I
\]
where \(c \geq 2\). Consider the rank- 1 approximation problem for this matrix \(A\). The optimal cost is at most \(n-1\).

Run their algorithm. The starting vectors are \(u(0) \sim N(0, I)\) and \(v(0) \sim N(0,1)\). We define two properties for a given vector \(y \in \mathbb{R}^{n}\). Property I is for all \(i \in[n],\left|y_{i}\right| \leq n / 8\), and Property II is there exist half of the \(i\) such that \(\left|y_{i}\right| \geq 1 / 2\). We can show that with probability \(1-2^{-\Omega(n)}\), both \(u(0)\) and \(v(0)\) satisfy Property I and II. After 1 iteration, we can show that \(u(1)_{1}=v(1)_{1}=0\)

Now let us use column vector \(u(0) \in \mathbb{R}^{n}\) to interpolate the first column of matrix \(A\). We simplify \(u(0)\) to be \(u\). Let \(u_{i}\) denote the \(i\)-th coordinate of vector \(u, \forall i \in[n]\). Let \(A_{1}\) denote the first column of matrix \(A\). We define \(\alpha=v(1)_{1}=\arg \min _{\alpha}\left\|\alpha u(0)-A_{1}\right\|_{1}\). For any scalar \(\alpha\), the cost we pay on the first column of matrix \(A\) is,
\[
\begin{equation*}
\left|\alpha \cdot u_{1}-n^{c}\right|+\sum_{i=2}^{n}\left|\alpha u_{i}\right| \geq\left|\alpha \cdot u_{1}-n^{c}\right|+\frac{n}{2}\left|\alpha \cdot \frac{1}{2}\right| \geq\left|\alpha \cdot u_{1}-n^{c}\right|+\frac{n}{4}|\alpha| \tag{17.44}
\end{equation*}
\]


Figure 17.4: The \(x\)-axis is \(n\) where \(A \in \mathbb{R}^{n \times n}\), and the \(y\)-axis is \(\left\|A^{\prime}-A\right\|_{1}\) where \(\operatorname{rank}\left(A^{\prime}\right)=k\). Algorithm [KK05] has two ways of initialization. The left figure shows the performance of both our algorithm and [KK05] with random vector initialization on input matrix \(A\) defined as Equation (17.43). The right figure shows the performance of both our algorithm and [KK05] with top singular vector initialization on input matrix \(A\) defined as Equation (17.45).

Notice that \(c \geq 2\). Then \(\left|\alpha \cdot u_{1}-n^{c}\right|+\frac{n}{4}|\alpha| \geq\left|\alpha \cdot \frac{n}{8}-n^{c}\right|+\frac{n}{4}|\alpha|\). If \(\alpha \leq 0\), then the cost is minimized when \(\alpha=0\). If \(\alpha \in\left[0,8 n^{c-1}\right]\), the cost is minimized when \(\alpha=0\). If \(\alpha \geq 8 n^{c-1}\), the cost is minimized when \(\alpha=8 n^{c-1}\). Putting it all together, to achieve the minimum cost, there is only one choice for \(\alpha\), which is \(\alpha=0\). The optimal cost is at least \(n^{c}\).

Then after \(T\) iterations(for any \(T \geq 1\) ), \(u(T)_{1}=v(T)_{1}=0\). Thus, we always pay at least \(n^{c}\) cost on the first entry.

Therefore, their algorithm cannot achieve any approximation ratio better than \(n^{c-1}\). Because \(c \geq 2\), we complete the proof.

Top singular vector initialization The counterexample input matrix \(A \in \mathbb{R}^{n \times n}\) is defined as,
\[
A=\left[\begin{array}{ll}
n & 0  \tag{17.45}\\
0 & B
\end{array}\right]
\]
where matrix \(B \in \mathbb{R}^{(n-1) \times(n-1)}\) contains all 1 s . Consider the rank-1 approximation problem for this matrix \(A\). The optimal cost is at most \(n\). Run their algorithm. The starting vectors \(u(0)\) and \(v(0)\) will be set to \((1,0, \cdots, 0) \in \mathbb{R}^{n}\). After \(T\) iterations(for any \(T>0\) ), the support of \(u(T)\) (resp. \(v(T)\) ) is the same as \(u(0)\) (resp. \(v(T)\) ). Thus, the cost is at least \(\|B\|_{1}=(n-1)^{2}\). Therefore, we can conclude that their algorithm cannot achieve any approximation ratio better than \((n-1)^{2} / n=\Theta(n)\).

\subsection*{17.10.6 Counterexample for all}


Figure 17.5: Let \(A\) be \((2 n+2) \times(2 n+2)\) input matirx. (a) shows the performance of all the algorithms when the matrix dimension is growing. The \(x\)-axis is \(n\), and the \(y\)-axis is \(\left\|A^{\prime}-A\right\|_{1}\) where \(A^{\prime}\) is the rank- 3 solution output by all the heuristic algorithms and also ours. The \(\ell_{1}\) residual cost of all the other algorithms is growing much faster than ours, which is consistent with our theoretical results. (b) shows the running time (in seconds) of all the algorithms when the matrix dimension \(n\) is growing. The \(x\)-axis is \(n\) and the \(y\)-axis is time (seconds). The running time of some of the algorithms is longer than 3 seconds. For most of the algorithms (including ours), the running time is always less than 3 seconds.

For any \(\epsilon \in(0,0.5)\) and \(\gamma>0\), we construct the input matrix \(A \in \mathbb{R}^{(2 n+2) \times(2 n+2)}\) as follows
\[
A=\left[\begin{array}{cccc}
n^{2+\gamma} & 0 & 0 & 0 \\
0 & n^{1.5+\epsilon} & 0 & 0 \\
0 & 0 & B & 0 \\
0 & 0 & 0 & B
\end{array}\right],
\]
where \(B\) is \(n \times n\) all 1 s matrix. We want to find a rank \(k=3\) solution for \(A\). Then any of those four heuristic algorithms [KK05, DZHZ06, Kwa08, BDB13] is not able to achieve better than \(n^{\min (\gamma, 0.5-\epsilon)}\) approximation ratio. We present our main experimental results in Figure 17.5.

Both [KK05] and [Kwa08] have two different ways of initialization. In Figure 17.5 we use KK05r(resp. Kwak08r) to denote the way that uses random vector as initialization, and use KK05s(resp. Kwak08s) to denote the way that uses top singular vector as initialization. Figure 17.5 (a) shows the performance of all the algorithms and Figure 17.5(b) presents the running time. The \(\ell_{1}\) residual cost of all the other algorithm is growing much faster than our algorithm. Most of the algorithms (including ours) are pretty efficient, i.e., the running time is always below 3 seconds. The running time of [BDB13, KK05] is increasing very fast when the matrix dimension \(n\) is growing.

In Figure 17.5(a), the cost of KK05r at \(\{82, \cdots, 142\}\) is in \(\left[10^{5}, 10^{6}\right]\), at \(\{162, \cdots, 302\}\) is in \(\left[10^{6}, 10^{7}\right]\), and at \(\{322, \cdots, 402\}\) is in \(\left[10^{7}, 10^{8}\right]\). In Figure \(17.5(\mathrm{~b})\), the time of KK05r at \(\{382,482\}\) is 64 s and 160 s . The running time of BDB 13 at \(\{82, \cdots, 222\}\) is between 1 minute and 1 hour. The running time of \(\operatorname{BDB} 13\) at \(\{242, \cdots, 322\}\) is between 1 hour and 20 hours. The running time of BDB 13 at \(\{342, \cdots, 402\}\) is more than 20 hours.

\subsection*{17.10.7 Discussion for Robust PCA [CLMW11]}

A popular method is robust PCA [CLMW11], which given a matrix \(A\), tries to find a matrix \(L\) for which \(\lambda\|A-L\|_{1}+\|L\|_{*}\) is minimized, where \(\lambda>0\) is a tuning parameter and \(\|L\|_{*}\) is the nuclear norm of \(L\). This is a convex program, but it need not return a low rank matrix \(L\) with relative error. As a simple example, suppose \(\lambda=1\), and the \(n \times n\) matrix \(A\) is a block-diagonal matrix of rank \(k\) and \(n=\frac{k}{2}(b+1)\). Further, the first \(k / 2\) blocks are \(b \times b\) matrices of all 1 s , while the next \(k / 2\) blocks are just a single value \(b\) on the diagonal.

Then the solution to the above problem may return \(L\) to be the first \(k / 2\) blocks of \(A\). The total cost of \(\lambda\|A-L\|_{1}+\|L\|_{*}\) is \((k / 2) b+(k / 2) b=k b\).

Also, the solution to the above problem may return \(L\) to be \(A\), which has cost \(0+\) \(\|A\|_{*}=k b\). Because this solution has the same cost, it means that their algorithm might output a rank- \(k\) solution, and also might output a rank- \(k / 2\) solution.

Therefore, the relative error of the output matrix may be arbitrarily bad for \(\ell_{1}\)-low rank approximation.

We also consider the following example. Suppose \(\lambda=1 / \sqrt{n}\), and let \(n \times n\) matrix \(A\) denote the Hadamard matrix \(H_{n}\). Recall that the Hadamard matrix \(H_{p}\) of size \(p \times p\) is defined recursively : \(\left[\begin{array}{cc}H_{p / 2} & H_{p / 2} \\ H_{p / 2} & -H_{p / 2}\end{array}\right]\) with \(H_{2}=\left[\begin{array}{cc}+1 & +1 \\ +1 & -1\end{array}\right]\). Notice that every singular values of \(A\) is \(\sqrt{n}\). We consider the objective function \(\lambda\|A-L\|_{1}+\|L\|_{*}\).

Then the solution to the above problem may return \(L\) to be the first \(n / 2\) rows of \(A\). The total cost of \(\lambda\|A-L\|_{1}+\|L\|_{*}\) is \((1 / \sqrt{n}) n^{2} / 2+(n / 2) \sqrt{n}=n^{1.5}\). Also, the solution to the above problem may return \(L\) to be \(A\), which has cost \(0+\|A\|_{*}=n \sqrt{n}=n^{1.5}\). For any \(i\), if the solution takes \(i\) rows of \(A\), the cost is \((1 / \sqrt{n})(n-i) n+i \sqrt{n}=n^{1.5}\). Because this
solution has the same cost, it means that their algorithm might output a rank- \(n\) solution, and also might output a rank- \(n / 2\) solution. Therefore, the relative error of the output matrix may be arbitrarily bad for \(\ell_{1}\)-low rank approximation.

\subsection*{17.11 Limited Independent Cauchy Random Variables}

This section presents the fundamental lemmas with limited independent Cauchy variables, which will be used in Section 17.12 and 17.13. In Section 17.11.1, we provide some notation, definitions and tools from previous work. Section 17.11.2 includes the main result.

\subsection*{17.11.1 Notations and tools}

For a function \(f: \mathbb{R} \rightarrow \mathbb{R}\) and nonnegative integer \(\ell, f^{(\ell)}\) denotes the \(\ell\) th derivative of \(f\), with \(f^{(0)}=f\). We also often use \(x \approx_{\epsilon} y\) to state that \(|x-y|=O(\epsilon)\). We use \(I_{[a, b]}\) to denote the indicator function of the interval \([a, b]\).

To optimize the communication complexity of our distributed algorithm, we show that instead of using fully independent Cauchy variables, poly \((k, d)\)-wise independent Cauchy variables suffice.

We start by stating two useful Lemmas from previous work [KNW10a].

Lemma 17.11.1 (Lemma 2.2 in [KNW10a]). There exists an \(\epsilon_{0}>0\) such that the following holds. Let \(n\) be a positive integer and \(0<\epsilon<\epsilon_{0}, 0<p<2\) be given. Let \(f: \mathbb{R} \rightarrow \mathbb{R}\) satisfy \(\left\|f^{(\ell)}\right\|_{\infty}=O\left(\alpha^{\ell}\right)\) for all \(\ell \geq 0\), for some \(\alpha\) satisfying \(\alpha^{p} \geq \log (1 / \epsilon)\). Let \(k=\alpha^{p}\). Let \(a \in \mathbb{R}^{n}\) satisfy \(\|a\|_{p}=O(1)\). Let \(X_{i}\) be a \(3 C k\)-independent family of p-stable random variables. Let \(X=\sum_{i} a_{i} X_{i}\) and \(Y=\sum_{i} a_{i} Y_{i}\). Then \(\mathbb{E}[f(x)]=\mathbb{E}[f(Y)]+O(\epsilon)\).

Lemma 17.11.2 (Lemma 2.5 in [KNW10a]). There exist constants \(c^{\prime}\), \(\epsilon_{0}>0\) such that for all \(c>0\) and \(0<\epsilon<\epsilon_{0}\), and for all \([a, b] \subseteq \mathbb{R}\), there exists a function \(J_{[a, b]}^{c}: \mathbb{R} \rightarrow \mathbb{R}\) satisfying:
i. \(\left\|\left(J_{[a, b]}^{c}\right)^{(\ell)}\right\|_{\infty}=O\left(c^{\ell}\right)\) for all \(\ell \geq 0\).
ii. For all \(x\) such that \(a, b \notin[x-\epsilon, x+\epsilon]\), and as long as \(c>c^{\prime} \epsilon^{-1} \log ^{3}(1 / \epsilon)\), \(\left|J_{[a, b]}^{c}(x)-I_{[a, b]}(x)\right|<\epsilon\).

\subsection*{17.11.2 Analysis of limited independent random Cauchy variables}

Lemma 17.11.3. Given a vector \(y \in \mathbb{R}^{n}\), choose \(Z\) to be the \(t \times n\) random Cauchy matrices with \(1 / t\) rescaling and \(t=O(k \log k)\). The variables from different rows are fully independent, and the variables from the same rows are \(O(1)\)-wise independent. Then, we have
\[
\|Z y\|_{1} \gtrsim\|y\|_{1}
\]
holds with probability at least \(1-2^{-\Omega(t)}\).

Proof. Let \(S\) denote the original fully independent matrix and \(Z\) denote the matrix for which the entries in the same row are \(w\)-wise independent, and the entries from different rows are fully independent. Notice we define the random matrices without rescaling by \(1 / t\) and it will be added back at the end. (We will decide \(w\) later)

We define random variable \(X\) such that \(X=1\) if \(\left|(Z y)_{i}\right| \leq \frac{1}{50}\) and \(X=0\) otherwise. We also define random variable \(Y\) such that \(Y=1\) if \(\left|(S y)_{i}\right| \leq \frac{1}{50}\) and \(Y=0\) otherwise. Then, we have
\[
\begin{aligned}
& \mathbb{E}[X]=\operatorname{Pr}\left[\left|(Z y)_{i}\right| \leq \frac{1}{50}\right]=\mathbb{E}\left[I_{\left[-\frac{1}{50}, \frac{1}{50}\right]}\left((Z y)_{i}\right)\right] \\
& \mathbb{E}[Y]=\operatorname{Pr}\left[\left|(S y)_{i}\right| \leq \frac{1}{50}\right]=\mathbb{E}\left[I_{\left[-\frac{1}{50}, \frac{1}{50}\right]}\left((S y)_{i}\right)\right]
\end{aligned}
\]

The goal is to show that \(\mathbb{E}[X] \approx_{\epsilon} \mathbb{E}[Y]\). Following the same idea from [KNW10a], we need to argue this chain of inequalities,
\[
\mathbb{E}\left[I_{[a, b]}(X)\right] \approx_{\epsilon} \mathbb{E}\left[J_{[a, b]}^{c}(X)\right] \approx_{\epsilon} \mathbb{E}\left[J_{[a, b]}^{c}(Y)\right] \approx_{\epsilon} \mathbb{E}\left[I_{[a, b]}(Y)\right]
\]

Using Lemma 2.2 and Lemma 2.5 from [KNW10a], choosing sufficiently small constant \(\epsilon(\) which implies \(w=O(1))\), it follows that for each \(i \in[t]\), we still have
\[
\operatorname{Pr}\left[\left|(Z y)_{i}\right|>\frac{1}{50}\|y\|_{1}\right] \gtrsim \operatorname{Pr}\left[\left|(S y)_{i}\right|>\frac{1}{50}\|y\|_{1}\right] \geq 0.9
\]

Because all rows of \(Z\) are fully independent, using the Chernoff bound we can get that
\[
\operatorname{Pr}\left[\|Z y\|_{1} \lesssim t\|y\|_{1}\right] \leq \exp (-\Omega(t))
\]
as we needed for the "no contraction" part of the net argument.

For the no dilation, we need to argue that

Lemma 17.11.4. Given a set of vectors \(\left\{y_{1}, y_{2}, \ldots, y_{d}\right\}\) where \(y_{i} \in \mathbb{R}^{n}, \forall i \in[d]\), choose \(Z\) to be the \(t \times n\) random Cauchy matrices with \(1 / t\) rescaling and \(t=O(k \log k)\), where the variables from different rows are fully independent, and the variables from the same rows are \(w\)-wise independent.
I. If \(w=\widetilde{O}(d k)\), we have
\[
\sum_{i=1}^{d}\left\|Z y_{i}\right\|_{1} \leq O(\log d) \sum_{i=1}^{d}\left\|y_{i}\right\|_{1}
\]
holds with probability at least . 999 .
II. If If \(w=\widetilde{O}(d)\), we have
\[
\sum_{i=1}^{d}\left\|Z y_{i}\right\|_{1} \leq O(k \log d) \sum_{i=1}^{d}\left\|y_{i}\right\|_{1}
\]
holds with probability at least . 999 .

Proof. Let \(m=t\). Let \(S \in \mathbb{R}^{m \times n}\) denote the original fully independent matrix and \(Z\) denote the matrix that for each entry in the same row are \(w\)-wise independent, where the entries from different rows are fully independent. (We will decide on \(w\) later)

Applying matrix \(S\) to those fixed set of vectors, we have
\[
\sum_{i=1}^{d}\left\|S y_{i}\right\|_{1}=\sum_{i=1}^{d} \sum_{j=1}^{m}\left|\sum_{l=1}^{n} \frac{1}{m} S_{j, l} \cdot\left(y_{i}\right)_{l}\right|=\frac{1}{m} \sum_{i=1}^{d} \sum_{j=1}^{m}\left|\sum_{l=1}^{n} S_{j, l} \cdot\left(y_{i}\right)_{l}\right|
\]

Applying matrix \(Z\) to those fixed set of vectors, we have a similar thing,
\[
\sum_{i=1}^{d}\left\|Z y_{i}\right\|_{1}=\sum_{i=1}^{d} \sum_{j=1}^{m}\left|\sum_{l=1}^{n} \frac{1}{m} Z_{j, l} \cdot\left(y_{i}\right)_{l}\right|=\frac{1}{m} \sum_{i=1}^{d} \sum_{j=1}^{m}\left|\sum_{l=1}^{n} Z_{j, l} \cdot\left(y_{i}\right)_{l}\right|
\]

The goal is to argue that, for any \(i \in[d], j \in[m]\),
\[
\mathbb{E}\left[\left|\sum_{l=1}^{n} Z_{j, l} \cdot\left(y_{i}\right)_{l}\right| \mid \xi\right] \lesssim \mathbb{E}\left[\left|\sum_{l=1}^{n} S_{j, l} \cdot\left(y_{i}\right)_{l}\right| \mid \xi\right]+\delta
\]

As long as \(\delta\) is small enough, we are in a good shape.
Let \(X=\frac{1}{\left\|y_{i}\right\|_{1}} \sum_{l=1}^{n} Z_{j, l}\left(y_{i}\right)_{l}\), and \(Y=\frac{1}{\left\|y_{i}\right\|_{1}} \sum_{l=1}^{m} S_{j, l}\left(y_{i}\right)_{l}\). Let \(D\) be the truncating threshold of each Cauchy random variable. Define \(T=O(\log D)\). On one hand, we have
\[
\begin{align*}
\mathbb{E}[|X| \mid \xi] & \leq 2\left(\operatorname{Pr}[X \in[0,1] \mid \xi] \cdot 1+\sum_{j=0}^{T} \operatorname{Pr}\left[X \in\left(2^{j}, 2^{j+1}\right] \mid \xi\right] \cdot 2^{j+1}\right) \\
& \leq 2\left(1+\sum_{j=0}^{T} \operatorname{Pr}\left[I_{\left(2^{j}, 2^{j+1}\right)}(X)=1 \mid \xi\right] \cdot 2^{j+1}\right) \\
& =2\left(1+\sum_{j=0}^{T} \mathbb{E}\left[I_{\left(2^{j}, 2^{j+1}\right)}(X) \mid \xi\right] \cdot 2^{j+1}\right) \tag{17.46}
\end{align*}
\]

On the other hand, we can show
\[
\begin{align*}
\mathbb{E}[|Y| \mid \xi] & \geq 2\left(\operatorname{Pr}[X \in[0,1] \mid \xi] \cdot 0+\sum_{j=0}^{T} \operatorname{Pr}\left[Y \in\left(2^{j}, 2^{j+1}\right] \mid \xi\right] \cdot 2^{j}\right) \\
& \geq 2 \sum_{j=0}^{T} \operatorname{Pr}\left[I_{\left(2^{j}, 2^{j+1}\right]}(Y)=1 \mid \xi\right] \cdot 2^{j} \\
& \geq 2 \sum_{j=0}^{T} \mathbb{E}\left[I_{\left(2^{j}, 2^{j+1}\right]}(Y) \mid \xi\right] \cdot 2^{j} \tag{17.47}
\end{align*}
\]

Thus, we need to show that, for each \(j\),
\[
\begin{equation*}
\mathbb{E}\left[I_{\left(2^{j}, 2^{j+1}\right)}(X) \mid \xi\right] \approx_{\epsilon} \mathbb{E}\left[I_{\left(2^{j}, 2^{j+1}\right]}(Y) \mid \xi\right] \tag{17.48}
\end{equation*}
\]

Following the same idea from [KNW10a], we need to argue this chain of inequalities,
\[
\mathbb{E}\left[I_{[a, b]}(X)\right] \approx_{\epsilon} \mathbb{E}\left[J_{[a, b]}^{c}(X)\right] \approx_{\epsilon} \mathbb{E}\left[J_{[a, b]}^{c}(Y)\right] \approx_{\epsilon} \mathbb{E}\left[I_{[a, b]}(Y)\right]
\]

We first show \(\mathbb{E}\left[I_{[a, b]}(X)\right] \approx_{\epsilon} \mathbb{E}\left[J_{[a, b]}^{c}(X)\right]\). Notice that \(I_{[a, b]}\) and \(J_{[a, b]}\) are within \(\epsilon\) everywhere except for two intervals of length \(O(\epsilon)\). Also the Cauchy distribution is anticoncentrated (any length- \(O(\epsilon)\) interval contains \(O(\epsilon)\) probability mass) and \(\left\|I_{[a, b]}\right\|_{\infty},\left\|J_{[a, b]}^{c}\right\|_{\infty}=O(1)\), these intervals contribute \(O(\epsilon)\) to the difference.

Second, we show \(\mathbb{E}\left[J_{[a, b]}^{c}(X)\right] \approx_{\epsilon} \mathbb{E}\left[J_{[a, b]}^{c}(Y)\right]\). This directly follows by Lemma 17.11.1 by choosing \(\alpha=O\left(\epsilon^{-1} \log ^{3}(1 / \epsilon)\right)\).

Third, we show \(\mathbb{E}\left[J_{[a, b]}^{c}(Y)\right] \approx_{\epsilon} \mathbb{E}\left[I_{[a, b]}(Y)\right]\). The argument is similar as the first step, but we need to show anticoncentration of \(Y\). Suppose for any \(t \in \mathbb{R}\) we had a nonnegative
function \(f_{\epsilon, t}: \mathbb{R} \rightarrow \mathbb{R}\) symmetric about \(t\) satisfying:
\[
\begin{aligned}
& \text { I. }\left\|f_{t, \epsilon}^{(\ell)}\right\|_{\infty}=O\left(\alpha^{\ell}\right) \text { for all } \ell \geq 0 \text {, with } \alpha=O(1 / \epsilon) \\
& \text { II. } \mathbb{E}\left[f_{t, \epsilon}(z)\right]=O(\epsilon) \text { for } z \sim \mathcal{D}_{1} \\
& \text { III. } f_{t, \epsilon}(t+\epsilon)=\Omega(1) \\
& \text { IV. } f_{t, \epsilon}(x) \text { is strictly decreasing as }|x-t| \rightarrow \infty
\end{aligned}
\]

By I, II and Lemma 17.11.1 we could have \(\mathbb{E}\left[f_{\epsilon, t}(Y)\right] \approx_{\epsilon} \mathbb{E}\left[f_{t, \epsilon}(z)\right]=O(\epsilon)\). Then, \(\mathbb{E}\left[f_{t, \epsilon}(Y)\right] \geq\) \(f_{t, \epsilon}(t+\epsilon) \cdot \operatorname{Pr}[Y \in[t-\epsilon, t+\epsilon]]=\Omega(\operatorname{Pr}[Y \in[t-\epsilon, t+\epsilon]])\) by III and IV, implying anticoncentration in \([t-\epsilon, t+\epsilon]\) as desired. For the details of function \(f_{t, \epsilon}\), we refer the readers to Section A. 4 in [KNW10a].

Now, combining Equation (17.46), (17.47) and (17.48) gives
\[
\begin{aligned}
\mathbb{E}\left[\left|\sum_{l=1}^{n} Z_{j, l} \cdot\left(y_{i}\right)_{l}\right| \mid \xi\right] & \lesssim \mathbb{E}\left[\left|\sum_{l=1}^{n} S_{j, l} \cdot\left(y_{i}\right)_{l}\right| \mid \xi\right]+\sum_{j=0}^{T} 2^{j} \cdot \epsilon \cdot\left\|y_{i}\right\|_{1} \\
& \lesssim \mathbb{E}\left[\left|\sum_{l=1}^{n} S_{j, l} \cdot\left(y_{i}\right)_{l}\right| \mid \xi\right]+D \cdot \epsilon \cdot\left\|y_{i}\right\|_{1}
\end{aligned}
\]

Overall, for the fixed \(j, S_{j, l}\) is \(\widetilde{O}(1 / \epsilon)\)-independent family of Cauchy random variable. Choosing \(D=O(d k)\) and \(\epsilon=O(1 / D)\), we can show
\[
\frac{1}{m} \sum_{i=1}^{d} \sum_{j=1}^{m} \mathbb{E}\left[\left|\sum_{l=1}^{n} S_{j, l} \cdot\left(y_{i}\right)_{l}\right| \mid \xi\right] \leq O(\log d) \sum_{i=1}^{d}\left\|y_{i}\right\|_{1}
\]
as before. Notice that \(D \epsilon\left\|y_{i}\right\|_{1}=O\left(\left\|y_{i}\right\|_{1}\right)\). Thus, we complete the proof of the first result.
Choosing \(D=O(d k)\) and \(\epsilon=O(1 / d)\), the dominant term becomes \(D \epsilon\left\|y_{i}\right\|_{1}=\) \(O\left(k\left\|y_{i}\right\|_{1}\right)\). Thus, we complete the proof of second result.

Corollary 17.11.5. Given \(U \in \mathbb{R}^{n \times k}\), let \(Z \in \mathbb{R}^{t \times n}\) be the same as the matrix stated in the Lemma 17.11.3, then with probability at least .95,
\[
\forall x \in \mathbb{R}^{k},\|Z U x\|_{1} \gtrsim\|U x\|_{1} .
\]

The proof is very similar to the proof of Lemma 17.5.14. Without loss of generality, we can suppose \(U\) is a well-conditioned basis. Due to Lemma 17.11.4, with arbitrarily high constant probability \(\|Z U\|_{1}\) is bounded by poly \((t, k)\). By simply applying the net argument and using Lemma 17.11.3 to take a union bound over net points, we can get the above corollary.

\subsection*{17.12 Streaming Setting}

Section 17.12.1 provides some notation and definitions about row-update streaming model and the turnstile streaming model. For some recent developments of row-update streaming and turnstile streaming models, we refer the readers to [CW09, KL11, GP14, Lib13, KLM \({ }^{+}\)14a, BWZ16] and the references therein. Section 17.12.2 presents our turnstile streaming algorithm. Section 17.12.3 presents our row-update streaming algorithm.

\subsection*{17.12.1 Definitions}

Definition 17.12.1 (Row-update model). Let matrix \(A \in \mathbb{R}^{n \times d}\) be a set of rows \(A_{1}, \cdots, A_{n}\). In the row-update streaming model, each row of \(A\) will occur in the stream exactly once. But the rows can be in arbitrary order. An algorithm in this model is only allowed a single pass over these rows. At the end of the stream, the algorithm stores some information of \(A\). The space of the algorithm is the total number of words required to store this information during the stream. Here, each word is \(O(\log (n d))\) bits.

Definition 17.12.2 (Turnstile model). At the beginning, let matrix \(A \in \mathbb{R}^{n \times d}\) be a zero matrix. In the turnstile streaming model, there is a stream of update operations, and the \(i^{\text {th }}\) operation has the form \(\left(x_{i}, y_{i}, c_{i}\right)\) which means that \(A_{x_{i}, y_{i}}\) should be incremented by \(c_{i}\). An algorithm in this model is only allowed a single pass over the stream. At the end of the stream, the algorithm stores some information of \(A\). The space complexity of the algorithm is the total number of words required to store this information during the stream. Here, each word is \(O(\log (n d))\) bits.

\subsection*{17.12.2 Turnstile model, \(\operatorname{poly}(k, \log (d), \log (n))\) approximation}

Definition 17.12.3 (Turnstile model \(\ell_{1}\)-low rank approximation - rank- \(k\) subspace version). Given matrix \(A \in \mathbb{R}^{n \times d}\) and \(k \in \mathbb{N}_{+}\), the goal is to propose an algorithm in the streaming model of Definition 17.12.2 such that
1. Upon termination, the algorithm outputs a matrix \(V^{*} \in \mathbb{R}^{k \times d}\).
2. \(V^{*}\) satisfies that
\[
\min _{U \in \mathbb{R}^{n \times k}}\left\|A-U V^{*}\right\|_{1} \leq \operatorname{poly}(k, \log (d), \log (n)) \cdot \min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|A-U V\|_{1}
\]
3. The space complexity is as small as possible

Theorem 17.12.1. Suppose \(A \in \mathbb{R}^{n \times d}\) is given in the turnstile streaming model (See Definition 17.12.2), there is an algorithm(in Algorithm 17.8 without decomposition) which solves the problem in Definition 17.12.3 with constant probability. Further, the space complexity of the algorithm is \(\operatorname{poly}(k)+\widetilde{O}(k d)\) words.

Proof. Correctness. The correctness is implied by (IV) of Lemma 17.5.5, and the proof of Theorem 17.4.3. Notice that \(L=T_{1} A R, N=S A T_{2}, M=T_{1} A T_{2}\), so \(\widehat{X} \in \mathbb{R}^{O(k \log k) \times O(k \log k)}\) minimizes
\[
\min _{\text {rank }-k X}\left\|T_{1} A R X S A T_{2}-T_{1} A T_{2}\right\|_{F}
\]

According to the proof of Theorem 17.4.3, \(A R \widehat{X} S A\) gives a \(\ell_{1}\) rank- \(k\) poly \((k, \log (d), \log (n))\) approximation to \(A\). Because \(\widehat{X}=\widehat{U} \widehat{\Sigma} \widehat{V}^{\top}, V^{*}=\widehat{\Sigma} \widehat{V}^{\top} S A\) satisfies:
\[
\min _{U \in \mathbb{R}^{n \times k}}\left\|A-U V^{*}\right\|_{1} \leq \operatorname{poly}(k, \log (d), \log (n)) \cdot \min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|A-U V\|_{1}
\]

Space complexity. Generating \(\widetilde{O}(k d)\)-wise independent random Cauchy variables needs \(\widetilde{O}(k d)\) bits. The size of \(L, N\) and \(M\) are \(k^{2} \log ^{2} k, k^{2} \log ^{3} k\) and \(k^{2} \log ^{3} k\) words separately. So the space of maintaining them is \(O\left(k^{2} \log ^{3} k\right)\) words. The size of \(D\) is \(O(k \log k) \times d\), so maintaining it needs \(O(k d \log k)\) words. Therefore, the total space complexity of the algorithm is poly \((k)+\widetilde{O}(k d)\) words.

It is easy to extend our algorithm to output a decomposition. The formal definition of the decomposition problem is as the following:

Definition 17.12.4 (Turnstile model \(\ell_{1}\)-low rank approximation - rank- \(k\) decomposition version). Given matrix \(A \in \mathbb{R}^{n \times d}\) and \(k \in \mathbb{N}_{+}\), the goal is to propose an algorithm in the streaming model of Definition 17.12.2 such that
1. Upon termination, the algorithm outputs a matrix \(U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times d}\).
2. \(U^{*}, V^{*}\) satisfies
\[
\left\|A-U^{*} V^{*}\right\|_{1} \leq \operatorname{poly}(k, \log (d), \log (n)) \cdot \min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|A-U V\|_{1} .
\]
3. The space complexity is as small as possible

Theorem 17.12.2. Suppose \(A \in \mathbb{R}^{n \times d}\) is given by the turnstile streaming model (See Definition 17.12.2). There is an algorithm( in Algorithm 17.8 with decomposition) which solves the problem in Definition 17.12.4 with constant probability. Further, the space complexity of the algorithm is \(\operatorname{poly}(k)+\widetilde{O}(k(d+n))\) words.

Proof. Correctness. The only difference from the Algorithm 17.8 (without decomposition) is that the algorithm maintains \(C\). Thus, finally it can compute \(U^{*}=A R \widehat{U}\). Notice that \(U^{*} V^{*}=A R \widehat{X} S A\), according to the proof of Theorem 17.4.3, \(U^{*} V^{*}\) gives a \(\ell_{1}\) rank- \(k\) \(\operatorname{poly}(k, \log (d), \log (n))\)-approximation to \(A\).

Space complexity. Since the size of \(C\) is \(O(n k \log k)\) words, the total space is \(\operatorname{poly}(k)+\widetilde{O}(k(d+n))\) words.

\subsection*{17.12.3 Row-update model, poly \((k) \log d\) approximation}

Definition 17.12.5 (Row-update model \(\ell_{1}\)-low rank approximation - rank- \(k\) subspace version). Given matrix \(A \in \mathbb{R}^{n \times d}\) and \(k \in \mathbb{N}_{+}\), the goal is to propose an algorithm in the streaming model of Definition 17.12 .1 such that
1. Upon termination, the algorithm outputs a matrix \(V^{*} \in \mathbb{R}^{k \times d}\).
2. \(V^{*}\) satisfies that
\[
\min _{U \in \mathbb{R}^{n \times k}}\left\|A-U V^{*}\right\|_{1} \leq \operatorname{poly}(k) \log (d) \cdot \min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|A-U V\|_{1} .
\]
3. The space complexity is as small as possible

Theorem 17.12.3. Suppose \(A \in \mathbb{R}^{n \times d}\) is given by the row-update streaming model (See Definition 17.12.1), there is an algorithm (in Algorithm 17.9 without decomposition) which solves the problem in Definition 17.12.5 with constant probability. Further, the space complexity of the algorithm is poly \((k)+\widetilde{O}(k d)\) words.

Proof. Correctness. Notice that \(L=T_{1} B R, N=S B T_{2}, M=T_{1} B T_{2}\). Thus, \(\widehat{X} \in\) \(\mathbb{R}^{O(k \log k) \times O(k \log k)}\) actually minimizes
\[
\min _{\text {rank }-k X}\left\|T_{1} B R X S B T_{2}-T_{1} B T_{2}\right\|_{F}
\]

Also notice that \(B\) is just taking each row of \(A\) and replacing it with its nearest point in the row span of \(S^{\prime} A\). According to the proof of Theorem 17.4.6 and (IV) of Lemma 17.5.5 \(B R \widehat{X} S B\) gives a poly \((k) \log d \ell_{1}\) norm rank- \(k\) approximation to \(A\). Since \(\widehat{X}=\widehat{U} \widehat{\Sigma} \widehat{V}^{\top}\), \(V^{*}=\widehat{\Sigma} \widehat{V}^{\top} S B\) satisfies:
\[
\min _{U \in \mathbb{R}^{n \times k}}\left\|A-U V^{*}\right\|_{1} \leq \operatorname{poly}(k, \log (d), \log (n)) \cdot \min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|A-U V\|_{1}
\]

Space complexity. Constructing sketching matrices needs \(\widetilde{O}(k d)\) bits to store random seeds. Maintaining \(L, N, M\) needs \(O\left(k^{2} \log ^{3} k\right)\) words. The cost of storing \(\widehat{X}\) is also \(O\left(k^{2} \log ^{2} k\right)\) words. Maintaining \(D\) needs \(O(k d \log k)\) words. Therefore, the total space cost of the algorithm is \(\operatorname{poly}(k)+\widetilde{O}(k d)\) words.

It is easy to extend our algorithm to output a decomposition. The formal definition of the decomposition problem is as the following:

Definition 17.12.6 (Row-update model \(\ell_{1}\)-low rank approximation - rank- \(k\) decomposition version). Given matrix \(A \in \mathbb{R}^{n \times d}\) and \(k \in \mathbb{N}_{+}\), the goal is to propose an algorithm in the streaming model of Definition 17.12 .1 such that
1. Upon termination, the algorithm outputs matrices \(U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times d}\).
2. \(U^{*}, V^{*}\) satisfies that
\[
\left\|A-U^{*} V^{*}\right\|_{1} \leq \operatorname{poly}(k) \log d \cdot \min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|A-U V\|_{1} .
\]
3. The space complexity is as small as possible.

Theorem 17.12.4. Suppose \(A \in \mathbb{R}^{n \times d}\) is given by the row-update streaming model (See Definition 17.12.1), there is an algorithm(in Algorithm 17.9 with decomposition ) which solves the problem in Definition 17.12.6 with constant probability. Further, the space complexity of the algorithm is \(\operatorname{poly}(k)+\widetilde{O}(k(n+d))\) words.

Proof. Correctness. The only difference is that the above algorithm maintains \(C\). Thus, We can compute \(U^{*}=C \widehat{U}\) in the end. Notice that \(U^{*} V^{*}=B R \widehat{X} S B\), according to the proof of Theorem 17.4.6, \(U^{*} V^{*}\) gives a poly \((k) \log d \ell_{1}\) norm rank- \(k\) approximation to \(A\).

Space complexity. Since the size of \(C\) is \(n k \log k\) words, the total space is poly \((k)+\) \(\widetilde{O}(k(n+d))\) words.

\subsection*{17.13 Distributed Setting}

Section 17.13 .1 provides some notation and definitions for the Row-partition distributed model and the Arbitrary-partition model. These two models were recently studied in a line of works such as [TD99, QOSG02, BCL05, BRB08, MBZ10, FEGK13, PMvdG \({ }^{+}\)13, KVW14, BKLW14, BLS \({ }^{+}\)16b, BWZ16, WZ16]. Section 17.13 .2 and 17.13.3 presents our distributed protocols for the Arbitrary-partition distributed model. Section 17.13.4 and 17.13.5 presents our distributed protocols for the Row-partition distributed model.

\subsection*{17.13.1 Definitions}

Definition 17.13.1 (Row-partition model [BWZ16]). There are \(s\) machines, and the \(i^{\text {th }}\) machine has a matrix \(A_{i} \in \mathbb{R}^{n_{i} \times d}\) as input. Suppose \(n=\sum_{i=1}^{s} n_{i}\), and the global data \(\operatorname{matrix} A \in \mathbb{R}^{n \times d}\) is denoted as
\[
\left(\begin{array}{c}
A_{1} \\
A_{2} \\
\cdots \\
A_{s}
\end{array}\right)
\]
we say \(A\) is row-partitioned into these \(s\) matrices distributed in \(s\) machines respectively. Furthermore, there is a machine which is a coordinator. The model only allows communication between the machines and the coordinator. The communication cost in this model is the total number of words transferred between machines and the coordinator. Each word is \(O(\log (s n d))\) bits.

Definition 17.13.2 (Arbitrary-partition model [BWZ16]). There are \(s\) machines, and the \(i^{\text {th }}\) machine has a matrix \(A_{i} \in \mathbb{R}^{n \times d}\) as input. Suppose the global data matrix \(A \in \mathbb{R}^{n \times d}\) is denoted as \(A=\sum_{i=1}^{s} A_{i}\). We say \(A\) is arbitrarily partitioned into these \(s\) matrices distributed in \(s\) machines respectively. Furthermore, there is a machine which is a coordinator. The model only allows communication between the machines and the coordinator. The communication cost in this model is the total number of words transferred between machines and the coordinator. Each word is \(O(\log (s n d))\) bits.
17.13.2 Arbitrary-partition model, subspace, \(\operatorname{poly}(k, \log (d), \log (n))\) approxima-
tion

Definition 17.13.3 (Arbitrary-partition model \(\ell_{1}\)-low rank approximation - rank- \(k\) subspace version). Given matrix \(A \in \mathbb{R}^{n \times d}\) arbitrarily partitioned into \(s\) matrices \(A_{1}, A_{2}, \cdots, A_{s}\) distributed in \(s\) machines respectively, and \(k \in \mathbb{N}_{+}\), the goal is to propose a protocol in the model of Definition 17.13.2 such that
1. Upon termination, the protocol leaves a matrix \(V^{*} \in \mathbb{R}^{k \times d}\) on the coordinator.
2. \(V^{*}\) satisfies that
\[
\min _{U \in \mathbb{R}^{n \times k}}\left\|A-U V^{*}\right\|_{1} \leq \operatorname{poly}(k, \log (d), \log (n)) \cdot \min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|A-U V\|_{1} .
\]
3. The communication cost is as small as possible

Theorem 17.13.1. Suppose \(A \in \mathbb{R}^{n \times d}\) is partitioned in the arbitrary partition model (See Definition 17.13.2). There is a protocol(in Algorithm 17.10) which solves the problem in Definition 17.13.3 with constant probability. Further, the communication complexity of the protocol is \(s(\operatorname{poly}(k)+\widetilde{O}(k d))\) words.

Proof. Correctness. The correctness is shown by the proof of Theorem 17.4.3 and (IV) of Lemma 17.5.5. Notice that \(\widehat{X} \in \mathbb{R}^{O(k \log k) \times O(k \log k)}\) minimizes
\[
\min _{\operatorname{rank}-k X}\|L X N-M\|_{F} .
\]
which is
\[
\min _{\operatorname{rank}-k X}\left\|T_{1} A R X S A T_{2}-T_{1} A T_{2}\right\|_{F}
\]

According to the proof of Theorem 17.4.3, \(A R \widehat{X} S A\) gives an \(\ell_{1} \operatorname{rank}-k \operatorname{poly}(k, \log (d), \log (n))\) approximation to \(A\). Because \(\widehat{X}=\widehat{U} \widehat{\Sigma} \widehat{V}^{\top}, V^{*}=\widehat{\Sigma} \widehat{V}^{\top} S A\) satisfies:
\[
\min _{U \in \mathbb{R}^{n \times k}}\left\|A-U V^{*}\right\|_{1} \leq \operatorname{poly}(k, \log (d), \log (n)) \cdot \min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|A-U V\|_{1}
\]

Communication complexity. Since the random seed generates \(\widetilde{O}(k d)\)-wise independent random Cauchy variables, the cost of line 5 is \(\widetilde{O}(s k d)\) bits. The size of \(L_{i}, N_{i}\) and \(M_{i}\) are \(k^{2} \log ^{2} k, k^{2} \log ^{3} k\) and \(k^{2} \log ^{3} k\) words separately. So the cost of line 15 is \(O\left(s k^{2} \log ^{3} k\right)\) words. Because the size of \(\widehat{X}\) is \(O\left(k^{2} \log ^{2} k\right)\), the cost of line 22 is \(O\left(s k^{2} \log ^{2} k\right)\) words. line 30 needs \(s k d\) words of communication. Therefore, the total communication of the protocol is \(s(\operatorname{poly}(k)+\widetilde{O}(k d))\) words.

\subsection*{17.13.3 Arbitrary-partition model, decomposition, \(\operatorname{poly}(k, \log (d), \log (n))\) approximation}

Definition 17.13.4 (Arbitrary-partition model \(\ell_{1}\)-low rank approximation - rank- \(k\) decomposition version). Given matrix \(A \in \mathbb{R}^{n \times d}\) arbitrarily partitioned into \(s\) matrices \(A_{1}, A_{2}, \cdots, A_{s}\) distributed in \(s\) machines respectively, and \(k \in \mathbb{N}_{+}\), the goal is to propose a protocol in the model of Definition 17.13.2 such that
1. Upon termination, the protocol leave matrices \(U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times d}\) on the coordinator.
2. \(U^{*}, V^{*}\) satisfies that
\[
\left\|A-U^{*} V^{*}\right\|_{1} \leq \operatorname{poly}(k, \log (d), \log (n)) \cdot \min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|A-U V\|_{1} .
\]
3. The communication cost is as small as possible.

Theorem 17.13.2. Suppose \(A \in \mathbb{R}^{n \times d}\) is partitioned in the arbitrary partition model (See Definition 17.13.2). There is a protocol(in Algorithm 17.10 with decomposition) which solves the problem in Definition 17.13.4 with constant probability. Further, the communication complexity of the protocol is \(s(\operatorname{poly}(k)+\widetilde{O}(k(d+n)))\) words.

Proof. Correctness. The only difference from the protocol (without decomposition) in Section 17.13.2 is that the protocol sends \(U_{i}\). Thus, the coordinator can compute \(U^{*}=A R \widehat{U}\). Notice that \(U^{*} V^{*}=A R \widehat{X} S A\). According to the proof of Theorem 17.4.3, \(U^{*} V^{*}\) gives a \(\ell_{1}\) rank- \(k \operatorname{poly}(k, \log (d), \log (n))\)-approximation to \(A\).

Communication complexity. Since the size of \(U_{i}\) is \(k n\) words, the total communication is \(s(\operatorname{poly}(k)+\widetilde{O}(k(d+n)))\) words.

\subsection*{17.13.4 Row-partition model, subspace, \(\operatorname{poly}(k) \log d\) approximation}

Definition 17.13.5 (Row-partition model \(\ell_{1}\)-low rank approximation - rank- \(k\) subspace version). Given matrix \(A \in \mathbb{R}^{n \times d}\) row-partitioned into \(s\) matrices \(A_{1}, A_{2}, \cdots, A_{s}\) distributed in \(s\) machines respectively, and \(k \in \mathbb{N}_{+}\), the goal is to propose a protocol in the model of Definition 17.13 .1 such that
1. Upon termination, the protocol leaves a matrix \(V^{*} \in \mathbb{R}^{k \times d}\) on the coordinator.
2. \(V^{*}\) satisfies that
\[
\min _{U \in \mathbb{R}^{n \times k}}\left\|A-U V^{*}\right\|_{1} \leq \operatorname{poly}(k) \log d \cdot \min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|A-U V\|_{1} .
\]
3. The communication cost is as small as possible

Theorem 17.13.3. Suppose \(A \in \mathbb{R}^{n \times d}\) is partitioned in the row partition model (See Definition 17.13.1). There is a protocol(in Algorithm 17.11 without decomposition) which solves the problem in Definition 17.13 .5 with constant probability. Further, the communication complexity of the protocol is \(s(\operatorname{poly}(k)+\widetilde{O}(k d))\) words.

Proof. Correctness. For convenience, we denote matrices \(B \in \mathbb{R}^{n \times d}, S \in \mathbb{R}^{O\left(k \log ^{2} k\right) \times n}\) and \(T_{1} \in \mathbb{R}^{O\left(k \log ^{2} k\right) \times n}\) as
\[
B=\left(\begin{array}{c}
B_{1} \\
B_{2} \\
\cdots \\
B_{s}
\end{array}\right) \quad S=\left(\begin{array}{llll}
S_{1} & S_{2} & \cdots & S_{s}
\end{array}\right) \quad T_{1}=\left(\begin{array}{llll}
T_{11} & T_{12} & \cdots & T_{1 s}
\end{array}\right) .
\]

Notice that \(L=T_{1} B R, N=S B T_{2}, M=T_{1} B T_{2}\). Thus, \(\widehat{X} \in \mathbb{R}^{O(k \log k) \times O(k \log k)}\) actually minimizes
\[
\min _{\text {rank }-k X}\left\|T_{1} B R X S B T_{2}-T_{1} B T_{2}\right\|_{F} .
\]

Also notice that \(B\) is just taking each row of \(A\) and replacing it with its nearest point in the row span of \(S^{\prime} A\). According to the proof of Theorem 17.4.6, \(B R \widehat{X} S B\) gives a \(\operatorname{poly}(k) \log d\) \(\ell_{1}\) norm rank- \(k\) approximation to \(A\). Since \(\widehat{X}=\widehat{U} \widehat{\Sigma} \widehat{V}^{\top}, V^{*}=\widehat{\Sigma} \widehat{V}^{\top} S B\) satisfies:
\[
\min _{U \in \mathbb{R}^{n \times k}}\left\|A-U V^{*}\right\|_{1} \leq \operatorname{poly}(k, \log (d), \log (n)) \cdot \min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|A-U V\|_{1} .
\]

Communication complexity. Since the \(R\) and \(T_{2}\) are \(\widetilde{O}(k d)\)-wise independent, line 6 needs \(O(s W)\) bits of communication. Line 16 needs \(O\left(s k^{2} \log ^{3} k\right)\) words. The cost of line 23 is \(O\left(s k^{2} \log ^{2} k\right)\) words. Line 32 needs \(s k d\) words of communication. Therefore, the total communication of the protocol is \(s(\operatorname{poly}(k)+\widetilde{O}(k d))\) words.

\subsection*{17.13.5 Row-partition model, decomposition, \(\operatorname{poly}(k) \log d\) approximation}

Definition 17.13.6 (Row-partition model \(\ell_{1}\)-low rank approximation - rank-k decomposition version). Given matrix \(A \in \mathbb{R}^{n \times d}\) row partitioned into \(s\) matrices \(A_{1}, A_{2}, \cdots, A_{s}\) distributed in \(s\) machines respectively, and a positive integer \(k<\operatorname{rank}(A)\), the goal is to propose a protocol in the model of Definition 17.13 .1 such that
1. Upon termination, the protocol leaves matrices \(U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times d}\) on the coordinator.
2. \(U^{*}, V^{*}\) satisfies that
\[
\left\|A-U^{*} V^{*}\right\|_{1} \leq \operatorname{poly}(k) \log d \cdot \min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}}\|A-U V\|_{1} .
\]
3. The communication cost is as small as possible.

Theorem 17.13.4. Suppose \(A \in \mathbb{R}^{n \times d}\) is partitioned in the row partition model (See Definition 17.13.1). There is a protocol(in Algorithm 17.11 with decomposition) which solves the problem in Definition 17.13 .6 with constant probability. Further, the communication complexity of the protocol is \(s(\operatorname{poly}(k)+\widetilde{O}(k(n+d)))\) words.

Proof. Correctness. The only difference is that the above protocol sends \(U_{i}\). Thus, the coordinator can compute \(U^{*}=B R \widehat{U}\). Notice that \(U^{*} V^{*}=B R \widehat{X} S B\), according to the proof of Theorem 17.4.6, \(U^{*} V^{*}\) gives a poly \((k) \log d \ell_{1}\) norm rank- \(k\) approximation to \(A\).

Communication complexity. Since the size of \(U_{i}\) is \(k n\) words, the total communication is \(s(\operatorname{poly}(k)+\widetilde{O}(k(n+d)))\) words.
```

Algorithm 17.8 Turnstile Streaming Algorithm
procedure TurnstileStreaming $(k, \mathcal{S})$
Construct sketching matrices $S \in \mathbb{R}^{O(k \log k) \times n}, R \in \mathbb{R}^{d \times O(k \log k)}, T_{1} \in \mathbb{R}^{O(k \log k) \times n}, T_{2} \in$
$\mathbb{R}^{d \times O\left(k \log ^{2} k\right)}$ where $R, T_{2}$ are fully independent random Cauchy matrices, and $S, T_{1}$ are
random Cauchy matrices with fully independent variables across different rows and $\widetilde{O}(d)$ -
wise independent variables from the same row.
Initialize matrices:
$L \leftarrow\{0\}^{O(k \log k) \times O(k \log k)}, N \leftarrow\{0\}^{O(k \log k) \times O\left(k \log ^{2} k\right)}$.
$M \leftarrow\{0\}^{O(k \log k) \times O\left(k \log ^{2} k\right)}, D \leftarrow\{0\}^{O(k \log k) \times d}$.
if need decomposition then
$C \leftarrow\{0\}^{n \times O(k \log k)}$.
end if
for $i \in[l]$ do
Receive update operation $\left(x_{i}, y_{i}, c_{i}\right)$ from the data stream $\mathcal{S}$.
for $r=1 \rightarrow O(k \log k), s=1 \rightarrow O(k \log k)$ do
$L_{r, s} \leftarrow L_{r, s}+T_{1_{r, x_{i}}} \cdot c_{i} \cdot R_{y_{i}, s}$.
end for
for $r=1 \rightarrow O(k \log k), s=1 \rightarrow O\left(k \log ^{2} k\right)$ do
$N_{r, s} \leftarrow N_{r, s}+S_{r, x_{i}} \cdot c_{i} \cdot T_{2 y_{i}, s}$.
end for
for $r=1 \rightarrow O(k \log k), s=1 \rightarrow O\left(k \log ^{2} k\right)$ do
$M_{r, s} \leftarrow M_{r, s}+T_{1 r, x_{i}} \cdot c_{i} \cdot T_{2 y_{i}, s}$.
end for
for $r=1 \rightarrow O(k \log k)$ do
$D_{r, y_{i}} \leftarrow D_{r, s}+S_{r, x_{i}} \cdot c_{i}$.
end for
if need decomposition then
for $s=1 \rightarrow O(k \log k)$ do
$C_{x_{i}, s} \leftarrow C_{x_{i}, s}+c_{i} \cdot R_{y_{i}, s}$.
end for
end if
end for
Compute the SVD of $L=U_{L} \Sigma_{L} V_{L}^{\top}$.
Compute the SVD of $N=U_{N} \Sigma_{N} V_{N}^{\top}$.
Compute $\widehat{X}=L^{\dagger}\left(U_{L} U_{L}^{\top} M V_{N} V_{N}^{\top}\right)_{k} N^{\dagger}$.
Compute the SVD of $\widehat{X}=\widehat{U} \widehat{\Sigma} \widehat{V}^{\top}$.
if need decomposition then
return $V^{*}=\widehat{\Sigma} \widehat{V}^{\top} D, U^{*}=C \widehat{U}$.
else
return $V^{*}=\widehat{\Sigma} \widehat{V}^{\top} D$.
1327
end if
end procedure

```
```

Algorithm 17.9 Row Update Streaming Algorithm
procedure RowUpdateStreaming $(k, \mathcal{S})$
Construct sketching matrices $S^{\prime} \in \mathbb{R}^{O(k \log k) \times n}, S \in \mathbb{R}^{O(k \log k) \times n}, R \in \mathbb{R}^{d \times O(k \log k)}, T_{1} \in$
$\mathbb{R}^{O(k \log k) \times n}, T_{2} \in \mathbb{R}^{d \times O\left(k \log ^{2} k\right)}$ where $R, T_{2}$ are fully independent random Cauchy vari-
ables, and $S, S^{\prime}, T_{1}$ are random Cauchy matrices with fully independent random vari-
ables from different rows and $\widetilde{O}(d)$-wise independent in the same row.
Initialize matrices:
$L \leftarrow\{0\}^{O(k \log k) \times O(k \log k)}, N \leftarrow\{0\}^{O(k \log k) \times O\left(k \log ^{2} k\right)}$.
$M \leftarrow\{0\}^{O(k \log k) \times O\left(k \log ^{2} k\right)}, D \leftarrow\{0\}^{O(k \log k) \times d}$.
if need decomposition then
$C \leftarrow\{0\}^{n \times O(k \log k)}$.
end if
for $i \in[n]$ do
Receive a row update $\left(i, A_{i}\right)$ from the data stream $\mathcal{S}$.
Compute $Y_{i}^{*} \in \mathbb{R}^{1 \times O(k \log k)}$ which minimizes $\min _{Y \in \mathbb{R}^{1 \times O(k \log k)}}\left\|Y S_{:, i}^{\prime} A_{i}-A_{i}\right\|_{1}$.
Compute $B_{i}=Y_{i}^{*} S_{:, i}^{\prime} A_{i}$.
for $r=1 \rightarrow O(k \log k), s=1 \rightarrow O(k \log k), j=1 \rightarrow d$ do
$L_{r, s} \leftarrow L_{r, s}+T_{1 r, i} \cdot B_{i, j} \cdot R_{j, s}$.
end for
for $r=1 \rightarrow O(k \log k), s=1 \rightarrow O\left(k \log ^{2} k\right), j=1 \rightarrow d$ do
$N_{r, s} \leftarrow N_{r, s}+S_{r, i} \cdot B_{i, j} \cdot T_{2 j, s}$.
end for
for $r=1 \rightarrow O(k \log k), s=1 \rightarrow O\left(k \log ^{2} k\right), j=1 \rightarrow d$ do
$M_{r, s} \leftarrow M_{r, s}+T_{1 r, i} \cdot B_{i, j} \cdot T_{2 j, s}$.
end for
for $r=1 \rightarrow O(k \log k), j=1 \rightarrow d$ do
$D_{r, j} \leftarrow D_{r, j}+S_{r, i} \cdot B_{i, j}$.
end for
if need decomposition then
for $s=1 \rightarrow O(k \log k), j=1 \rightarrow d$ do
$C_{i, s}:=C_{i, s}+B_{i, j} \cdot R_{j, s}$.
end for
end if
end for
Compute the SVD of $L=U_{L} \Sigma_{L} V_{L}^{\top}$.
Compute the SVD of $N=U_{N} \Sigma_{N} V_{N}^{\top}$.
Compute $\widehat{X}=L^{\dagger}\left(U_{L} U_{L}^{\top} M V_{N} V_{N}^{\top}\right)_{k} N^{\dagger}$.
Compute the SVD of $\widehat{X}=\widehat{U} \widehat{\Sigma} \widehat{V}^{\top}$.
if need decomposition then
return $V^{*}=\widehat{\Sigma} \widehat{V}^{\top} D, U^{*}=C \widehat{U} .1328$
else
return $V^{*}=\widehat{\Sigma} \widehat{V}^{\top} D$.
end if
end procedure

```
```

Algorithm 17.10 Arbitrary Partition Distributed Protocol
procedure ArbitraryPartitionDistributedProtocol $(k, s, A)$
$A \in \mathbb{R}^{n \times d}$ was arbitrarily partitioned into $s$ matrices $A_{1}, \cdots, A_{s} \in \mathbb{R}^{n \times d}$ distributed
in $s$ machines.

```

4: Chooses a random seed.
5: \(\quad\) Sends it to all machines.
```

6:

```
6:
7:
8 :
9:
    Cauchy
10:
11:
    indepen-
12:
13:
14:
\(15:\)
17: \(\quad\) Computes \(L=\sum_{i=1}^{s} L_{i}, N=\sum_{i=1}^{s} N_{i}\).
18: \(\quad\) Computes \(M=\sum_{i=1}^{s} M_{i}\).
            Computes the SVD of \(L=U_{L} \Sigma_{L} V_{L}^{\top}\).
            Computes the SVD of \(N=U_{N} \Sigma_{N} V_{N}^{\top}\).
            Computes \(\widehat{X}=L^{\dagger}\left(U_{L} U_{L}^{\top} M V_{N} V_{N}^{\top}\right)_{k} N^{\dagger}\).
            Sends \(\widehat{X}\) to machines.
                                    \(-------->\)
                                    Computes the SVD of \(\widehat{X}=\widehat{U} \widehat{\Sigma} \widehat{V}^{\top}\).
                    Computes \(V_{i}^{*}=\widehat{\Sigma} \widehat{V}^{\top} S A_{i}\).
                        If need decomposition
                                    \(U_{i}^{*}=A_{i} \widehat{R U}\).
                                    Sends \(U_{i}^{*}, V_{i}^{*}\) to the coordinator.
                                    Else
                                    Sends \(V_{i}^{*}\) to the coordinator.
                                    Endif
                                    <------- - -
            If need decomposition,
            return \(V^{*}=\sum_{i=1}^{s} V_{i}^{*}, U^{*}=\sum 329 U_{i=1}^{*}\).
            Else
            return \(V^{*}=\sum_{i=1}^{s} V_{i}^{*}\).
            Endif
    end procedure
```

```
Algorithm 17.11 Row Partition Distributed Protocol
    procedure RowPartitionDistributedProtocol \((k, s, A)\)
        \(A \in \mathbb{R}^{n \times d}\) was row partitioned into \(s\) matrices \(A_{1} \in \mathbb{R}^{n_{1} \times d}, \cdots, A_{s} \in \mathbb{R}^{n_{s} \times d}\) distributed
    in \(s\) machines.
                    Coordinator
Machines \(i\)
Chooses a random seed.
Sends it to all machines.
```



```
Agrees on \(R, T_{2}\) which are fully
independent random Cauchy variables.
Generates random Cauchy matrices
\(S_{i}^{\prime} \in \mathbb{R}^{O(k \log k) \times n_{i}}, \quad S_{i}, T_{1 i} \in\)
\(\mathbb{R}^{O\left(k \log ^{2} k\right) \times n_{i}}\).
11:
\(Y_{i}^{*}=\underset{Y \in \mathbb{R}^{n_{i} \times O(k \log k)}}{\arg \min }\left\|Y S_{i}^{\prime} A_{i}-A_{i}\right\|_{1}\).
```

12:
13:
14:
15:
16:

22: $\quad$ Sends $\widehat{X}$ to machines.

## Chapter 18

## L1 Low Rank Approximation with Regularization

Low rank matrix approximation with respect to the entrywise $\ell_{1}$ norm error is a fundamental problem in both theoretical computer science and machine learning community. At the same time, $\ell_{1}$ regularization (a.k.a. the LASSO) has been used for many statistical learning techniques on finding solutions with good performance and interpretability. In this paper, we consider how to combine the two techniques and study the $\ell_{1}$ regularized version of $\ell_{1}$-low rank approximation problem. In particular, given matrix $A \in \mathbb{R}^{n \times n}$, parameter $k \geq 1, \lambda>0$, the problem we interested in is

$$
\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times n}}\|U V-A\|_{1}+\lambda\left(\|U\|_{1}+\|V\|_{1}\right)
$$

where $\|B\|_{1}:=\sum_{i=1}^{n} \sum_{j=1}^{n}\left|B_{i, j}\right|$ for any matrix $B$.
We introduce the first algorithm for solving this problem with probable guarantee. Precisely, we compute two matrices $U \in \mathbb{R}^{n \times k}$ and $V \in \mathbb{R}^{k \times n}$ such that

$$
\|U V-A\|_{1}+\lambda\left(\|U\|_{1}+\|V\|_{1}\right) \leq \alpha \mathrm{OPT}
$$

holds with probability at least $9 / 10$, where OPT is the cost of the optimal solution, and $\alpha$ is approximation ratio.

In addition, our algorithm shows a trade-off between running time and approximation ratio:

- $O(\operatorname{nnz}(A))+n \operatorname{poly}(\log n, k)$ time with $\alpha=\operatorname{poly}(k, \log n)$.
- $\widetilde{O}(k) \operatorname{nnz}(A)+n$ poly $(\log n, k)$ time with $\alpha=\widetilde{O}\left(k^{2}\right) \log ^{3} n$.
where $n n z(A)$ denotes the number of non-zero entries in $A$.


### 18.1 Introduction

The low-rank matrix approximation is a core problem in machine learning and data mining. It approximating a matrix by one whose rank is less than that of the original matrix. The goal of low-rank approximation is to obtain more condensed representations of the data with limited loss of information. Given a matrix $A \in \mathbb{R}^{m \times n}$ and a desired rank $k$, the low-rank matrix approximation problem can be stated as

$$
\begin{equation*}
\min _{\widehat{A} \in \mathbb{R}^{m \times n}}\|\widehat{A}-A\| \quad \text { subject to } \quad \operatorname{rank}(\widehat{A}) \leq k \tag{18.1}
\end{equation*}
$$

where $\|\cdot\|$ is a matrix norm measuring the approximation error.
Generally speaking, a rank- $k$ matrix $\widehat{A} \in \mathbb{R}^{n \times n}$ can be expressed by the matrix product of two matrices $U, V$, where $U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times n}$. The above equation can be rewritten as

$$
\begin{equation*}
\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times n}}\|U V-A\| . \tag{18.2}
\end{equation*}
$$

The key point in low-rank approximation problem is how to choose the matrix norm $\|\cdot\|$. The most widely used norm is the Frobenius norm, which is defined as $\|B\|_{F}=$ $\left(\sum_{i=1}^{n} \sum_{j=1}^{n} B_{i, j}^{2}\right)^{\frac{1}{2}}$ for matrix $B \in \mathbb{R}^{n \times n}$, The analytical optimal solution for Frobenius norm can be obtained by using the singular value decomposition (SVD) of $B$.

The Frobenius norm fits well with Gaussian noise, but it is sensitive to sparse outliers. Thus, an interesting extension for solving (18.2) is when the noise is sparse. This problem can be written as

$$
\begin{equation*}
\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times n}}\|U V-A\|_{0} \tag{18.3}
\end{equation*}
$$

where $\|\cdot\|_{0}$ is $\ell_{0}$ norm which measures the sparsity of the matrix.

Nevertheless, solving $\ell_{0}$ low-rank approximation has been proven to be NP-hard [GG11]. In practice, instead of solving (18.3), a more feasible way is to solve

$$
\begin{equation*}
\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times n}}\|U V-A\|_{1}, \tag{18.4}
\end{equation*}
$$

where we replace the $\ell_{0}$ by $\ell_{1}$, and $\|B\|_{1}:=\sum_{i=1}^{n} \sum_{j=1}^{n}\left|B_{i, j}\right|$ for matrix $B \in \mathbb{R}^{n \times n}$. This problem sometimes is referred to as robust PCA [CLMW11]. Finding solution to (18.4) has been widely studied by machine learning and theoretical computer science community for a long time. It has been validated both empirically and theoretically that $\ell_{1}$-norm is usually more robust than $\ell_{2}$-norm and tends to preserve the sparsity of the original matrix [CLMW11, XCS10, WY13, XCS11]. Also, $\ell_{1}$ norm has been extensively studied in many other problems such as sparse recovery [PW11], Fourier transform $\left[\mathrm{BCG}^{+} 12\right]$, furthest pair [Yao82, GBT84], closest pair [PS85, DHKP97] and so on.

Although solving (18.4) has been proven to be NP-hard [GV15], there are a lot of efforts for the $\ell_{1}$ low-rank approximation problem in a heuristic way. [KK05] introduced a method based on alternating convex minimization. [WYWY12] devise a parallelizable expectation-maximization (EM) algorithm which can potentially be applied to large-scale applications. [BDB13] exploit the efficient calculation of the optimal solution of the $\ell_{1}$-norm best-fit hyperplane problem. Unfortunately, those heuristic method do not come with any performance guarantees. Recently, [SWZ17, CGK $\left.{ }^{+} 17 \mathrm{~b}\right]$ introduce efficient algorithms for $\ell_{1}$-low rank approximation with provable approximation guarantees.
$\ell_{1}$ regularization can be very effective for many statistical learning techniques [LZWW16, Tib96]. However, none of the previous $\ell_{1}$ low rank approximation algorithms can handle the regularized version of $\ell_{1}$ low rank approximation problem. Though [LZWW16] studies $\ell_{1}$
principal component analysis problem with $\ell_{1}$ regularization, their algorithm does not have any provable guarantee. [ACW17] studied regularized low rank approximation problem, but they can only deal with the regularized Frobenious norm low rank approximation, i.e. $\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times n}}\|U V-A\|_{F}^{2}+\lambda\left(\|U\|_{F}^{2}+\|V\|_{F}^{2}\right)$, and their techniques cannot be extended to regularized $\ell_{1}$ low rank approximation case. It is natural to ask whether there is an efficient algorithm for solving $\ell_{1}$-regularized $\ell_{1}$ low-rank approximation with a good approximation ratio.

In this work, we study the $\ell_{1}$ regularized version of $\ell_{1}$ low rank approximation problem. The problem can be formulated as

$$
\begin{equation*}
\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times n}}\|U V-A\|_{1}+\lambda\left(\|U\|_{1}+\|V\|_{1}\right) \tag{18.5}
\end{equation*}
$$

where $\lambda$ is a non-negative tunable parameter which controls the weight between approximation error and the sparsity of $U, V$. When $\lambda=0$, (18.5) degenerates to (18.4). By solving (18.5), our model provides the flexibility that can find the balance between the reconstruction error and the sparsity of each low-rank component. Notably, when $\lambda=0$, our results exactly match the results shown in [SWZ17].

### 18.1.1 Our results

We give the first efficient algorithms for $\ell_{1}$ low-rank approximation with $\ell_{1}$ regularization problem with both provable running time and approximation ratio guarantees. A crucial tool in our algorithms is sketching, i.e. compressing the large input data into a small size summary which is enough to provide a good solution to the original input. In particular, we use the Cauchy transformation as our sketching tool. The Cauchy transformation is widely used in $\ell_{1}$ numerical linear algebra problems such as subspace embedding
and linear regression (see e.g. [SW11, MM13]). However, people did not know how to use the Cauchy transformation technique to solve low-rank matrix approximation problem until recent work [SWZ17]. In our work, we show how to extend this technique to handle $\ell_{1}$ regularization. But the analysis of the algorithms is very different from the previous works [SWZ17, SW11, MM13]. Algorithm 18.1 shows the framework of our algorithms.

There are different kind of the Cauchy transformations given by previous works [SW11, MM13, CDMI ${ }^{+}$13]. They have different trade-offs between distortion ratio and the running time of applying the transformation. By using different Cauchy transformation, we show algorithms with different running time/approximation ratio tradeoffs.

By using sparse Cauchy transform [MM13] (see Section 18.2.2), we can get the following result,

Theorem 18.1.1 (Input sparsity time algorithm). Given matrix $A \in \mathbb{R}^{n \times n}$, parameters $k \geq 1, \lambda>0$, let OPT denote $\min _{U^{\prime} \in \mathbb{R}^{n \times k}, V^{\prime} \in \mathbb{R}^{k \times n}}\left\|U^{\prime} V^{\prime}-A\right\|_{1}+\lambda\left(\left\|U^{\prime}\right\|_{1}+\left\|V^{\prime}\right\|_{1}\right)$, there is an algorithm that runs in $O(\mathrm{nnz}(A)+n \operatorname{poly}(k, \log n))$ time to output two matrices $U \in \mathbb{R}^{n \times k}$, $V \in \mathbb{R}^{k \times n}$ such that,

$$
\|U V-A\|_{1}+\lambda\left(\|U\|_{1}+\|V\|_{1}\right) \leq \operatorname{poly}(k, \log n) \cdot \mathrm{OPT}
$$

holds with probability at least $9 / 10$, where $\mathrm{nnz}(A)$ is the number of nonzeros in matrix $A$.

By using dense Cauchy transform [SW11] (see Section 18.2.2), we can get a different trade-off between approximation ration and running time. Theorem 18.1.1 provides a nnz $(A)$ running time with $\operatorname{poly}(k, \log n)$-approximation. Here we provide a better approximation ratio under polynomial running time.

Theorem 18.1.2 $\left(\widetilde{O}\left(k^{2}\right)\right.$ poly $(\log n)$-approximation). Given matrix $A \in \mathbb{R}^{n \times n}$, parameters $k \geq 1, \lambda>0$, let OPT denote $\min _{U^{\prime} \in \mathbb{R}^{n \times k}, V^{\prime} \in \mathbb{R}^{k \times n}}\left\|U^{\prime} V^{\prime}-A\right\|_{1}+\lambda\left(\left\|U^{\prime}\right\|_{1}+\left\|V^{\prime}\right\|_{1}\right)$, there is an algorithm that runs in $\widetilde{O}(k) \cdot \mathrm{nnz}(A)+n \cdot \operatorname{poly}(k, \log n)$ time to output two matrices $U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times n}$ such that,

$$
\|U V-A\|_{1}+\lambda\left(\|U\|_{1}+\|V\|_{1}\right) \leq \widetilde{O}\left(k^{2}\right) \log ^{3} n \cdot \mathrm{OPT}
$$

holds with probability at least 9/10.

Notice that when $\lambda=0$, all of our results exactly match the $\ell_{1}$ low-rank approximation results shown by [SWZ17]. Since [SWZ17] also showed $k^{\Omega(1)}$ factor is necessary in the approximation ratio, our results are almost tight.

```
Algorithm 18.1 A Meta Algorithm
    procedure L1REguLARIzEDL1LowRAnkApprox \((A, n, k, \lambda) \quad \triangleright\) Theorem 18.1.1 and
    18.1.2
        Choose sketching matrices \(S \in \mathbb{R}^{s \times n}, R \in \mathbb{R}^{n \times r}, T_{1} \in \mathbb{R}^{t_{1} \times n}, T_{2} \in \mathbb{R}^{n \times t_{2}}\) to be random
    Cauchy transformation matrices
        Compute \(S A\) and \(A R\)
        Compute \(S A T_{2}, T_{1} A R\)
        Compute \(T_{1} A T_{2}\)
        Solve \(\min _{X, Y}\left\|T_{1} A R X Y S A T_{2}-T_{1} A T_{2}\right\|_{F}^{2}+\lambda^{2}\left(\left\|T_{1} A R X\right\|_{F}^{2}+\left\|Y S A T_{2}\right\|_{F}^{2}\right)\)
        \(U \leftarrow A R X, V \leftarrow Y S A\)
        return \(U, V\)
    end procedure
```


### 18.1.2 Technique overview

Sketching is used widely in many important numerical linear algebra problems such as overconstrained regression [CW13, DMMS11, MM13, NN13a, Sar06], low rank matrix
approximation [AM07, CW13, DV06, DKM06, MM13, NN13a, Sar06, CW15b, SWZ17, BWZ16] computation of leverage scores [CW13, DMIMW12, MM13, NN13a], regularized data fitting [ACW17]. Roughly speaking, the main idea of sketching is to highly compress the original data into a small size summary, and then to recover properties of the original data from the such small summary.
[SWZ17] proposed the first sketching based algorithm for solving $\ell_{1}$ norm low rank matrix approximation problem. In that paper, they focused on the problem that given $A \in \mathbb{R}^{n \times n}, k>0$, the goal is to solve

$$
\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times n}}\|U V-A\|_{1} .
$$

This problem is $\ell_{1}$ norm low rank approximation problem, and it is a special case of our problem, i.e. when $\lambda=0$. The main idea of [SWZ17] is as the following: for a given matrix $A \in \mathbb{R}^{n \times n}$, and a rank parameter $k$, they firstly chose two small random sketching matrices $S \in \mathbb{R}^{s \times n}$ and $R \in \mathbb{R}^{n \times r}$ and argued that there is a rank- $k$ matrix $B$ in the both small column span of $A R$ and small row span of $S A$, and $\|B-A\|_{1}$ has a very good approximation ratio to the optimal cost, i.e. $\operatorname{poly}(k, \log n)$. Notice that $B$ is in the both column space of $A R$ and the row space of $S A$, it means that

$$
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\|A R X \cdot Y S A-A\|_{1}
$$

is a good approximation of $\min _{\text {rank }-k U V}\|U V-A\|_{1}$, and since the dimension of the column space of $A R$ and the row space of $S A$ are very small, then they further sketched down the problem by using oblivious $\ell_{1}$ subspace embedding [MM13, SW11] or Lewis weights sampling [CP15]. Precisely, they chose additional two small random sketching matrices
$T_{1} \in \mathbb{R}^{t_{1} \times n}, T_{2} \in \mathbb{R}^{n \times t_{2}},\left(t_{1}, t_{2} \leq \operatorname{poly}(k)\right)$ and further sketched down the optimization problem to get a new minimization problem i.e.,

$$
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\left\|T_{1} A R \cdot X \cdot Y \cdot S A T_{2}-T_{1} A T_{2}\right\|_{1}
$$

Since the new optimization problem is very small, then it can be solved by relaxing to $\ell_{2}$ minimization problem with losing mild additional poly $(k)$ factors in approximation ratio or solved by using exponential running time optimizer with losing mild additional $2^{\text {poly }(k)}$ running time, e.g. polynomial system verifier [Ren92a, Ren92b, BPR96, RSW16, SWZ17]. Now come back to our regularized $\ell_{1}$ norm low rank approximation problem. We can find that we cannot extend the technique of [SWZ17] directly. Because our objective function is

$$
\min _{\text {rank }-k U V}\|U V-A\|_{1}+\lambda\left(\|U\|_{1}+\|V\|_{1}\right)
$$

if $\lambda$ is large, then the cost induced by $\|U\|_{1}+\|V\|_{1}$ may dominate the cost induced by $\|U V-A\|_{1}$. Thus, the main issue we are facing is that it is not clear whether there is still a $\widehat{U}$ in the column space of $A R$ such that

$$
\min _{\operatorname{rank}-k V}\|\widehat{U} V-A\|_{1}+\lambda\left(\|\widehat{U}\|_{1}+\|V\|_{1}\right)
$$

has a good approximation ratio to the optimal cost (same issue appears for $V$, i.e. we are not sure whether there is a good $\widehat{V}$ in the row space of $S A$.) If we want to sketch $\|U V-A\|_{1},\|U\|_{1}$ and $\|V\|_{1}$ separately, we can not get a poly $(k)$ size sketch, since the freedom of $U$ and $V$ itself is too large, i.e. we cannot restrict them on any small subspace. Thus, it motivated us to look at the cost induced by $\|U V-A\|_{1},\|U\|_{1}$ and $\|V\|_{1}$ at the same time. Suppose $U^{*}$ is the optimal solution of

$$
\min _{\text {rank }-k U V}\|U V-A\|_{1}+\lambda\left(\|U\|_{1}+\|V\|_{1}\right)
$$

Then we have $\min _{\text {rank }-k V}\left\|U^{*} V-A\right\|_{1}+\lambda\left(\left\|U^{*}\right\|_{1}+\|V\|_{1}\right)=$ OPT. Notice that it is actually an $\ell_{1}$ regression problem:

$$
\min _{\operatorname{rank}-k}\left\|\left[\begin{array}{c}
U^{*} \\
\lambda I
\end{array}\right] V-\left[\begin{array}{c}
A \\
0
\end{array}\right]\right\|_{1}+\lambda\left\|U^{*}\right\|_{1}
$$

We choose a small random sketching matrix $S$, and we proved that a good approximation for

$$
\min _{\operatorname{rank}-k V}\left\|\left[\begin{array}{c}
S U^{*}  \tag{18.6}\\
\lambda I
\end{array}\right] V-\left[\begin{array}{c}
S A \\
0
\end{array}\right]\right\|_{1}
$$

is a good approximation for

$$
\min _{\operatorname{rank}-k V}\left\|\left[\begin{array}{c}
U^{*} \\
\lambda I
\end{array}\right] V-\left[\begin{array}{c}
A \\
0
\end{array}\right]\right\|_{1} .
$$

Notice that since $S$ has only poly $(k)$ number of rows, the number of rows of

$$
\left[\begin{array}{c}
S U^{*} \\
\lambda I
\end{array}\right] V-\left[\begin{array}{c}
S A \\
0
\end{array}\right]
$$

is at most $\operatorname{poly}(k)$. Thus, we can relax optimization problem (18.6) to

$$
\min _{\operatorname{rank}-k} \sum_{i=1}^{n}\left\|\left[\begin{array}{c}
S U^{*} \\
\lambda I
\end{array}\right] V_{i}-\left[\begin{array}{c}
S A \\
0
\end{array}\right]_{i}\right\|_{2}
$$

with at most $\operatorname{poly}(k)$ factor loss in approximation ratio. For the above optimization problem, we can get the optimal solution which has form

$$
\widehat{V}=\left[\begin{array}{c}
S U^{*} \\
\lambda I
\end{array}\right]^{\dagger} \cdot\left[\begin{array}{c}
S A \\
0
\end{array}\right]=\left[\begin{array}{c}
S U^{*} \\
\lambda I
\end{array}\right]^{\dagger} \cdot\left[\begin{array}{c}
I \\
0
\end{array}\right] \cdot S A .
$$

Thus, we can prove that

$$
\min _{\text {rank }-k U}\|U \widehat{V}-A\|_{1}+\lambda\left(\|U\|_{1}+\|\widehat{V}\|_{1}\right) .
$$

is a good approximation to OPT. Then, we choose a small random sketching matrix $R$, we can apply the "iterative existential proof" [SWZ19b] argument to show that there is a matrix $\widehat{U}$ in the column space of $A R$ such that

$$
\|\widehat{U} \widehat{V}-A\|_{1}+\lambda\left(\|\widehat{U}\|_{1}+\|\widehat{V}\|_{1}\right)
$$

is also a good approximation to OPT. Since $\widehat{U}$ is in the column space of $A R$ and $\widehat{V}$ is in the row space of $S A$, we can get that

$$
\min _{\text {rank }-k X Y}\|A R X Y S A-A\|_{1}+\lambda\left(\|A R X\|_{1}+\|Y S A\|_{1}\right)
$$

is a good approximation to OPT. Because we shrink the solution space to a very small subspace, then we can use the oblivious subspace embedding [MM13, SW11] or Lewis weights sampling [CP15] to further choose two random sketching $T_{1}, T_{2}$ to sketch down the optimization problem to get

$$
\min _{\text {rank }-k X Y}\left\|T_{1} A R X Y S A T_{2}-T_{1} A T_{2}\right\|_{1}+\lambda\left(\left\|T_{1} A R X\right\|_{1}+\left\|Y S A T_{2}\right\|_{1}\right)
$$

It is equivalent to solve

$$
\min _{\operatorname{rank}-k X Y}\left\|\left[\begin{array}{c}
T_{1} A R X \\
\lambda I
\end{array}\right]\left[\begin{array}{ll}
Y S A T_{2} & \lambda I
\end{array}\right]-\left[\begin{array}{cc}
T_{1} A T_{2} & 0 \\
0 & \lambda^{2} I
\end{array}\right]\right\|_{1} .
$$

To solve the above small problem, we can relax it into $\ell_{2}$ version by losing poly $(k)$ approximation factor or use some powerful optimizers by losing $2^{\text {poly(k) }}$ running time. If we relax it into $\ell_{2}$ version, then we need to solve

$$
\min _{\operatorname{rank}-k X Y}\left\|\left[\begin{array}{c}
T_{1} A R X \\
\lambda I
\end{array}\right]\left[\begin{array}{ll}
Y S A T_{2} & \lambda I
\end{array}\right]-\left[\begin{array}{cc}
T_{1} A T_{2} & 0 \\
0 & \lambda^{2} I
\end{array}\right]\right\|_{F}^{2}
$$

which is

$$
\min _{\operatorname{rank}-k X Y}\left\|T_{1} A R X Y S A T_{2}-T_{1} A T_{2}\right\|_{F}^{2}+\lambda^{2}\left(\left\|T_{1} A R X\right\|_{F}^{2}+\left\|Y S A T_{2}\right\|_{F}^{2}\right) .
$$

It is exactly a generalized version of the Frobenius norm ridge low rank matrix approximation problem (see Section 18.2.1 for more details), and we can directly get the closed form of the optimal solution of this problem [ACW17, UHZ $\left.{ }^{+} 16\right]$.

Surprisingly, we further showed that all the sketching matrices $T_{1}, S, R, T_{2}$ used in [SWZ17] satisfy the property we needed our problem. It means that we can actually fully generalize almost all the algorithms ${ }^{1}$ in [SWZ17] to the $\ell_{1}$ regularizer version.

[^64]
### 18.2 Preliminaries

In this Section, we mainly introduce the notations and some background knowledge.
For an $n \in \mathbb{N}_{\geq 0}$, let $[n]$ denote the set $\{1,2, \cdots, n\}$.
For any function $f$, we define $\widetilde{O}(f)$ to be $f \cdot \log ^{O(1)}(f)$. In addition to $O(\cdot)$ notation, for two functions $f, g$, we use the shorthand $f \lesssim g$ (resp. $\gtrsim$ ) to indicate that $f \leq C g$ (resp. $\geq)$ for an absolute constant $C$. We use $f \approx g$ to mean $c f \leq g \leq C f$ for constants $c, C$.

Given a matrix $A \in \mathbb{R}^{n \times d}$, we say the set $\left\{y \mid x \in \mathbb{R}^{d}, y=A x\right\}$ is the column space of $A$. Similarly, we say the set $\left\{y \mid x \in \mathbb{R}^{n}, y=x^{\top} A\right\}$ is the row space of $A$.

Let $n n z(A)$ denote the number of nonzero entries of $A$. Let $\operatorname{det}(A)$ denote the determinant of a square matrix $A$. Let $A^{\top}$ denote the transpose of $A$. Let $A^{\dagger}$ denote the Moore-Penrose pseudoinverse of $A$. Let $A^{-1}$ denote the inverse of a full rank square matrix.

Let $\|A\|_{F}$ denote the Frobenius norm of matrix $A$ and $\|A\|_{1}$ denote the entry-wise 1-norm of matrix $A$.

### 18.2.1 Frobenius norm ridge low-rank approximation and $\ell_{2}$ relaxation

In this Section, we introduce a highly related problem called Frobenius norm ridge low rank approximation problem. It is just an $\ell_{2}$ variation of the problem we studied in this work:

$$
\begin{equation*}
\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times n}}\|U V-A\|_{F}^{2}+\lambda\left(\|U\|_{F}^{2}+\|V\|_{F}^{2}\right) . \tag{18.7}
\end{equation*}
$$

In [ACW17], they have an efficient algorithm for more generalized problem:

Theorem 18.2.1 (Lemma 27 in [ACW17]). For $T_{1} \in \mathbb{R}^{t_{1} \times n}, T_{2} \in \mathbb{R}^{n \times t_{2}}, B \in \mathbb{R}^{t_{1} \times t_{2}}$, the problem of finding

$$
\min _{\widehat{X} \in \mathbb{R}^{n \times k}, \widehat{Y} \in \mathbb{R}^{k \times n}}\left\|T_{1} \widehat{X} \widehat{Y} T_{2}-B\right\|_{F}^{2}+\lambda\left(\left\|T_{1} \widehat{X}\right\|_{F}^{2}+\left\|\widehat{Y} T_{2}\right\|_{F}^{2}\right)
$$

and the minimizing $T_{1} \widehat{X} \in \mathbb{R}^{t_{1} \times k}$ and $\widehat{Y} T_{2} \in \mathbb{R}^{k \times t_{2}}$, can be solved in $O\left(t_{1} n \operatorname{rank}\left(T_{1}\right)+\right.$ $\left.t_{2} n \operatorname{rank}\left(T_{2}\right)+\operatorname{rank}\left(T_{2}\right) t_{1}\left(t_{2}+\operatorname{rank}\left(T_{1}\right)\right)\right)$ time.

The following Claim shows the relation between the solution of $\ell_{2}$ regression and the solution of $\ell_{1}$ regression.

Claim 18.2.2 ( $\ell_{2}$ relaxation). Let $A \in \mathbb{R}^{n \times d}$. We have $\|A\|_{F} \leq\|A\|_{1} \leq \sqrt{n d}\|A\|_{F}$.

Roughly speaking, our algorithms firstly sketch down the size of the input, and then relax the original problem to $\ell_{2}$ case. The above lemma and the claim are fundamental building blocks in the reduction from the original problem to the $\ell_{2}$ relaxation.

### 18.2.2 Cauchy transforms

In this section, we introduce the main sketching tool used in this work: the Cauchy transformation. We mainly introduce two kinds of Cauchy transformations: Dense Cauchy transform and Sparse Cauchy transform where each of them is introduced by [SW11] and [MM13] separately.

Definition 18.2.1 (Dense Cauchy transform [SW11]). Let $S=\sigma \cdot C \in \mathbb{R}^{m \times n}$ where $\sigma$ is a scalar, and each entry of $C \in \mathbb{R}^{m \times n}$ is chosen independently from the standard Cauchy distribution. $S$ is called Dense Cauchy transform matrix.

It is easy to see that for any $A \in \mathbb{R}^{n \times d}, S A$ can be computed in $O(m \cdot n n z(A))$ time [SW11].

Definition 18.2.2 (Sparse Cauchy transform [MM13]). Let $\Pi=\sigma \cdot S C \in \mathbb{R}^{m \times n}$, where $\sigma$ is a scalar, $S \in \mathbb{R}^{m \times n}$ has each column chosen independently and uniformly from the $m$ standard basis vectors of $\mathbb{R}^{m}$, and $C \in \mathbb{R}^{n \times n}$ is a diagonal matrix with diagonals chosen independently from the standard Cauchy distribution. $\Pi$ is called Sparse Cauchy transform matrix.

For any matrix $A \in \mathbb{R}^{n \times d}, \Pi A$ can be computed in $O(\operatorname{nnz}(A))$ time. For any matrix $A \in \mathbb{R}^{n \times d_{1} \times d_{2}}, \Pi A$ can be computed in $O(\mathrm{nnz}(A))$ time [MM13].

The properties of Cauchy transformation needed in this work is the following Lemma:
Lemma 18.2.3 (Lemma D. 11 in [SWZ17]). Given $M \in \mathbb{R}^{n \times d}, U \in \mathbb{R}^{n \times t}, d \geq t=\operatorname{rank}(U), n \geq$ $d \geq r=\operatorname{rank}(M)$, if sketching matrix $S \in \mathbb{R}^{m \times n}$ is drawn from any of the following probability distributions of matrices, with .99 probability $S$ satisfies the following properties:

1. $\|S M\|_{1} \leq c_{1}\|M\|_{1}$,
2. $\forall x \in \mathbb{R}^{t},\|S U x\|_{1} \geq \frac{1}{c_{2}}\|U x\|_{1}$,
where $c_{1}, c_{2}$ are parameters which depend on the distribution over $S$.

- $S \in \mathbb{R}^{m \times n}$ is a dense Cauchy transform matrix: a matrix with i.i.d. entries drawn from the standard Cauchy distribution. If $m=O(t \log t)$, then $c_{1} c_{2}=O(\log d)$. If $m=O((t+r) \log (t+r))$, then $c_{1} c_{2}=O(\min (\log d, r \log r))$.
- $S \in \mathbb{R}^{m \times n}$ is a sparse Cauchy transform matrix: $S=T D$, where $T \in \mathbb{R}^{m \times n}$ has each column i.i.d. from the uniform distribution on standard basis vector of $\mathbb{R}^{m}$, and $D \in \mathbb{R}^{n \times n}$ is a diagonal matrix with i.i.d. diagonal entries from standard Cauchy distribution. If $m=O\left(t^{5} \log ^{5} t\right)$, then $c_{1} c_{2}=O\left(t^{2} \log ^{2} t \log d\right)$. If $m=O\left((t+r)^{5} \log ^{5}(t+\right.$ $r)$ ), then $c_{1} c_{2}=O\left(t^{2} \log ^{2} t \log d, r^{3} \log ^{3} r\right)$.

In the above, if we replace $S$ with $\sigma \cdot S$ where $\sigma \in \mathbb{R} \backslash\{0\}$ is any scalar, then the relation between $m$ and $c_{1} c_{2}$ is still preserved.

### 18.3 The Algorithm

Our algorithm is shown in Algorithm 18.1. If we choose all the sketching matrices $S, R, T_{1}, T_{2}$ as dense Cauchy transformation matrices (see Definition 18.2.1), then we are able to finally get Theorem 18.1.2. If we choose all the sketching matrices $S, R, T_{1}, T_{2}$ as sparse Cauchy transformation matrices (see Definition 18.2.2), then we are able to finally get Theorem 18.1.1.

For convenience, we define $\operatorname{OPT}(A, k, \lambda)$ as follows

$$
\operatorname{OPT}(A, k, \lambda)=\min _{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times n}}\|U V-A\|_{1}+\lambda\left(\|U\|_{1}+\|V\|_{1}\right) .
$$

If $A, k, \lambda$ is clear from the context, we just use $\operatorname{OPT}$ to denote $\operatorname{OPT}(A, k, \lambda)$.
Then we present the analysis of our algorithm in the following way:

1. In Section 18.3.1, we show how to reduce the original problem to the following problem,

$$
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\|A R X Y S A-A\|_{1}+\lambda\left(\|A R X\|_{1}+\|Y S A\|_{1}\right)
$$

2. In Section 18.3.2, we show how to further reduce the above problem to:

$$
\begin{aligned}
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}} & \left\|T_{1} A R X Y S A T_{2}-T_{1} A T_{2}\right\|_{1} \\
& +\lambda\left(\left\|T_{1} A R X\right\|_{1}+\left\|Y S A T_{2}\right\|_{1}\right) .
\end{aligned}
$$

3. Finally, in Section 18.3.3, we show how to solve the above problem with a small approximation ratio.

### 18.3.1 Sketching Lemmas

In this section, we show how to use sketching techniques to prove Corollary 18.3.2 which means that there is a good approximate solution whose columns are in the column space of $A R$ and rows are in the row space of $S A$ where $S, A$ are two sketching matrices (shown in the algorithm 18.1) which satisfy desired properties. Before proving Corollary 18.3.2, we first prove the following lemma which is a "one-side" version of the corollary: it shows that there is a good approximate solution in the row space of $S A$.

Lemma 18.3.1. (Existence of a good solution in a small column subspace). Given matrix $A \in \mathbb{R}^{n \times n}, \alpha \geq 1, \beta \geq 1, \gamma \geq 1, \lambda>0$ and $k \geq 1$. Let $\widehat{U} \in \mathbb{R}^{n \times k}$ satisfy

$$
\min _{V \in \mathbb{R}^{k \times n}}\|\widehat{U} V-A\|_{1}+\lambda\left(\|\widehat{U}\|_{1}+\|V\|_{1}\right) \leq \gamma \mathrm{OPT} .
$$

Let $V^{*} \in \mathbb{R}^{k \times n}$ denote the optimal solution of

$$
\min _{V \in \mathbb{R}^{k \times n}}\|\widehat{U} V-A\|_{1}+\lambda\left(\|\widehat{U}\|_{1}+\|V\|_{1}\right)
$$

Let matrix $S \in \mathbb{R}^{m \times n}$ satisfy:

$$
\begin{aligned}
& \|S \widehat{U} x\|_{1} \geq \frac{1}{\alpha}\|\widehat{U} x\|_{1} \text { for all } x \in \mathbb{R}^{k}, \\
& \text { and }\left\|S \widehat{U} V^{*}-S A\right\|_{1} \leq \beta\left\|\widehat{U} V^{*}-A\right\|_{1} .
\end{aligned}
$$

Let $\widehat{V} \in \mathbb{R}^{k \times n}$ denote $\left[\begin{array}{c}S \widehat{U} \\ \lambda I\end{array}\right]^{\dagger} \cdot\left[\begin{array}{l}I \\ 0\end{array}\right] \cdot S A$. Then

$$
\|\widehat{U} \widehat{V}-A\|_{1}+\lambda\left(\|\widehat{U}\|_{1}+\|\widehat{V}\|_{1}\right) \leq \alpha \beta \gamma(\sqrt{m+k}+2) \mathrm{OPT} .
$$

Proof. Due to the space limitation, we provide the proofs in Appendix 18.4.

Corollary 18.3.2 (Existence of a good solution in both row and column subspace). Given matrix $A \in \mathbb{R}^{n \times n}, \alpha_{1}, \beta_{1}, \alpha_{2}, \beta_{2} \geq 1, \lambda>0$ and $k \geq 1$. Let $U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times n}$ denote the optimal solution of the problem,

$$
\min _{U, V}\|U V-A\|_{1}+\lambda\left(\|U\|_{1}+\|V\|_{1}\right) .
$$

Let matrix $S \in \mathbb{R}^{s \times n}$ satisfy :

$$
\begin{aligned}
& \left\|S U^{*} x\right\|_{1} \geq \frac{1}{\alpha_{1}}\left\|U^{*} x\right\|_{1} \text { for all } x \in \mathbb{R}^{k} \\
& \text { and }\left\|S U^{*} V^{*}-S A\right\|_{1} \leq \beta_{1}\left\|U^{*} V^{*}-A\right\|_{1} .
\end{aligned}
$$

Let $\widehat{V} \in \mathbb{R}^{k \times n}$ denote $\left[\begin{array}{c}S U^{*} \\ \lambda I\end{array}\right]^{\dagger} \cdot\left[\begin{array}{l}I \\ 0\end{array}\right] \cdot$ SA. Let $\widetilde{U} \in \mathbb{R}^{n \times k}$ denote the optimal solution of

$$
\min _{U \in \mathbb{R}^{n \times k}}\|U \widehat{V}-A\|_{1}+\lambda\left(\|U\|_{1}+\|\widehat{V}\|_{1}\right)
$$

Let $R^{\top} \in \mathbb{R}^{r \times n}$ satisfy:

$$
\begin{aligned}
& \left\|R^{\top} \widehat{V}^{\top} x\right\|_{1} \geq \frac{1}{\alpha_{2}}\left\|\widehat{V}^{\top} x\right\|_{1} \text { for all } x \in \mathbb{R}^{k} \\
& \text { and }\left\|R^{\top} \widehat{V}^{\top} \widetilde{U}^{\top}-R^{\top} A^{\top}\right\|_{1} \leq \beta_{2}\left\|\widehat{V}^{\top} \widetilde{U}^{\top}-A^{\top}\right\|_{1} .
\end{aligned}
$$

Let $\widehat{U}^{\top} \in \mathbb{R}^{k \times n}$ denote $\left[\begin{array}{c}R^{\top} \widehat{V}^{\top} \\ \lambda I\end{array}\right]^{\dagger} \cdot\left[\begin{array}{l}I \\ 0\end{array}\right] \cdot R^{\top} A^{\top}$. Then

$$
\begin{aligned}
& \|\widehat{U} \widehat{V}-A\|_{1}+\lambda\left(\|\widehat{U}\|_{1}+\|\widehat{V}\|_{1}\right) \\
\leq & 4\left(\alpha_{1} \beta_{1} \alpha_{2} \beta_{2}\right)(\sqrt{s+k}+2)(\sqrt{r+k}+2) \mathrm{OPT}
\end{aligned}
$$

and there exists $X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}$ such that

$$
\|A R X Y S A-A\|_{1}+\lambda\left(\|A R X\|_{1}+\|Y S A\|_{1}\right) \leq\|\widehat{U} \widehat{V}-A\|_{1}+\lambda\left(\|\widehat{U}\|_{1}+\|\widehat{V}\|_{1}\right)
$$

Proof. It is actually applying Lemma 18.3.1 twice. Due to the property of sketching matrix $S$, we can apply Lemma 18.3.1 to get

$$
\begin{aligned}
& \min _{U \in \mathbb{R}^{n \times k}}\|U \widehat{V}-A\|_{1}+\lambda\left(\|U\|_{1}+\|\widehat{V}\|_{1}\right) \\
\leq & \left\|U^{*} \widehat{V}-A\right\|_{1}+\lambda\left(\left\|U^{*}\right\|_{1}+\|\widehat{V}\|_{1}\right) \\
\leq & 2 \alpha_{1} \beta_{1}(\sqrt{s+k}+2) \mathrm{OPT} .
\end{aligned}
$$

We set $\gamma=2 \alpha_{1} \beta_{1}(\sqrt{s+k}+2)$. Due to the property of sketching matrix $R$, we can apply Lemma 18.3.1 again to get

$$
\begin{aligned}
& \|\widehat{U} \widehat{V}-A\|_{1}+\lambda\left(\|\widehat{U}\|_{1}+\|\widehat{V}\|_{1}\right) \\
\leq & 2 \alpha_{2} \beta_{2} \gamma(\sqrt{r+k}+2) \mathrm{OPT} \\
= & 4 \alpha_{1} \beta_{1} \alpha_{2} \beta_{2}(\sqrt{s+k}+2)(\sqrt{r+k}+2) \mathrm{OPT} .
\end{aligned}
$$

Notice that $\widehat{U}$ is in the column space of $A R, \widehat{V}$ is in the row space of $S A$, thus,

$$
\begin{aligned}
& \min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\|A R X Y S A-A\|_{1} \\
& +\lambda\left(\|A R X\|_{1}+\|Y S A\|_{1}\right) \\
\leq & \|\widehat{U} \widehat{V}-A\|_{1}+\lambda\left(\|\widehat{U}\|_{1}+\|\widehat{V}\|_{1}\right) \\
\leq & 4 \alpha_{1} \beta_{1} \alpha_{2} \beta_{2}(\sqrt{s+k}+2)(\sqrt{r+k}+2) \mathrm{OPT} .
\end{aligned}
$$

The above corollary shows that if we can find a good approximate solution $\widehat{X}, \widehat{Y}$ to

$$
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\|A R X Y S A-A\|_{1}+\lambda\left(\|A R X\|_{1}+\|Y S A\|_{1}\right),
$$

then we can have a good approximation $U=A R \widehat{X}, V=\widehat{Y} S A$ to the original problem.

### 18.3.2 Reducing to small problems

In the last section, we already restrict the solution in a low dimensional column space and in a low dimensional row space. But the problem size is till too large since the size of $A$ is still $n$-by- $n$. In this section, we further sketch down the size of the input. The Corollary 18.3.4 shows how to further reduce the problem to solving

$$
\begin{aligned}
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}} & \left\|T_{1} A R X Y S A T_{2}-T_{1} A T_{2}\right\|_{1} \\
& +\lambda\left(\left\|T_{1} A R X\right\|_{1}+\left\|Y S A T_{2}\right\|_{1}\right)
\end{aligned}
$$

where $T_{1}, T_{2}$ are two sketching matrices with some desired properties. Before proving Corollary 18.3.4, let us look at the following lemma which provides a "one-side" sketch version of that corollary:

Lemma 18.3.3. (Reducing the problem size by sketching on one side). Given three matrices $L \in \mathbb{R}^{d_{1} \times l_{1}}, M^{\top} \in \mathbb{R}^{d_{2} \times l_{2}}, N \in \mathbb{R}^{d_{1} \times d_{2}}$ and $\alpha \geq 1, \beta \geq 1, \gamma \geq 1, \lambda>0$. Let OPT $=$ $\min _{X, Y}\|L X Y M-N\|_{1}+\lambda\left(\|L X\|_{1}+\|Y M\|_{1}\right)$. Let $X^{*}, Y^{*}$ denote the optimal solution of

$$
\min _{X, Y}\|L X Y M-N\|_{1}+\lambda\left(\|L X\|_{1}+\|Y M\|_{1}\right)
$$

Let $T \in \mathbb{R}^{t \times d_{1}}$ satisfy :

$$
\|T L x\|_{1} \geq \frac{1}{\alpha}\|L x\|_{1}, \text { for all } x \in \mathbb{R}^{l_{1}}
$$

and

$$
\begin{aligned}
& \left\|T L X^{*} Y^{*} M-T N\right\|_{1}+\lambda \cdot\left\|T L X^{*}\right\|_{1} \\
\leq & \beta\left(\left\|L X^{*} Y^{*} M-N\right\|_{1}+\lambda \cdot\left\|L X^{*}\right\|_{1}\right)
\end{aligned}
$$

Let $\widehat{X}, \widehat{Y}$ satisfy

$$
\begin{aligned}
& \|T L \widehat{X} \widehat{Y} M-T N\|_{1}+\lambda \cdot\left(\|T L \widehat{X}\|_{1}+\|\widehat{Y} M\|_{1}\right) \\
\leq & \gamma \min _{X, Y}\|T L X Y M-T N\|_{1}+\lambda \cdot\left(\|T L X\|_{1}+\|Y M\|_{1}\right)
\end{aligned}
$$

Then

$$
\begin{aligned}
& \|L \widehat{X} \widehat{Y} M-N\|_{1}+\lambda\left(\|L \widehat{X}\|_{1}+\|\widehat{Y} M\|_{1}\right) \\
\leq & (\alpha \beta \gamma+\alpha \beta+1) \mathrm{OPT}
\end{aligned}
$$

Proof. Due to space limitation, we provide the proofs in Appendix 18.5.

Corollary 18.3.4 (Reducing the problem size by sketching on both sides). Given three matrices $L \in \mathbb{R}^{d_{1} \times l_{1}}, M^{\top} \in \mathbb{R}^{d_{2} \times l_{2}}, N \in \mathbb{R}^{d_{1} \times d_{2}}$ and $\alpha_{1} \geq 1, \alpha_{2} \geq 1, \beta_{1} \geq 1, \beta_{2} \geq 1, \gamma \geq$ $1, \lambda>0$. Let $\mathrm{OPT}=\min _{X, Y}\|L X Y M-N\|_{1}+\lambda\left(\|L X\|_{1}+\|Y M\|_{1}\right)$. Let $X^{*}, Y^{*}$ denote the optimal solution of

$$
\min _{X, Y}\|L X Y M-N\|_{1}+\lambda\left(\|L X\|_{1}+\|Y M\|_{1}\right)
$$

Let $T_{1} \in \mathbb{R}^{t_{1} \times d_{1}}$ satisfy : for all $x \in \mathbb{R}^{l_{1}}$

$$
\left\|T_{1} L x\right\|_{1} \geq \frac{1}{\alpha_{1}}\|L x\|_{1}
$$

and

$$
\begin{aligned}
& \left\|T_{1} L X^{*} Y^{*} M-T_{1} N\right\|_{1}+\lambda \cdot\left\|T_{1} L X^{*}\right\|_{1} \\
\leq & \beta_{1}\left(\left\|L X^{*} Y^{*} M-N\right\|_{1}+\lambda \cdot\left\|L X^{*}\right\|_{1}\right)
\end{aligned}
$$

Let $\widehat{X}, \widehat{Y}$ the optimal solution

$$
\min _{X, Y}\left\|T_{1} L X Y M-T_{1} N\right\|_{1}+\lambda \cdot\left(\left\|T_{1} L X\right\|_{1}+\|Y M\|_{1}\right)
$$

Let $T_{2}^{\top} \in \mathbb{R}^{t_{2} \times d_{2}}$ satisfy : for all $x \in \mathbb{R}^{l_{2}}$

$$
\left\|T_{2}^{\top} M^{\top} x\right\|_{1} \geq \frac{1}{\alpha_{2}}\left\|M^{\top} x\right\|_{1}
$$

and

$$
\begin{aligned}
& \left\|T_{1} L \widehat{X} \widehat{Y} M T_{2}-T_{1} N T_{2}\right\|_{1}+\lambda \cdot\left\|\widehat{Y} M T_{2}\right\|_{1} \\
\leq & \beta_{2}\left(\left\|T_{1} L \widehat{X} \widehat{Y} M-T_{1} N\right\|_{1}+\lambda \cdot\|\widehat{Y} M\|_{1}\right)
\end{aligned}
$$

Let $\widetilde{X}, \widetilde{Y}$ satisfy

$$
\begin{aligned}
& \left\|T_{1} L \tilde{X} \tilde{Y} M L_{2}-T_{1} N T_{2}\right\|_{1}+\lambda\left(\left\|T_{1} L \widetilde{X}\right\|_{1}+\left\|M \tilde{Y} T_{2}\right\|_{1}\right) \\
\leq & \gamma \min _{X, Y}\left\|T_{1} L X Y M L_{2}-T_{1} N T_{2}\right\|_{1}+\lambda\left(\left\|T_{1} L X\right\|_{1}+\left\|M Y T_{2}\right\|_{1}\right)
\end{aligned}
$$

Then

$$
\begin{aligned}
& \|L \tilde{X} \tilde{Y} M-N\|_{1}+\lambda\left(\|L \tilde{X}\|_{1}+\|M \tilde{Y}\|_{1}\right) \\
\leq & 9 \alpha_{1} \alpha_{2} \beta_{1} \beta_{2} \gamma \mathrm{OPT}
\end{aligned}
$$

Proof. This follows by applying Lemma 18.3.3 twice. Due to the property of $T_{2}$ and the definition of $\tilde{X}, \tilde{Y}$, by applying Lemma 18.3.3, we have that

$$
\begin{aligned}
& \left\|T_{1} L \tilde{X} \tilde{Y} M-T_{1} N\right\|_{1}+\lambda \cdot\left(\left\|T_{1} L \widetilde{X}\right\|_{1}+\|\widetilde{Y} M\|_{1}\right) \\
\leq & \left(\alpha_{2} \beta_{2} \gamma+\alpha_{2} \beta_{2}+1\right)\left(\left\|T_{1} L \widehat{X} \widehat{Y} M-T_{1} N\right\|_{1}+\lambda \cdot\left(\left\|T_{1} L \widehat{X}\right\|_{1}+\|\widehat{Y} M\|_{1}\right)\right) .
\end{aligned}
$$

Now we set the new $\gamma^{\prime}=\left(\alpha_{2} \beta_{2} \gamma+\alpha_{2} \beta_{2}+1\right)$. Then, due to the property of $T_{1}$ and the definition of $\widehat{X}, \widehat{Y}$, we can apply Lemma 18.3.3 again, we then have

$$
\begin{aligned}
& \|L \widetilde{X} \widetilde{Y} M-N\|_{1}+\lambda\left(\|L \widetilde{X}\|_{1}+\|M \widetilde{Y}\|_{1}\right) \\
\leq & \left(\alpha_{1} \beta_{1} \gamma^{\prime}+\alpha_{1} \beta_{1}+1\right) \mathrm{OPT}
\end{aligned}
$$

Since $\alpha_{1}, \alpha_{2}, \beta_{1}, \beta_{2} \geq 1$, we then have $\left(\alpha_{1} \beta_{1} \gamma^{\prime}+\alpha_{1} \beta_{1}+1\right) \leq 3 \alpha_{1} \beta_{1} \gamma^{\prime} \leq 9 \alpha_{1} \beta_{1} \alpha_{2} \beta_{2} \gamma$.

Let us set $L=A R, M=S A, N=A$, then due to the above corollary, we show that once we get a good approximation $\widetilde{X}, \widetilde{Y}$ to $\min _{X, Y}\left\|T_{1} A R X Y S A T_{2}-T_{1} A T_{2}\right\|_{1}+\lambda\left(\left\|T_{1} A R X\right\|_{1}+\right.$ $\left\|Y S A T_{2}\right\|_{1}$ ), we get a good approximation to $\min _{X, Y}\|A R X Y S A-A\|_{1}+\lambda\left(\|A R X\|_{1}+\right.$ $\left.\|Y S A\|_{1}\right)$. Since the number of rows of $T_{1}$ is small and the number of columns of $T_{2}$ is also small, the problem size becomes very small. In the next section, we will see how to use the $\ell_{2}$ relaxation to solve the small problem.

### 18.3.3 Solving small problems

In this section, we focus on solving

$$
\min _{X, Y}\left\|T_{1} A R X Y S A T_{2}-T_{1} A T_{2}\right\|_{1}+\lambda\left(\left\|T_{1} A R X\right\|_{1}+\left\|Y S A T_{2}\right\|_{1}\right)
$$

Firstly, we have the following easy fact.

Fact 18.3.5. Given five matrices $L, X, Y, M$ and $N$ with matching dimensions, we have

$$
\begin{aligned}
& \|L X Y M-N\|_{1}+\lambda\left(\|L X\|_{1}+\|Y M\|_{1}\right) \\
= & \left\|\left[\begin{array}{c}
L X \\
\lambda I
\end{array}\right] \cdot\left[\begin{array}{cc}
Y M & \lambda I
\end{array}\right]-\left[\begin{array}{cc}
N & 0 \\
0 & \lambda^{2} I
\end{array}\right]\right\|_{1} .
\end{aligned}
$$

We define $f(X, Y)$ as follows

$$
f(X, Y)=\left[\begin{array}{c}
T_{1} A R X  \tag{18.8}\\
\lambda I
\end{array}\right] \cdot\left[\begin{array}{ll}
Y S A T_{2} & \lambda I
\end{array}\right]-\left[\begin{array}{cc}
T_{1} A T_{2} & 0 \\
0 & \lambda^{2} I
\end{array}\right]
$$

Due to the above fact, we can relax the $\ell_{1}$-norm minimization problem

$$
\begin{equation*}
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\|f(X, Y)\|_{1} \tag{18.9}
\end{equation*}
$$

to Frobenius norm minimization problem

$$
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\|f(X, Y)\|_{F}
$$

which is equivalent to

$$
\begin{equation*}
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\|f(X, Y)\|_{F}^{2} \tag{18.10}
\end{equation*}
$$

Lemma 18.3.6 (Relaxing $\ell_{1}$ norm to Frobenius norm). Given $A \in \mathbb{R}^{n \times n}, T_{1} \in \mathbb{R}^{t_{1} \times n}$, $T_{2} \in \mathbb{R}^{n \times t_{2}}, R \in \mathbb{R}^{n \times r}$ and $S \in \mathbb{R}^{s \times n}$. Let $f_{T_{1}, T_{2}, R, S, A, \lambda}(X, Y) \in \mathbb{R}^{\left(t_{1}+k\right) \times\left(t_{2}+k\right)}$ be defined in Eq. (18.8), if $X^{\prime}, Y^{\prime}$ satisfies that

$$
\left\|f\left(X^{\prime}, Y^{\prime}\right)\right\|_{F}^{2} \leq(1+\epsilon) \min _{X, Y}\|f(X, Y)\|_{F}^{2}
$$

for some $\epsilon>0$, then we have

$$
\left\|f\left(X^{\prime}, Y^{\prime}\right)\right\|_{1} \leq \sqrt{t_{1}+k} \cdot \sqrt{t_{2}+k} \cdot \sqrt{1+\epsilon} \min _{X, Y}\|f(X, Y)\|_{1} .
$$

Proof. Let $X^{*}, Y^{*}$ satisfy $\left\|f\left(X^{*}, Y^{*}\right)\right\|_{1}=\min _{X, Y}\|f(X, Y)\|_{1}$. We have:

$$
\begin{aligned}
& \left\|f\left(X^{\prime}, Y^{\prime}\right)\right\|_{1} \\
\leq & \sqrt{t_{1}+k} \cdot \sqrt{t_{2}+k} \cdot\left\|f\left(X^{\prime}, Y^{\prime}\right)\right\|_{F} \\
\leq & \sqrt{t_{1}+k} \cdot \sqrt{t_{2}+k} \cdot \sqrt{1+\epsilon} \cdot \min _{X, Y}\|f(X, Y)\|_{F} \\
\leq & \sqrt{t_{1}+k} \cdot \sqrt{t_{2}+k} \cdot \sqrt{1+\epsilon} \cdot\left\|f\left(X^{*}, Y^{*}\right)\right\|_{F} \\
\leq & \sqrt{t_{1}+k} \cdot \sqrt{t_{2}+k} \cdot \sqrt{1+\epsilon} \cdot\left\|f\left(X^{*}, Y^{*}\right)\right\|_{1}
\end{aligned}
$$

where the first inequality and the last inequality follow the Claim 18.2.2.

Therefore, we can reduce the entry-wise $\ell_{1}$ problem (18.9) to Frobenius norm problem (18.10).

Using the Theorem 18.2.1, we can solve the Frobenius norm problem (18.10). Now, we are ready to present our main results.

Theorem 18.3.7 (Restatement of Theorem 18.1.1). Given matrix $A \in \mathbb{R}^{n \times n}$, parameters $k \geq 1, \lambda>0$, let OPT denote $\min _{U^{\prime} \in \mathbb{R}^{n \times k}, V^{\prime} \in \mathbb{R}^{k \times n}}\left\|U^{\prime} V^{\prime}-A\right\|_{1}+\lambda\left(\left\|U^{\prime}\right\|_{1}+\left\|V^{\prime}\right\|_{1}\right)$, there is an algorithm that runs in $O(\mathrm{nnz}(A)+n \operatorname{poly}(k, \log n))$ time to output two matrices $U \in \mathbb{R}^{n \times k}$, $V \in \mathbb{R}^{k \times n}$ such that,

$$
\|U V-A\|_{1}+\lambda\left(\|U\|_{1}+\|V\|_{1}\right) \leq \operatorname{poly}(k, \log n) \cdot \mathrm{OPT}
$$

holds with probability at least 9/10.

Proof. Due to the space limits, we put the whole proof into Appendix 18.6.

Theorem 18.3.8 (Restatement of Theorem 18.1.2). Given matrix $A \in \mathbb{R}^{n \times n}$, parameters $k \geq 1, \lambda>0$, let OPT denote $\min _{U^{\prime} \in \mathbb{R}^{n \times k}, V^{\prime} \in \mathbb{R}^{k \times n}}\left\|U^{\prime} V^{\prime}-A\right\|_{1}+\lambda\left(\left\|U^{\prime}\right\|_{1}+\left\|V^{\prime}\right\|_{1}\right)$, there is an algorithm that runs in $\widetilde{O}(k) \cdot \mathrm{nnz}(A)+n \cdot \operatorname{poly}(k, \log n)$ time to output two matrices $U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times n}$ such that,

$$
\|U V-A\|_{1}+\lambda\left(\|U\|_{1}+\|V\|_{1}\right) \leq \widetilde{O}\left(k^{2}\right) \log ^{3} n \cdot \mathrm{OPT}
$$

holds with probability at least 9/10.

Proof. Due to the space limits, we put the whole proof into Appendix 18.7.

Remark 18.3.1. Combining our new technique with $p$-stable transform (e.g. Section E of [SWZ17] and [SWZ19b]), all of our results can be extended to $\ell_{p}$-norm, where $1<p<2$.

### 18.4 Proof of Lemma 18.3.1

Lemma 18.3.1. (Existence of a good solution in a small column subspace). Given matrix $A \in \mathbb{R}^{n \times n}, \alpha \geq 1, \beta \geq 1, \gamma \geq 1, \lambda>0$ and $k \geq 1$. Let $\widehat{U} \in \mathbb{R}^{n \times k}$ satisfy

$$
\min _{V \in \mathbb{R}^{k \times n}}\|\widehat{U} V-A\|_{1}+\lambda\left(\|\widehat{U}\|_{1}+\|V\|_{1}\right) \leq \gamma \mathrm{OPT} .
$$

Let $V^{*} \in \mathbb{R}^{k \times n}$ denote the optimal solution of

$$
\min _{V \in \mathbb{R}^{k \times n}}\|\widehat{U} V-A\|_{1}+\lambda\left(\|\widehat{U}\|_{1}+\|V\|_{1}\right)
$$

Let matrix $S \in \mathbb{R}^{m \times n}$ satisfy:

$$
\begin{aligned}
& \|S \widehat{U} x\|_{1} \geq \frac{1}{\alpha}\|\widehat{U} x\|_{1} \text { for all } x \in \mathbb{R}^{k}, \\
& \text { and }\left\|S \widehat{U} V^{*}-S A\right\|_{1} \leq \beta\left\|\widehat{U} V^{*}-A\right\|_{1} .
\end{aligned}
$$

Let $\widehat{V} \in \mathbb{R}^{k \times n}$ denote $\left[\begin{array}{c}S \widehat{U} \\ \lambda I\end{array}\right]^{\dagger} \cdot\left[\begin{array}{l}I \\ 0\end{array}\right] \cdot S A$. Then

$$
\|\widehat{U} \widehat{V}-A\|_{1}+\lambda\left(\|\widehat{U}\|_{1}+\|\widehat{V}\|_{1}\right) \leq \alpha \beta \gamma(\sqrt{m+k}+2) \mathrm{OPT} .
$$

Proof. We can show the following Claim,

## Claim 18.4.1.

$$
\|\widehat{U} \widehat{V}-A\|_{1}+\lambda\left(\|\widehat{U}\|_{1}+\|\widehat{V}\|_{1}\right) \leq \alpha\left\|\left[\begin{array}{c}
S \widehat{U} \\
\lambda I
\end{array}\right] \widehat{V}-\left[\begin{array}{c}
S A \\
0
\end{array}\right]\right\|_{1}+\alpha\left\|\left[\begin{array}{c}
S \widehat{U} \\
\lambda I
\end{array}\right] V^{*}-\left[\begin{array}{c}
S A \\
0
\end{array}\right]\right\|_{1}+\gamma \mathrm{OPT} .
$$

Proof.

$$
\begin{aligned}
& \|\widehat{U} \widehat{V}-A\|_{1}+\lambda\left(\|\widehat{U}\|_{1}+\|\widehat{V}\|_{1}\right) \\
= & \left\|\left[\begin{array}{c}
\widehat{U} \\
\lambda I
\end{array}\right] \widehat{V}-\left[\begin{array}{c}
A \\
0
\end{array}\right]\right\|_{1}+\lambda\|\widehat{U}\|_{1} \\
\leq & \left\|\left[\begin{array}{c}
\widehat{U} \\
\lambda I
\end{array}\right]\left(\widehat{V}-V^{*}\right)\right\|_{1}+\left\|\left[\begin{array}{c}
\widehat{U} \\
\lambda I
\end{array}\right] V^{*}-\left[\begin{array}{c}
A \\
0
\end{array}\right]\right\|_{1}+\lambda\|\widehat{U}\|_{1} \\
\leq & \left\|\left[\begin{array}{c}
\widehat{U} \\
\lambda I
\end{array}\right]\left(\widehat{V}-V^{*}\right)\right\|_{1}+\gamma \mathrm{OPT} \\
= & \left\|\widehat{U}\left(\widehat{V}-V^{*}\right)\right\|_{1}+\lambda\left\|\widehat{V}-V^{*}\right\|_{1}+\gamma \mathrm{OPT} \\
\leq & \alpha\left\|S \widehat{U}\left(\widehat{V}-V^{*}\right)\right\|_{1}+\lambda\left\|\widehat{V}-V^{*}\right\|_{1}+\gamma \mathrm{OPT} \\
\leq & \alpha\left\|S \widehat{U}\left(\widehat{V}-V^{*}\right)\right\|_{1}+\lambda \alpha\left\|\widehat{V}-V^{*}\right\|_{1}+\gamma \mathrm{OPT} \\
= & \alpha\left\|\left[\begin{array}{c}
S \widehat{U} \\
\lambda I
\end{array}\right]\left(\widehat{V}-V^{*}\right)\right\|_{1}+\gamma \mathrm{OPT} \\
\leq & \alpha\left\|\left[\begin{array}{c}
S \widehat{U} \\
\lambda I
\end{array}\right] \widehat{V}-\left[\begin{array}{c}
S A \\
0
\end{array}\right]\right\|_{1}+\alpha\left\|\left[\begin{array}{c}
S \widehat{U} \\
\lambda I
\end{array}\right] V^{*}-\left[\begin{array}{c}
S A \\
0
\end{array}\right]\right\|_{1}+\gamma \mathrm{OPT}
\end{aligned}
$$

where the second step follows by triangle inequality, the third step follows by

$$
\left\|\left[\begin{array}{c}
\widehat{U} \\
\lambda I
\end{array}\right] V^{*}-\left[\begin{array}{c}
A \\
0
\end{array}\right]\right\|_{1}+\lambda\|\widehat{U}\|_{1} \leq \gamma \mathrm{OPT}
$$

the fifth step follows by $\left\|\widehat{U}\left(\widehat{V}-V^{*}\right)\right\|_{1} \leq \alpha\left\|S \widehat{U}\left(\widehat{V}-V^{*}\right)\right\|_{1}$, the sixth step follows by $\alpha \geq 1$, the eighth step follows by triangle inequality.

We now show how to bound the second term on RHS of the equation shown in Claim 18.4.1.

## Claim 18.4.2.

$$
\alpha\left\|\left[\begin{array}{c}
S \widehat{U} \\
\lambda I
\end{array}\right] V^{*}-\left[\begin{array}{c}
S A \\
0
\end{array}\right]\right\|_{1} \leq \alpha \beta \gamma \mathrm{OPT} .
$$

Proof. We have

$$
\begin{aligned}
& \alpha\left\|\left[\begin{array}{c}
S \hat{U} \\
\lambda I
\end{array}\right] V^{*}-\left[\begin{array}{c}
S A \\
0
\end{array}\right]\right\|_{1} \\
= & \alpha\left\|S\left(\widehat{U} V^{*}-A\right)\right\|_{1}+\alpha \lambda\left\|V^{*}\right\|_{1} \\
\leq & \alpha \beta\left\|\widehat{U} V^{*}-A\right\|_{1}+\alpha \lambda\left\|V^{*}\right\|_{1} \\
\leq & \alpha \beta\left\|\widehat{U} V^{*}-A\right\|_{1}+\alpha \beta \lambda\left\|V^{*}\right\|_{1}+\alpha \beta \lambda\|\widehat{U}\|_{1} \\
\leq & \alpha \beta \gamma \mathrm{OPT},
\end{aligned}
$$

the second step follows by $\left\|S\left(\widehat{U} V^{*}-A\right)\right\|_{1} \leq \beta\left\|\widehat{U} V^{*}-A\right\|_{1}$.

We show how to bound the first term on RHS of Claim 18.4.1,

## Claim 18.4.3.

$$
\alpha\left\|\left[\begin{array}{c}
S \widehat{U} \\
\lambda I
\end{array}\right] \widehat{V}-\left[\begin{array}{c}
S A \\
0
\end{array}\right]\right\|_{1} \leq \alpha \beta \sqrt{m+k} \mathrm{OPT}
$$

Proof. We have

$$
\begin{aligned}
& \alpha\left\|\left[\begin{array}{c}
S \widehat{U} \\
\lambda I
\end{array}\right] \widehat{V}-\left[\begin{array}{c}
S A \\
0
\end{array}\right]\right\|_{1} \\
= & \alpha \sum_{i=1}^{n}\left\|\left[\begin{array}{c}
S \widehat{U} \\
\lambda I
\end{array}\right] \widehat{V}_{i}-\left[\begin{array}{c}
S A_{i} \\
0
\end{array}\right]\right\|_{1} \\
\leq & \alpha \sqrt{m+k} \sum_{i=1}^{n}\left\|\left[\begin{array}{c}
S \widehat{U} \\
\lambda I
\end{array}\right] \widehat{V}_{i}-\left[\begin{array}{c}
S A_{i} \\
0
\end{array}\right]\right\|_{2} \\
\leq & \alpha \sqrt{m+k} \sum_{i=1}^{n}\left\|\left[\begin{array}{c}
S \widehat{U} \\
\lambda I
\end{array}\right] V_{i}^{*}-\left[\begin{array}{c}
S A_{i} \\
0
\end{array}\right]\right\|_{2} \\
\leq & \alpha \sqrt{m+k} \sum_{i=1}^{n}\left\|\left[\begin{array}{c}
S \widehat{U} \\
\lambda I
\end{array}\right] V_{i}^{*}-\left[\begin{array}{c}
S A_{i} \\
0
\end{array}\right]\right\|_{1} \\
= & \alpha \sqrt{m+k}\left\|\left[\begin{array}{c}
S \widehat{U} \\
\lambda I
\end{array}\right] V^{*}-\left[\begin{array}{c}
S A \\
0
\end{array}\right]\right\|_{1} \\
= & \alpha \sqrt{m+k}\left(\left\|S \widehat{U} V^{*}-S A\right\|_{1}+\lambda\left\|V^{*}\right\|_{1}\right) \\
\leq & \alpha \sqrt{m+k}\left(\beta\left\|\widehat{U} V^{*}-A\right\|_{1}+\lambda\left\|V^{*}\right\|_{1}\right) \\
\leq & \alpha \beta \sqrt{m+k} \mathrm{OPT}
\end{aligned}
$$

where the second step follows by $\forall d \in \mathbb{Z}_{>0}, x \in \mathbb{R}^{d},\|x\|_{1} \leq \sqrt{d}\|x\|_{2}$, the third step follows by definition of $\widehat{V}$, the fourth step follows by $\forall d \in \mathbb{Z}_{>0}, x \in \mathbb{R}^{d}\|x\|_{2} \leq\|x\|_{1}$, the seventh step follows by $\left\|S \widehat{U} V^{*}-S A\right\|_{1} \leq \beta\left\|\widehat{U} V^{*}-A\right\|_{1}$, and the last step follows by $\beta \geq 1, \lambda\|\widehat{U}\|_{1} \geq 0$.

Putting it all together, we

$$
\begin{aligned}
& \|\widehat{U} \widehat{V}-A\|_{1}+\lambda\left(\|\widehat{U}\|_{1}+\|\widehat{V}\|_{1}\right) \\
\leq & \alpha \beta \sqrt{m+k} \mathrm{OPT}+\alpha \beta \gamma \mathrm{OPT}+\gamma \mathrm{OPT} \\
\leq & \alpha \beta \sqrt{m+k} \mathrm{OPT}+2 \alpha \beta \gamma \mathrm{OPT} \\
\leq & \alpha \beta \gamma(\sqrt{m+k}+2) \mathrm{OPT}
\end{aligned}
$$

where the first step follows by Claim 18.4.1, Claim 18.4.2 and Claim 18.4.3, the second step follows by $\alpha \beta \geq 1$, the last step follows by $\gamma \geq 1$.

Thus, we complete the proof.

### 18.5 Proof of Lemma 18.3.3

Lemma 18.3.3. (Reducing the problem size by sketching on one side). Given three matrices $L \in \mathbb{R}^{d_{1} \times l_{1}}, M^{\top} \in \mathbb{R}^{d_{2} \times l_{2}}, N \in \mathbb{R}^{d_{1} \times d_{2}}$ and $\alpha \geq 1, \beta \geq 1, \gamma \geq 1, \lambda>0$. Let OPT $=$ $\min _{X, Y}\|L X Y M-N\|_{1}+\lambda\left(\|L X\|_{1}+\|Y M\|_{1}\right)$. Let $X^{*}, Y^{*}$ denote the optimal solution of

$$
\min _{X, Y}\|L X Y M-N\|_{1}+\lambda\left(\|L X\|_{1}+\|Y M\|_{1}\right) .
$$

Let $T \in \mathbb{R}^{t \times d_{1}}$ satisfy :

$$
\|T L x\|_{1} \geq \frac{1}{\alpha}\|L x\|_{1}, \text { for all } x \in \mathbb{R}^{l_{1}}
$$

and

$$
\begin{aligned}
& \left\|T L X^{*} Y^{*} M-T N\right\|_{1}+\lambda \cdot\left\|T L X^{*}\right\|_{1} \\
\leq & \beta\left(\left\|L X^{*} Y^{*} M-N\right\|_{1}+\lambda \cdot\left\|L X^{*}\right\|_{1}\right)
\end{aligned}
$$

Let $\widehat{X}, \widehat{Y}$ satisfy

$$
\begin{aligned}
& \|T L \widehat{X} \widehat{Y} M-T N\|_{1}+\lambda \cdot\left(\|T L \widehat{X}\|_{1}+\|\widehat{Y} M\|_{1}\right) \\
\leq & \gamma \min _{X, Y}\|T L X Y M-T N\|_{1}+\lambda \cdot\left(\|T L X\|_{1}+\|Y M\|_{1}\right)
\end{aligned}
$$

Then

$$
\begin{aligned}
& \|L \widehat{X} \widehat{Y} M-N\|_{1}+\lambda\left(\|L \widehat{X}\|_{1}+\|\widehat{Y} M\|_{1}\right) \\
\leq & (\alpha \beta \gamma+\alpha \beta+1) \mathrm{OPT} .
\end{aligned}
$$

## Proof.

$$
\begin{aligned}
& \|L \widehat{X} \widehat{Y} M-N\|_{1}+\lambda\left(\|L \widehat{X}\|_{1}+\|\widehat{Y} M\|_{1}\right) \\
\leq & \left\|L\left(\widehat{X} \widehat{Y}-X^{*} Y^{*}\right) M\right\|_{1}+\left\|L X^{*} Y^{*} M-N\right\|_{1}+\lambda\left(\|L \widehat{X}\|_{1}+\|\widehat{Y} M\|_{1}\right) \\
\leq & \left\|L\left(\widehat{X} \widehat{Y}-X^{*} Y^{*}\right) M\right\|_{1}+\mathrm{OPT}+\lambda\left(\|L \widehat{X}\|_{1}+\|\widehat{Y} M\|_{1}\right) \\
\leq & \alpha\left\|T L\left(\widehat{X} \widehat{Y}-X^{*} Y^{*}\right) M\right\|_{1}+\mathrm{OPT}+\lambda\left(\alpha\|T L \widehat{X}\|_{1}+\|\widehat{Y} M\|_{1}\right) \\
\leq & \alpha\|T L \widehat{X} \widehat{Y} M-T N\|_{1}+\alpha\left\|T L X^{*} Y^{*} M-T N\right\|_{1}+\lambda\left(\alpha\|T L \widehat{X}\|_{1}+\|\widehat{Y} M\|_{1}\right)+\mathrm{OPT} \\
\leq & \alpha\left(\|T L \widehat{X} \widehat{Y} M-T N\|_{1}+\lambda\left(\|T L \widehat{X}\|_{1}+\|\widehat{Y} M\|_{1}\right)\right)+\alpha\left\|T L X^{*} Y^{*} M-T N\right\|_{1}+\mathrm{OPT} \\
\leq & \alpha \gamma\left(\left\|T L X^{*} Y^{*} M-T N\right\|_{1}+\lambda\left(\left\|T L X^{*}\right\|_{1}+\left\|Y^{*} M\right\|_{1}\right)\right)+\alpha\left\|T L X^{*} Y^{*} M-T N\right\|_{1}+\mathrm{OPT} \\
\leq & \alpha \beta \gamma \mathrm{OPT}+\alpha\left\|T L X^{*} Y^{*} M-T N\right\|_{1}+\mathrm{OPT} \\
\leq & \alpha \beta \gamma \mathrm{OPT}+\alpha \beta\left\|L X^{*} Y^{*} M-N\right\|_{1}+\mathrm{OPT} \\
= & (\alpha \beta \gamma+\alpha \beta+1) \mathrm{OPT},
\end{aligned}
$$

where the first step follows by triangle inequality, the second step follows by $\| L X^{*} Y^{*} M-$ $N \|_{1} \leq$ OPT, the third step follows by the property of $T$, the forth step follows by triangle inequality, the fifth step follows by $\alpha \geq 1$, the sixth step follows by the definition of $\widehat{X}, \widehat{Y}$, the seventh step follows by the property of $T$, the eighth step follows by the property of $T$.

### 18.6 Proof of Theorem 18.3.7

Proof. We choose all the matrices $S \in \mathbb{R}^{s \times n}, R \in \mathbb{R}^{n \times r}, T_{1} \in \mathbb{R}^{t_{1} \times n}, T_{2} \in \mathbb{R}^{n \times t_{2}}$ in Algorithm 18.1 as sparse Cauchy transformation matrices (see Definition 18.2.2) which have the same scalar $\sigma=1$. We set $s=O\left(k^{5} \log ^{5} k\right), r=O\left(k^{5} \log ^{5} k\right), t_{1}=O\left(r^{5} \log ^{5} r\right), t_{2}=$ $O\left(s^{5} \log ^{5} s\right)$.

Since computing $T_{1} A, A T_{2}$ can be done in $\mathrm{nnz}(A)$ time, and the size of $T_{1} A, A T_{2}$ is at most $n \cdot \operatorname{poly}(k)$, the total running time of computing $T_{1} A R, S A T_{2}, T_{1} A T_{2}$ is at most $n n z(A)+n \cdot \operatorname{poly}(k)$. Since the size of the final small Frobenius norm ridge low-rank matrix approximation problem is poly $(k)$, the running time of the whole algorithm is at most $n n z(A)+n \cdot \operatorname{poly}(k)$.

Now let us consider the correctness of the algorithm.
Let $U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times n}$ satisfy $\left\|U^{*} V^{*}-A\right\|_{1}+\lambda\left(\left\|U^{*}\right\|_{1}+\left\|V^{*}\right\|_{1}\right)=$ OPT. Then due to Lemma 18.2.3, with probability at least $0.99, S$ satisfies

$$
\left\|S U^{*} x\right\|_{1} \geq \frac{1}{\alpha_{1}}\left\|U^{*} x\right\|_{1} \text { for all } x \in \mathbb{R}^{k} \text { and }\left\|S U^{*} V^{*}-S A\right\|_{1} \leq \beta_{1}\left\|U^{*} V^{*}-A\right\|_{1}
$$

where $\alpha_{1} \beta_{1}=O\left(k^{2} \log ^{2} k \log n\right)$. Let $\widehat{V} \in \mathbb{R}^{k \times n}$ denote $\left[\begin{array}{c}S U^{*} \\ \lambda I\end{array}\right]^{\dagger} \cdot\left[\begin{array}{l}I \\ 0\end{array}\right] \cdot S A$. Let $\widetilde{U} \in \mathbb{R}^{n \times k}$ denote the optimal solution of

$$
\min _{U \in \mathbb{R}^{n \times k}}\|U \widehat{V}-A\|_{1}+\lambda\left(\|U\|_{1}+\|\widehat{V}\|_{1}\right)
$$

Due to Lemma 18.2.3, with probability at least $0.99, R^{\top} \in \mathbb{R}^{r \times n}$ satisfies:

$$
\begin{aligned}
& \left\|R^{\top} \widehat{V}^{\top} x\right\|_{1} \geq \frac{1}{\alpha_{2}}\left\|\widehat{V}^{\top} x\right\|_{1} \text { for all } x \in \mathbb{R}^{k}, \text { and } \\
& \left\|R^{\top} \widehat{V}^{\top} \widetilde{U}^{\top}-R^{\top} A^{\top}\right\|_{1} \leq \beta_{2}\left\|\widehat{V}^{\top} \widetilde{U}^{\top}-A^{\top}\right\|_{1}
\end{aligned}
$$

where $\alpha_{2} \beta_{2}=O\left(k^{2} \log ^{2} k \log n\right)$. Then, by applying Corollary 18.3.2, we have that

$$
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\|A R X Y S A-A\|_{1}+\lambda\left(\|A R X\|_{1}+\|Y S A\|_{1}\right) \leq \operatorname{poly}(k, \log n) \text { OPT } .
$$

Let $X^{*} \in \mathbb{R}^{r \times k}, Y^{*} \in \mathbb{R}^{k \times s}$ denote the optimal solution of

$$
\min _{X, Y}\|A R X Y S A-A\|_{1}+\lambda\left(\|A R X\|_{1}+\|Y S A\|_{1}\right)
$$

Due to Lemma 18.2.3, with probability at least $0.99, T_{1}$ satisfies that $\forall x \in \mathbb{R}^{r},\left\|T_{1} A R x\right\|_{1} \geq$ $\frac{1}{\alpha_{1}^{\prime}}\|A R x\|_{1}$, and $\left\|T_{1}\left(\left[A R X^{*} Y^{*} S A-A \quad \lambda A R X^{*}\right]\right)\right\|_{1} \leq \beta_{1}^{\prime} \|\left[\begin{array}{ll}A R X^{*} Y^{*} S A-A & \left.\lambda A R X^{*}\right] \|_{1}\end{array}\right.$ where $\alpha_{1}^{\prime} \beta_{1}^{\prime}=O\left(r^{2} \log ^{2} r \log n\right)$. Due to Lemma 18.2.3, with probability at least $0.99, T_{2}$ satisfies that $\forall x \in \mathbb{R}^{s},\left\|T_{2}^{\top} A^{\top} S^{\top} x\right\|_{1} \geq \frac{1}{\alpha_{2}^{\prime}}\left\|A^{\top} S^{\top} x\right\|_{1}$ and

$$
\begin{aligned}
&\left\|T_{2}^{\top}\left(\left[A^{\top} S^{\top}\left(Y^{*}\right)^{\top}\left(X^{*}\right)^{\top} R^{\top} A^{\top} T_{1}^{\top}-A^{\top} T_{1}^{\top} \quad \lambda A^{\top} S^{\top}\left(Y^{*}\right)^{\top}\right]\right)\right\|_{1} \\
& \leq \beta_{2}^{\prime} \|\left[A^{\top} S^{\top}\left(Y^{*}\right)^{\top}\left(X^{*}\right)^{\top} R^{\top} A^{\top} T_{1}^{\top}-A^{\top} T_{1}^{\top}\right. \\
&\left.\lambda A^{\top} S^{\top}\left(Y^{*}\right)^{\top}\right] \|_{1}
\end{aligned}
$$

where $\alpha_{2}^{\prime} \beta_{2}^{\prime} \leq O\left(s^{2} \log ^{2} s \log n\right)$. Thus, by Corollary 18.3.4, for any $\gamma \geq 1$, if $\tilde{X}, \tilde{Y}$ can achieve

$$
\begin{aligned}
& \left\|T_{1} A R \tilde{X} \tilde{Y} S A T_{2}-T_{1} A T_{2}\right\|_{1}+\lambda\left(\left\|T_{1} A R \widetilde{X}\right\|_{1}+\left\|\widetilde{Y} S A T_{2}\right\|_{1}\right) \\
\leq & \gamma \min _{X, Y}\left(\left\|T_{1} A R X Y S A T_{2}-T_{1} A T_{2}\right\|_{1}+\lambda\left(\left\|T_{1} A R X\right\|_{1}+\left\|Y S A T_{2}\right\|_{1}\right)\right)
\end{aligned}
$$

then

$$
\begin{aligned}
& \|A R \widetilde{X} \tilde{Y} S A-A\|_{1}+\lambda\left(\|A R \widetilde{X}\|_{1}+\|\widetilde{Y} S A\|_{1}\right) \\
\leq & \operatorname{poly}(k, \log n) \gamma \min _{X, Y}\left(\|A R X Y S A-A\|_{1}+\lambda\left(\|A R X\|_{1}+\|Y S A\|_{1}\right)\right)
\end{aligned}
$$

Finally, due to Lemma 18.3 .6 , we can find such $\widetilde{X} \widetilde{Y}$ with $\gamma=O\left(\sqrt{t_{1}+k} \sqrt{t_{2}+k}\right)$. Thus, we can conclude that $\|A R \widetilde{X} \widetilde{Y} S A-A\|_{1}+\lambda\left(\|A R \widetilde{X}\|_{1}+\|\widetilde{Y} S A\|_{1}\right) \leq \operatorname{poly}(k, \log n) \cdot$ OPT .

By taking union bound over all the success event of $S, R, T_{1}, T_{2}$, the algorithm will succeed with probability at least $9 / 10$.

### 18.7 Proof of Theorem 18.3.8

Proof. We choose all the matrices $S \in \mathbb{R}^{s \times n}, R \in \mathbb{R}^{n \times r}, T_{1} \in \mathbb{R}^{t_{1} \times n}, T_{2} \in \mathbb{R}^{n \times t_{2}}$ in Algorithm 18.1 as dense Cauchy transformation matrices (see Definition 18.2.1) which have the same scalar $\sigma=1$. We set $s=O(k \log k), r=O(k \log k), t_{1}=O(r \log r), t_{2}=O(s \log s)$.

Since computing $T_{1} A, A T_{2}$ can be done in $t_{1} \mathrm{nnz}(A), t_{2} \mathrm{nnz}(A)$ time respectiviely, and the size of $T_{1} A, A T_{2}$ is at most $n \cdot \operatorname{poly}(k)$, the total running time of computing

$$
T_{1} A R, S A T_{2}, T_{1} A T_{2}
$$

is at most $\widetilde{O}(k) \cdot \mathrm{nnz}(A)+n \cdot \operatorname{poly}(k)$. Since the size of the final small Frobenius norm ridge low-rank matrix approximation problem is $\operatorname{poly}(k)$, the running time of the whole algorithm is at most $\widetilde{O}(k) \cdot \operatorname{nnz}(A)+n \cdot \operatorname{poly}(k)$.

Now let us consider the correctness of the algorithm.
Let $U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{k \times n}$ satisfy $\left\|U^{*} V^{*}-A\right\|_{1}+\lambda\left(\left\|U^{*}\right\|_{1}+\left\|V^{*}\right\|_{1}\right)=$ OPT . Then due to Lemma 18.2.3, with probability at least $0.99, S$ satisfies

$$
\left\|S U^{*} x\right\|_{1} \geq \frac{1}{\alpha_{1}}\left\|U^{*} x\right\|_{1} \text { for all } x \in \mathbb{R}^{k} \text { and }\left\|S U^{*} V^{*}-S A\right\|_{1} \leq \beta_{1}\left\|U^{*} V^{*}-A\right\|_{1}
$$

where $\alpha_{1} \beta_{1}=O(\log n)$. Let $\widehat{V} \in \mathbb{R}^{k \times n}$ denote $\left[\begin{array}{c}S U^{*} \\ \lambda I\end{array}\right]^{\dagger} \cdot\left[\begin{array}{c}I \\ 0\end{array}\right] \cdot S A$. Let $\widetilde{U} \in \mathbb{R}^{n \times k}$ denote the optimal solution of $\min _{U \in \mathbb{R}^{n \times k}}\|U \widehat{V}-A\|_{1}+\lambda\left(\|U\|_{1}+\|\widehat{V}\|_{1}\right)$. Due to Lemma 18.2.3, with probability at least $0.99, R^{\top} \in \mathbb{R}^{r \times n}$ satisfies: $\left\|R^{\top} \widehat{V}^{\top} x\right\|_{1} \geq \frac{1}{\alpha_{2}}\left\|\widehat{V}^{\top} x\right\|_{1}$ for all $x \in$ $\mathbb{R}^{k}$, and $\left\|R^{\top} \widehat{V}^{\top} \widetilde{U}^{\top}-R^{\top} A^{\top}\right\|_{1} \leq \beta_{2}\left\|\widehat{V}^{\top} \widetilde{U}^{\top}-A^{\top}\right\|_{1}$ where $\alpha_{2} \beta_{2}=O(\log n)$. Then, by applying Corollary 18.3.2, we have that

$$
\min _{X \in \mathbb{R}^{r \times k}, Y \in \mathbb{R}^{k \times s}}\|A R X Y S A-A\|_{1}+\lambda\left(\|A R X\|_{1}+\|Y S A\|_{1}\right) \leq \widetilde{O}(k) \operatorname{poly}(\log n) \mathrm{OPT}
$$

Let $X^{*} \in \mathbb{R}^{r \times k}, Y^{*} \in \mathbb{R}^{k \times s}$ denote the optimal solution of $\min _{X, Y} \| A R X Y S A-$ $A \|_{1}+\lambda\left(\|A R X\|_{1}+\|Y S A\|_{1}\right)$. Due to Lemma 18.2.3, with probability at least $0.99, T_{1}$ satisfies that $\forall x \in \mathbb{R}^{r},\left\|T_{1} A R x\right\|_{1} \geq \frac{1}{\alpha_{1}^{\prime}}\|A R x\|_{1}$, and $\left\|T_{1}\left(\left[A R X^{*} Y^{*} S A-A \quad \lambda A R X^{*}\right]\right)\right\|_{1} \leq$ $\beta_{1}^{\prime}\left\|\left[\begin{array}{lll}A R X^{*} Y^{*} S A-A & \lambda A R X^{*}\end{array}\right]\right\|_{1}$ where $\alpha_{1}^{\prime} \beta_{1}^{\prime}=O(\log n)$. Due to Lemma 18.2.3, with probability at least $0.99, T_{2}$ satisfies that

$$
\forall x \in \mathbb{R}^{s},\left\|T_{2}^{\top} A^{\top} S^{\top} x\right\|_{1} \geq \frac{1}{\alpha_{2}^{\prime}}\left\|A^{\top} S^{\top} x\right\|_{1}
$$

and

$$
\begin{aligned}
& \left\|T_{2}^{\top}\left(\left[A^{\top} S^{\top}\left(Y^{*}\right)^{\top}\left(X^{*}\right)^{\top} R^{\top} A^{\top} T_{1}^{\top}-A^{\top} T_{1}^{\top} \quad \lambda A^{\top} S^{\top}\left(Y^{*}\right)^{\top}\right]\right)\right\|_{1} \\
\leq & \beta_{2}^{\prime}\left\|\left[A^{\top} S^{\top}\left(Y^{*}\right)^{\top}\left(X^{*}\right)^{\top} R^{\top} A^{\top} T_{1}^{\top}-A^{\top} T_{1}^{\top} \quad \lambda A^{\top} S^{\top}\left(Y^{*}\right)^{\top}\right]\right\|_{1}
\end{aligned}
$$

where $\alpha_{2}^{\prime} \beta_{2}^{\prime} \leq O(\log n)$. Thus, by Corollary 18.3.4, for any $\gamma \geq 1$, if $\widetilde{X}, \tilde{Y}$ can achieve

$$
\begin{aligned}
& \left\|T_{1} A R \tilde{X} \tilde{Y} S A T_{2}-T_{1} A T_{2}\right\|_{1}+\lambda\left(\left\|T_{1} A R \tilde{X}\right\|_{1}+\left\|\tilde{Y} S A T_{2}\right\|_{1}\right) \\
\leq & \gamma \min _{X, Y}\left(\left\|T_{1} A R X Y S A T_{2}-T_{1} A T_{2}\right\|_{1}+\lambda\left(\left\|T_{1} A R X\right\|_{1}+\left\|Y S A T_{2}\right\|_{1}\right)\right),
\end{aligned}
$$

then

$$
\begin{aligned}
& \|A R \widetilde{X} \widetilde{Y} S A-A\|_{1}+\lambda\left(\|A R \widetilde{X}\|_{1}+\|\widetilde{Y} S A\|_{1}\right) \\
\leq & \widetilde{O}(k) \cdot \operatorname{poly}(\log n) \gamma \min _{X, Y}\left(\|A R X Y S A-A\|_{1}+\lambda\left(\|A R X\|_{1}+\|Y S A\|_{1}\right)\right) .
\end{aligned}
$$

Finally, due to Lemma 18.3.6, we can find such $\tilde{X} \widetilde{Y}$ with $\gamma=O\left(\sqrt{t_{1}+k} \sqrt{t_{2}+k}\right)$. Thus, we can conclude that

$$
\|A R \widetilde{X} \widetilde{Y} S A-A\|_{1}+\lambda\left(\|A R \widetilde{X}\|_{1}+\|\widetilde{Y} S A\|_{1}\right) \leq \widetilde{O}\left(k^{2}\right) \operatorname{poly}(\log n) \cdot \mathrm{OPT} .
$$

By taking union bound over all the success event of $S, R, T_{1}, T_{2}$, the algorithm will succeed with probability at least $9 / 10$.

## Chapter 19

## Average Case L1 Low Rank Approximation

We study the column subset selection problem with respect to the entrywise $\ell_{1}$-norm loss. It is known that in the worst case, to obtain a good rank- $k$ approximation to a matrix, one needs an arbitrarily large $n^{\Omega(1)}$ number of columns to obtain a $(1+\epsilon)$-approximation to the best entrywise $\ell_{1}$-norm low rank approximation of an $n \times n$ matrix. Nevertheless, we show that under certain minimal and realistic distributional settings, it is possible to obtain a $(1+\epsilon)$-approximation with a nearly linear running time and poly $(k / \epsilon)+O(k \log n)$ columns. Namely, we show that if the input matrix $A$ has the form $A=B+E$, where $B$ is an arbitrary rank- $k$ matrix, and $E$ is a matrix with i.i.d. entries drawn from any distribution $\mu$ for which the $(1+\gamma)$-th moment exists, for an arbitrarily small constant $\gamma>0$, then it is possible to obtain a $(1+\epsilon)$-approximate column subset selection to the entrywise $\ell_{1}$-norm in nearly linear time. Conversely we show that if the first moment does not exist, then it is not possible to obtain a $(1+\epsilon)$-approximate subset selection algorithm even if one chooses any $n^{o(1)}$ columns. This is the first algorithm of any kind for achieving a $(1+\epsilon)$-approximation for entrywise $\ell_{1}$-norm loss low rank approximation.

### 19.1 Introduction

Numerical linear algebra algorithms are fundamental building blocks in many machine learning and data mining tasks. A well-studied problem is low rank matrix approximation. The most common version of the problem is also known as Principal Component Analysis (PCA), in which the goal is to find a low rank matrix to approximate a given matrix such that the Frobenius norm of the error is minimized. The optimal solution of this objective can be obtained via the singular value decomposition (SVD). Hence, the problem can be solved in polynomial time. If approximate solutions are allowed, then the running time can be made almost linear in the number of non-zero entries of the given matrix [Sar06, CW13, MM13, NN13a, BDN15, Coh16a].

An important variant of the PCA problem is the entrywise $\ell_{1}$-norm low rank matrix approximation problem. In this problem, instead of minimizing the Frobenius norm of the error, we seek to minimize the $\ell_{1}$-norm of the error. In particular, given an $n \times n$ input matrix $A$, and a rank parameter $k$, we want to find a matrix $B$ with rank at most $k$ such that $\|A-B\|_{1}$ is minimized, where for a matrix $C,\|C\|_{1}$ is defined to be $\sum_{i, j}\left|C_{i, j}\right|$. There are several reasons for using the $\ell_{1}$-norm as the error measure. For example, solutions with respect to the $\ell_{1}$-norm loss are usually more robust than solutions with Frobenius norm loss [Hub64, CLMW11]. Further, the $\ell_{1}$-norm loss is often used as a relaxation of the $\ell_{0}$-loss, which has wide applications including sparse recovery, matrix completion, and robust PCA; see e.g., [XCS10, CLMW11]. Although a number of algorithms have been proposed for the $\ell_{1}$-norm loss $\left[\mathrm{KK} 03, \mathrm{KK} 05, \mathrm{KLC}^{+}\right.$15, Kwa08, ZLS ${ }^{+}$12, BJ12, BD13, BDB13, MXZZ13, MKP13, MKP14, MKCP16, PK16], the problem is known to be NP-hard [GV15]. The first $\ell_{1}$-low rank approximation with provable guarantees was proposed by [SWZ17]. To cope
with NP-hardness, the authors gave a solution with a $\operatorname{poly}(k \log n)$-approximation ratio, i.e., their algorithm outputs a rank- $k$ matrix $B^{\prime} \in \mathbb{R}^{n \times n}$ for which

$$
\left\|A-B^{\prime}\right\|_{1} \leq \alpha \cdot \min _{\text {rank }-k}\|A-B\|_{1}
$$

for $\alpha=\operatorname{poly}(k \log n)$. The approximation ratio $\alpha$ was further improved to $O(k \log k)$ by allowing $B^{\prime}$ to have a slightly larger $k^{\prime}=O(k \log n)$ rank [CGK ${ }^{+} 17$ a]. Such $B^{\prime}$ with larger rank is referred to as a bicriteria solution. However, in high precision applications, such approximation factors are too large. A natural question is if one can compute a $(1+\epsilon)$ approximate solution efficiently for $\ell_{1}$-norm low rank approximation. In fact, a $(1+\epsilon)$ approximation algorithm was given in $\left[\mathrm{BBB}^{+} 19 \mathrm{a}\right]$, but the running time of their algorithm is a prohibitive $n^{\operatorname{poly}(k / \epsilon)}$. Unfortunately, $\left[\mathrm{BBB}^{+} 19 \mathrm{a}\right]$ shows in the worst case that a $2^{k^{\Omega(1)}}$ running time is necessary for any constant approximation given a standard conjecture in complexity theory.

### 19.1.1 Notation

To describe our results, let us first introduce some notation. We will use $[n]$ to denote the set $\{1,2, \cdots, n\}$. We use $A_{i}$ to denote the $i^{\text {th }}$ column of $A$. We use $A^{j}$ to denote the $j^{\text {th }}$ row of $A$. Let $Q \subseteq[n]$. We use $A_{Q}$ to denote the matrix which is comprised of the columns of $A$ with column indices in $Q$. Similarly, we use $A^{Q}$ to denote the matrix which is comprised of the rows of $A$ with row indices in $Q$. We use $\binom{[n]}{t}$ to denote the set of all the size- $t$ subsets of $[n]$. Let $\|A\|_{F}$ denote the Frobenius norm of a matrix $A$, i.e., $\|A\|_{F}$ is the square root of the sum of squares of all the entries in $A$. For $1 \leq p<2$, we use $\|A\|_{p}$ to denote the entry-wise $\ell_{p}$-norm of a matrix $A$, i.e., $\|A\|_{p}$ is the $p$-th root of the sum of $p$-th powers of the absolute values of the entries of $A .\|A\|_{1}$ is an important special case of $\|A\|_{p}$,
which corresponds to the sum of absolute values of the entries in $A$. A random variable $X$ has the Cauchy distribution if its probability density function is $f(z)=\frac{1}{\pi\left(1+z^{2}\right)}$.

### 19.1.2 Our Results

We propose an efficient bicriteria $(1+\epsilon)$-approximate column subset selection algorithm for the $\ell_{1}$-norm. We bypass the running time lower bound mentioned above by making a mild assumption on the input data, and also show that our assumption is necessary in a certain sense.

Our main algorithmic result is described as follows.
Theorem 19.1.1 (Informal version of Theorem 19.2.9). Suppose we are given a matrix $A=$ $A^{*}+\Delta \in \mathbb{R}^{n \times n}$, where $\operatorname{rank}\left(A^{*}\right)=k$ for $k=n^{o(1)}$, and $\Delta$ is a random matrix for which the $\Delta_{i, j}$ are i.i.d. symmetric random variables with $\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]=O\left(\mathbb{E}\left[\left|\Delta_{i, j}\right|\right]^{p}\right)$ for some constant $p>1$. Let $\epsilon \in(0,1 / 2)$ satisfy $1 / \epsilon=n^{o(1)}$. There is an $\widetilde{O}\left(n^{2}+n \operatorname{poly}(k / \epsilon)\right)^{1}$ time algorithm (Algorithm 19.1) which can output a subset $S \subseteq[n]$ with $|S| \leq \operatorname{poly}(k / \epsilon)+O(k \log n)$ for which

$$
\min _{X \in \mathbb{R}^{|S| \times n}}\left\|A_{S} X-A\right\|_{1} \leq(1+\epsilon)\|\Delta\|_{1},
$$

holds with probability at least 99/100.

Note the running time in Theorem 19.1.1 is nearly linear in the number of nonzero entries of $A$, since for an $n \times n$ matrix with i.i.d. noise drawn from any continuous distribution, the number of non-zero entries of $A$ will be $n^{2}$ with probability 1 . We also show the moment assumption of Theorem 19.1.1 is necessary in the following precise sense.

[^65]Theorem 19.1.2 (Hardness, informal version of Theorem 19.4.17). Let $n>0$ be sufficiently large. Let $A=\eta \cdot \mathbf{1} \cdot \mathbf{1}^{\top}+\Delta \in \mathbb{R}^{n \times n}$ be a random matrix where $\eta=n^{c_{0}}$ for some sufficiently large constant $c_{0}, \mathbf{1} \in \mathbb{R}^{n}$ is the all-ones vector, and $\forall i, j \in[n], \Delta_{i, j} \sim C(0,1)$ are i.i.d. standard Cauchy random variables. Let $r=n^{o(1)}$. Then with probability at least $1-O(1 / \log \log n), \forall S \subseteq[n]$ with $|S|=r$,

$$
\min _{X \in \mathbb{R}^{r} \times n}\left\|A_{S} X-A\right\|_{1} \geq 1.002\|\Delta\|_{1}
$$

### 19.1.3 Our Techniques

For an overview of our hardness result, we refer readers to the supplementary material, namely, Appendix 19.4. In the following, we will outline the main techniques used in our algorithm.
$(1+\epsilon)$-Approximate $\ell_{1}$-Low Rank Approximation. We make the following distributional assumption on the input matrix $A \in \mathbb{R}^{n \times n}$ : namely, $A=A^{*}+\Delta$ where $A^{*}$ is an arbitrary rank- $k$ matrix and the entries of $\Delta$ are i.i.d. from any symmetric distribution with $\mathbb{E}\left[\left|\Delta_{i, j}\right|\right]=1$ and $\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]=O(1)$ for any real number $p$ strictly greater than 1 , e.g., $p=1.000001$ would suffice. Note that such an assumption is mild compared to typical noise models which require the noise be Gaussian or have bounded variance; in our case the random variables may even be heavy-tailed with infinite variance. In this setting we show it is possible to obtain a subset of poly $\left(k\left(\epsilon^{-1}+\log n\right)\right)$ columns spanning a $(1+\epsilon)$-approximation. This provably overcomes the column subset selection lower bound of [SWZ17] which shows for entrywise $\ell_{1}$-low rank approximation that there are matrices for which any subset of poly $(k)$ columns spans at best a $k^{\Omega(1)}$-approximation.

Consider the following algorithm: sample poly $(k / \epsilon)$ columns of $A$, and try to cover as many of the remaining columns as possible. Here, by covering a column $i$, we mean that if $A_{I}$ is the subset of columns sampled, then $\min _{y}\left\|A_{I} y-A_{i}\right\|_{1} \leq(1+O(\epsilon)) n$. The reason for this notion of covering is that we are able to show in Lemma 19.2.1 that in this noise model, $\|\Delta\|_{1} \geq(1-\epsilon) n^{2}$ w.h.p., and so if we could cover every column $i$, our overall cost would be $(1+O(\epsilon)) n^{2}$, which would give a $(1+O(\epsilon))$-approximation to the overall cost.

We will not be able to cover all columns, unfortunately, with our initial sample of $\operatorname{poly}(k / \epsilon)$ columns of $A$. Instead, though, we will show that we will be able to cover all but a set $T$ of $\epsilon n /(k \log k)$ of the columns. Fortunately, we show in Lemma 19.2.4 another property of the noise matrix $\Delta$ is that all subsets $S$ of columns of size at most $n / r$, for $r \geq(1 / \gamma)^{1+1 /(p-1)}$ satisfy $\sum_{j \in S}\left\|\Delta_{j}\right\|_{1}=O\left(\gamma n^{2}\right)$. Thus, for the above set $T$ that we do not cover, we can apply this lemma to it with $\gamma=\epsilon /(k \log k)$, and then we know that $\sum_{j \in T}\left\|\Delta_{j}\right\|_{1}=O\left(\epsilon n^{2} /(k \log k)\right)$, which then enables us to run a previous $\widetilde{O}(k)$-approximate $\ell_{1}$ low rank approximation algorithm $\left[\mathrm{CGK}^{+} 17 \mathrm{a}\right]$ on the set $T$, which will only incur total cost $O\left(\epsilon n^{2}\right)$, and since by Lemma 19.2.1 above the overall cost is at least $(1-\epsilon) n^{2}$, we can still obtain a $(1+O(\epsilon))$-approximation overall.

The main missing piece of the algorithm to describe is why we are able to cover all but a small fraction of the columns. One thing to note is that our noise distribution may not have a finite variance, and consequently, there can be very large entries $\Delta_{i, j}$ in some columns. In Lemma 19.2.3, we show the number of columns in $\Delta$ for which there exists an entry larger than $n^{1 / 2+1 /(2 p)}$ in magnitude is $O\left(n^{(2-p) / 2}\right)$, which since $p>1$ is a constant bounded away from 1, is sublinear. Let us call this set with entries larger than $n^{1 / 2+1 /(2 p)}$ in magnitude the set $H$ of "heavy" columns; we will not make any guarantees about $H$,
rather, we will stuff it into the small set $T$ of columns above on which we will run our earlier $O(k \log k)$-approximation.

For the remaining, non-heavy columns, which constitute almost all of our columns, we show in Lemma 19.2.5 that $\left\|\Delta_{i}\right\|_{1} \leq(1+\epsilon) n$ w.h.p. The reason this is important is that recall to cover some column $i$ by a sample set $I$ of columns, we need $\min _{y}\left\|A_{I} y-A_{i}\right\|_{1} \leq(1+O(\epsilon)) n$. It turns out, as we now explain, that we will get

$$
\min _{y}\left\|A_{I} y-A_{i}\right\|_{1} \leq\left\|\Delta_{i}\right\|_{1}+e_{i}
$$

where $e_{i}$ is a quantity which we can control and make $O(\epsilon n)$ by increasing our sample size $I$. Consequently, since $\left\|\Delta_{i}\right\|_{1} \leq(1+\epsilon) n$, overall we will have $\min _{y}\left\|A_{I} y-A_{i}\right\|_{1} \leq(1+O(\epsilon)) n$, which means that $i$ will be covered. We now explain what $e_{i}$ is, and why $\min _{y}\left\|A_{I} y-A_{i}\right\|_{1} \leq$ $\left\|\Delta_{i}\right\|_{1}+e_{i}$.

Towards this end, we first explain a key insight in this model. Since the $p$-th moment exists for some real number $p>1$ (e.g., $p=1.000001$ suffices), averaging helps reduce the noise of fitting a column $A_{i}$ by subsets of other columns. Namely, we show in Lemma 19.2.2 that for any $t$ non-heavy column $\Delta_{i_{1}}, \ldots, \Delta_{i_{t}}$ of $\Delta$, and any coefficients $\alpha_{1}, \alpha_{2}, \ldots, \alpha_{t} \in$ $[-1,1],\left\|\sum_{j=1}^{t} \alpha_{j} \Delta_{i_{j}}\right\|_{1}=O\left(t^{1 / p} n\right)$, that is, since the individual coordinates of the $\Delta_{i_{j}}$ are zero-mean random variables, their sum concentrates as we add up more columns. We do not need bounded variance for this property.

How can we use this averaging property for subset selection? The idea is, instead of sampling a single subset $I$ of $O(k)$ columns and trying to cover each remaining column with this subset as shown in $\left[\mathrm{CGK}^{+} 17 \mathrm{a}\right]$, we will sample multiple independent subsets $I_{1}, I_{2}, \ldots, I_{t}$. By a similar argument of $\left[\mathrm{CGK}^{+} 17 \mathrm{a}\right]$, for any given column index $i \in[n]$, for most of these
subset $I_{j}$, we have that $A_{i}^{*} /\left\|\Delta_{i}\right\|_{1}$ can be expressed as a linear combination of columns $A_{\ell}^{*} /\left\|\Delta_{\ell}\right\|_{1}, \ell \in I_{j}$, via coefficients of absolute value at most 1 . Note that this is only true for most $i$ and most $j$; we develop terminology for this in Definitions 19.2.1, 19.2.2, 19.2.3, and 19.2.4, referring to what we call a good core. We quantify what we mean by most $i$ and most $j$ having this property in Lemma 19.2.7 and Lemma 19.2.8.

The key though, that drives the analysis, is Lemma 19.2.6, which shows that $\min _{y} \| A_{i} y-$ $A_{i}\left\|_{1} \leq\right\| \Delta_{i} \|_{1}+e_{i}$, where $e_{i}=O\left(q^{1 / p} / t^{1-1 / p} n\right)$, where $q$ is the size of each $I_{j}$, and $t$ is the number of different $I_{j}$. We need $q$ to be at least $k$, just as before, so that we can be guaranteed that when we adjoin a column index $i$ to $I_{j}$, there is some positive probability that $A_{i}^{*} /\left\|\Delta_{i}\right\|_{1}$ can be expressed as a linear combination of columns $A_{\ell}^{*} /\left\|\Delta_{\ell}\right\|_{1}, \ell \in I_{j}$, with coefficients of absolute value at most 1 . What is different in our noise model though is the division by $t^{1-1 / p}$. Since $p>1$, if we set $t$ to be a large enough $\operatorname{poly}(k / \epsilon)$, then $e_{i}=O(\epsilon n)$, and then we will have covered $A_{i}$, as desired. This captures the main property that averaging the linear combinations for expression $A_{i}^{*} /\left\|\Delta_{i}\right\|_{1}$ using different subsets $I_{j}$ gives us better and better approximations to $A_{i}^{*} /\left\|\Delta_{i}\right\|_{1}$. Of course we need to ensure several properties such as not sampling a heavy column (the averaging in Lemma 19.2.2 does not apply when this happens), we need to ensure most of the $I_{j}$ have small-coefficient linear combinations expressing $A_{i}^{*} /\left\|\Delta_{i}\right\|_{1}$, etc. This is handled in our main theorem, Theorem 19.2.9.

## $19.2 \ell_{1}$-Norm Column Subset Selection

We first present two subroutines.
Linear regression with $\ell_{1}$ loss. The first subroutine needed is an approximate $\ell_{1}$ linear regression solver. In particular, given a matrix $M \in \mathbb{R}^{n \times d}$, $n$ vectors $b_{1}, b_{2}, \cdots, b_{n} \in \mathbb{R}^{n}$, and an error parameter $\epsilon \in(0,1)$, we want to compute $x_{1}, x_{2}, \cdots, x_{n} \in \mathbb{R}^{d}$ for which $\forall i \in[n]$, we have

$$
\left\|M x_{i}-b_{i}\right\|_{1} \leq(1+\epsilon) \cdot \min _{x \in \mathbb{R}^{d}}\left\|M x-b_{i}\right\|_{1}
$$

Furthermore, we also need an estimate $v_{i}$ of the regression cost $\left\|M x_{i}-b_{i}\right\|_{1}$ for each $i \in[n]$ such that $\left\|M x_{i}-b_{i}\right\|_{1} \leq v_{i} \leq(1+\epsilon)\left\|M x_{i}-b_{i}\right\|_{1}$. Such an $\ell_{1}$-regression problem can be solved efficiently (see [Woo14b] for a survey). The total running time to solve these $n$ regression problems simultaneously is at most $\widetilde{O}\left(n^{2}\right)+n \cdot \operatorname{poly}(d \log n)$, and the success probability is at least 0.999.
$\ell_{1}$ Column subset selection for general matrices. The second subroutine needed is an $\ell_{1}$-low rank approximation solver for general input matrices, though we allow a large approximation ratio. We use the algorithm proposed by [CGK+ 17 a] for this purpose. In particular, given an $n \times d(d \leq n)$ matrix $M$ and a rank parameter $k$, the algorithm can output a small set $S \subset[n]$ with size at most $O(k \log n)$, such that

$$
\min _{X \in \mathbb{R}^{|S| \times d}}\left\|M_{S} X-M\right\|_{1} \leq O(k \log k) \cdot \min _{\text {rank }-k}\|M-B\|_{1} .
$$

Furthermore, the running time is at most $\widetilde{O}\left(n^{2}\right)+n \cdot \operatorname{poly}(k \log n)$, and the success probability is at least 0.999. Now we can present our algorithm, Algorithm 19.1.

```
Algorithm \(19.1 \ell_{1}\)-Low Rank Approximation with Input Assumption
    procedure L1NoisyLowRAnkApprox \(\left(A \in \mathbb{R}^{n \times n}, k, \epsilon\right) \quad \triangleright\) Theorem 19.2.9
        Sample a set \(I\) from \(\binom{[n]}{s}\) uniformly at random, where \(s=\operatorname{poly}(k / \epsilon)\).
        Solve the approximate \(\ell_{1}\)-regression problem \(\min _{x \in \mathbb{R}^{I I I}}\left\|A_{I} x-A_{i}\right\|_{1}\) for each \(i \in[n]\),
    and let \(v_{i}\) be the estimated regression cost.
        Compute the set \(T=\left\{i \in[n] \mid v_{i}\right.\) is one of the top \(l\) largest values among \(\left.v_{1}, v_{2}, \cdots, v_{n}\right\}\),
    where \(l=n / \operatorname{poly}(k / \epsilon)\).
        Solve \(\ell_{1}\)-column subset selection for \(A_{T}\). Let the solution be \(A_{Q}\).
        Solve the approximate \(\ell_{1}\)-regression problem \(\min _{X \in \mathbb{R}^{(|I|+|Q|) \times n}}\left\|A_{(I \cup S)} X-A\right\|_{1}\), and let
    \(\widehat{X}\) be the solution. Return \(A_{(I \cup S)}\) and \(\widehat{X} . \triangleright A_{(I \cup S)} \widehat{X}\) is a good low rank approximation
    to \(A\)
    end procedure
```

Running time. Uniformly sampling a set $I$ can be done in poly $(k / \epsilon)$ time. According to our $\ell_{1}$-regression subroutine, solving $\min _{x}\left\|A_{I} x-A_{i}\right\|_{1}$ for all $i \in[n]$ can be finished in $\widetilde{O}\left(n^{2}\right)+n \cdot \operatorname{poly}(k \log (n) / \epsilon)$ time. We only need sorting to compute the set $T$ which takes $O(n \log n)$ time. By our second subroutine, the $\ell_{1}$-column subset selection for $A_{T}$ will take $\widetilde{O}\left(n^{2}\right)+n \cdot \operatorname{poly}(k \log n)$. The last step only needs an $\ell_{1}$-regression solver, which takes $\widetilde{O}\left(n^{2}\right)+n \cdot \operatorname{poly}(k \log (n) / \epsilon)$ time. Thus, the overall running time is $\widetilde{O}\left(n^{2}\right)+n \cdot \operatorname{poly}(k \log (n) / \epsilon)$.

The remaining parts in this section will focus on analyzing the correctness of the algorithm.

### 19.2.1 Properties of the Noise Matrix

Recall that the input matrix $A \in \mathbb{R}^{n \times n}$ can be decomposed as $A^{*}+\Delta$, where $A^{*}$ is the ground truth, and $\Delta$ is a random noise matrix. In particular, $A^{*}$ is an arbitrary rank- $k$ matrix, and $\Delta$ is a random matrix where each entry is an i.i.d. sample drawn from an unknown symmetric distribution. The only assumption on $\Delta$ is that each entry $\Delta_{i, j}$
satisfies $\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]=O\left(\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]\right)$ for some constant $p>1$, i.e., the $p$-th moment of the noise distribution is bounded. Without loss of generality, we will suppose $\mathbb{E}\left[\left|\Delta_{i, j}\right|\right]=1$, $\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]=O(1)$, and $p \in(1,2)$ throughout the paper. In this section, we will present some key properties of the noise matrix.

The following lemma provides a lower bound on $\|\Delta\|_{1}$. Once we have the such lower bound, we can focus on finding a solution for which the approximation cost is at most that lower bound.

Lemma 19.2.1 (Lower bound on the noise matrix). Let $\Delta \in \mathbb{R}^{n \times n}$ be a random matrix where $\Delta_{i, j}$ are i.i.d. samples drawn from a symmetric distribution. Suppose $\mathbb{E}\left[\left|\Delta_{i, j}\right|\right]=1$ and $\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]=O(1)$ for some constant $p \in(1,2)$. Then, $\forall \epsilon \in(0,1)$ which satisfies $1 / \epsilon=n^{o(1)}$, we have

$$
\operatorname{Pr}\left[\|\Delta\|_{1} \geq(1-\epsilon) n^{2}\right] \geq 1-e^{-\Theta(n)}
$$

The next lemma shows the main reason why we are able to get a small fitting cost when running regression. Consider a toy example. Suppose we have a target number $a \in \mathbb{R}$, and another $t$ numbers $a+g_{1}, a+g_{2}, \cdots, a+g_{t} \in \mathbb{R}$, where $g_{i}$ are i.i.d. samples drawn from the standard Gaussian distribution $N(0,1)$. If we use $a+g_{i}$ to fit $a$, then the expected cost is $\mathbb{E}\left[\left|a+g_{i}-a\right|\right]=\mathbb{E}\left[\left|g_{i}\right|\right]=\sqrt{2 / \pi}$. However, if we use the average of $a+g_{1}, a+g_{2}, \cdots, a+g_{t}$ to fit $a$, then the expected cost is $\mathbb{E}\left[\left|\sum_{i=1}^{t} g_{i}\right| / t\right]$. Since the $g_{i}$ are independent, $\sum_{i=1}^{t} g_{i}$ is a random Gaussian variable with variance $t$, which means that the above expected cost is $\sqrt{2 / \pi} / \sqrt{t}$. Thus the fitting cost is reduced by a factor $\sqrt{t}$. By generalizing the above argument, we obtain the following lemma.

Lemma 19.2.2 (Averaging reduces the noise). Let $\Delta_{1}, \Delta_{2}, \cdots, \Delta_{t} \in \mathbb{R}^{n}$ be $t$ random vectors. The $\Delta_{i, j}$ are i.i.d. symmetric random variables with $\mathbb{E}\left[\left|\Delta_{i, j}\right|\right]=1$ and $\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]=O(1)$ for some constant $p \in(1,2)$. Let $\alpha_{1}, \alpha_{2}, \cdots, \alpha_{t} \in[-1,1]$ be $t$ real numbers. Conditioned on $\forall i \in[n], j \in[t],\left|\Delta_{i, j}\right| \leq n^{1 / 2+1 /(2 p)}$, with probability at least $1-2^{-n^{\Theta(1)}}$,

$$
\left\|\sum_{i=1}^{t} \alpha_{i} \Delta_{i}\right\|_{1} \leq O\left(t^{1 / p} n\right)
$$

The above lemma needs a condition that each entry in the noise column should not be too large. Fortunately, we can show that most of the (noise) columns do not have any large entry.

Lemma 19.2.3 (Only a small number of columns have large entries). Let $\Delta \in \mathbb{R}^{n \times n}$ be $a$ random matrix where the $\Delta_{i, j}$ are i.i.d. symmetric random variables with $\mathbb{E}\left[\left|\Delta_{i, j}\right|\right]=1$ and $\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]=O(1)$ for some constant $p \in(1,2)$. Let

$$
H=\left\{j \in[n]\left|\exists i \in[n],\left|\Delta_{i, j}\right|>n^{1 / 2+1 /(2 p)}\right\} .\right.
$$

Then with probability at least 0.999 ,

$$
|H| \leq O\left(n^{1-(p-1) / 2}\right)
$$

The following lemma shows that any small subset of the columns of the noise matrix $\Delta$ cannot contribute too much to the overall error. By combining with the previous lemma, the entrywise $\ell_{1}$ cost of all columns containing large entries can be bounded.

Lemma 19.2.4. Let $\Delta \in \mathbb{R}^{n \times n}$ be a random matrix where $\Delta_{i, j}$ are i.i.d. symmetric random variables with $\mathbb{E}\left[\left|\Delta_{i, j}\right|\right]=1$ and $\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]=O(1)$ for some constant $p \in(1,2)$. Let $\epsilon \in(0,1)$
satisfy $1 / \epsilon=n^{o(1)}$. Let $r \geq(1 / \epsilon)^{1+1 /(p-1)}$. Then, with probability at least $.999, \forall S \subset[n]$ with $|S| \leq n / r$,

$$
\sum_{j \in S}\left\|\Delta_{j}\right\|_{1}=O\left(\epsilon n^{2}\right)
$$

We say a (noise) column is good if it does not have a large entry. We can show that, with high probability, the entry-wise $\ell_{1}$ cost of a good (noise) column is small.

Lemma 19.2.5 (Cost of good noise columns). Let $\Delta \in \mathbb{R}^{n}$ be a random vector where $\Delta_{i}$ are i.i.d. symmetric random variables with $\mathbb{E}\left[\left|\Delta_{i}\right|\right]=1$ and $\mathbb{E}\left[\left|\Delta_{i}\right|^{p}\right]=O(1)$ for some constant $p \in(1,2)$. Let $\epsilon \in(0,1)$ satisfy $1 / \epsilon=n^{o(1)}$. If $\forall i \in[n],\left|\Delta_{i}\right| \leq n^{1 / 2+1 /(2 p)}$, then with probability at least $1-2^{-n^{\Theta(1)}}$,

$$
\|\Delta\|_{1} \leq(1+\epsilon) n
$$

### 19.2.2 Definition of Tuples and Cores

In this section, we provide some basic definitions, e.g., of a tuple, a good tuple, the core of a tuple, and a coefficients tuple. These definitions will be heavily used later when we analyze the correctness of our algorithm.

Before we present the definitions, we introduce a notion $R_{A^{*}}(S)$. Given a matrix $A^{*} \in \mathbb{R}^{n_{1} \times n_{2}}$, for a set $S \subseteq\left[n_{2}\right]$, we define

$$
R_{A^{*}}(S):=\arg \max _{P: P \subseteq S}\left\{\left|\operatorname{det}\left(\left(A^{*}\right)_{P}^{Q}\right)\right|| | P\left|=|Q|=\operatorname{rank}\left(A_{S}^{*}\right), Q \subseteq\left[n_{1}\right]\right\}\right.
$$

where for a squared matrix $C$, $\operatorname{det}(C)$ denotes the determinant of $C$. Roughly speaking, by Cramer's rule, if we use the columns of $A^{*}$ with index in the set $R_{A^{*}}(S)$ to fit any column
of $A^{*}$ with index in the set $S$, the absolute value of any fitting coefficient will be at most 1 . Small fitting coefficients are good since they will not increase the noise by too much. For example, suppose $A_{i}^{*}=A_{S}^{*} x$ and $\|x\|_{\infty} \leq 1$, i.e., the $i$-th column can be fit by the columns with indices in the set $S$ and the fitting coefficients $x \in \mathbb{R}^{|S|}$ are small. If we use the noisy columns of $A_{S}^{*}+\Delta_{S}$ to fit the noisy column $A_{i}^{*}+\Delta_{i}$, then the fitting cost is at most

$$
\left\|\left(A_{S}^{*}+\Delta_{S}\right) x-\left(A_{i}^{*}+\Delta_{i}\right)\right\|_{1} \leq\left\|\Delta_{i}\right\|_{1}+\left\|\Delta_{S} x\right\|_{1} .
$$

Since $\|x\|_{\infty} \leq 1$, it is possible to give a good upper bound for $\left\|\Delta_{S} x\right\|_{1}$.

Definition 19.2.1 (Tuple). A ( $q, t, n$ )-tuple is defined to be $\left(S_{1}, S_{2}, \cdots, S_{t}, i\right)$, where $\forall j \in$ $[t], S_{j} \subset[n]$ with $\left|S_{j}\right|=q$. Let $S=\bigcup_{j=1}^{t} S_{j}$. Then $|S|=q t$, i.e., $S_{1}, S_{2}, \cdots, S_{t}$ are disjoint. Furthermore, $i \in[n]$ and $i \notin S$. For simplicity, we use ( $S_{[t]}, i$ ) to denote ( $S_{1}, S_{2}, \cdots, S_{t}, i$ ).

We next provide the definition of a good tuple.

Definition 19.2.2 (Good tuple). Given a rank- $k$ matrix $A^{*} \in \mathbb{R}^{n \times n}$, an $\left(A^{*}, q, t, \alpha\right)$-good tuple is a $(q, t, n)$-tuple $\left(S_{[t]}, i\right)$ which satisfies

$$
\left|\left\{j \in[t] \mid i \notin R_{A^{*}}\left(S_{j} \cup\{i\}\right)\right\}\right| \geq \alpha \cdot t
$$

We need the definition of the core of a tuple.

Definition 19.2.3 (Core of a tuple). The core of $\left(S_{[t]}, i\right)$ is defined to be the set

$$
\left\{j \in[t] \mid i \notin R_{A^{*}}\left(S_{j} \cup\{i\}\right)\right\}
$$

We define a coefficients tuple as follows.

Definition 19.2.4 (Coefficients tuple). Given a rank- $k$ matrix $A^{*} \in \mathbb{R}^{n \times n}$, let $\left(S_{[t]}, i\right)$ be an ( $\left.A^{*}, q, t, \alpha\right)$-good tuple. Let $C$ be the core of $\left(S_{[t]}, i\right)$. A coefficients tuple corresponding to $\left(S_{[t]}, i\right)$ is defined to be $\left(x_{1}, x_{2}, \cdots, x_{t}\right)$ where $\forall j \in[t], x_{j} \in \mathbb{R}^{q}$. The vector $x_{j} \in \mathbb{R}^{q}$ satisfies: $x_{j}=0$ if $j \in[t] \backslash C$, while $A_{S_{j}}^{*} x_{j}=A_{i}^{*}$ and $\left\|x_{j}\right\|_{\infty} \leq 1$, if $j \in C$. To guarantee the coefficients tuple is unique, we restrict each vector $x_{j} \in \mathbb{R}^{q}$ to be one that has the minimum lexicographic order.

### 19.2.3 Properties of a Good Tuple and a Coefficients Tuple

Consider a good tuple $\left(S_{1}, S_{2}, \cdots, S_{t}, i\right)$. By the definition of a good tuple, the size of the core $C$ of the tuple is large. For each $j \in C$, the coefficients $x_{j}$ of using $A_{S_{j}}^{*}$ to fit $A_{i}^{*}$ should have absolute value at most 1. Now consider the noisy setting. As discussed in the previous section, using $A_{S_{j}}$ to fit $A_{i}$ has cost at most $\left\|\Delta_{i}\right\|_{1}+\left\|\Delta_{S_{j}} x_{j}\right\|_{1}$. Although $\left\|\Delta_{S_{j}} x_{j}\right\|_{1}$ has a good upper bound, it is not small enough. To further reduce the $\ell_{1}$ fitting cost, we can now apply the averaging argument (Lemma 19.2.2) over all the fitting choices corresponding to $C$. Formally, we have the following lemma.

Lemma 19.2.6 (Good tuples imply low fitting cost). Suppose we are given a matrix $A \in$ $\mathbb{R}^{n \times n}$ which satisfies $A=A^{*}+\Delta$, where $A^{*} \in \mathbb{R}^{n \times n}$ has rank $k$. Here $\Delta \in \mathbb{R}^{n \times n}$ is a random matrix where $\Delta_{i, j}$ are i.i.d. symmetric random variables with $\mathbb{E}\left[\left|\Delta_{i, j}\right|\right]=1$ and $\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]=O(1)$ for some constant $p \in(1,2)$. Let $H \subset[n]$ be defined as follows:

$$
H=\left\{j \in[n]\left|\exists i \in[n],\left|\Delta_{i, j}\right|>n^{1 / 2+1 /(2 p)}\right\} .\right.
$$

Let $q, t \leq n^{o(1)}$. Then, with probability at least $1-2^{-n^{\ominus(1)}}$, for all $\left(A^{*}, q, t, 1 / 2\right)$-good tuples
$\left(S_{1}, S_{2}, \cdots, S_{t}, i\right)$ which satisfy $H \cap\left(\bigcup_{j=1}^{t} S_{j}\right)=\emptyset$, we have

$$
\min _{y \in \mathbb{R}^{q t}}\left\|A_{\left\{\cup_{j=1}^{t} S_{j}\right\}} y-A_{i}\right\|_{1} \leq\left\|\frac{1}{|C|} \sum_{j=1}^{t} A_{S_{j}} x_{j}-A_{i}\right\|_{1} \leq\left\|\Delta_{i}\right\|_{1}+O\left(q^{1 / p} / t^{1-1 / p} n\right),
$$

where $C$ is the core of $\left(S_{1}, S_{2}, \cdots, S_{t}, i\right)$, and $\left(x_{1}, x_{2}, \cdots, x_{t}\right)$ is the coefficients tuple corresponding to $\left(S_{1}, S_{2}, \cdots, S_{t}, i\right)$.

We next show that if we choose columns randomly, it is easy to find a good tuple.
Lemma 19.2.7. Given a rank-k matrix $A^{*} \in \mathbb{R}^{n \times n}$, let $q>10 k, t>0$. Let $I=\left\{i_{1}, i_{2}, \cdots, i_{q t+1}\right\}$ be a subset drawn uniformly at random from $\binom{[n]}{q t+1}$. Let $\pi: I \rightarrow I$ be a random permutation of $q t+1$ elements. $\forall j \in[t]$, let

$$
S_{j}=\left\{i_{\pi((j-1) q+1)}, i_{\pi((j-1) q+2)}, \cdots, i_{\pi((j-1) q+q)}\right\} .
$$

We use $i$ to denote $i_{\pi(q t+1)}$. With probability $\geq 1-2 k / q,\left(S_{1}, S_{2}, \cdots, S_{t}, i\right)$ is an $\left(A^{*}, q, t, 1 / 2\right)$-good tuple.

Lemma 19.2.7 implies that if we randomly choose $S_{1}, S_{2}, \cdots, S_{t}$, then with high probability, there are many choices of $i \in[n]$, such that $\left(S_{1}, S_{2}, \cdots, S_{t}, i\right)$ is a good tuple. Precisely, we can show the following.

Lemma 19.2.8. Given a rank-k matrix $A^{*} \in \mathbb{R}^{n \times n}$, let $q>10 k, t>0$. Let $I=\left\{i_{1}, i_{2}, \cdots, i_{q t}\right\}$ be a random subset uniformly drawn from $\binom{[n]}{q t}$. Let $\pi$ be a random permutation of $q t$ elements. $\forall j \in[t]$, we define $S_{j}$ as follows:

$$
S_{j}=\left\{i_{\pi((j-1) q+1)}, i_{\pi((j-1) q+2)}, \cdots, i_{\pi((j-1) q+q)}\right\} .
$$

Then with probability at least $2 k / q$,

$$
\mid\left\{i \in[n] \backslash I \mid\left(S_{1}, S_{2}, \cdots, S_{t}, i\right) \text { is an }\left(A^{*}, q, t, 1 / 2\right) \text {-good tuple }\right\} \mid \geq(1-4 k / q)(n-q t) .
$$

### 19.2.4 Main Result

Now we are able to put all ingredients together to prove our main theorem, Theorem 19.2.9.

Theorem 19.2.9 (Formal version of Theorem 19.1.1). Suppose we are given a matrix $A=$ $A^{*}+\Delta \in \mathbb{R}^{n \times n}$, where $\operatorname{rank}\left(A^{*}\right)=k$ for $k=n^{o(1)}$, and $\Delta$ is a random matrix for which the $\Delta_{i, j}$ are i.i.d. symmetric random variables with $\mathbb{E}\left[\left|\Delta_{i, j}\right|\right]=1$ and $\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]=O(1)$ for some constant $p \in(1,2)$. Let $\epsilon \in(0,1 / 2)$ satisfy $1 / \epsilon=n^{o(1)}$. There is an $\widetilde{O}\left(n^{2}+n \operatorname{poly}(k / \epsilon)\right)$ time algorithm (Algorithm 19.1) which can output a subset $S \in[n]$ with $|S| \leq \operatorname{poly}(k / \epsilon)+$ $O(k \log n)$ for which

$$
\min _{X \in \mathbb{R}^{|S| \times n}}\left\|A_{S} X-A\right\|_{1} \leq(1+\epsilon)\|\Delta\|_{1},
$$

holds with probability at least 99/100.

Proof. We discussed the running time at the beginning of Section 19.2. Next, we turn to correctness. Let

$$
q=\Omega\left(\frac{k(k \log k)^{1+\frac{1}{p-1}}}{\epsilon^{1+\frac{1}{p-1}}}\right), \quad t=\frac{q^{\frac{1}{p-1}}}{\epsilon^{1+\frac{1}{p-1}}} .
$$

Let $r=\Theta(q / k)$. Let

$$
I_{1}=\left\{i_{1}^{(1)}, i_{2}^{(1)}, \cdots, i_{q t}^{(1)}\right\}, I_{2}=\left\{i_{1}^{(2)}, i_{2}^{(2)}, \cdots, i_{q t}^{(2)}\right\}, \cdots, I_{r}=\left\{i_{1}^{(r)}, i_{2}^{(r)}, \cdots, i_{q t}^{(r)}\right\}
$$

be $r$ independent subsets drawn uniformly at random from $\binom{[n]}{q t}$. Let $I=\bigcup_{s \in[r]} I_{s}$, which is the same as that in Algorithm 19.1. Let $\pi_{1}, \pi_{2}, \cdots, \pi_{r}$ be $r$ independent random permutations
of $q t$ elements. Due to Lemma 19.2.8 and a Chernoff bound, with probability at least .999, $\exists s \in[r]$,

$$
\mid\left\{i \in[n] \backslash I_{s} \mid\left(S_{1}, S_{2}, \cdots, S_{t}, i\right) \text { is an }\left(A^{*}, q, t, 1 / 2\right)-\operatorname{good} \text { tuple }\right\} \mid \geq(1-4 k / q)(n-q t)
$$

where

$$
S_{j}=\left\{i_{\pi_{s}((j-1) q+1)}^{(s)}, i_{\pi_{s}((j-1) q+2)}^{(s)}, \cdots, i_{\pi_{s}((j-1) q+q)}^{(s)}\right\}, \forall j \in[t] .
$$

Let set $H \subset[n]$ be defined as follows:

$$
H=\left\{j \in[n]\left|\exists i \in[n],\left|\Delta_{i, j}\right|>n^{1 / 2+1 /(2 p)}\right\} .\right.
$$

Then due to Lemma 19.2.3, with probability at least $0.999,|H| \leq O\left(n^{1-(p-1) / 2}\right)$. Thus, for $j \in[r]$, the probability that $H \cap I_{j} \neq \emptyset$ is at most $O\left(q t \cdot n^{1-(p-1) / 2} /(n-q t)\right)=1 / n^{\Omega(1)}$. By taking a union bound over all $j \in[r]$, with probability at least $1-1 / n^{\Omega(1)}, \forall j \in[r], I_{j} \cap H=\emptyset$. Thus, we can condition on $I_{s} \cap H=\emptyset$. Due to Lemma 19.2 .6 and $q^{1 / p} / t^{1-1 / p}=\epsilon$,

$$
\left|\left\{i \in[n] \backslash I_{s} \mid \min _{y \in \mathbb{R}^{q t}}\left\|A_{I_{s}} y-A_{i}\right\|_{1} \leq\left\|\Delta_{i}\right\|_{1}+O(\epsilon n)\right\}\right| \geq(1-4 k / q)(n-q t) .
$$

Due to Lemma 19.2.5 and a union bound over all $i \in[n] \backslash H$, with probability at least .999, $\forall i \notin H,\left\|\Delta_{i}\right\| \leq(1+\epsilon) n$. Thus,

$$
\left|\left\{i \in[n] \backslash I_{s} \mid \min _{y \in \mathbb{R}^{q t}}\left\|A_{I_{s}} y-A_{i}\right\|_{1} \leq(1+O(\epsilon)) n\right\}\right| \geq(1-4 k / q)(n-q t)-|H| .
$$

Let

$$
T^{\prime}=[n] \backslash\left\{i \in[n] \mid \min _{y \in \mathbb{R}^{q t}}\left\|A_{I_{s}} y-A_{i}\right\|_{1} \leq(1+O(\epsilon)) n\right\}
$$

Then $\left|T^{\prime}\right| \leq O\left(k n / q+n^{1-(p-1) / 2}\right)=O(k n / q)=O\left((\epsilon /(k \log k))^{1+1 /(p-1)} n\right)$. By our selection of $T$ in algorithm 19.1, $T^{\prime}$ should be a subset of $T$. Due to Lemma 19.2.4, with probability at least .999, $\left\|\Delta_{T}\right\|_{1} \leq O\left(\epsilon n^{2} /(k \log k)\right)$. By our second subroutine mentioned at the beginning of Section 19.2 it can find a set $Q \subset[n]$ with $|Q|=O(k \log n)$ such that

$$
\min _{X \in \mathbb{R}^{|Q| \times|T|} \mid}\left\|A_{Q} X-A_{T}\right\|_{1} \leq O(k \log k)\left\|\Delta_{T}\right\|_{1} \leq O\left(\epsilon n^{2}\right) .
$$

Thus, we have

$$
\begin{aligned}
\min _{X \in \mathbb{R}(|Q|+q \cdot t \cdot r) \times n}\left\|A_{(Q \cup I)} X-A\right\|_{1} & \leq \min _{X_{1} \in \mathbb{R}^{(q \cdot t \cdot r) \times n}}\left\|A_{I} X_{1}-A_{[n] \backslash T}\right\|_{1}+\min _{X_{2} \in \mathbb{R}|Q| \times n}\left\|A_{Q} X_{2}-A_{T}\right\|_{1} \\
& \leq(1+O(\epsilon)) n^{2} .
\end{aligned}
$$

Due to Lemma 19.2.1, with probability at least .999, $\|\Delta\|_{1} \geq(1-\epsilon) n^{2}$, and thus

$$
\min _{X \in \mathbb{R}(|Q|+q \cdot t \cdot r) \times n}\left\|A_{(Q \cup I)} X-A\right\|_{1} \leq(1+O(\epsilon))\|\Delta\|_{1} .
$$

### 19.3 Missing Proofs in Section 19.2

### 19.3.1 Proof of Lemma 19.2.1

Proof. Let $Z \in \mathbb{R}^{n \times n}$ be a random matrix. For each $i, j \in[n]$, define random variable $Z_{i, j}$ as

$$
Z_{i, j}= \begin{cases}\left|\Delta_{i, j}\right|, & \text { if }\left|\Delta_{i, j}\right| \leq n ; \\ n, & \text { otherwise } .\end{cases}
$$

For $i, j \in[n]$, by Markov's inequality, we have

$$
\begin{equation*}
\operatorname{Pr}\left[\left|\Delta_{i, j}\right| \geq n\right]=\operatorname{Pr}\left[\left|\Delta_{i, j}\right|^{p} \geq n^{p}\right] \leq \mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right] / n^{p}=O\left(1 / n^{p}\right) . \tag{19.1}
\end{equation*}
$$

Notice that

$$
\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]=\int_{0}^{n} x^{p} f(x) \mathrm{d} x+\int_{n}^{\infty} x^{p} f(x) \mathrm{d} x=O(1)
$$

where $f(x)$ is the probability density function of $\left|\Delta_{i, j}\right|$. Thus we have

$$
\int_{n}^{\infty} x f(x) \mathrm{d} x \leq \int_{n}^{\infty} x^{p} / n^{p-1} \cdot f(x) \mathrm{d} x=O\left(1 / n^{p-1}\right) .
$$

Because $\mathbb{E}\left[\left|\Delta_{i, j}\right|\right]=1$, we have

$$
\begin{equation*}
\int_{0}^{\infty} x f(x) \mathrm{d} x=\mathbb{E}\left[\left|\Delta_{i, j}\right|\right]-\int_{n}^{\infty} x f(x) \mathrm{d} x \geq 1-O\left(1 / n^{p-1}\right) \tag{19.2}
\end{equation*}
$$

By Equation (19.2), we have

$$
\mathbb{E}\left[Z_{i, j}\right]=\int_{0}^{n} x f(x) d x+n \cdot \operatorname{Pr}\left[\left|\Delta_{i, j}\right| \geq n\right] \geq \int_{0}^{n} x f(x) d x \geq 1-O\left(1 / n^{p-1}\right)
$$

By Equation (19.1) and $E\left[\left|\Delta_{i, j}\right|^{p}\right] \leq O(1)$, we have

$$
\mathbb{E}\left[Z_{i, j}^{2}\right]=\int_{0}^{n} x^{2} f(x) \mathrm{d} x+n^{2} \operatorname{Pr}\left[\left|\Delta_{i, j}\right| \geq n\right] \leq O\left(n^{2-p}\right)+O\left(n^{2-p}\right)=O\left(n^{2-p}\right)
$$

By the inequality of [Mau03],

$$
\begin{aligned}
\operatorname{Pr}\left[\mathbb{E}\left[\|Z\|_{1}\right]-\|Z\|_{1} \geq \epsilon \mathbb{E}\left[\|Z\|_{1}\right] / 2\right] & \leq \exp \left(\frac{-\epsilon^{2} \mathbb{E}\left[\|Z\|_{1}\right]^{2} / 4}{2 \sum_{i, j} \mathbb{E}\left[Z_{i, j}^{2}\right]}\right) \\
& \leq \exp \left(\frac{-\epsilon^{2}\left(n^{2}-O\left(n^{3-p}\right)\right)^{2} / 4}{2 n^{2} \cdot O\left(n^{2-p}\right)}\right) \\
& \leq e^{-\Theta(n)}
\end{aligned}
$$

Thus with probability at least $1-e^{-\Theta(n)},\|Z\|_{1} \geq(1-\epsilon / 2) \mathbb{E}\left[\|Z\|_{1}\right] \geq(1-\epsilon) n^{2}$ where the last inequality follows by $\mathbb{E}\left[\|Z\|_{1} \geq n^{2}-O\left(n^{3-p}\right)\right]$ and $1 / \epsilon=n^{o(1)}$. Since $\|\Delta\|_{1} \geq\|Z\|_{1}$, we complete the proof.

### 19.3.2 Proof of Lemma 19.2.2

Proof. Let $Z \in \mathbb{R}^{n \times t}$ be a random matrix where $Z_{i, j}$ are i.i.d. random variables with probability density function:

$$
g(x)= \begin{cases}f(x) / \operatorname{Pr}\left[\left|\Delta_{1,1}\right| \leq n^{1 / 2+1 /(2 p)}\right], & \text { if }|x| \leq n^{1 / 2+1 /(2 p)} ; \\ 0, & \text { otherwise }\end{cases}
$$

where $f(x)$ is the probability density function of $\Delta_{1,1}$. (Note that in the above equation, $\operatorname{Pr}\left[\left|\Delta_{1,1}\right| \leq n^{1 / 2+1 /(2 p)}\right]>0$.) Now, we have $\forall a \geq 0$,

$$
\operatorname{Pr}\left[\left\|\sum_{j=1}^{t} \alpha_{j} \Delta_{j}\right\|_{1} \leq a\left|\forall i \in[n], j \in[t],\left|\Delta_{i, j}\right| \leq n^{1 / 2+1 /(2 p)}\right]=\operatorname{Pr}\left[\left\|\sum_{j=1}^{t} \alpha_{j} Z_{j}\right\|_{1} \leq a\right] .\right.
$$

Now we look at the $i$-th row of $\sum_{j=1}^{t} \alpha_{j} Z_{j}$. We have

$$
\begin{align*}
\mathbb{E}\left[\left|\sum_{j=1}^{t} \alpha_{j} Z_{i, j}\right|\right] & =\left(\mathbb{E}\left[\left|\sum_{j=1}^{t} \alpha_{j} Z_{i, j}\right|\right]^{p}\right)^{1 / p} \\
& \leq \mathbb{E}\left[\left|\sum_{j=1}^{t} \alpha_{j} Z_{i, j}\right|^{p}\right]^{1 / p} \\
& \leq \mathbb{E}\left[\left(\left(\sum_{j=1}^{t} \alpha_{j}^{2} Z_{i, j}^{2}\right)^{1 / 2}\right)^{p}\right]^{1 / p} \\
& \leq \mathbb{E}\left[\sum_{j=1}^{t}\left|\alpha_{j} Z_{i, j}\right|^{p}\right]^{1 / p} \\
& \leq\left(\sum_{j=1}^{t} \mathbb{E}\left[\left|\alpha_{j} Z_{i, j}\right|^{p}\right]\right)^{1 / p} \\
& \leq\left(\sum_{j=1}^{t} \mathbb{E}\left[\left|Z_{i, j}\right|^{p}\right]\right)^{1 / p} \\
& \leq O\left(t^{1 / p}\right) \tag{19.3}
\end{align*}
$$

where the first inequality follows by Jensen's inequality, the second inequality follows by Remark 3 of [Lat97], the third inequality follows by $\|x\|_{2} \leq\|x\|_{p}$ for $p<2$, the fourth inequality follows by $\left|\alpha_{j}\right| \leq 1$, the fifth inequality follows by $\mathbb{E}\left[\left|Z_{i, j}\right|^{p}\right]=\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}| | \Delta_{1,1} \mid \leq\right.$
$\left.n^{1 / 2+1 /(2 p)}\right] \leq \mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]=O(1)$. For the second moment, we have

$$
\begin{align*}
\mathbb{E}\left[\left|\sum_{j=1}^{t} \alpha_{j} Z_{i, j}\right|^{2}\right] & =\sum_{j=1}^{t} \mathbb{E}\left[\alpha_{j}^{2} Z_{i, j}^{2}\right]+\sum_{j \neq k} \mathbb{E}\left[\alpha_{j} \alpha_{k} Z_{i, j} Z_{i, k}\right] \\
& =\sum_{j=1}^{t} \alpha_{j}^{2} \mathbb{E}\left[Z_{i, j}^{2}\right]+\sum_{j \neq k} \alpha_{j} \alpha_{k} \mathbb{E}\left[Z_{i, j}\right] \mathbb{E}\left[Z_{i, k}\right] \\
& \leq \sum_{j=1}^{t} \mathbb{E}\left[Z_{i, j}^{2}\right] \\
& =t \cdot 2 \int_{0}^{n^{1 / 2+1 /(2 p)}} x^{2} f(x) / \operatorname{Pr}\left[\left|\Delta_{i, j}\right| \leq n^{1 / 2+1 /(2 p)}\right] \mathrm{d} x \\
& \leq 2 t / \operatorname{Pr}\left[\left|\Delta_{i, j}\right| \leq n^{1 / 2+1 /(2 p)}\right] \cdot\left(n^{1 / 2+1 /(2 p)}\right)^{2-p} \int_{0}^{n^{1 / 2+1 /(2 p)}} x^{p} f(x) \mathrm{d} x \\
& \leq O\left(\operatorname{tn}^{2-p}\right) \tag{19.4}
\end{align*}
$$

where the second inequality follows by independence of $Z_{i, j}$ and $Z_{i, k}$. The first inequality follows by $\left|\alpha_{j}\right| \leq 1$ and $\mathbb{E}\left[Z_{i, j}\right]=\mathbb{E}\left[Z_{i, k}\right]=0$. The third equality follows by the probability density function of $Z_{i, j}$. The second inequality follows by $x^{2-p} \leq\left(n^{1 / 2+1 /(2 p)}\right)^{2-p}$ when $0 \leq$ $x \leq n^{1 / 2+1 /(2 p)}$. The last inequality follows by $\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]=O(1), p>1$ and $\operatorname{Pr}\left[\left|\Delta_{i, j}\right| \leq\right.$ $\left.n^{1 / 2+1 /(2 p)}\right] \geq 1-\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right] /\left(n^{1 / 2+1 /(2 p)}\right)^{p}=1-O\left(1 / n^{p / 2+1 / 2}\right) \geq 1 / 2$.

For $i \in[n]$, define $X_{i}=\left|\sum_{j=1}^{t} \alpha_{j} Z_{i, j}\right|$. Then, by Bernstein's inequality

$$
\begin{aligned}
& \operatorname{Pr}\left[\left\|\sum_{j=1}^{t} \alpha_{j} Z_{j}\right\|_{1}-\mathbb{E}\left[\left\|\sum_{j=1}^{t} \alpha_{j} Z_{j}\right\|_{1}\right] \geq 0.5 t^{1 / p} n\right] \\
= & \operatorname{Pr}\left[\sum_{i=1}^{n} X_{i}-\mathbb{E}\left[\sum_{i=1}^{n} X_{i}\right] \geq 0.5 t^{1 / p} n\right] \\
\leq & \exp \left(-\frac{0.5 \cdot 0.5^{2} t^{2 / p} n^{2}}{\sum_{i=1}^{n} \mathbb{E}\left[X_{i}^{2}\right]+\frac{1}{3} n^{1 / 2+1 /(2 p)} \cdot 0.5 t^{1 / p} n}\right) \\
\leq & e^{-n^{\ominus(1)}} .
\end{aligned}
$$

The last inequality follows by Equation (19.4). According to Equation (19.3), with probability at least $1-e^{-n^{\ominus(1)}}$,

$$
\left\|\sum_{j=1}^{t} \alpha_{j} Z_{j}\right\|_{1} \leq \mathbb{E}\left[\left\|\sum_{j=1}^{t} \alpha_{j} Z_{j}\right\|_{1}\right]+0.5 t^{1 / p} n \leq O\left(t^{1 / p} n\right)
$$

### 19.3.3 Proof of Lemma 19.2.3

Proof. For $i, j \in[n]$, we have

$$
\begin{aligned}
\operatorname{Pr}\left[\left|\Delta_{i, j}\right|>n^{1 / 2+1 /(2 p)}\right] & =\operatorname{Pr}\left[\left|\Delta_{i, j}\right|^{p}>n^{p / 2+1 / 2}\right] \\
& \leq \mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right] / n^{p / 2+1 / 2} \\
& \leq O\left(1 / n^{p / 2+1 / 2}\right)
\end{aligned}
$$

For column $j$, by taking a union bound,

$$
\operatorname{Pr}[j \in H]=\operatorname{Pr}\left[\exists i \in[n],\left|\Delta_{i, j}\right|>n^{1 / 2+1 /(2 p)}\right] \leq O\left(1 / n^{p / 2-1 / 2}\right)
$$

Thus, $\mathbb{E}[|H|] \leq O\left(n^{1-(p-1) / 2}\right)$. By applying Markov's inequality, we complete the proof.

### 19.3.4 Proof of Lemma 19.2.4

Proof. For $l \in \mathbb{N}_{\geq 0}$, define $G_{l}=\left\{j \mid\left\|\Delta_{j}\right\|_{1} \in\left(n \cdot 2^{l}, n \cdot 2^{l+1}\right]\right\}$. We have

$$
\begin{aligned}
\mathbb{E}\left[\left|G_{l}\right|\right] & \leq \sum_{j=1}^{n} \operatorname{Pr}\left[\left\|\Delta_{j}\right\|_{1} \geq n \cdot 2^{l}\right] \\
& =n \operatorname{Pr}\left[\left\|\Delta_{1}\right\|_{1} \geq n \cdot 2^{l}\right] \\
& \leq n \operatorname{Pr}\left[n^{1-1 / p}\left\|\Delta_{1}\right\|_{p} \geq n \cdot 2^{l}\right] \\
& =n \operatorname{Pr}\left[n^{p-1}\left\|\Delta_{1}\right\|_{p}^{p} \geq n^{p} \cdot 2^{l p}\right] \\
& \leq n \mathbb{E}\left[n^{p-1}\left\|\Delta_{1}\right\|_{p}^{p}\right] /\left(n^{p} \cdot 2^{l p}\right) \\
& \leq O\left(n / 2^{l p}\right) .
\end{aligned}
$$

The first inequality follows by the definition of $G_{l}$. The second inequality follows since $\forall x \in \mathbb{R}^{n},\|x\|_{1} \leq n^{1-1 / p}\|x\|_{p}$. The third inequality follows by Markov's inequality. The last inequality follows since $\forall i, j \in[n], \mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]=O(1)$.

Let $l^{*} \in \mathbb{N}_{\geq 0}$ satisfy $2^{l^{*}}<\epsilon r$ and $2^{l^{*}+1} \geq \epsilon r$. We have

$$
\begin{aligned}
\mathbb{E}\left[\sum_{j:\left\|\Delta_{j}\right\|_{1} \geq n 2^{l^{*}}}\left\|\Delta_{j}\right\|_{1}\right] & \leq \mathbb{E}\left[\sum_{l=l^{*}}^{\infty}\left|G_{l}\right| \cdot n 2^{l+1}\right] \\
& =\sum_{l=l^{*}}^{\infty} \mathbb{E}\left[\left|G_{l}\right|\right] \cdot n 2^{l+1} \\
& \leq \sum_{l=l^{*}}^{\infty} O\left(n / 2^{l p}\right) \cdot n 2^{l+1} \\
& =\sum_{l=l^{*}}^{\infty} O\left(n^{2} / 2^{l(p-1)}\right) \\
& =O\left(n^{2} / 2^{l^{*}(p-1)}\right) \\
& =O\left(n^{2} /(\epsilon r)^{p-1}\right) \\
& =O\left(\epsilon n^{2}\right)
\end{aligned}
$$

By Markov's inequality, with probability at least .999, $\sum_{j:\left\|\Delta_{j}\right\|_{1} \geq n 2^{2^{*}}}\left\|\Delta_{j}\right\|_{1} \leq O\left(\epsilon n^{2}\right)$. Conditioned on $\sum_{j:\left\|\Delta_{j}\right\|_{1} \geq n 2^{*}}\left\|\Delta_{j}\right\|_{1} \leq O\left(\epsilon n^{2}\right)$, for any $S \subset[n]$ with $|S| \leq n / r$, we have

$$
\sum_{j \in S}\left\|\Delta_{j}\right\|_{1} \leq|S| \cdot n 2^{l^{*}}+\sum_{j:\left\|\Delta_{j}\right\|_{1} \geq n 2^{* *}}\left\|\Delta_{j}\right\|_{1} \leq \epsilon n^{2}+O\left(\epsilon n^{2}\right)=O\left(\epsilon n^{2}\right)
$$

The second inequality follows because $|S| \leq n / r, 2^{l^{*}} \leq \epsilon r$ and $\sum_{j:\left\|\Delta_{j}\right\|_{1} \geq n 2^{2^{*}}}\left\|\Delta_{j}\right\|_{1} \leq O\left(\epsilon n^{2}\right)$.

### 19.3.5 Proof of Lemma 19.2.5

Proof. Let $M=n^{1 / 2+1 /(2 p)}$. Let $Z \in \mathbb{R}^{n}$ be a random vector where $Z_{i}$ are i.i.d. random variables with probability density function

$$
g(x)= \begin{cases}f(x) / \operatorname{Pr}\left[\left|\Delta_{1}\right| \leq M\right] & \text { if } 0 \leq x \leq M \\ 0 & \text { otherwise }\end{cases}
$$

where $f(x)$ is the probability density function of $\left|\Delta_{1}\right|$. Then $\forall a>0$

$$
\operatorname{Pr}\left[\|\Delta\|_{1} \leq a\left|\forall i \in[n],\left|\Delta_{i}\right| \leq M\right]=\operatorname{Pr}\left[\|Z\|_{1} \leq a\right] .\right.
$$

For $i \in[n]$, because $\mathbb{E}\left[\left|\Delta_{i}\right|\right]=1$, it holds that $\mathbb{E}\left[Z_{i}\right] \leq 1$. We have $\mathbb{E}\left[\sum_{i=1}^{n} Z_{i}\right] \leq n$. For the second moment, we have

$$
\begin{aligned}
\mathbb{E}\left[Z_{i}^{2}\right] & =\int_{0}^{M} x^{2} f(x) / \operatorname{Pr}\left[\left|\Delta_{1}\right| \leq M\right] \mathrm{d} x \\
& \leq M^{2-p} / \operatorname{Pr}\left[\left|\Delta_{1}\right| \leq M\right] \int_{0}^{M} x^{p} f(x) \mathrm{d} x \\
& \leq O\left(M^{2-p}\right) \\
& \leq O\left(n^{2-p}\right)
\end{aligned}
$$

where the second inequality follows by $\mathbb{E}\left[\left|\Delta_{1}\right|^{p}\right]=O(1)$, and

$$
\operatorname{Pr}\left[\left|\Delta_{1}\right| \leq M\right] \geq 1-\mathbb{E}\left[\left|\Delta_{1}\right|^{p}\right] / M^{p} \geq 1 / 2
$$

Then by Bernstein's inequality, we have

$$
\begin{aligned}
& \operatorname{Pr}\left[\sum_{i=1}^{n} Z_{i}-E\left[\sum_{i=1}^{n} Z_{i}\right] \geq \epsilon n\right] \\
\leq & \exp \left(\frac{-0.5 \epsilon^{2} n^{2}}{\sum_{i=1}^{n} \mathbb{E}\left[Z_{i}^{2}\right]+\frac{1}{3} M \cdot \epsilon n}\right) \\
\leq & e^{-n^{\ominus(1)}}
\end{aligned}
$$

Thus,

$$
\operatorname{Pr}\left[\|\Delta\|_{1} \leq(1+\epsilon) n\left|\forall i \in[n],\left|\Delta_{i}\right| \leq M\right]=\operatorname{Pr}\left[\|Z\|_{1} \leq(1+\epsilon) n\right] \geq 1-e^{-n^{\ominus(1)}}\right.
$$

### 19.3.6 Proof of Lemma 19.2.6

Proof. Recall that $\left(S_{1}, S_{2}, \cdots, S_{t}, i\right)$ is equivalent to $\left(S_{[t]}, i\right)$. Let $\left(S_{[t]}, i\right)$ be an $\left(A^{*}, q, t, 1 / 2\right)$ good tuple which satisfies $H \cap\left(\bigcup_{j=1}^{t} S_{j}\right)=\emptyset$. Let $C$ be the core of $\left(S_{[t]}, i\right)$. Let $\left(x_{1}, x_{2}, \cdots, x_{t}\right)$ be the coefficients tuple corresponding to $\left(S_{[t]}, i\right)$. Then we have that

$$
\begin{aligned}
& \left\|\frac{1}{|C|} \sum_{j=1}^{t} A_{S_{j}} x_{j}-A_{i}\right\|_{1} \\
= & \left\|\frac{1}{|C|} \sum_{j=1}^{t}\left(A_{S_{j}}^{*}+\Delta_{S_{j}}\right) x_{j}-\left(A_{i}^{*}+\Delta_{i}\right)\right\|_{1} \\
\leq & \left\|\frac{1}{|C|} \sum_{j=1}^{t} A_{S_{j}}^{*} x_{j}-A_{i}^{*}\right\|_{1}+\left\|\Delta_{i}\right\|_{1}+\frac{1}{|C|}\left\|\sum_{j=1}^{t} \Delta_{S_{j}} x_{j}\right\|_{1} \\
= & \left\|\Delta_{i}\right\|_{1}+\frac{1}{|C|}\left\|\sum_{j=1}^{t} \Delta_{S_{j}} x_{j}\right\|_{1} \\
\leq & \left\|\Delta_{i}\right\|_{1}+\frac{2}{t}\left\|\sum_{j=1}^{t} \Delta_{S_{j}} x_{j}\right\|_{1} \\
\leq & \left\|\Delta_{i}\right\|_{1}+O\left(\frac{1}{t} \cdot(q t)^{1 / p} n\right) \\
= & \left\|\Delta_{i}\right\|_{1}+O\left(q^{1 / p} / t^{1-1 / p} n\right)
\end{aligned}
$$

holds with probability at least $1-2^{-n^{\ominus(1)}}$. The first equality follows using $A=A^{*}+\Delta$. The first inequality follows using the triangle inequality. The second equality follows using the definition of the core and the coefficients tuple (see Definition 19.2.2 and Definition 19.2.4). The second inequality follows using Definition 19.2.2. The third inequality follows by Lemma 19.2.2 and the condition that $H \cap\left(\bigcup_{j=1}^{t} S_{j}\right)=\emptyset$.

Since the size of $\left|\{i\} \cup\left(\bigcup_{j=1}^{t} S_{j}\right)\right|=q t+1$, the total number of $\left(A^{*}, q, t, 1 / 2\right)$-good tuples is upper bounded by $n^{q t+1} \leq 2^{n^{o(1)}}$. By taking a union bound, we complete the
proof.

### 19.3.7 Proof of Lemma 19.2.7

Proof. For $j \in[t]$, by symmetry of the choices of $S_{j}$ and $i$, we have $\operatorname{Pr}\left[i \in R_{A^{*}}\left(S_{j} \cup\{i\}\right)\right] \leq$ $k /(q+1)$. Thus, by Markov's inequality,

$$
\begin{aligned}
& \operatorname{Pr}\left[\left|\left\{j \in[t] \mid i \in R_{A^{*}}\left(S_{j} \cup\{i\}\right)\right\}\right|>0.5 t\right] \\
\leq & \mathbb{E}\left[\left|\left\{j \in[t] \mid i \in R_{A^{*}}\left(S_{j} \cup\{i\}\right)\right\}\right|\right] /(0.5 t) \\
\leq & 2 k / q .
\end{aligned}
$$

Thus,

$$
\operatorname{Pr}\left[\left|\left\{j \in[t] \mid i \notin R_{A^{*}}\left(S_{j} \cup\{i\}\right)\right\}\right| \geq 0.5 t\right] \geq 1-2 k / q .
$$

### 19.3.8 Proof of Lemma 19.2.8

Proof. For $S_{1}, S_{2}, \cdots, S_{t} \in\binom{[n]}{q}$ with $\sum_{j=1}^{t}\left|S_{j}\right|=q t$, define

$$
P_{\left(S_{1}, S_{2}, \cdots, S_{t}\right)}=\operatorname{Pr}_{i \in[n] \backslash\left(\cup_{j=1}^{t} S_{j}\right)}\left[\left(S_{1}, S_{2}, \cdots, S_{t}, i\right) \text { is an }\left(A^{*}, q, t, 1 / 2\right)-\text { good tuple }\right] .
$$

Let set $T$ be defined as follows:

$$
\left\{\left(S_{1}, S_{2}, \cdots, S_{t}\right) \mid S_{1}, S_{2}, \cdots, S_{t} \in\binom{[n]}{q} \text { with } \sum_{j=1}^{t}\left|S_{j}\right|=q t\right\} \text {. }
$$

Let $G$ be the set of all the $\left(A^{*}, q, t, 1 / 2\right)$-good tuples. Then, we have

$$
\begin{array}{rl} 
& \operatorname{Pr}_{\left(S_{1}, S_{2}, \cdots, S_{t}\right) \sim T}\left[\left|\left\{i \in[n] \backslash\left(\cup_{j=1}^{t} S_{j}\right) \mid\left(S_{1}, S_{2}, \cdots, S_{t}, i\right) \in G\right\}\right| \geq(1-4 k / q)(n-q t)\right] \\
= & \left.\left.\frac{1}{|T|} \right\rvert\,\left\{\left(S_{1}, S_{2}, \cdots, S_{t}\right) \mid\left(S_{1}, S_{2}, \cdots, S_{t}\right) \in T \text { and } P_{\left(S_{1}, S_{2}, \cdots, S_{t}\right)} \geq 1-4 k / q\right\} \right\rvert\, \\
= & \frac{1}{|T|} \sum_{\substack{\left(S_{1}, S_{2}, \cdots, S_{t}\right) \in T \\
P_{\left(S_{1}, S_{2}, \cdots, S_{t}\right) \geq 1-4 k / q}}} 1 \\
\geq & \frac{1}{|T|} \sum_{\left(S_{1}, S_{2}, \cdots, S_{t}\right) \in T} \sum_{\left(S_{1}, S_{2}, \cdots, S_{t}\right)} \\
\quad P_{\left(S_{1}, S_{2}, \cdots, S_{t}\right) \geq 1-4 k / q} & 1-2 k / q-\frac{1}{|T|} \sum_{\left(S_{1}, S_{2}, \cdots, S_{t}\right) \in T} P_{\left(S_{1}, S_{2}, \cdots, S_{t}\right)} \\
\geq & 1-2 k / q-(1-4 k / q) \\
\geq & 2 k / q .
\end{array}
$$

The second inequality follows from Lemma 19.2.7

$$
\frac{1}{|T|} \sum_{\substack{\left(S_{1}, S_{2}, \cdots, S_{t}\right) \in T \\ P_{\left(S_{1}, S_{2}, \cdots, S_{t}\right)<1-4 k / q}}} P_{\left(S_{1}, S_{2}, \cdots, S_{t}\right)}+\frac{1}{|T|} \sum_{\substack{\left(S_{1}, S_{2}, \cdots, S_{t}\right) \in T \\ P_{\left(S_{1}, S_{2}, \cdots, S_{t}\right) \geq 1-4 k / q}}} P_{\left(S_{1}, S_{2}, \cdots, S_{t}\right)} \geq 1-2 k / q .
$$

### 19.4 Hardness Result

An overview of the hardness result. Recall that we overcame the column subset selection lower bound of [SWZ17], which shows for entrywise $\ell_{1}$-low rank approximation that there are matrices for which any subset of poly $(k)$ columns spans at best a $k^{\Omega(1)}$-approximation. Indeed, we came up with a column subset of size poly $\left(k\left(\epsilon^{-1}+\log n\right)\right)$ spanning a $(1+\epsilon)$ approximation. To do this, we assumed $A=A^{*}+\Delta$, where $A^{*}$ is an arbitrary rank- $k$ matrix, and the entries are i.i.d. from a distribution with $\mathbb{E}\left[\left|\Delta_{i, j}\right|\right]=1$ and $\mathbb{E}\left[\left|\Delta_{i, j}\right|^{p}\right]=O(1)$ for any real number $p$ strictly greater than 1 .

Here we show an assumption on the moments is necessary, by showing if instead $\Delta$ were drawn from a matrix of i.i.d. Cauchy random variables, for which the $p$-th moment is undefined or infinite for all $p \geq 1$, then for any subset of $n^{o(1)}$ columns, it spans at best a 1.002 approximation. The input matrix $A=n^{C} 1 \cdot 1^{\top}+\Delta$, where $C>0$ is a constant and we show that $n^{\Omega(1)}$ columns need to be chosen to obtain a 1.001-approximation, even for $k=1$. Note that this result is stronger than that in [SWZ17] in that it rules out column subset selection even if one were to choose $n^{o(1)}$ columns; the result in [SWZ17] requires at most $\operatorname{poly}(k)$ columns, which for $k=1$, would just rule out $O(1)$ columns. Our main goal here is to show that a moment assumption on our distribution is necessary, and our result also applies to a symmetric noise distribution which is i.i.d. on all entries, whereas the result of [SWZ17] requires a specific deterministic pattern (namely, the identity matrix) on certain entries.

Our main theorem is given in Theorem 19.4.17. The outline of the proof is as follows. We first condition on the event that $\|\Delta\|_{1} \leq \frac{4.0002}{\pi} n^{2} \ln n$, which is shown in Lemma 19.4.2 and follows form standard analysis of sums of absolute values of Cauchy random variables.

Thus, it is sufficient to show if we choose any subset $S$ of $r=n^{o(1)}$ columns, denoted by the submatrix $A_{S}$, then $\min _{X \in \mathbb{R}^{r \times n}}\left\|A_{S} X-A\right\|_{1} \geq \frac{4.01}{\pi} \cdot n^{2} \ln n$, as indeed then $\min _{X \in \mathbb{R}^{r \times n}} \| A_{S} X-$ $A\left\|_{1} \geq 1.002\right\| \Delta \|_{1}$ and we rule out a $(1+\epsilon)$-approximation for $\epsilon$ a sufficiently small constant. To this end, we instead show for a fixed $S$, that $\min _{X \in \mathbb{R}^{r \times n}}\left\|A_{S} X-A\right\|_{1} \geq \frac{4.01}{\pi} \cdot n^{2} \ln n$ with probability $1-2^{-n^{\ominus(1)}}$, and then apply a union bound over all $S$. To prove this for a single subset $S$, we argue that for every "coefficient matrix" $X$, that $\left\|A_{S} X-A\right\|_{1} \geq \frac{4.01}{\pi} \cdot n^{2} \ln n$.

We show in Lemma 19.4.5, that with probability $1-(1 / n)^{\Theta(n)}$ over $\Delta$, simultaneously for all $X$, if $X$ has a column $X_{j}$ with $\left\|X_{j}\right\|_{1} \geq n^{c}$ for a constant $c>0$, then $\left\|A_{S} X_{j}-A_{j}\right\|_{1} \geq$ $.9 n^{3}$, which is already too large to provide an $O(1)$-approximation. Note that we need such a high probability bound to later union bound over all $S$. Lemma 19.4.5 is in turn shown via a net argument on all $X_{j}$ (it suffices to prove this for a single $j \in[n]$, since there are only $n$ different $j$, so we can union bound over all $j$ ). The net bounds are given in Definition 19.4.1 and Definition 19.4.4, and the high probability bound for a given coefficient vector $X_{j}$ is shown in Lemma 19.4.3, where we use properties of the Cauchy distribution. Thus, we can assume $\left\|X_{j}\right\|_{1}<n^{c}$ for all $j \in[n]$. We also show in Fact 19.4.1, conditioned on the fact that $\|\Delta\|_{1} \leq \frac{4.002}{\pi} n^{2} \ln n$, it holds that for any vector $X_{j}$, if $\left\|X_{j}\right\|_{1}<n^{c}$ and $\left|1-1^{\top} X_{j}\right|>1-10^{-20}$, then $\left\|A_{S} X-A\right\|_{1} \geq\left\|A_{S} X_{j}-A_{j}\right\|_{1}>n^{3}$. The intuition here is $A=n^{c_{0}} \mathbf{1} \cdot \mathbf{1}^{\top}+\Delta$ for a large constant $c_{0}$, and $X_{j}$ does not have enough norm $\left(\left\|X_{j}\right\|_{1} \leq n^{c}\right)$ or correlation with the vector $\mathbf{1}\left(\left|1-\mathbf{1}^{\top} X_{j}\right|>1-10^{-20}\right)$ to make $\left\|A_{S} X_{j}-A_{j}\right\|_{1}$ small.

Given the above, we can assume both that $\left\|X_{j}\right\|_{1} \leq n^{c}$ and $\left|1-\mathbf{1}^{\top} X_{j}\right| \leq 1-10^{-20}$ for all columns $j$ of our coefficient matrix $X$. We can also assume that $\left\|A_{S} X-A\right\|_{1} \leq 4 n^{2} \ln n$, as otherwise such an $X$ already satisfies $\left\|A_{S} X-A\right\|_{1} \geq \frac{4.01}{\pi} \cdot n^{2} \ln n$ and we are done. To analyze $\left\|A_{S} X-A\right\|_{1}=\sum_{i, j}\left|\left(A_{S} X-A_{[n] \backslash S}\right)_{i, j}\right|$ in Theorem 19.4.17, we then split the sum over
"large coordinates" $(i, j)$ for which $\left|\Delta_{i, j}\right|>n^{1.0002}$, and "small coordinates" $(i, j)$ for which $\left|\Delta_{i, j}\right|<n^{.9999}$, and since we seek to lower bound $\left\|A_{S} X-A_{[n] \backslash S}\right\|_{1}$, we drop the remaining coordinates $(i, j)$. To handle large coordinates, we observe that since the column span of $A_{S}$ is only $r=n^{o(1)}$-dimensional, as one ranges over all vectors $y$ in its span of 1-norm, say, $O\left(n^{2} \ln n\right)$, there is only a small subset $T$, of size at most $n^{.99999}$ of coordinates $i \in[n]$ for which we could ever have $\left|y_{i}\right| \geq n^{1.0001}$. We show this in Lemma 19.4.7. This uses the property of vectors in low-dimensional subspaces, and has been exploited in earlier works in the context of designing so-called subspace embeddings [CW13, MM13]. We call $T$ the "bad region" for $A_{S}$. While the column span of $A_{S}$ depends on $\Delta_{S}$, it is independent of $\Delta_{[n] \backslash S}$, and thus it is extremely unlikely that the large coordinate of $\Delta_{S}$ "match up" with the bad region of $A_{S}$. This is captured in Lemma 19.4.10, where we show that if $\left\|A_{S} X-A_{[n] \backslash S}\right\|_{1} \leq 4 n^{2} \ln n$ (as we said we could assume above), then $\sum_{\text {large coordinates } i, j}\left|\left(A_{S} X-A_{[n \backslash \backslash S}\right)_{i, j}\right|$ is at least $\frac{1.996}{\pi} n^{2} \ln n$. Intuitively, the heavy coordinates make up about $\frac{2}{\pi} n^{2} \ln n$ of the total mass of $\|\Delta\|_{1}$, by tail bounds of the Cauchy distribution, and for any set $S$ of size $n^{o(1)}, A_{S}$ fits at most a small portion of this, still leaving us left with $\frac{1.996}{\pi} n^{2} \ln n$ in cost. Our goal is to show that $\left\|A_{S} X-A_{[n] \backslash S}\right\|_{1} \geq \frac{4.01}{\pi} \cdot n^{2} \ln n$, so we still have a way to go.

We next analyze $\sum_{\text {small coordinates } i, j}\left|\left(A_{S} X-A_{[n] \backslash S}\right)_{i, j}\right|$. Via Bernstein's inequality, in Lemma 19.4.11 we argue that for any fixed vector $y$ and random vector $\Delta_{j}$ of i.i.d. Cauchy entries, roughly half of the contribution of coordinates to $\left\|\Delta_{j}\right\|_{1}$ will come from coordinates $j$ for which $\operatorname{sign}\left(y_{j}\right)=\operatorname{sign}\left(\Delta_{j}\right)$ and $\left|\Delta_{j}\right| \leq n^{.9999}$, giving us a contribution of roughly $\frac{.9998}{\pi} n \ln n$ to the cost. The situation we will actually be in, when analyzing a column of $A_{S} X-A_{[n] \backslash S}$, is that of taking the sum of two independent Cauchy vectors, shifted by a multiple of $\mathbf{1}^{\top}$. We analyze this setting in Lemma 19.4.13, after first conditioning on certain level sets having
typical behavior in Lemma 19.4.12. This roughly doubles the contribution, gives us roughly a contribution of $\frac{1.996}{\pi} n^{2} \ln n$ from coordinates $j$ for which $(i, j)$ is a small coordinate and we look at coordinates $i$ on which the sum of two independent Cauchy vectors have the same sign. Combined with the contribution from the heavy coordinates, this gives us a cost of roughly $\frac{3.992}{\pi} n^{2} \ln n$, which still falls short of the $\frac{4.01}{\pi} \cdot n^{2} \ln n$ total cost we are aiming for. Finally, if we sum up two independent Cauchy vectors and look at the contribution to the sum from coordinates which disagree in sign, due to the anti-concentration of the Cauchy distribution we can still "gain a little bit of cost" since the values, although differing in sign, are still likely not to be very close in magnitude. We formalize this in Lemma 19.4.14. We combine all of the costs from small coordinates in Lemma 19.4.15, where we show we obtain a contribution of at least $\frac{2.025}{\pi} n \ln n$. This is enough, when combined with our earlier $\frac{1.996}{\pi} n^{2} \ln n$ contribution from the heavy coordinates, to obtain an overall $\frac{4.01}{\pi} \cdot n^{2} \ln n$ lower bound on the cost, and conclude the proof of our main theorem in Theorem 19.4.17.

In the remaining sections, we will present detailed proofs.

### 19.4.1 A Useful Fact

Fact 19.4.1. Let $c_{0}>0$ be a sufficiently large constant. Let $u=n^{c_{0}} \cdot \mathbf{1} \in \mathbb{R}^{n}$ and $\Delta \in$ $\mathbb{R}^{n \times(d+1)}$. If $\sum_{i=1}^{d+1}\left\|\Delta_{i}\right\|_{1} \leq n^{3}$ and if $\alpha \in \mathbb{R}^{d}$ satisfies $\left|1-\mathbf{1}^{\top} \alpha\right|>1 / n^{c_{1}}$ and $\|\alpha\|_{1} \leq n^{c}$, where $0<c<c_{0}-10$ is a constant and $c_{1}>3$ is another constant depending on $c_{0}, c$, then

$$
\left\|u-u \mathbf{1}^{\top} \alpha+\Delta_{d+1}-\Delta_{[d]} \alpha\right\|_{1}>n^{3} .
$$

Proof.

$$
\begin{aligned}
& \left\|u-u \mathbf{1}^{\top} \alpha+\Delta_{d+1}-\Delta_{[d]} \alpha\right\|_{1} \\
\geq & \left|1-\mathbf{1}^{\top} \alpha\right| \cdot\|u\|_{1}-\left\|\Delta_{d+1}\right\|_{1}-\left\|\Delta_{[d]} \alpha\right\|_{1} \\
\geq & \left|1-\mathbf{1}^{\top} \alpha\right| \cdot n \cdot n^{c_{0}}-n^{3}-n^{4}\|\alpha\|_{1} \\
\geq & \left|1-\mathbf{1}^{\top} \alpha\right| \cdot n \cdot n^{c_{0}}-n^{5+c} \\
\geq & n^{c_{0}+1-c_{1}}-n^{5+c} \\
\geq & n^{3} .
\end{aligned}
$$

The first inequality follows by the triangle inequality. The second inequality follows since $u=n^{c_{0}} \cdot \mathbf{1} \in \mathbb{R}^{n}$ and $\sum_{i=1}^{d+1}\left\|\Delta_{i}\right\|_{1} \leq n^{3}$. The third inequality follows since $\|\alpha\|_{1} \leq n^{c}$. The fourth inequality follows since $\left|1-\mathbf{1}^{\top} \alpha\right|>1 / n^{c_{1}}$. The last inequality follows since $c_{0}-c_{1}>c+5$.

### 19.4.2 One-Sided Error Concentration Bound for a Random Cauchy Matrix

Lemma 19.4.2 (Lower bound on the cost). If $n$ is sufficiently large, then

$$
\operatorname{Pr}_{\Delta \sim\{C(0,1)\}^{n \times n}}\left[\|\Delta\|_{1} \leq \frac{4.0002}{\pi} n^{2} \ln n\right] \geq 1-O(1 / \log \log n) .
$$

Proof. Let $\Delta \in \mathbb{R}^{n \times n}$ be a random matrix such that each entry is an i.i.d. $C(0,1)$ random Cauchy variable. Let $B=n^{2} \ln \ln n$. Let $Z \in \mathbb{R}^{n \times n}$ and $\forall i, j \in[n]$,

$$
Z_{i, j}=\left\{\begin{array}{ll}
\left|\Delta_{i, j}\right| & \left|\Delta_{i, j}\right|<B \\
B & \text { Otherwise }
\end{array} .\right.
$$

For fixed $i, j \in[n]$, we have

$$
\begin{aligned}
\mathbb{E}\left[Z_{i, j}\right] & =\frac{2}{\pi} \int_{0}^{B} \frac{x}{1+x^{2}} \mathrm{~d} x+\operatorname{Pr}\left[\left|\Delta_{i, j}\right| \geq B\right] \cdot B \\
& =\frac{1}{\pi} \ln \left(B^{2}+1\right)+\operatorname{Pr}\left[\left|\Delta_{i, j}\right| \geq B\right] \cdot B \\
& \leq \frac{1}{\pi} \ln \left(B^{2}+1\right)+1
\end{aligned}
$$

where the first inequality follows by the cumulative distribution function of a half Cauchy random variable. We also have $\mathbb{E}\left[Z_{i, j}\right] \geq \frac{1}{\pi} \ln \left(B^{2}+1\right)$. For the second moment, we have

$$
\begin{aligned}
\mathbb{E}\left[Z_{i, j}^{2}\right] & =\frac{2}{\pi} \int_{0}^{B} \frac{x^{2}}{1+x^{2}} \mathrm{~d} x+\operatorname{Pr}\left[\left|\Delta_{i, j}\right| \geq B\right] \cdot B^{2} \\
& =\frac{2}{\pi}\left(B-\tan ^{-1} B\right)+\operatorname{Pr}\left[\left|\Delta_{i, j}\right| \geq B\right] \cdot B^{2} \\
& \leq \frac{2}{\pi} B+B \\
& \leq 2 B
\end{aligned}
$$

where the first inequality follows by the cumulative distribution function of a half Cauchy random variable. By applying Bernstein's inequality, we have

$$
\begin{align*}
& \operatorname{Pr}\left[\|Z\|_{1}-\mathbb{E}\left[\|Z\|_{1}\right]>0.0001 \mathbb{E}\left[\|Z\|_{1}\right]\right] \\
\leq & \exp \left(-\frac{0.5 \cdot 0.0001^{2} \mathbb{E}\left[\|Z\|_{1}\right]^{2}}{n^{2} \cdot 2 B+\frac{1}{3} B \cdot 0.0001 \mathbb{E}\left[\|Z\|_{1}\right]}\right) \\
\leq & \exp (-\Omega(\ln n / \ln \ln n)) \\
\leq & O(1 / \ln n) \tag{19.5}
\end{align*}
$$

The first inequality follows by the definition of $Z$ and the second moment of $Z_{i, j}$. The second
inequality follows from $\mathbb{E}\left[\|Z\|_{1}\right]=\Theta\left(n^{2} \ln n\right)$ and $B=\Theta\left(n^{2} \ln \ln n\right)$. Notice that

$$
\begin{aligned}
& \operatorname{Pr}\left[\|\Delta\|_{1}>\frac{4.0002}{\pi} n^{2} \ln n\right] \\
= & \operatorname{Pr}\left[\|\Delta\|_{1}>\frac{4.0002}{\pi} n^{2} \ln n\left|\forall i, j,\left|\Delta_{i, j}\right|<B\right] \operatorname{Pr}\left[\forall i, j,\left|\Delta_{i, j}\right|<B\right]\right. \\
& +\operatorname{Pr}\left[\|\Delta\|_{1}>\frac{4.0002}{\pi} n^{2} \ln n\left|\exists i, j,\left|\Delta_{i, j}\right| \geq B\right] \operatorname{Pr}\left[\exists i, j,\left|\Delta_{i, j}\right| \geq B\right]\right. \\
\leq & \operatorname{Pr}\left[\|\Delta\|_{1}>\frac{4.0002}{\pi} n^{2} \ln n\left|\forall i, j,\left|\Delta_{i, j}\right|<B\right]+\operatorname{Pr}\left[\exists i, j,\left|\Delta_{i, j}\right| \geq B\right]\right. \\
\leq & \operatorname{Pr}\left[\|Z\|_{1}>\frac{4.0002}{\pi} n^{2} \ln n\right]+\operatorname{Pr}\left[\exists i, j,\left|\Delta_{i, j}\right| \geq B\right] \\
\leq & \operatorname{Pr}\left[\|Z\|_{1}>\frac{4.0002}{\pi} n^{2} \ln n\right]+n^{2} \cdot 1 / B \\
\leq & \operatorname{Pr}\left[\|Z\|_{1}>1.0001 \mathbb{E}\left[\|Z\|_{1}\right]\right]+n^{2} \cdot 1 / B \\
\leq & O(1 / \log (n))+O(1 / \log \log n) \\
\leq & O(1 / \log \log n)
\end{aligned}
$$

The second inequality follows by the definition of $Z$. The third inequality follows by the union bound and the cumulative distribution function of a half Cauchy random variable. The fourth inequality follows from $\mathbb{E}\left[\|Z\|_{1}\right] \leq n^{2}\left(1 / \pi \cdot \ln \left(B^{2}+1\right)+1\right) \leq 4.0000001 / \pi \cdot n^{2} \ln n$ when $n$ is sufficiently large.

### 19.4.3 "For Each" Guarantee

In the following Lemma, we show that, for each fixed coefficient vector $\alpha$, if the entry of $\alpha$ is too large, the fitting cost cannot be small.

Lemma 19.4.3 (For each fixed $\alpha$, the entry cannot be too large). Let $c>0$ be a sufficiently
large constant, $n \geq d \geq 1, u \in \mathbb{R}^{n}$ be any fixed vector and $\Delta \in \mathbb{R}^{n \times d}$ be a random matrix where $\forall i \in[n], j \in[d], \Delta_{i, j} \sim C(0,1)$ independently. For any fixed $\alpha \in \mathbb{R}^{d}$ with $\|\alpha\|_{1}=n^{c}$,

$$
\operatorname{Pr}_{\Delta \sim\{C(0,1)\}^{n \times d}}\left[\left\|\left(u \cdot \mathbf{1}^{\top}+\Delta\right) \alpha\right\|_{1}>n^{3}\right]>1-(1 / n)^{\Theta(n)}
$$

Proof. Let $c$ be a sufficiently large constant. Let $\alpha \in \mathbb{R}^{d}$ with $\|\alpha\|_{1}=n^{c}$. Let $u \in \mathbb{R}^{n}$ be any fixed vector. Let $\Delta \in \mathbb{R}^{n \times d}$ be a random matrix where $\forall i \in[n], j \in[d], \Delta_{i, j} \sim C(0,1)$. Then $\Delta \alpha \in \mathbb{R}^{n}$ is a random vector with each entry drawn independently from $C\left(0,\|\alpha\|_{1}\right)$. Due to the probability density function of standard Cauchy random variables,

$$
\operatorname{Pr}\left[\|\Delta \alpha\|_{1}<n^{3}\right] \geq \operatorname{Pr}\left[\left\|\Delta \alpha+u \cdot \mathbf{1}^{\top} \alpha\right\|_{1}<n^{3}\right] .
$$

It suffices to upper bound $\operatorname{Pr}\left[\|\Delta \alpha\|_{1}<n^{3}\right]$. If $c>10$, then due to the cumulative distribution function of Cauchy random variables, for a fixed $i \in[n], \operatorname{Pr}\left[\left|(\Delta \alpha)_{i}\right|<n^{3}\right]<1 / n$. Thus, $\operatorname{Pr}\left[\|\Delta \alpha\|_{1}<n^{3}\right]<\left(\frac{1}{n}\right)^{n}$. Thus,

$$
\operatorname{Pr}_{\Delta \sim\{C(0,1)\}^{n \times d}}\left[\left\|\left(u \cdot \mathbf{1}^{\top}+\Delta\right) \alpha\right\|_{1}>n^{3}\right]>1-(1 / n)^{n} .
$$

### 19.4.4 From "For Each" to "For All" via an $\epsilon$-Net

Definition 19.4.1 ( $\epsilon$-net for the $\ell_{1}$-norm ball). Let $A \in \mathbb{R}^{n \times d}$ have rank $d$, and let $L=$ $\left\{y \in \mathbb{R}^{n} \mid y=A x, x \in \mathbb{R}^{d}\right\}$ be the column space of $A$. An $\epsilon$-net of the $\ell_{1}$-unit sphere $\mathcal{S}^{d-1}=\left\{y \mid\|y\|_{1}=1, y \in L\right\} \subset L$ is a set $N \subset \mathcal{S}^{d-1}$ of points for which $\forall y \in \mathcal{S}^{d-1}, \exists y^{\prime} \in N$ for which $\left\|y-y^{\prime}\right\| \leq \epsilon$.
$\left[\mathrm{DDH}^{+} 09\right]$ proved an upper bound on the size of an $\epsilon$-net.

Lemma 19.4.4 (See, e.g., the ball $B$ on page 2068 of $\left.\left[\mathrm{DDH}^{+} 09\right]\right)$. Let $A \in \mathbb{R}^{n \times d}$ have rank $d$, and let $L=\left\{y \in \mathbb{R}^{n} \mid y=A x, x \in \mathbb{R}^{d}\right\}$ be the column space of $A$. For $\epsilon \in(0,1)$, an $\epsilon$-net (Definition 19.4.1) $N$ of the $\ell_{1}$-unit sphere $\mathcal{S}^{d-1}=\left\{y \mid\|y\|_{1}=1, y \in L\right\} \subset L$ exists. Furthermore, the size of $N$ is at most $(3 / \epsilon)^{d}$.

Lemma 19.4.5 (For all possible $\alpha$, the entry cannot be too large). Let $n \geq 1, d=n^{o(1)}$. Let $u=n^{c_{0}} \cdot \mathbf{1} \in \mathbb{R}^{n}$ denote a fixed vector where $c_{0}$ is a constant. Let $\Delta \in \mathbb{R}^{n \times d}$ be a random matrix where $\forall i \in[n], j \in[d], \Delta_{i, j} \sim C(0,1)$ independently. Let $c>0$ be a sufficiently large constant. Conditioned on $\|\Delta\|_{1} \leq n^{3}$, with probability at least $1-(1 / n)^{\Theta(n)}$, for all $\alpha \in \mathbb{R}^{d}$ with $\|\alpha\|_{1} \geq n^{c}$, we have $\left\|\left(u \cdot \mathbf{1}^{\top}+\Delta\right) \alpha\right\|_{1}>0.9 n^{3}$.

Proof. Due to Lemma 19.4.4, there is a set $N \subset\left\{\alpha \in \mathbb{R}^{d} \mid\|\alpha\|_{1}=n^{c}\right\} \subset \mathbb{R}^{d}$ with $|N| \leq$ $2^{\Theta(d \log n)}$ such that $\forall \alpha \in \mathbb{R}^{d}$ with $\|\alpha\|_{1}=n^{c}, \exists \alpha^{\prime} \in N$ such that $\left\|\alpha-\alpha^{\prime}\right\|_{1} \leq 1 / n^{c^{\prime}}$ where $c^{\prime}>c_{0}+100$ is a constant. By applying Lemma 19.4.3 and union bounding over all the points in $N$, with probability at least $1-(1 / n)^{n} \cdot|N| \geq 1-(1 / n)^{n} \cdot 2^{n^{\circ(1)}}=1-(1 / n)^{\Theta(n)}$, $\forall \alpha^{\prime} \in N,\left\|\left(u \cdot \mathbf{1}^{\top}+\Delta\right) \alpha^{\prime}\right\|_{1}>n^{3} . \forall \alpha \in \mathbb{R}^{d}$ with $\|\alpha\|_{1}=n^{c}$, we can find $\alpha^{\prime} \in N$ such that
$\left\|\alpha-\alpha^{\prime}\right\|_{1} \leq 1 / n^{c^{\prime}}$. Let $\gamma=\alpha-\alpha^{\prime}$. Then,

$$
\begin{aligned}
& \left\|\left(u \cdot \mathbf{1}^{\top}+\Delta\right) \alpha\right\|_{1} \\
= & \left\|\left(u \cdot \mathbf{1}^{\top}+\Delta\right)\left(\alpha^{\prime}+\gamma\right)\right\|_{1} \\
\geq & \left\|\left(u \cdot \mathbf{1}^{\top}+\Delta\right) \alpha^{\prime}\right\|_{1}-\left\|\left(u \cdot \mathbf{1}^{\top}+\Delta\right) \gamma\right\|_{1} \\
\geq & n^{3}-\sqrt{n}\left\|\left(u \cdot \mathbf{1}^{\top}+\Delta\right) \gamma\right\|_{2} \\
\geq & n^{3}-\sqrt{n}\left(\left\|u \cdot \mathbf{1}^{\top}\right\|_{2}+\|\Delta\|_{2}\right)\|\gamma\|_{2} \\
\geq & n^{3}-n^{c_{0}+50} / n^{c^{\prime}} \\
\geq & 0.9 n^{3}
\end{aligned}
$$

The first equality follows from $\alpha=\alpha^{\prime}+\gamma$. The first inequality follows by the triangle inequality. The second inequality follows by the relaxation from the $\ell_{1}$ norm to the $\ell_{2}$ norm. The third inequality follows from the operator norm and the triangle inequality. The fourth inequality follows using $\|\Delta\|_{2} \leq\|\Delta\|_{1} \leq n^{3},\|u\|_{2} \leq n^{c_{0}+10},\|\gamma\|_{2} \leq\|\gamma\|_{1} \leq(1 / n)^{c^{\prime}}$. The last inequality follows since $c^{\prime}>c_{0}+100$.

For $\alpha \in \mathbb{R}^{n}$ with $\|\alpha\|_{1}>n^{c}$, let $\alpha^{\prime}=\alpha /\|\alpha\|_{1} \cdot n^{c}$. Then

$$
\left\|\left(u \cdot \mathbf{1}^{\top}+\Delta\right) \alpha\right\|_{1} \geq\left\|\left(u \cdot \mathbf{1}^{\top}+\Delta\right) \alpha^{\prime}\right\|_{1} \geq 0.9 n^{3} .
$$

### 19.4.5 Bounding the Cost from the Large-Entry Part via "Bad" Regions

In this section, we will use the concept of well-conditioned basis in our analysis.
Definition 19.4.2 (Well-conditioned basis $\left.\left[\mathrm{DDH}^{+} 09\right]\right)$. Let $A \in \mathbb{R}^{n \times m}$ have rank $d$. Let $p \in[1, \infty)$, and let $\|\cdot\|_{q}$ be the dual norm of $\|\cdot\|_{p}$, i.e., $1 / p+1 / q=1$. If $U \in \mathbb{R}^{n \times d}$ satisfies

1. $\|U\|_{p} \leq \alpha$,
2. $\forall z \in \mathbb{R}^{d},\|z\|_{q} \leq \beta\|U z\|_{p}$,
then $U$ is an $(\alpha, \beta, p)$ well-conditioned basis for the column space of $A$.

The following theorem gives an existence result of a well-conditioned basis.

Theorem 19.4.6 ( $\ell_{1}$ well-conditioned basis $\left[\mathrm{DDH}^{+} 09\right]$ ). Let $A \in \mathbb{R}^{n \times m}$ have rank $d$. There exists $U \in \mathbb{R}^{n \times d}$ such that $U$ is a $(d, 1,1)$ well-conditioned basis for the column space of $A$.

In the following lemma, we consider vectors from low-dimensional subspaces. For a coordinate, if there is a vector from the subspace for which this entry is large, but the norm of the vector is small, then this kind of coordinate is pretty "rare". More formally,

Lemma 19.4.7. Given a matrix $U \in \mathbb{R}^{n \times r}$ for a sufficiently large $n \geq 1$, let $r=n^{o(1)}$. Let $S=\left\{y \mid y=U x, x \in \mathbb{R}^{r}\right\}$. Let the set $T$ denote $\left\{i \in[n]\left|\exists y \in S,\left|y_{i}\right| \geq n^{1.0001}\right.\right.$ and $\|y\|_{1}<$ $\left.8 n^{2} \ln n\right\}$. Then we have

$$
|T| \leq n^{0.99999}
$$

Proof. Due to Theorem 19.4.6, let $U \in \mathbb{R}^{n \times r}$ be the $(r, 1,1)$ well-conditioned basis of the column space of $U$. If $i \in T$, then $\exists x \in \mathbb{R}^{r}$ such that $\left|(U x)_{i}\right| \geq n^{1.0001}$ and $\|U x\|_{1}<8 n^{2} \ln n$. Thus, we have

$$
n^{1.0001} \leq\left|(U x)_{i}\right| \leq\left\|U^{i}\right\|_{1}\|x\|_{\infty} \leq\left\|U^{i}\right\|_{1}\|U x\|_{1} \leq\left\|U^{i}\right\|_{1} \cdot 8 n^{2} \ln n
$$

The first inequality follows using $n^{1.0001} \leq\left|(U x)_{i}\right|$. The second inequality follows by Hölder's inequality. The third inequality follows by the second property of the well-conditioned basis. The fourth inequality follows using $\|U x\|_{1}<8 n^{2} \ln n$. Thus, we have

$$
\left\|U^{i}\right\|_{1} \geq n^{1.0001} / n^{2+o(1)} \geq 1 / n^{0.9999-o(1)} .
$$

Notice that $\sum_{j=1}^{n}\left\|U^{j}\right\|_{1}=\|U\|_{1} \leq r$. Thus,

$$
|T| \leq r /\left(1 / n^{0.9999-o(1)}\right)=n^{0.9999+o(1)} \leq n^{0.99999} .
$$

Definition 19.4.3 (Bad region). Given a matrix $U \in \mathbb{R}^{n \times r}$, we say $\mathcal{B}(U)=\{i \in[n] \mid \exists y \in$ $\operatorname{colspan}(U) \subset \mathbb{R}^{n}$ s.t. $y_{i} \geq n^{1.0001}$ and $\left.\|y\|_{1} \leq 8 n^{2} \ln n\right\}$ is a bad region for $U$.

Next we state a lower and an upper bound on the probability that a Cauchy random variable is in a certain range,

Claim 19.4.8. Let $X \sim C(0,1)$ be a standard Cauchy random variable. Then for any $x>1549$,

$$
\frac{2}{\pi} \cdot \frac{\ln (1.001)}{x} \geq \operatorname{Pr}[|X| \in(x, 1.001 x]] \geq \frac{1.999}{\pi} \cdot \frac{\ln (1.001)}{x}
$$

Proof. When $x>1549, \frac{2}{\pi} \cdot \frac{\ln (1.001)}{x} \geq \frac{2}{\pi} \cdot\left(\tan ^{-1}(1.001 x)-\tan ^{-1}(x)\right) \geq \frac{1.999}{\pi} \cdot \frac{\ln (1.001)}{x}$.

We build a level set for the "large" noise values, and we show the bad region cannot cover much of the large noise. The reason is that the bad region is small, and for each row, there is always some large noise.

Lemma 19.4.9. Given a matrix $U \in \mathbb{R}^{n \times r}$ with $n$ sufficiently large, let $r=n^{o(1)}$, and consider a random matrix $\Delta \in \mathbb{R}^{n \times(n-r)}$ with $\Delta_{i, j} \sim C(0,1)$ independently. Let $L_{t}=$ $\left\{(i, j)\left|(i, j) \in[n] \times[n-r],\left|\Delta_{i, j}\right| \in\left(1.001^{t}, 1.001^{t+1}\right]\right\}\right.$. With probability at least $1-1 / 2^{n^{\Theta(1)}}$, for all $t \in\left(\frac{1.0002 \ln n}{\ln 1.001}, \frac{1.9999 \ln n}{\ln 1.001}\right) \cap \mathbb{N}$,

$$
\left|L_{t} \backslash(\mathcal{B}(U) \times[n-r])\right| \geq n(n-r) \cdot 1.998 \cdot \ln (1.001) /\left(\pi \cdot 1.001^{t}\right)
$$

Proof. Let $N=n \cdot(n-r)$. Then according to Claim 19.4.8, $\forall t \in\left(\frac{1.0002 \ln n}{\ln 1.001}, \frac{1.9999 \ln n}{\ln 1.001}\right) \cap \mathbb{N}$, $\mathbb{E}\left(\left|L_{t}\right|\right) \geq N \cdot 1.999 \cdot \ln (1.001) /\left(\pi \cdot 1.001^{t}\right) \geq n^{\Theta(1)}$. For a fixed $t$, by a Chernoff bound, with probability at least $1-1 / 2^{n^{\Theta(1)}},\left|L_{t}\right| \geq N \cdot 1.9989 \cdot \ln (1.001) /\left(\pi \cdot 1.001^{t}\right)$. Due to Lemma 19.4.7, $|\mathcal{B}(U) \times[n-r]| \leq n^{0.99999}(n-r)=N / n^{0.00001}$. Due to the Chernoff bound, with probability at least $1-1 / 2^{n^{\ominus(1)}},\left|L_{t} \cap(\mathcal{B}(U) \times[n-r])\right|<N / n^{0.00001} \cdot 2.0001 \cdot \ln (1.001) /\left(\pi \cdot 1.001^{t}\right)$. Thus, with probability at least $1-1 / 2^{n^{\Theta(1)}},\left|L_{t} \backslash(\mathcal{B}(U) \times[n-r])\right| \geq N \cdot 1.998 \cdot \ln (1.001) /\left(\pi \cdot 1.001^{t}\right)$. By taking a union bound over all $t \in\left(\frac{1.0002 \ln n}{\ln 1.001}, \frac{1.9999 \ln n}{\ln 1.001}\right) \cap \mathbb{N}$, we complete the proof.

Lemma 19.4.10 (The cost of the large noise part). Let $n \geq 1$ be sufficiently large, and let $r=n^{o(1)}$. Given a matrix $U \in \mathbb{R}^{n \times r}$, and a random matrix $\Delta \in \mathbb{R}^{n \times(n-r)}$ with $\Delta_{i, j} \sim C(0,1)$ independently, let $\mathcal{J}=\left\{(i, j) \in[n] \times[n-r]| | \Delta_{i, j} \mid \geq n^{1.0002}\right\}$. If $\|\Delta\|_{1} \leq 4 n^{2} \ln n$, then with probability at least $1-1 / 2^{n^{\Theta(1)}}$, for all $X \in \mathbb{R}^{r \times n}$, either

$$
\sum_{(i, j) \in \mathcal{J}}\left|(U X-\Delta)_{i, j}\right|>\frac{1.996}{\pi} n^{2} \ln n,
$$

or

$$
\|U X-\Delta\|_{1}>4 n^{2} \ln n
$$

Proof.

$$
\begin{align*}
& \sum_{(i, j) \in \mathcal{J}}\left|(U X-\Delta)_{i, j}\right| \\
\geq & \sum_{(i, j) \in \mathcal{J} \backslash \mathcal{B}(U)}\left|(U X-\Delta)_{i, j}\right| \\
\geq & \sum_{(i, j) \in \mathcal{J} \backslash \mathcal{B}(U)}\left|(\Delta)_{i, j}\right|-\sum_{(i, j) \in \mathcal{J} \backslash \mathcal{B}(U)}\left|(U X)_{i, j}\right| \tag{19.6}
\end{align*}
$$

Let $N=n(n-r)$. By a Chernoff bound and the cumulative distribution function of a Cauchy random variable, with probability at least $1-1 / 2^{n^{\Theta(1)}}$, $|\mathcal{J}| \leq 1.1 \cdot N / n^{1.0002}$. If $\exists(i, j) \in \mathcal{J} \backslash \mathcal{B}(U)$ which has $\left|(U X)_{i, j}\right|>n^{1.0001}$, then according to the definition of $\mathcal{B}(U)$, $\|U X\|_{1} \geq\left\|(U X)_{j}\right\|_{1} \geq 8 n^{2} \ln n$. Due to the triangle inequality, $\|U X-\Delta\|_{1} \geq\|U X\|_{1}-$ $\|\Delta\|_{1} \geq 4 n^{2} \ln n$. If $\forall(i, j) \in \mathcal{J} \backslash \mathcal{B}(U)$ we have $\left|(U X)_{i, j}\right| \leq n^{1.0001}$, then

$$
\begin{equation*}
\sum_{(i, j) \in \mathcal{J} \backslash \mathcal{B}(U)}\left|(U X)_{i, j}\right| \leq|\mathcal{J}| \cdot n^{1.0001} \leq 1.1 \cdot N / n^{0.0001} \tag{19.7}
\end{equation*}
$$

Due to Lemma 19.4.9, with probability at least $1-1 / 2^{n^{\ominus(1)}}$,

$$
\begin{align*}
& \sum_{(i, j) \in \mathcal{J} \backslash \mathcal{B}(U)}\left|(\Delta)_{i, j}\right| \\
\geq & \sum_{t \in\left(\frac{1.0002 \ln n}{\ln 1.001}, \frac{1.9999 \ln n}{\ln 1.001}\right) \cap \mathbb{N}} 1.001^{t} \cdot N \cdot 1.998 \cdot \ln (1.001) /\left(\pi \cdot 1.001^{t}\right) \\
\geq & \frac{1.997}{\pi} \cdot N \ln n . \tag{19.8}
\end{align*}
$$

We plug (19.7) and (19.8) into (19.6), from which we have

$$
\sum_{(i, j) \in \mathcal{J}}\left|(U X-\Delta)_{i, j}\right| \geq \frac{1.996}{\pi} n^{2} \ln n
$$

### 19.4.6 Cost from the Sign-Agreement Part of the Small-Entry Part

We use $-y$ to fit $\Delta$ (we think of $A_{S} \alpha=A_{S}^{*} \alpha-y$, and want to minimize $\|-y-\Delta\|_{1}$ ). If the sign of $y_{j}$ is the same as the sign of $\Delta_{j}$, then both coordinate values will collectively contribute.

Lemma 19.4.11 (The contribution from $\Delta_{i}$ when $\Delta_{i}$ and $y_{i}$ have the same sign). Suppose we are given a vector $y \in \mathbb{R}^{n}$ and a random vector $\Delta \in \mathbb{R}^{n}$ with $\Delta_{j} \sim C(0,1)$ independently. Then with probability at least $1-1 / 2^{n^{\ominus(1)}}$,

$$
\sum_{j: \operatorname{sign}\left(y_{j}\right)=\operatorname{sign}\left(\Delta_{j}\right) \text { and }\left|\Delta_{j}\right| \leq n^{0.9999}}\left|\Delta_{j}\right|>\frac{0.9998}{\pi} n \ln n .
$$

Proof. For $j \in[n]$, define the random variable

$$
Z_{j}= \begin{cases}\Delta_{j} & 0<\Delta_{j} \leq n^{0.9999} \\ 0 & \text { otherwise }\end{cases}
$$

Then, we have

$$
\operatorname{Pr}\left[\sum_{j: \operatorname{sign}\left(y_{j}\right)=\operatorname{sign}\left(\Delta_{i, j}\right) \text { and }\left|\Delta_{j}\right| \leq n^{0.9999}}\left|\Delta_{j}\right|>\frac{0.9998}{\pi} n \ln n\right]=\operatorname{Pr}\left[\sum_{j=1}^{n} Z_{j}>\frac{0.9998}{\pi} n \ln n\right] .
$$

Let $B=n^{0.9999}$. For $j \in[n]$,

$$
\mathbb{E}\left[Z_{j}\right]=\frac{1}{\pi} \int_{0}^{B} \frac{x}{1+x^{2}} \mathrm{~d} x=\frac{1}{2 \pi} \ln \left(B^{2}+1\right) .
$$

Also,

$$
\mathbb{E}\left[Z_{j}^{2}\right]=\frac{1}{\pi} \int_{0}^{B} \frac{x^{2}}{1+x^{2}} \mathrm{~d} x=\frac{B-\tan ^{-1}(B)}{\pi} \leq B
$$

By Bernstein's inequality,

$$
\begin{aligned}
& \operatorname{Pr}\left[\mathbb{E}\left[\sum_{j=1}^{n} Z_{j}\right]-\sum_{j=1}^{n} Z_{j}>10^{-5} \mathbb{E}\left[\sum_{j=1}^{n} Z_{j}\right]\right] \\
\leq & \exp \left(-\frac{0.5 \cdot\left(10^{-5} \mathbb{E}\left[\sum_{j=1}^{n} Z_{j}\right]\right)^{2}}{\sum_{j=1}^{n} \mathbb{E}\left[Z_{j}^{2}\right]+\frac{1}{3} B \cdot 10^{-5} \mathbb{E}\left[\sum_{j=1}^{n} Z_{j}\right]}\right) \\
\leq & \exp \left(-\frac{5 \cdot 10^{-11} n^{2} \ln ^{2}\left(B^{2}+1\right) /\left(4 \pi^{2}\right)}{n B+\frac{1}{3} B \cdot 10^{-5} n \ln \left(B^{2}+1\right) /(2 \pi)}\right) \\
\leq & e^{-n^{\ominus(1)}}
\end{aligned}
$$

The last inequality follows since $B=n^{0.9999}$. Thus, we have

$$
\operatorname{Pr}\left[\sum_{j=1}^{n} Z_{j}<0.9998 / \pi \cdot n \ln n\right] \leq \operatorname{Pr}\left[\sum_{j=1}^{n} Z_{j}<0.99999 n \ln \left(B^{2}+1\right) /(2 \pi)\right] \leq e^{-n^{\Theta(1)}}
$$

Lemma 19.4.12 (Bound on level sets of a Cauchy vector). Suppose we are given a random vector $y \in \mathbb{R}^{n}$ with $y_{i} \sim C(0,1)$ chosen independently. Let

$$
L_{t}^{-}=\left\{i \in[n] \mid-y_{i} \in\left(1.001^{t}, 1.001^{t+1}\right]\right\} \text { and } L_{t}^{+}=\left\{i \in[n] \mid y_{i} \in\left(1.001^{t}, 1.001^{t+1}\right]\right\}
$$

With probability at least $1-1 / 2^{n^{\ominus(1)}}$, for all $t \in\left(\frac{\ln 1549}{\ln 1.001}, \frac{0.9999 \ln n}{\ln 1.001}\right) \cap \mathbb{N}$,

$$
\min \left(\left|L_{t}^{-}\right|,\left|L_{t}^{+}\right|\right) \geq 0.999 n \cdot \frac{1}{\pi} \frac{\ln 1.001}{1.001^{t}}
$$

Proof. For $i \in[n], t \geq \frac{\ln 1549}{\ln 1.001}$, according to Claim 19.4.8, $\operatorname{Pr}\left[y_{i} \in\left(1.001^{t}, 1.001^{t+1}\right]\right] \geq$ $0.9995 / \pi \cdot \ln (1.0001) / 1.001^{t}$. Thus, $\mathbb{E}\left[\left|L_{t}^{+}\right|\right]=\mathbb{E}\left[\left|L_{t}^{-}\right|\right]=n \cdot 0.9995 / \pi \cdot \ln (1.0001) / 1.001^{t}$.

Since $t \leq \frac{0.9999 \ln n}{\ln 1.001}, 1.001^{t} \leq n^{0.9999}$, we have $\mathbb{E}\left[\left|L_{t}^{+}\right|\right]=\mathbb{E}\left[\left|L_{t}^{-}\right|\right] \geq n^{\Theta(1)}$. By applying a Chernoff bound,

$$
\operatorname{Pr}\left[\left|L_{t}^{+}\right|>0.999 n / \pi \cdot \ln (1.0001) / 1.001^{t}\right] \geq 1-1 / 2^{n^{\ominus(1)}}
$$

Similarly, we have

$$
\operatorname{Pr}\left[\left|L_{t}^{-}\right|>0.999 n / \pi \cdot \ln (1.0001) / 1.001^{t}\right] \geq 1-1 / 2^{n^{\ominus(1)}}
$$

By taking a union bound over all the $L_{t}^{+}$and $L_{t}^{-}$, we complete the proof.

Lemma 19.4.13 (The contribution from $y_{i}$ when $\Delta_{i}$ and $y_{i}$ have the same sign). Let $u=$ $\eta \cdot \mathbf{1} \in \mathbb{R}^{n}$ where $\eta \in \mathbb{R}$ is an arbitrary real number. Let $y \in \mathbb{R}^{n}$ be a random vector with $y_{i} \sim C(0, \beta)$ independently for some $\beta>0$. Let $\Delta \in \mathbb{R}^{n}$ be a random vector with $\Delta_{i} \sim C(0,1)$ independently. With probability at least $1-1 / 2^{n^{\ominus(1)}}$,

$$
\sum_{i: \operatorname{sign}\left((u+y)_{i}\right)=\operatorname{sign}\left(\Delta_{i}\right) \text { and }\left|\Delta_{i}\right| \leq n^{0.9999}}\left|(u+y)_{i}\right| \geq \beta \cdot \frac{0.997}{\pi} n \ln n .
$$

Proof. For all $t \in\left(\frac{\ln 1549}{\ln 1.001}, \frac{0.9999 \ln n}{\ln 1.001}\right) \cap \mathbb{N}$, define

$$
\begin{aligned}
& L_{t}^{-}=\left\{i \in[n] \mid-y_{i} \in\left(\beta \cdot 1.001^{t}, \beta \cdot 1.001^{t+1}\right]\right\} \\
& L_{t}^{+}=\left\{i \in[n] \mid+y_{i} \in\left(\beta \cdot 1.001^{t}, \beta \cdot 1.001^{t+1}\right]\right\}
\end{aligned}
$$

Define

$$
G=\left\{i \in[n] \mid \operatorname{sign}\left((u+y)_{i}\right)=\operatorname{sign}\left(\Delta_{i}\right) \text { and }\left|\Delta_{i}\right| \leq n^{0.9999}\right\}
$$

Then $\forall i \in[n], \operatorname{Pr}[i \in G] \geq 0.5-1 / n^{0.9999} \geq 0.4999999999$. Due to Lemma 19.4.12,

$$
\min \left(\left|L_{t}^{-}\right|,\left|L_{t}^{+}\right|\right) \geq 0.999 n \cdot \frac{1}{\pi} \frac{\ln 1.001}{1.001^{t}} \geq n^{\Theta(1)}
$$

By the Chernoff bound and a union bound, with probability at least $1-1 / 2^{n^{\ominus(1)}}$, $\forall t \in$ $\left(\frac{\ln 1549}{\ln 1.001}, \frac{0.9999 \ln n}{\ln 1.001}\right) \cap \mathbb{N}$,

$$
\begin{equation*}
\min \left(\left|L_{t}^{-} \cap G\right|,\left|L_{t}^{+} \cap G\right|\right) \geq 0.499 n \cdot \frac{1}{\pi} \frac{\ln 1.001}{1.001^{t}} \tag{19.9}
\end{equation*}
$$

Then we have

$$
\begin{aligned}
& \sum_{i \in G}\left|(u+y)_{i}\right| \\
\geq & \sum_{t \in\left(\frac{\ln 1599}{\ln 1.001}, \frac{0.9999 \ln n}{\ln 1.001}\right) \cap \mathbb{N}}\left(\sum_{i \in L_{t}^{+}, i \in G}\left|y_{i}+\eta\right|+\sum_{i \in L_{t}^{-}, i \in G}\left|-y_{i}-\eta\right|\right) \\
\geq & \sum_{t \in\left(\frac{\ln 1549}{\ln 1.001}, \frac{0.9999 \ln n}{\ln 1.001}\right) \cap \mathbb{N}} 0.499 n \cdot \frac{1}{\pi} \frac{\ln 1.001}{1.001^{t}} \cdot 2 \cdot 1.001^{t} \cdot \beta \\
\geq & \beta \cdot \frac{0.997}{\pi} n \ln n
\end{aligned}
$$

The second inequality follows by Equation (19.9) and the triangle inequality, i.e., $\forall a, b, c \in$ $\mathbb{R},|a+c|+|b-c| \geq|a+b|$.

### 19.4.7 Cost from the Sign-Disagreement Part of the Small-Entry Part

Lemma 19.4.14. Given a vector $y \in \mathbb{R}^{n}$ and a random vector $\Delta \in \mathbb{R}^{n}$ with $\Delta_{i} \sim C(0,1)$ independently, with probability at least $1-1 / 2^{n^{\ominus(1)}}$,

$$
\sum_{i: \operatorname{sign}\left(y_{i}\right) \neq \operatorname{sign}\left(\Delta_{i}\right) \text { and }\left|\Delta_{i}\right|<n^{0.9999}}\left|y_{i}+\Delta_{i}\right|>\frac{0.03}{\pi} n \ln n .
$$

Proof. For $t \in\left[0, \frac{0.9999 \ln n}{\ln 4}\right) \cap \mathbb{N}$ define

$$
L_{t}=\left\{i \in[n]\left|\operatorname{sign}\left(y_{i}\right) \neq \operatorname{sign}\left(\Delta_{i}\right),\left|\Delta_{i}\right| \in\left(4^{t}, 4^{t+1}\right],\left|\Delta_{i}\right| \notin\left[\left|y_{i}\right|-4^{t},\left|y_{i}\right|+4^{t}\right]\right\} .\right.
$$

$\forall x \geq 1, y>0$, we have

$$
\begin{aligned}
& \operatorname{Pr}_{X \sim C(0,1)}[|X| \in(x, 4 x],|X| \notin[y-x, y+x]] \\
\geq & \operatorname{Pr}_{X \sim C(0,1)}[|X| \in(3 x, 4 x]] \\
= & \frac{2}{\pi} \cdot\left(\tan ^{-1}(4 x)-\tan ^{-1}(3 x)\right) \\
\geq & \frac{0.1}{\pi} \cdot \frac{\ln (4)}{x}
\end{aligned}
$$

Thus, $\forall i \in[n], t \in\left[0, \frac{0.9999 \ln n}{\ln 4}\right) \cap \mathbb{N}$,

$$
\operatorname{Pr}\left[i \in L_{t}\right] \geq \frac{0.05}{\pi} \cdot \frac{\ln (4)}{4^{t}}
$$

Thus, $\forall t \in\left[0, \frac{0.9999 \ln n}{\ln 4}\right) \cap \mathbb{N}$,

$$
\mathbb{E}\left[\left|L_{t}\right|\right] \geq 0.05 n / \pi \cdot \ln (4) / 4^{t} \geq n^{\Theta(1)}
$$

By a Chernoff bound and a union bound, with probability at least $1-1 / 2^{n^{\ominus(1)}} \forall t \in$ $\left[0, \frac{0.9999 \ln n}{\ln 4}\right) \cap \mathbb{N}$,

$$
\left|L_{t}\right| \geq 0.04 n / \pi \cdot \ln (4) / 4^{t}
$$

Thus, we have, with probability at least $1-1 / 2^{n^{\ominus(1)}}$,

$$
\begin{aligned}
& \sum_{i: \operatorname{sign}\left(y_{i}\right) \neq \operatorname{sign}\left(\Delta_{i}\right) \text { and }\left|\Delta_{i}\right|<n^{0.9999}}\left|y_{i}+\Delta_{i}\right| \\
\geq & \sum_{t \in\left[0, \frac{0.999 \ln n}{\ln 4}\right) \cap \mathbb{N}}\left|L_{t}\right| \cdot 4^{t} \\
\geq & \frac{0.03}{\pi} n \ln n .
\end{aligned}
$$

### 19.4.8 Overall Cost of the Small-Entry Part

Lemma 19.4.15 (For each). Let $u=\eta \cdot \mathbf{1} \in \mathbb{R}^{n}$ where $\eta \in \mathbb{R}$ is an arbitrary real number. Let $\alpha \in \mathbb{R}^{d}$ where $\|\alpha\|_{1} \geq 1-10^{-20}$. Let $\Delta \in \mathbb{R}^{n \times(d+1)}$ and $\forall(i, j) \in[n] \times[d+1], \Delta_{i, j} \sim C(0,1)$ are i.i.d. standard Cauchy random variables. Then with probability at least $1-1 / 2^{n^{\ominus(1)}}$,

$$
\sum_{j \in[n],\left|\Delta_{j, d+1}\right|<n^{0.9999}}\left|\left(u+\Delta_{d+1}-\left(u \mathbf{1}^{\top}+\Delta_{[d]}\right) \alpha\right)_{j}\right| \geq \frac{2.025}{\pi} n \ln n
$$

Proof. Let $G_{1}$ and $G_{2}$ be defined as

$$
\begin{aligned}
& \left.G_{1}=\left\{j \in[n]| | \Delta_{j, d+1} \mid<n^{0.9999}, \operatorname{sign}\left(\left(u\left(1-1^{\top} \alpha\right)-\Delta_{[d]} \alpha\right)_{j}\right)=\operatorname{sign}\left(\Delta_{d+1}\right)_{j}\right)\right\}, \\
& \left.G_{2}=\left\{j \in[n]| | \Delta_{j, d+1} \mid<n^{0.9999}, \operatorname{sign}\left(\left(u\left(1-\mathbf{1}^{\top} \alpha\right)-\Delta_{[d]} \alpha\right)_{j}\right) \neq \operatorname{sign}\left(\Delta_{d+1}\right)_{j}\right)\right\} .
\end{aligned}
$$

Notice that $\Delta_{[d]} \alpha$ is a random vector with each entry independently drawn from $C\left(0,\|\alpha\|_{1}\right)$. Then with probability at least $1-1 / 2^{n^{\ominus(1)}}$,

$$
\begin{aligned}
& \sum_{j \in[n],\left|\Delta_{j, d+1}\right|<n^{0.9999}}\left|\left(u+\Delta_{d+1}-\left(u \mathbf{1}^{\top}+\Delta_{[d]}\right) \alpha\right)_{j}\right| \\
= & \sum_{j \in[n],\left|\Delta_{j, d+1}\right|<n^{0.9999}}\left|\left(u\left(1-\mathbf{1}^{\top} \alpha\right)-\Delta_{[d]} \alpha+\Delta_{d+1}\right)_{j}\right| \\
= & \sum_{j \in G_{1}}\left|\left(u\left(1-\mathbf{1}^{\top} \alpha\right)-\Delta_{[d]} \alpha+\Delta_{d+1}\right)_{j}\right|+\sum_{j \in G_{2}}\left|\left(u\left(1-\mathbf{1}^{\top} \alpha\right)-\Delta_{[d]} \alpha+\Delta_{d+1}\right)_{j}\right| \\
= & \sum_{j \in G_{1}}\left|\left(u\left(1-\mathbf{1}^{\top} \alpha\right)-\Delta_{[d]} \alpha\right)_{j}\right|+\sum_{j \in G_{1}}\left|\left(\Delta_{d+1}\right)_{j}\right|+\sum_{j \in G_{2}}\left|\left(u\left(1-\mathbf{1}^{\top} \alpha\right)-\Delta_{[d]} \alpha+\Delta_{d+1}\right)_{j}\right| \\
\geq & \|\alpha\|_{1} \cdot \frac{0.997}{\pi} \cdot n \ln n+\frac{0.9998}{\pi} n \ln n+\frac{0.03}{\pi} n \ln n \\
\geq & \frac{2.025}{\pi} n \ln n
\end{aligned}
$$

The first inequality follows by Lemma 19.4.13, Lemma 19.4.11 and Lemma 19.4.14. The second inequality follows by $\|\alpha\|_{1} \geq 1-10^{-20}$.

Lemma 19.4.16 (For all). Let $c>0, c_{0}>0$ be two arbitrary constants. Let $u=\eta \cdot \mathbf{1} \in \mathbb{R}^{n}$ where $\eta \in \mathbb{R}$ satisfies $|\eta| \leq n^{c_{0}}$. Consider a random matrix $\Delta \in \mathbb{R}^{n \times(d+1)}$ with $d=n^{o(1)}$ and $\forall(i, j) \in[n] \times[d+1], \Delta_{i, j} \sim C(0,1)$ are i.i.d. standard Cauchy random variables. Conditioned on $\|\Delta\|_{1} \leq n^{3}$, with probability at least $1-1 / 2^{n^{\Theta(1)}}, \forall \alpha \in \mathbb{R}^{d}$ with $1-10^{-20} \leq\|\alpha\|_{1} \leq n^{c}$,

$$
\sum_{j \in[n],\left|\Delta_{j, d+1}\right|<n^{0.9999}}\left|\left(u+\Delta_{d+1}-\left(u \mathbf{1}^{\top}+\Delta_{[d]}\right) \alpha\right)_{j}\right| \geq \frac{2.024}{\pi} n \ln n
$$

Proof. Let $\mathcal{N}$ be a set of points:

$$
\mathcal{N}=\left\{\alpha \in \mathbb{R}^{d} \mid 1-10^{-20} \leq\|\alpha\|_{1} \leq n^{c} \text { and } \exists q \in \mathbb{Z}^{d}, \text { such that } \alpha=q / n^{c+c_{0}+1000}\right\}
$$

Since $d=n^{o(1)}$, we have $|\mathcal{N}| \leq\left(n^{2 c+c_{0}+2000}\right)^{d}=2^{n^{o(1)}}$. By Lemma 19.4.15 and a union bound, with probability at least $1-1 / 2^{n^{\Theta(1)}} \cdot|\mathcal{N}| \geq 1-1 / 2^{n^{\ominus(1)}}, \forall \alpha \in \mathcal{N}$, we have

$$
\sum_{j \in[n],\left|\Delta_{j, d+1}\right|<n^{0.9999}}\left|\left(u+\Delta_{d+1}-\left(u \mathbf{1}^{\top}+\Delta_{[d]}\right) \alpha\right)_{j}\right| \geq \frac{2.025}{\pi} n \ln n
$$

Due to the construction of $\mathcal{N}$, we have $\forall \alpha \in \mathbb{R}^{d}$ with $1-10^{-20} \leq\|\alpha\|_{1} \leq n^{c}, \exists \alpha^{\prime} \in \mathcal{N}$ such that $\left\|\alpha-\alpha^{\prime}\right\|_{\infty} \leq 1 / n^{c+c_{0}+1000}$. Let $\gamma=\alpha-\alpha^{\prime}$. Then

$$
\begin{aligned}
& \sum_{j \in[n],\left|\Delta_{j, d+1}\right|<n^{0.9999}}\left|\left(u+\Delta_{d+1}-\left(u \mathbf{1}^{\top}+\Delta_{[d]}\right) \alpha\right)_{j}\right| \\
= & \sum_{j \in[n],\left|\Delta_{j, d+1}\right|<n^{0.9999}}\left|\left(u+\Delta_{d+1}-\left(u \mathbf{1}^{\top}+\Delta_{[d]}\right)\left(\alpha^{\prime}+\gamma\right)\right)_{j}\right| \\
\geq & \sum_{j \in[n],\left|\Delta_{j, d+1}\right|<n^{0.9999}}\left|\left(u+\Delta_{d+1}-\left(u \mathbf{1}^{\top}+\Delta_{[d]}\right) \alpha^{\prime}\right)_{j}\right|-\sum_{j \in[n],\left|\Delta_{j, d+1}\right|<n^{0.9999}}\left|\left(\left(u \mathbf{1}^{\top}+\Delta_{[d]}\right) \gamma\right)_{j}\right| \\
\geq & \sum_{j \in[n],\left|\Delta_{j, d+1}\right|<n^{0.9999}}\left|\left(u+\Delta_{d+1}-\left(u \mathbf{1}^{\top}+\Delta_{[d]}\right) \alpha^{\prime}\right)_{j}\right|-\left\|\left(u \mathbf{1}^{\top}+\Delta_{[d]}\right) \gamma\right\|_{1} \\
\geq & \frac{2.025}{\pi} n \ln n-1 / n^{500} \\
\geq & \frac{2.024}{\pi} n \ln n
\end{aligned}
$$

The first equality follows from $\alpha=\alpha^{\prime}+\gamma$. The first inequality follows by the triangle inequality. The third inequality follows from $\|\gamma\|_{1} \leq 1 / n^{c+c_{0}+800},\left\|u \mathbf{1}^{\top}\right\|_{1} \leq n^{c_{0}+10},\|\Delta\|_{1} \leq$ $n^{3}$, and $\forall \alpha^{\prime} \in \mathcal{N}$,

$$
\sum_{j \in[n],\left|\Delta_{j, d+1}\right|<n^{0.9999}}\left|\left(u+\Delta_{d+1}-\left(u \mathbf{1}^{\top}+\Delta_{[d]}\right) \alpha^{\prime}\right)_{j}\right| \geq \frac{2.025}{\pi} n \ln n .
$$

### 19.4.9 Main result

Theorem 19.4.17 (Formal version of Theorem 19.1.2). Let $n>0$ be sufficiently large. Let $A=\eta \cdot \mathbf{1} \cdot \mathbf{1}^{\top}+\Delta \in \mathbb{R}^{n \times n}$ be a random matrix where $\eta=n^{c_{0}}$ for some sufficiently large constant $c_{0}$, and $\forall i, j \in[n], \Delta_{i, j} \sim C(0,1)$ are i.i.d. standard Cauchy random variables. Let $r=n^{o(1)}$. Then with probability at least $1-O(1 / \log \log n), \forall S \subset[n]$ with $|S|=r$,

$$
\min _{X \in \mathbb{R}^{r \times n}}\left\|A_{S} X-A\right\|_{1} \geq 1.002\|\Delta\|_{1}
$$

Proof. We first argue that for a fixed set $S$, conditioned on $\|\Delta\|_{1} \leq 100 n^{2} \ln n$, with probability at least $1-1 / 2^{n^{\Theta(1)}}$,

$$
\min _{X \in \mathbb{R}^{r \times n}}\left\|A_{S} X-A\right\|_{1} \geq 1.002\|\Delta\|_{1}
$$

Then we can take a union bound over the at most $n^{r}=2^{n^{o(1)}}$ possible choices of $S$. It suffices to show for a fixed set $S, \min _{X \in \mathbb{R}^{r \times n}}\left\|A_{S} X-A\right\|_{1}$ is not small.

Without loss of generality, let $S=[r]$, and we want to argue the cost

$$
\min _{X \in \mathbb{R}^{r \times n}}\left\|A_{S} X-A\right\|_{1} \geq \min _{X \in \mathbb{R}^{r} \times n}\left\|A_{S} X_{[n] \backslash S}-A_{[n] \backslash S}\right\|_{1} \geq 1.002\|\Delta\|_{1} .
$$

Due to Lemma 19.4.2, with probability at least $1-O(1 / \log \log n)$,

$$
\|\Delta\|_{1} \leq 4.0002 / \pi \cdot n^{2} \ln n
$$

Now, we can condition on $\|\Delta\|_{1} \leq 4.0002 / \pi \cdot n^{2} \ln n$.
Consider $j \in[n] \backslash S$. Due to Lemma 19.4.5, with probability at least $1-(1 / n)^{\Theta(n)}$, for all $X_{j} \in \mathbb{R}^{r}$ with $\left\|X_{j}\right\|_{1} \geq n^{c}$ for some constant $c>0$, we have

$$
\left\|A_{S} X_{j}-A_{j}\right\|_{1}=\left\|\left(\eta \cdot \mathbf{1} \cdot \mathbf{1}^{\top}+\left[\Delta_{S} \Delta_{j}\right]\right)\left[X_{j}^{\top}-1\right]^{\top}\right\|_{1} \geq 0.9 n^{3}
$$

By taking a union bound over all $j \in[n] \backslash S$, with probability at least $1-(1 / n)^{\Theta(n)}$, for all $X \in \mathbb{R}^{r \times n}$ with $\exists j \in[n] \backslash S,\left\|X_{j}\right\|_{1} \geq n^{c}$, we have

$$
\left\|A_{S} X-A\right\|_{1} \geq 0.9 n^{3}
$$

Thus, we only need to consider the case $\forall j \in[n] \backslash S,\left\|X_{j}\right\|_{1} \leq n^{c}$. Notice that we condition on $\|\Delta\|_{1} \leq 4.0002 / \pi \cdot n^{2} \ln n$. By Fact 19.4.1, we have that if $\left\|X_{j}\right\|_{1} \leq n^{c}$ and $\left|1-\mathbf{1}^{\top} X_{j}\right|>$ $1-10^{-20}$, then

$$
\left\|A_{S} X-A\right\|_{1} \geq\left\|A_{S} X_{j}-A_{j}\right\|_{1}>n^{3}
$$

Thus, we only need to consider the case $\forall j \in[n] \backslash S,\left\|X_{j}\right\|_{1} \leq n^{c},\left|1-\mathbf{1}^{\top} X_{j}\right| \leq 1-10^{-20}$. $\forall X \in \mathbb{R}^{r \times n}$ with $\forall j \in[n] \backslash S,\left\|X_{j}\right\|_{1} \leq n^{c},\left|1-\mathbf{1}^{\top} X_{j}\right| \leq 1-10^{-20}$, if $\left\|A_{S} X_{[n] \backslash S}-A_{[n] \backslash S}\right\|_{1} \leq$
$4 n^{2} \ln n$, then

$$
\begin{aligned}
& \left\|A_{S} X_{[n] \backslash S}-A_{[n] \backslash S}\right\|_{1} \\
= & \left\|\left(\eta \cdot \mathbf{1} \cdot \mathbf{1}^{\top}+\Delta_{S}\right) X_{[n] \backslash S}-\left(\eta \cdot \mathbf{1} \cdot \mathbf{1}^{\top}+\Delta_{[n] \backslash S}\right)\right\|_{1} \\
\geq & \sum_{i \in[n], j \in[n] \backslash S,\left|\Delta_{i, j}\right| \geq n^{1.0002}}\left|\left(\left(\left(\eta \cdot \mathbf{1} \cdot \mathbf{1}^{\top}+\Delta_{S}\right) X_{[n] \backslash S}-\eta \cdot \mathbf{1} \cdot \mathbf{1}^{\top}\right)-\Delta_{[n] \backslash S}\right)_{i, j}\right| \\
& +\sum_{i \in[n], j \in[n] \backslash S,\left|\Delta_{i, j}\right|<n^{0.9999}}\left|\left(\left(\eta \cdot \mathbf{1} \cdot \mathbf{1}^{\top}+\Delta_{S}\right) X_{[n] \backslash S}-\left(\eta \cdot \mathbf{1} \cdot \mathbf{1}^{\top}+\Delta_{[n] \backslash S}\right)\right)_{i, j}\right| \\
\geq & \frac{1.996}{\pi} \cdot n^{2} \ln n+\sum_{i \in[n], j \in[n] \backslash S,\left|\Delta_{i, j}\right|<n^{0.9999}}\left|\left(\left(\eta \cdot \mathbf{1} \cdot \mathbf{1}^{\top}+\Delta_{S}\right) X_{[n] \backslash S}-\left(\eta \cdot \mathbf{1} \cdot \mathbf{1}^{\top}+\Delta_{[n] \backslash S}\right)\right)_{i, j}\right| \\
= & \frac{1.996}{\pi} \cdot n^{2} \ln n+\sum_{j \in[n] \backslash S} \sum_{i \in[n],\left|\Delta_{i, j}\right|<n^{0.9999}}\left|\left(\left(\eta \cdot \mathbf{1} \cdot \mathbf{1}^{\top}+\Delta_{S}\right) X_{j}-\eta \cdot \mathbf{1}-\Delta_{j}\right)_{i}\right| \\
\geq & \frac{1.996}{\pi} \cdot n^{2} \ln n+\sum_{j \in[n] \backslash S} \frac{2.024}{\pi} n \ln n \\
\geq & \frac{1.996}{\pi} \cdot n^{2} \ln n+\frac{2.023}{\pi} n^{2} \ln n \\
\geq & \frac{4.01}{\pi} \cdot n^{2} \ln n
\end{aligned}
$$

holds with probability at least $1-1 / 2^{n^{\ominus}(1)}$. The first equality follows by the definition of $A$. The first inequality follows by the partition by $\left|\Delta_{i, j}\right|$. Notice that $\left[\mathbf{1} \Delta_{S}\right]$ has rank at most $r+1=n^{o(1)}$. Then, due to Lemma 19.4.10, and the condition $\left\|A_{S} X_{[n] \backslash S}-A_{[n] \backslash S}\right\|_{1} \leq 4 n^{2} \ln n$, the second inequality holds with probability at least $1-1 / 2^{n^{\ominus(1)}}$. The second equality follows by grouping the cost by each column. The third inequality holds with probability at least $1-1 / 2^{n^{\Theta(1)}}$ by Lemma 19.4.16, and a union bound over all the columns in $[n] \backslash S$. The fourth inequality follows by $n-r=n-n^{o(1)} \geq\left(1-10^{-100}\right) n$.

Thus, conditioned on $\|\Delta\|_{1} \leq 4.0002 / \pi \cdot n^{2} \ln n$, with probability at least $1-1 / 2^{n^{\ominus(1)}}$,
we have

$$
\min _{X \in \mathbb{R}^{r \times n}}\left\|A_{S} X-A\right\|_{1} \geq \frac{4.02}{\pi} \cdot n^{2} \ln n .
$$

By taking a union bound over all the $\binom{n}{r}=2^{n^{o(1)}}$ choices of $S$, we have that conditioned on $\|\Delta\|_{1} \leq \frac{4.0002}{\pi} n^{2} \ln n$, with probability at least $1-1 / 2^{n^{\ominus(1)}}, \forall S \subset[n]$ with $|S|=r=n^{o(1)}$, $\min _{X \in \mathbb{R}^{r \times n}}\left\|A_{S} X-A\right\|_{1} \geq \frac{4.02}{\pi} \cdot n^{2} \ln n$. Since $4.01 / 4.0002>1.002$,

$$
\min _{X \in \mathbb{R}^{r \times n}}\left\|A_{S} X-A\right\|_{1} \geq 1.002\|\Delta\|_{1} .
$$

Since $\|\Delta\|_{1} \leq \frac{4.0002}{\pi} n^{2} \ln n$ happens with probability at least $1-O(1 / \log \log n)$, this completes the proof.

## Chapter 20

## Towards a Zero-One Law for Low Rank Approximation

There are a number of approximation algorithms for NP-hard versions of low rank approximation, such as finding a rank- $k$ matrix $B$ minimizing the sum of absolute values of differences to a given $n$-by- $n$ matrix $A, \min _{\text {rank-k }}\|A-B\|_{1}$, or more generally finding a rank- $k$ matrix $B$ which minimizes the sum of $p$-th powers of absolute values of differences, $\min _{\text {rank-k } B}\|A-B\|_{p}^{p}$. Many of these algorithms are linear time columns subset selection algorithms, returning a subset of $\operatorname{poly}(k \log n)$ columns whose cost is no more than a poly $(k)$ factor larger than the cost of the best rank- $k$ matrix. The above error measures are special cases of the following general entrywise low rank approximation problem: given an arbitrary function $g: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$, find a rank- $k$ matrix $B$ which minimizes $\|A-B\|_{g}=\sum_{i, j} g\left(A_{i, j}-B_{i, j}\right)$. A natural question is which functions $g$ admit efficient approximation algorithms? Indeed, this is a central question of recent work studying generalized low rank models. In this work we give approximation algorithms for every function $g$ which is approximately monotone and satisfies an approximate triangle inequality, and we show both of these conditions are necessary. Further, our algorithm is efficient if the function $g$ admits an efficient approximate regression algorithm. Our approximation algorithms handle functions which are not even scale-invariant, such as the Huber loss function, which we show have very different structural properties than $\ell_{p}$-norms, e.g., one can show the lack of scale-invariance causes any column subset selection algorithm to provably require a $\sqrt{\log n}$ factor larger number of columns
than $\ell_{p}$-norms; nevertheless we design the first efficient column subset selection algorithms for such error measures.

### 20.1 Introduction

A well-studied problem in machine learning and numerical linear algebra, with applications to recommendation systems, text mining, and computer vision, is that of computing a low-rank approximation of a matrix. Such approximations reveal low-dimensional structure, provide a compact way of storing a matrix, and can quickly be applied to a vector.

A commonly used version of the problem is to compute a near optimal low-rank approximation with respect to the Frobenius norm. That is, given an $n \times n$ input matrix $A$ and an accuracy parameter $\epsilon>0$, output a rank- $k$ matrix $B$ with large probability so that $\|A-B\|_{F}^{2} \leq(1+\epsilon)\left\|A-A_{k}\right\|_{F}^{2}$, where for a matrix $C,\|C\|_{F}^{2}=\sum_{i, j} C_{i, j}^{2}$ is its squared Frobenius norm, and $A_{k}=\operatorname{argmin}_{\text {rank-k }}{ }_{B}\|A-B\|_{F} . A_{k}$ can be computed exactly using the singular value decomposition (SVD), but takes $O\left(n^{3}\right)$ time in practice and $n^{\omega}$ time in theory, where $\omega \approx 2.373$ is the exponent of matrix multiplication [Str69, CW87, Wil12, LG14].

Sárlos [Sar06] showed how to achieve the above guarantee with constant probability in $\widetilde{O}(\mathrm{nnz}(A) \cdot k / \epsilon)+n \cdot \operatorname{poly}(k / \epsilon)$ time, where $\mathrm{nnz}(A)$ denotes the number of non-zero entries of $A$. This was improved in [CW13, MM13, NN13a, BDN15, Coh16a] using sparse random projections in $O(\mathrm{nnz}(A))+n \cdot \operatorname{poly}(k / \epsilon)$ time. Large sparse datasets in recommendation systems are common, such as the Bookcrossing ( $100 \mathrm{~K} \times 300 \mathrm{~K}$ with $10^{6}$ observations) [ZMKL05] and Yelp datasets ( $40 \mathrm{~K} \times 10 \mathrm{~K}$ with $10^{5}$ observations) [Yel14], and this is a substantial improvement over the SVD.

Robust Low Rank Approximation. To understand the role of the Frobenius norm in the algorithms above, we recall a standard motivation for this error measure. Suppose one has $n$ data points in a $k$-dimensional subspace of $\mathbb{R}^{d}$, where $k \ll d$. We can write these points
as the rows of an $n \times d$ matrix $A^{*}$ which has rank $k$. The matrix $A^{*}$ is often called the ground truth matrix. In a number of settings, due to measurement noise or other kinds of noise, we only observe the matrix $A=A^{*}+\Delta$, where each entry of the noise matrix $\Delta \in \mathbb{R}^{n \times n}$ is an i.i.d. random variable from a certain mean-zero noise distribution $\mathcal{D}$. One method for approximately recovering $A^{*}$ from $A$ is maximum likelihood estimation. Here one tries to find a matrix $B$ maximizing the $\log$-likelihood: $\max _{\text {rank-k } B} \sum_{i, j} \log p\left(A_{i, j}-B_{i, j}\right)$, where $p(\cdot)$ is the probability density function of the underlying noise distribution $\mathcal{D}$. For example, when the noise distribution is Gaussian with mean zero and variance $\sigma^{2}$, denoted by $N\left(0, \sigma^{2}\right)$, then the optimization problem is $\max _{\text {rank-k } B} \sum_{i, j}\left(\log \left(1 / \sqrt{2 \pi \sigma^{2}}\right)-\left(A_{i, j}-B_{i, j}\right)^{2} /\left(2 \sigma^{2}\right)\right)$, which is equivalent to solving the Frobenius norm loss low rank approximation problem defined above.

The Frobenius norm loss, while having nice statistical properties for Gaussian noise, is well-known to be sensitive to outliers. Applying the same maximum likelihood framework above to other kinds of noise distributions results in minimizing other kinds of loss functions. In general, if the density function of the underlying noise $\mathcal{D}$ is $p(z)=c \cdot e^{-g(z)}$, where $c$ is a normalization constant, then the maximum likelihood estimation problem for this noise distribution becomes the following generalized entry-wise loss low rank approximation problem: $\min _{\text {rank-k }} \sum_{i, j} g\left(A_{i, j}-B_{i, j}\right)=\min _{\text {rank-k } B}\|A-B\|_{g}$, which is a central topic of recent work on generalized low-rank models $\left[\mathrm{UHZ}^{+} 16\right]$. For example, when the noise is Laplacian, the entrywise $\ell_{1}$ loss is the maximum likelihood estimation, which is also robust to sparse outliers. A natural setting is when the noise is a mixture of small Gaussian noise and sparse outliers; this noise distribution is referred to as the Huber density. In this case the Huber loss function gives the maximum likelihood estimate $\left[\mathrm{UHZ}^{+} 16\right]$, where the Huber
function [Hub64] is defined to be: $g(x)=x^{2} /(2 \tau)$ if $|x|<\tau / 2$, and $g(x)=|x|-\tau / 2$ if $|x| \geq \tau$. Another nice property of the Huber error measure is that it is differentiable everywhere, unlike the $\ell_{1}$-norm, yet still enjoys the robustness properties as one moves away from the origin, making it less sensitive to outliers than the $\ell_{2}$-norm. There are many other kinds of loss functions, known as $M$-estimators [Zha97], which are widely used as loss functions in robust statistics [HRRS11].

Although several specific cases have been studied, such as entry-wise $\ell_{p}$ loss [CLMW11, SWZ17, CGK+17a, BKW17, BBB ${ }^{+}$19a], weighted entry-wise $\ell_{2}$ loss [RSW16], and cascaded $\ell_{p}\left(\ell_{2}\right)$ loss [DVTV09, CW15a], the landscape of general entry-wise loss functions remains elusive. There are no results known for any loss function which is not scale-invariant, much less any kind of characterization of which loss functions admit efficient algorithms. This is despite the importance of these loss functions; we refer the reader to [UHZ $\left.{ }^{+} 16\right]$ for a survey of generalized low rank models. This motivates the main question in our work:

Question 20.1.1 (General Loss Functions). For a given approximation factor $\alpha>1$, which functions $g$ allow for efficient low-rank approximation algorithms? Formally, given an $n \times d$ matrix $A$, can we find a rank-k matrix $B$ for which $\|A-B\|_{g} \leq \alpha \min _{\text {rank }-k B^{\prime}}\left\|A-B^{\prime}\right\|_{g}$, where for a matrix $C,\|C\|_{g}=\sum_{i \in[n], j \in[d]} g\left(C_{i, j}\right)$ ? What if we also allow $B$ to have rank $\operatorname{poly}(k \log n)$ ?

For Question 20.1.1, one has $g(x)=|x|^{p}$ for $p$-norms, and note the Huber loss function also fits into this framework. Allowing $B$ to have slightly larger rank than $k$, namely, $\operatorname{poly}(k \log n)$, is often sufficient for applications as it still allows for the space savings and computational gains outlined above. These are referred to as bicriteria approximations and are the focus of our work.

Notation. Before we present our results, let us briefly introduce the notation. For $n \in \mathbb{Z}_{\geq 0}$, let $[n]$ denote the set $\{1,2, \cdots, n\}$. Let $A \in \mathbb{R}^{n \times m}$. $A_{i}$ and $A^{j}$ denote the $i^{\text {th }}$ column and the $j^{\text {th }}$ row of $A$ respectively. Let $P \subseteq[m], Q \subseteq[n]$. $A_{P}$ denotes the matrix which is composed by the columns of $A$ with column indices in $P$. Similarly, $A^{Q}$ denotes the matrix composed by the rows of $A$ with row indices in $Q$. Let $S$ be a set and $s \in \mathbb{Z}_{\geq 0}$. We use $\binom{S}{s}$ to denote the set of all the size- $s$ subsets of $S$.

### 20.1.1 Our Results

We give necessary and sufficient conditions for low rank approximation with respect to general error measures. Our algorithm is a column subset selection algorithm, returning a small subset of columns which span a good low rank approximation. Column subset selection has the benefit of preserving sparsity and interpretability, as described above. We describe two properties on the function $g$ that we need to obtain our low rank approximation algorithms, and show that if $g$ does not have either property, then there are matrices for which no small subset of columns spanning a good low rank approximation with respect to the error measure $g$ exists. We make these terms precise in the theorem statements below.

Since we obtain column subset selection algorithms for a wide class of functions, our algorithms must necessarily be bicriteria and have approximation factor at least poly $(k)$. Indeed, a special case of our class of functions includes entrywise $\ell_{1}$-low rank approximation, for which it was shown in Theorem G. 27 of [SWZ17] that any subset of poly $(k)$ columns incurs an approximation error of at least $k^{\Omega(1)}$. We also show that for the entrywise Huberlow rank approximation, already for $k=1, \sqrt{\log n}$ columns are needed to obtain any constant factor approximation, thus showing that for some of the functions we consider, a dependence
on $n$ in our column subset size is necessary.
We note that previously for almost all such functions, it was not known how to obtain any non-trivial approximation factor with any sublinear number of columns.

### 20.1.1.1 A Zero-One Law

We first state three general properties, the first two of which are structural properties and are necessary and sufficient for obtaining a good approximation from a small subset of columns. The third property is needed for efficient running time.

Approximate triangle inequality. For $t \in \mathbb{Z}_{>0}$, we say a function $g(x): \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ satisfies the ati $_{g, t}$-approximate triangle inequality if for any $x_{1}, x_{2}, \cdots, x_{t} \in \mathbb{R}, g\left(\sum x_{i}\right) \leq$ $\operatorname{ati}_{g, t} \cdot \sum g\left(x_{i}\right)$.

Monotone property. For any parameter $\operatorname{mon}_{g} \geq 1$, we say function $g(x): \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ is $\operatorname{mon}_{g}$-monotone if for any $x, y \in \mathbb{R}$ with $0 \leq|x| \leq|y|$, we have $g(x) \leq \operatorname{mon}_{g} \cdot g(y)$.

Many functions including most $M$-estimators [Zha97] and the quantile function [KBJ78] satisfy the above two properties. See Table 20.1 for several examples. We refer the reader to the later sections for the necessity of these two properties. Our next property is not structural, but rather states that if the loss function has an efficient regression algorithm, then that suffices to efficiently find a small subset of columns spanning a good low rank approximation.

Regression property. We say function $g(x): \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ has the $\left(\operatorname{reg}_{g, d}, \mathcal{T}_{\text {reg }, g, n, d, m}\right)$ regression property if the following holds: given two matrices $A \in \mathbb{R}^{n \times d}$ and $B \in \mathbb{R}^{n \times m}$, for each $i \in[m]$, let $\mathrm{OPT}_{i}$ denote $\min _{x \in \mathbb{R}^{d}}\left\|A x-B_{i}\right\|_{g}$. There is an algorithm that runs in

Table 20.1: Example functions satisfying both structural properties.

|  | $g(x)$ | $\mathrm{ati}_{g, t}$ | $\mathrm{mon}_{g}$ |
| :---: | :---: | :---: | :---: |
| Huber | $\begin{cases}x^{2} / 2 & \|x\| \leq \tau \\ \tau(\|x\|-\tau / 2) & \|x\|>\tau\end{cases}$ | $O(t)$ | 1 |
| $\ell_{p}(p \geq 1)$ | $\|x\|^{p} / p$ | $O\left(t^{p-1}\right)$ | 1 |
| $\ell_{1}-\ell_{2}$ | $2\left(\sqrt{1+x^{2} / 2}-1\right)$ | $O(t)$ | 1 |
| Geman-McClure | $x^{2} /\left(2+2 x^{2}\right)$ | $O(t)$ | 1 |
| "FAIR" | $\tau^{2}(\|x\| / \tau-\log (1+\|x\| / \tau))$ | $O(t)$ | 1 |
| Tukey | $\begin{cases}\tau^{2} / 6 \cdot\left(1-\left(1-(x / \tau)^{2}\right)^{3}\right) & \|x\| \leq \tau \\ \tau^{2} / 6 & \|x\|>\tau\end{cases}$ | $O(t)$ | 1 |
| CaUCHY | $\tau^{2} / 2 \cdot \log \left(1+(x / \tau)^{2}\right)$ | $O(t)$ | 1 |
| Quantile ( $\tau \in(0,1)$ ) | $\begin{cases}\tau x & x \geq 0 \\ (\tau-1) x & x<0\end{cases}$ | 1 | $\max \left(\frac{\tau}{1-\tau}, \frac{1-\tau}{\tau}\right)$ |

$\mathcal{T}_{\text {reg }, g, n, d, m}$ time and outputs a matrix $X^{\prime} \in \mathbb{R}^{d \times m}$ such that $\left\|A X_{i}^{\prime}-B_{i}\right\|_{g} \leq \operatorname{reg}_{g, d} . \mathrm{OPT}_{i}, \forall i \in$ $[m]$ and outputs a vector of estimated regression costs $v \in \mathbb{R}^{d}$ such that $\mathrm{OPT}_{i} \leq v_{i} \leq$ $\operatorname{reg}_{g, d} \cdot \mathrm{OPT}_{i}, \forall i \in[m]$. The success probability is at least $1-1 / \operatorname{poly}(n m)$.

Some functions for which regression itself is non-trivial are e.g., the $\ell_{0}$-loss function and Tukey function. The $\ell_{0}$-loss function corresponds to the nearest codeword problem over the reals and has slightly better than an $O(d)$-approximation ([BK02, APY09], see also [BKW17]). For the Tukey function, [CWW19] shows that Tukey regression is NP-hard, and it also gives approximation algorithms. For discussion on regression solvers, we refer the reader to Appendix 20.5.

For any function, as long as the above general three properties hold, we can provide an efficient algorithm, as our following main theorem shows.

Theorem 20.1.2. Given a matrix $A \in \mathbb{R}^{n \times n}$, let $k \geq 1, k^{\prime}=2 k+1$. Let $g: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$
denote a function satisfying the ati ${ }_{g, k^{\prime}}$-approximate triangle inequality, the mon $_{g}$-monotone property, and the $\left(\operatorname{reg}_{g, k^{\prime}}, \mathcal{T}_{\text {reg }, g, n, k^{\prime}, n}\right)$-regression property. Let $\mathrm{OPT}=\min _{\text {rank }-k A^{\prime}}\left\|A^{\prime}-A\right\|_{g}$. There is an algorithm that runs in $\widetilde{O}\left(n+\mathcal{T}_{\text {reg, }, n, k^{\prime}, n}\right)$ time and outputs a set $S \subseteq[n]$ with $|S|=O(k \log n)$ such that with probability at least 0.99,

$$
\min _{X \in \mathbb{R}^{|S| \times n}}\left\|A_{S} X-A\right\|_{g} \leq \operatorname{ati}_{g, k^{\prime}} \cdot \operatorname{mon}_{g} \cdot \operatorname{reg}_{g, k^{\prime}} \cdot O(k \log k) \cdot \mathrm{OPT} .
$$

Although the input matrix $A$ in the above statement is a square matrix, it is straightforward to extend the result to the rectangular case. By the above theorem, we can obtain a good subset of columns. To further get a low rank matrix $B$ which is a good low rank approximation to $A$, it is sufficient to take an additional $\mathcal{T}_{\text {reg, }, g, n,|S|, n}$ time to solve the regression problem.

### 20.1.1.2 Lower Bound on the Number of Columns

One may wonder if the $\log n$ blowup in rank is necessary in our theorem. We show some dependence on $n$ is necessary by showing that for the important Huber loss function, at least $\sqrt{\log n}$ columns are required in order to obtain a constant factor approximation for $k=1$ :

Theorem 20.1.3. Let $H(x)$ denote the following function: $H(x)= \begin{cases}x^{2}, & \text { if }|x|<1 ; \\ |x|, & \text { if }|x| \geq 1 .\end{cases}$
For any $n \geq 1$, there is a matrix $A \in \mathbb{R}^{n \times n}$ such that, if we select o $(\sqrt{\log n})$ columns to fit the entire matrix, there is no $O(1)$-approximation, i.e., for any subset $S \subseteq[n]$ with $|S|=o(\sqrt{\log n})$,

$$
\min _{X \in \mathbb{R}^{|S| \times n}}\left\|A_{S} X-A\right\|_{H} \geq \omega(1) \cdot \min _{\operatorname{rank}-1 A^{\prime}}\left\|A^{\prime}-A\right\|_{H} .
$$

Notice that the above function $H(x)$ is always a constant approximation to the Huber function (see Table 20.1) with $\tau=1$. Thus, the hardness also holds for the Huber function. For more discussion on our lower bound, we refer the reader to Appendix 20.6.

### 20.1.2 Overview of our Approach and Related Work

Low Rank Approximation for General Functions. A natural approach to low rank approximation is "column subset selection", which has been extensively studied in numerical linear algebra [DMM06c, DMM06b, DMM08, BMD09, BDM11, FEGK13, BW14, WS15, SWZ17, SWZ19b]. One can take the column subset selection algorithm for $\ell_{p}$-low rank approximation in $\left[\mathrm{CGK}^{+} 17 \mathrm{a}\right]$ and try to adapt it to general loss functions. Namely, their argument shows that for any matrix $A \in \mathbb{R}^{n \times n}$ there exists a subset $S$ of $k$ columns of $A$, denoted by $A_{S} \in \mathbb{R}^{n \times k}$, for which there exists a $k \times n$ matrix $V$ for which $\left\|A_{S} V-A\right\|_{p}^{p} \leq$ $(k+1)^{p} \min _{\text {rank-k } B^{\prime}}\left\|A-B^{\prime}\right\|_{p}^{p}$; we refer the reader to Theorem 3 of [CGK $\left.{ }^{+} 17 \mathrm{a}\right]$. Given the existence of such a subset $S$, a natural next idea is to then sample a set $T$ of $k$ columns of $A$ uniformly at random. It is then likely the case that if we look at a random column $A_{i}$, (1) with probability $1 /(k+1), i$ is not among the subset $S$ of $k$ columns out of the $k+1$ columns $T \cup\{i\}$ defining the optimal rank- $k$ approximation to the submatrix $A_{T \cup\{i\}}$, and (2) with probability at least $1 / 2$, the best rank- $k$ approximation to $A_{T \cup\{i\}}$ has cost at most

$$
\begin{equation*}
(2(k+1) / n) \cdot \min _{\text {rank-k } B^{\prime}}\left\|A-B^{\prime}\right\|_{p}^{p} . \tag{20.1}
\end{equation*}
$$

Indeed, (1) follows from $T \cup\{i\}$ being a uniformly random subset of $k+1$ columns, while (2) follows from a Markov bound. The argument in Theorem 7 of [CGK ${ }^{+}$17a] is then able to "prune" a $1 /(k+1)$ fraction of columns (this can be optimized to a constant fraction) in expectation, by "covering" them with the random set $T$. Recursing on the remaining columns,
this procedure stops after $k \log n$ iterations, giving a column subset of size $O\left(k^{2} \log n\right)$ (which can be optimized to $O(k \log n))$ and an $O(k)$-approximation.

The proof in [CGK $\left.{ }^{+} 17 \mathrm{a}\right]$ of the existence of a subset $S$ of $k$ columns of $A$ spanning a $(k+1)$-approximation above is quite general, and one might suspect it generalizes to a large class of error functions. Suppose, for example, that $k=1$. The idea there is to write $A=A^{*}+\Delta$, where $A^{*}=U \cdot V$ is the optimal rank-1 $\ell_{p}$-low rank approximation to $A$. One then "normalizes" by the error, defining $\widetilde{A}_{i}^{*}=A_{i}^{*} /\left\|\Delta_{i}\right\|_{p}$ and letting $s$ be such that $\left\|\widetilde{A}_{s}^{*}\right\|_{p}$ is largest. The rank-1 subset $S$ is then just $A_{s}$. Note that since $\widetilde{A}^{*}$ has rank-1 and $\left\|\widetilde{A}_{s}^{*}\right\|_{p}$ is largest, one can write $\widetilde{A}_{j}^{*}$ for every $j \neq s$ as $\alpha_{j} \cdot \widetilde{A}_{s}^{*}$ for $\left|\alpha_{j}\right| \leq 1$. The fact that $\left|\alpha_{j}\right| \leq 1$ is crucial; indeed, consider what happens when we try to "approximate" $A_{j}$ by $A_{s} \cdot \frac{\alpha_{j}\left\|\Delta_{j}\right\|_{p}}{\left\|\Delta_{s}\right\|_{p}}$. Then $\left\|A_{j}-A_{s} \alpha_{j}\right\| \Delta_{j}\left\|_{p} /\right\| \Delta_{s}\left\|_{p}\right\|_{p} \leq\left\|A_{j}-A_{j}^{*}\right\|_{p}+\left\|A_{j}^{*}-A_{s} \alpha_{j}\right\| \Delta_{j}\left\|_{p} /\right\| \Delta_{s}\left\|_{p}\right\|_{p}=$ $\left\|\Delta_{j}\right\|_{p}+\left\|A_{j}^{*}-\left(A_{s}^{*}+\Delta_{s}\right) \alpha_{j}\right\| \Delta_{j}\left\|_{p} /\right\| \Delta_{s}\left\|_{p}\right\|_{p}=\left\|\Delta_{j}\right\|_{p}+\left\|\Delta_{s} \alpha_{j}\right\| \Delta_{j}\left\|_{p} /\right\| \Delta_{s}\left\|_{p}\right\|_{p}$, and since the $p$-norm is monotonically increasing and $\alpha_{j} \leq 1$, the latter is at most $\left\|\Delta_{j}\right\|_{p}+\left\|\Delta_{s} \frac{\left\|\Delta_{j}\right\|_{p}}{\left\|\Delta_{s}\right\|_{p}}\right\|_{p}$. So far, all we have used about the $p$-norm is the monotone increasing property, so one could hope that the argument could be generalized to a much wider class of functions.

However, at this point the proof uses that the p-norm has scale-invariance, and so $\left\|\Delta_{s} \frac{\left\|\Delta_{j}\right\|_{p}}{\left\|\Delta_{s}\right\|_{p}}\right\|_{p}=\left\|\Delta_{j}\right\|_{p} \cdot\left\|\frac{\Delta_{s}}{\left\|\Delta_{s}\right\|_{p}}\right\|_{p}=\left\|\Delta_{j}\right\|_{p}$, and it follows that $\left\|A_{j}-A_{s} \frac{\alpha_{j}\left\|\Delta_{j}\right\|_{p}}{\left\|\Delta_{s}\right\|_{p}}\right\|_{p} \leq 2\left\|\Delta_{j}\right\|_{p}$, giving an overall 2-approximation (recall $k=1$ ). But what would happen for a general, not necessarily scale-invariant function $g$ ? We need to bound $\left\|\Delta_{s} \frac{\left\|\Delta_{j}\right\|_{g}}{\left\|\Delta_{s}\right\|_{g}}\right\|_{g}$. If we could bound this by $O\left(\left\|\Delta_{j}\right\|_{g}\right)$, we would obtain the same conclusion as before, up to constant factors. Consider, though, the "reverse Huber function": $g(x)=x^{2}$ if $x \geq 1$ and $g(x)=|x|$ for $x \leq 1$. Suppose that $\Delta_{s}$ and $\Delta_{j}$ were just 1-dimensional vectors, i.e., real numbers, so we need to bound $g\left(\Delta_{s} g\left(\Delta_{j}\right) / g\left(\Delta_{s}\right)\right)$ by $O\left(g\left(\Delta_{j}\right)\right)$. Suppose $\Delta_{s}=1$. Then $g\left(\Delta_{s}\right)=1$ and
$g\left(\Delta_{s} g\left(\Delta_{j}\right) / g\left(\Delta_{s}\right)\right)=g\left(g\left(\Delta_{j}\right)\right)$ and if $\Delta_{j}=n$, then $g\left(g\left(\Delta_{j}\right)\right)=n^{4}=g\left(\Delta_{j}\right)^{2}$, much larger than the $O\left(g\left(\Delta_{j}\right)\right)$ we were aiming for.

Maybe the analysis can be slightly changed to correct for these normalization issues? This is not the case, as we show that unlike for $\ell_{p}$-low rank approximation, for the reverse Hu ber function there is no subset of 2 columns of $A$ obtaining better than an $n^{1 / 4}$-approximation factor. (See Section 20.6.2 for more details). Further, the lack of scale invariance not only breaks the argument in $\left[\mathrm{CGK}^{+} 17 \mathrm{a}\right]$, it shows that combinatorially such functions $g$ behave very differently than $\ell_{p}$-norms. We show more generally there exist functions, in particular the Huber function, for which one needs to choose $\Omega(\sqrt{\log n})$ columns to obtain a constant factor approximation; we describe this more below. Perhaps more surprisingly, we show a subset of $O(\log n)$ columns suffice to obtain a constant factor approximation to the best rank-1 approximation for any function $g(x)$ which is approximately monotone and has the approximate triangle inequality, the latter implying for any constant $C>0$ and any $x \in \mathbb{R}_{\geq 0}$, $g(C x)=O(g(x))$. For $k>1$, these conditions become: (1) $g(x)$ is monotone non-decreasing in $x$, (2) $g(x)$ is within a $\operatorname{poly}(k)$ factor of $g(-x)$, and (3) for any real number $x \in \mathbb{R}^{\geq 0}$, $g(O(k x)) \leq \operatorname{poly}(k) \cdot g(x)$. We show it is possible to obtain an $O\left(k^{2} \log k\right)$ approximation with $O(k \log n)$ columns. We give the intuition and main lemma statements for our result in Section 20.2, deferring proofs to the supplementary material.

Even for $\ell_{p}$-low rank approximation, our algorithms slightly improve and correct a minor error in $\left[\mathrm{CGK}^{+} 17 \mathrm{a}\right]$ which claims in Theorem 7 an $O(k)$-approximation with $O(k \log n)$ columns for $\ell_{p}$-low rank approximation. However, their algorithm actually gives an $O(k \log n)$ approximation with $O(k \log n)$ columns. In [CGK ${ }^{+} 17$ a] it was argued that one expects to pay a cost of $O(k / n) \cdot \min _{\text {rank-k } B^{\prime}}\left\|A-B^{\prime}\right\|_{p}^{p}$ per column as in (20.1), and since each column is only
counted in one iteration, summing over the columns gives $O(k) \cdot \min _{\text {rank-k } B^{\prime}}\left\|A-B^{\prime}\right\|_{p}$ total cost. The issue is that the value of $n$ is changing in each iteration, so if in the $i$-th iteration it is $n_{i}$, then we could pay $n_{i} \cdot O\left(k / n_{i}\right) \cdot \min _{\text {rank-k } B^{\prime}}\left\|A-B^{\prime}\right\|_{p}=O(k) \cdot \min _{\text {rank-k } B^{\prime}}\left\|A-B^{\prime}\right\|_{p}$ in each of $O(\log n)$ iterations, giving $O(k \log n)$ approximation ratio. In contrast, our algorithm achieves an $O(k \log k)$ approximation ratio for $\ell_{p}$-low rank approximation as a special case, which gives the first $O(1)$ approximation in nearly linear time for any constant $k$ for $\ell_{p}$ norms. Our analysis is finer in that we show not only do we expect to pay a cost of $O\left(k / n_{i}\right) \cdot \min _{\text {rank-k } B^{\prime}}\left\|A-B^{\prime}\right\|_{p}^{p}$ per column in iteration $i$, we pay $O\left(k / n_{i}\right)$ times the cost of the best rank- $k$ approximation to $A$ after the most costly $n / k$ columns have been removed; thus we pay $O\left(k / n_{i}\right)$ times a residual cost with the top $n / k$ columns removed. This ultimately implies any column's cost can contribute in at most $O(\log k)$ of $O(\log n)$ recursive calls, replacing an $O(\log n)$ factor with an $O(\log k)$ factor in the approximation ratio. This also gives the first $\operatorname{poly}(k)$-approximation for $\ell_{0}$-low rank approximation, studied in [BKW17], improving the $O\left(k^{2} \log (n / k)\right)$-approximation there to $O\left(k^{2} \log k\right)$ and giving the first constant approximation for constant $k$.

```
Algorithm 20.1 Low Rank Approximation Algorithm for General Functions
    procedure GeneralFunctionLowRankApprox \(\left(A \in \mathbb{R}^{n \times n}, k \in \mathbb{Z}_{\geq 1}, g: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}\right)\)
        Initialization: \(r \leftarrow O(\log n), T_{0} \leftarrow[n]\)
        for \(i=1 \rightarrow r\) do
            for \(j=1 \rightarrow \log n\) do
                    Sample \(S_{i}^{(j)}\) from \(\binom{T_{i-1}}{2 k}\) uniformly at random
                    Solve the \(\operatorname{reg}_{g, 2 k}\)-approximate regression \(\min _{x \in \mathbb{R}^{2 k}}\left\|A_{S_{i}^{(j)}} x-A_{t}\right\|_{g}\) for each \(t \in\)
    \(T_{i-1} \backslash S_{i}^{(j)}\), and let \(v_{i, t}^{(j)}\) be the reg \(_{g, 2 k}\)-estimated regression cost \(\triangleright\) See Section 20.1.1.1 for
    regression property
                        \(R_{i}^{(j)} \leftarrow\left\{t \mid v_{i, t}^{(j)}\right.\) is the bottom \(\left|T_{i-1} \backslash S_{i}^{(j)}\right| / 20\) largest value in \(\left\{v_{i, t^{\prime}}^{(j)} \mid t^{\prime} \in T_{i-1} \backslash\right.\)
    \(\left.\left.S_{i}^{(j)}\right\}\right\}\)
            \(c_{i}^{(j)} \leftarrow \sum_{t \in R_{i}^{(j)}} v_{i, t}^{(j)}\)
            end for
            \(j^{*} \leftarrow \arg \min _{j \in[\log n]}\left\{c_{i}^{(j)}\right\}\)
            \(S_{i} \leftarrow S_{i}^{\left(j^{*}\right)}, R_{i} \leftarrow R_{i}^{\left(j^{*}\right)}, T_{i} \leftarrow T_{i-1} \backslash\left(S_{i} \cup R_{i}\right)\)
            end for
            return \(S=T_{r} \cup \bigcup_{i \in[r]} S_{i}\)
    end procedure
```


### 20.2 Algorithm for General Loss Low Rank Approximation

Our algorithm is presented in Algorithm 20.1. First, let us briefly analyze the running time. Consider fixed $i \in[r], j \in[\log n]$. Sampling $S_{i}^{(j)}$ takes $O(k)$ time. Solving $\operatorname{reg}_{g, 2 k^{-}}$ approximate regression $\min _{x}\left\|A_{S_{i}^{(j)}} x-A_{t}\right\|_{g}$ for all $t \in T_{i-1} \backslash S_{i}^{(j)}$ takes $\mathcal{T}_{\text {reg,g,n,2k,|Ti-1 } \backslash S_{i}^{(j)} \mid} \leq$ $\mathcal{T}_{\text {reg }, g, n, 2 k+1, n}$ time. Since finding $\left|T_{i-1} \backslash S_{i}^{(j)}\right| / 20$ smallest element can be done in $O(n)$ time, $R_{i}^{(j)}$ can be computed in $O(n)$ time. Thus the inner loop takes $O\left(n+\mathcal{T}_{\text {reg, }, g, n, 2 k+1, n}\right)$ time. Since $r=O(\log n)$, the total running time over all $i, j$ is $O\left(\left(n+\mathcal{T}_{\text {reg, }, n, 2 k+1, n}\right) \log ^{2} n\right)$. In the remainder of the section, we will sketch the proof of the correctness. For the missing proofs, we refer the reader to Appendix 20.3.

### 20.2.1 Properties of Uniform Column Sampling

Let us first introduce some useful notation. Consider a rank- $k$ matrix $M^{*} \in \mathbb{R}^{n \times m}$. For a set $H \subseteq[m]$, let $R_{M^{*}}(H) \subseteq H$ be a set such that

$$
R_{M^{*}}(H)=\arg \max _{P: P \subseteq H}\left\{\operatorname{det}\left(\left(M^{*}\right)_{P}^{Q}\right)| | P\left|=|Q|=\operatorname{rank}\left(M_{H}^{*}\right), Q \subseteq[n]\right\}\right.
$$

where $\operatorname{det}(C)$ denotes the determinant of a square matrix $C$. By Cramer's rule, if we use a linear combination of the columns of $M_{R_{M^{*}}(H)}^{*}$ to express any column of $M_{H}^{*}$, the absolute value of every fitting coefficient will be at most 1 .

Consider an arbitrary matrix $M \in \mathbb{R}^{n \times m}$. We can write $M=M^{*}+N$, where $M^{*} \in \mathbb{R}^{n \times m}$ is an arbitrary rank- $k$ matrix, and $N \in \mathbb{R}^{n \times m}$ is the residual matrix. The following lemma shows that, if we randomly choose a subset $H \subseteq[m]$ of $2 k$ columns, and we randomly look at another column $i$, then with constant probability, the absolute values of all the coefficients of using a linear combination of the columns of $M_{H}^{*}$ to express $M_{i}^{*}$ are at most 1, and furthermore, if we use the same coefficients to use columns of $M_{H}$ to fit $M_{i}$, then the fitting cost is proportional to $\left\|N_{H}\right\|_{g}+\left\|N_{i}\right\|_{g}$.

Lemma 20.2.1. Given a matrix $M \in \mathbb{R}^{n \times m}$ and a parameter $k \geq 1$, let $M^{*} \in \mathbb{R}^{n \times m}$ be an arbitrary rank-k matrix. Let $N=M-M^{*}$. Let $H \subseteq[m]$ be a uniformly random subset of $[m]$, and let $i$ denote a uniformly random index sampled from $[m] \backslash H$. Then (I) $\operatorname{Pr}\left[i \notin R_{M^{*}}(H \cup\{i\})\right] \geq 1 / 2$; (II) If $i \notin R_{M^{*}}(H \cup\{i\})$, then there exist $|H|$ coefficients $\alpha_{1}, \alpha_{2}, \cdots, \alpha_{|H|}$ for which $M_{i}^{*}=\sum_{j=1}^{|H|} \alpha_{j}\left(M_{H}^{*}\right)_{j}, \forall j \in[|H|],\left|\alpha_{j}\right| \leq 1$, and $\min _{x \in \mathbb{R}^{|H|}} \| M_{H} x-$ $M_{i} \|_{g} \leq \operatorname{ati}_{g,|H|+1} \cdot \operatorname{mon}_{g} \cdot\left(\left\|N_{i}\right\|_{g}+\sum_{j=1}^{|H|}\left\|\left(N_{H}\right)_{j}\right\|_{g}\right)$.

Notice that part (II) of the above lemma does not depend on any randomness of $H$ or $i$. By applying part (I) of the above lemma, it is enough to prove that if we randomly
choose a subset $H$ of $2 k$ columns, there is a constant fraction of columns that each column $M_{i}^{*}$ can be expressed by a linear combination of columns in $M_{H}^{*}$, and the absolute values of all the fitting coefficients are at most 1. Because of Cramer's rule, it thus suffices to prove the following lemma.

## Lemma 20.2.2.

$$
\operatorname{Pr}_{H \sim\binom{[m])}{2 k}}\left[\left|\left\{i \mid i \in[m] \backslash H, i \notin R_{M^{*}}(H \cup\{i\})\right\}\right| \geq(m-2 k) / 4\right] \geq 1 / 4 .
$$

### 20.2.2 Correctness of the Algorithm

We write the input matrix $A$ as $A^{*}+\Delta$, where $A^{*} \in \mathbb{R}^{n \times n}$ is the best rank- $k$ approximation to $A$, and $\Delta \in \mathbb{R}^{n \times n}$ is the residual matrix with respect to $A^{*}$. Then $\|\Delta\|_{g}=\sum_{i=1}^{n}\left\|\Delta_{i}\right\|_{g}$ is the optimal cost. As shown in Algorithm 20.1, our approach iteratively eliminates all the columns. In each iteration, we sample a subset of columns, and use these columns to fit other columns. We drop a constant fraction of columns which have a good fitting cost. Suppose the indices of the columns surviving after the $i$-th outer iteration are $T_{i}=\left\{t_{i, 1}, t_{i, 2}, \cdots, t_{i, m_{i}}\right\} \subseteq[n]$. Without loss of generality, we can assume $\left\|\Delta_{t_{i, 1}}\right\|_{g} \geq\left\|\Delta_{t_{i, 2}}\right\|_{g} \geq \cdots \geq\left\|\Delta_{t_{i, m_{i}}}\right\|_{g}$. The following claim shows that if we randomly sample $2 k$ column indices $H$ from $T_{i}$, then the cost of $\Delta_{H}$ will not be large.

Claim 20.2.3. If $\left|T_{i}\right|=m_{i} \geq 1000 k$,

$$
\operatorname{Pr}_{H \sim\binom{T_{i}}{2 k}}\left[\sum_{j \in H}\left\|\Delta_{j}\right\|_{g} \leq 400 \frac{k}{m_{i}} \sum_{j=\frac{m_{i}}{100 k}}^{m_{i}}\left\|\Delta_{t_{i, j}}\right\|_{g}\right] \geq \frac{19}{20} .
$$

By an averaging argument, in the following claim, we can show that there is a constant fraction of columns in $T_{i}$ whose optimal cost is also small.

Claim 20.2.4. If $\left|T_{i}\right|=m_{i} \geq 1000 k$,

$$
\left|\left\{t_{i, j} \mid t_{i, j} \in T_{i},\left\|\Delta_{t_{i, j}}\right\|_{g} \geq \frac{20}{m_{i}} \sum_{j^{\prime}=\frac{m_{i}}{100 k}}^{m_{i}}\left\|\Delta_{t_{i, j^{\prime}}}\right\|_{g}\right\}\right| \leq \frac{1}{5} m_{i} .
$$

By combining Lemma 20.2.2, part (II) of Lemma 20.2.1 with the above two claims, it is sufficient to prove the following core lemma. It says that if we randomly choose a subset of $2 k$ columns from $T_{i}$, then we can fit a constant fraction of the columns from $T_{i}$ with a small cost.

Lemma 20.2.5. If $\left|T_{i}\right|=m_{i} \geq 1000 k$,

$$
\operatorname{Pr}_{H \sim\left(\begin{array}{c}
T_{i} \\
2 k \\
2
\end{array}\right]}\left[\left|\left\{j \mid j \in T_{i}, \min _{x \in \mathbb{R}^{|H|}}\left\|A_{H} x-A_{j}\right\|_{g} \leq C_{1} \cdot \frac{1}{m_{i}} \cdot \sum_{j^{\prime}=\frac{m_{i}}{100 k}}^{m_{i}}\left\|\Delta_{t_{i, j^{\prime}}}\right\|_{g}\right\}\right| \geq \frac{1}{20} m_{i}\right] \geq \frac{1}{5}
$$

where $C_{1}=500 \cdot k \cdot \operatorname{ati}_{g,|S|+1} \cdot \operatorname{mon}_{g}$.

Let us briefly explain why the above lemma is enough to prove the correctness of our algorithm. For each column $j \in[m]$, either the column $j$ is in $T_{r}$ and is selected by the end of the algorithm, or $\exists i<r$ such that $j \in T_{i} \backslash T_{i+1}$. If $j \in T_{i} \backslash T_{i+1}$, then by the above lemma, we can show that with high probability, $\min _{x}\left\|A_{S_{i+1}} x-A_{j}\right\|_{g} \leq O\left(C_{1}\|\Delta\|_{1} /\left|T_{i}\right|\right)$. Thus, $\min _{X}\left\|A_{S_{i+1}} X-A_{T_{i} \backslash T_{i+1}}\right\|_{g} \leq O\left(C_{1}\|\Delta\|_{1}\right)$. It directly gives a $O\left(r C_{1}\right)=O\left(C_{1} \log n\right)$ approximation. For the detailed proof of Theorem 20.1.2, we refer the reader to Appendix 20.3.

### 20.3 Missing Proofs in Section 20.2

### 20.3.1 Proof of Lemma 20.2.1

Proof. (I) Since $\operatorname{rank}\left(M_{H \cup\{i\}}^{*}\right) \leq \operatorname{rank}\left(M^{*}\right)=k,\left|R_{M^{*}}(H \cup\{i\})\right| \leq k$. Note that $H$ is sampled from $\binom{[m]}{2 k}$ uniformly at random, $i$ is sampled from $[m] \backslash H$ uniformly at random, and $|H \cup\{i\}|=2 k+1$. By symmetry we have

$$
\operatorname{Pr}_{H \sim\binom{[m b}{2 k}, i \sim[m] \backslash H}\left[i \notin R_{M^{*}}(H \cup\{i\})\right] \geq 1-\frac{k}{2 k+1} \geq 1 / 2 .
$$

(II) Since $i \notin R_{M^{*}}(H \cup\{i\})$, by Cramer's rule, there exist $\alpha_{1}, \alpha_{2}, \cdots, \alpha_{|H|}$ such that $M_{i}^{*}=\sum_{j=1}^{|H|} \alpha_{j} \cdot\left(M_{H}^{*}\right)_{j}$ and $\forall j \in[|H|],\left|\alpha_{j}\right| \leq 1$. Then we have

$$
\begin{aligned}
\min _{x \in \mathbb{R}^{|H|}}\left\|M_{H} x-M_{i}\right\|_{g} & \leq\left\|\sum_{j=1}^{|H|}\left(M_{H}\right)_{j} \alpha_{j}-M_{i}\right\|_{g} \\
& =\left\|\sum_{j=1}^{|H|}\left(M_{H}^{*}\right)_{j} \alpha_{j}-M_{i}^{*}+\sum_{j=1}^{|H|}\left(N_{H}\right)_{j} \alpha_{j}-N_{i}\right\|_{g} \\
& =\left\|\sum_{j=1}^{|H|}\left(N_{H}\right)_{j} \alpha_{j}-N_{i}\right\|_{g} \\
& \leq \operatorname{ati}_{g,|H|+1} \cdot\left(\left\|-N_{i}\right\|_{g}+\sum_{j=1}^{|H|}\left\|\left(N_{H}\right)_{j} \alpha_{j}\right\|_{g}\right) \\
& \leq \operatorname{ati}_{g,|H|+1} \cdot\left(\left\|-N_{i}\right\|_{g}+\operatorname{mon}_{g} \cdot \sum_{j=1}^{|H|}\left\|\left(N_{H}\right)_{j}\right\|_{g}\right) \\
& \leq \operatorname{ati}_{g,|H|+1} \cdot\left(\operatorname{mon}_{g} \cdot\left\|N_{i}\right\|_{g}+\operatorname{mon}_{g} \cdot \sum_{j=1}^{|H|}\left\|\left(N_{H}\right)_{j}\right\|_{g}\right) \\
& \leq \operatorname{ati}_{g,|H|+1} \cdot \operatorname{mon}_{g} \cdot\left(\left\|N_{i}\right\|_{g}+\sum_{j=1}^{|H|}\left\|\left(N_{H}\right)_{j}\right\|_{g}\right)
\end{aligned}
$$

where the second step follows from $M=M^{*}+N$, the third step follows from $\sum_{j=1}^{|H|}\left(M_{H}^{*}\right)_{j} \alpha_{j}-$ $M_{i}^{*}=0$, the fourth step follows from the approximate triangle inequality, the fifth step follows from the fact that $\left|\alpha_{j}\right| \leq 1$ and $g$ is $\operatorname{mon}_{g}$-monotone, and the sixth step follows from that $g$ is mon $_{g}$-monotone.

### 20.3.2 Proof of Lemma 20.2.2

Proof. Using Part (I) of Lemma 20.2.1, we have

$$
\operatorname{Pr}_{H \sim\binom{[m]}{2 k}, i \sim[m] \backslash H}\left[i \notin R_{M^{*}}(H \cup\{i\})\right] \geq 1 / 2 .
$$

For each set $H$, we define $P_{H}=\operatorname{Pr}_{i \sim[m] \backslash H}\left[i \notin R_{M^{*}}(H \cup\{i\})\right]$. We have

$$
\begin{equation*}
1 \geq \frac{1}{\binom{m}{2 k}} \sum_{H \in\binom{[m]}{2 k}} P_{H} \geq 1 / 2 \tag{20.2}
\end{equation*}
$$

We can show

$$
\begin{aligned}
& \frac{1}{\binom{m}{2 k}}\left|\left\{H \left\lvert\, H \in\binom{[m]}{2 k}\right., P_{H} \geq 1 / 4\right\}\right| \\
= & \frac{1}{\binom{m}{2 k}} \sum_{H \in\binom{[m]}{2 k}, P_{H} \geq 1 / 4} 1 \\
\geq & \frac{1}{\binom{m}{2 k}} \sum_{H \in\binom{m n}{2 k}, P_{H} \geq 1 / 4} P_{H} \\
\geq & \frac{1}{2}-\frac{1}{\binom{m}{2 k}} \sum_{H \in\binom{[m]}{2 k}, P_{H}<1 / 4} P_{H} \\
\geq & \frac{1}{2}-\frac{1}{\binom{m}{2 k}} \sum_{H \in\binom{[m]}{2 k}, P_{H}<1 / 4} \frac{1}{4} \\
\geq & \frac{1}{2}-\frac{1}{\binom{m}{2 k}}\binom{m}{2 k} \frac{1}{4} \\
= & \frac{1}{4},
\end{aligned}
$$

where the second step follows since $1 \geq P_{H}$, the third step follows since Eq. (20.2), the fourth step follows since $P_{H}<1 / 4$.

Thus, we have

$$
\left|\left\{H \left\lvert\, H \in\binom{[m]}{2 k}\right., P_{H} \geq 1 / 4\right\}\right| \geq\binom{ m}{2 k} / 4 .
$$

Recall the definition of $P_{H}$, we have

$$
\left|\left\{H \left\lvert\, H \in\binom{[m]}{2 k}\right., \underset{i \sim[m] \backslash H}{\operatorname{Pr}}\left[i \notin R_{M^{*}}(H \cup\{i\})\right] \geq 1 / 4\right\}\right| \geq\binom{ m}{2 k} / 4,
$$

which implies

$$
\operatorname{Pr}_{H \sim\binom{[m]}{2 k}}\left[\left|\left\{i \mid i \in[m] \backslash H, i \notin R_{M^{*}}(H \cup\{i\})\right\}\right| \geq(m-2 k) / 4\right] \geq 1 / 4 .
$$

### 20.3.3 Proof of Claim 20.2.3

Proof. For simplicity, we omit $i$ in all the subscripts in this proof.

$$
\begin{aligned}
& \operatorname{Pr}_{H \sim\binom{T}{2 k}}\left[\sum_{j \in H}\left\|\Delta_{j}\right\|_{g} \leq 400 \frac{k}{m} \sum_{j=\frac{m}{100 k}}^{m}\left\|\Delta_{t_{j}}\right\|_{g}\right] \\
= & \operatorname{Pr}_{H \sim\binom{T}{2 k}}\left[\left.\sum_{j \in H}\left\|\Delta_{j}\right\|_{g} \leq 400 \frac{k}{m} \sum_{j=\frac{m}{100 k}}^{m}\left\|\Delta_{t_{j}}\right\|_{g} \right\rvert\, \exists j \leq \frac{m}{100 k}, t_{j} \in H\right] \cdot \operatorname{Pr}_{H \sim\binom{T}{2 k}}\left[\exists j \leq \frac{m}{100 k}, t_{j} \in H\right] \\
+ & \operatorname{Pr}_{H \sim\binom{T}{2 k}}\left[\left.\sum_{j \in H}\left\|\Delta_{j}\right\|_{g} \leq 400 \frac{k}{m} \sum_{j=\frac{m}{100 k}}^{m}\left\|\Delta_{t_{j}}\right\|_{g} \right\rvert\, \forall j \leq \frac{m}{100 k}, t_{j} \notin H\right] \cdot \operatorname{Pr}_{H \sim\left(\begin{array}{c}
T h
\end{array}\right)}^{\operatorname{Pr}_{2 k}}\left[\forall j \leq \frac{m}{100 k}, t_{j} \notin H\right] \\
\leq & \underbrace{\operatorname{Pr}_{H \sim\binom{T}{2 k}}\left[\exists j \leq \frac{m}{100 k}, t_{j} \in H\right]}_{C_{1}}+\underbrace{\operatorname{Pr}_{H \sim\left(\frac{T}{2 k}\right)}\left[\left.\sum_{j \in H}\left\|\Delta_{j}\right\|_{g} \leq 400 \frac{k}{m} \sum_{j=\frac{m}{100 k}}^{m}\left\|\Delta_{t_{j}}\right\|_{g} \right\rvert\, \forall j \leq \frac{m}{100 k}, t_{j} \notin H\right]}_{C_{2}}
\end{aligned}
$$

It remains to upper bound the terms $C_{1}$ and $C_{2}$. We can upper bound $C_{1}$ :

$$
\begin{aligned}
C_{1} & =1-\left(1-\frac{m / 100 k}{m}\right) \cdot\left(1-\frac{m / 100 k}{m-1}\right) \cdots \cdots\left(1-\frac{m / 100 k}{m-2 k+1}\right) \\
& \leq 1-\left(1-\frac{m / 100 k}{m / 2}\right)^{2 k} \\
& \leq 1-\left(1-\frac{1}{25}\right) \\
& =\frac{1}{25},
\end{aligned}
$$

where the second step follows since $m \geq 1000 k$.

Using Markov's inequality,

$$
C_{2} \leq \frac{\underset{H \sim\left(T 2_{2 k}^{T}\right)}{\mathbb{E}}\left[\left.\sum_{j \in H}\left\|\Delta_{j}\right\|_{g} \leq 400 \frac{k}{m} \sum_{j=\frac{m}{100 k}}^{m}\left\|\Delta_{t_{j}}\right\|_{g} \right\rvert\, \forall j \leq \frac{m}{100 k}, t_{j} \notin H\right]}{400 \frac{k}{m} \sum_{j=\frac{m}{100 k}}^{m}\left\|\Delta_{t_{j}}\right\|_{g}} \leq 1 / 100
$$

where the second step follows since

$$
\begin{aligned}
& \underset{H \sim(2 k}{\mathbb{E}}) \\
\leq & \frac{2 k}{m-m / 100 k} \sum_{j=\frac{m}{100 k}}^{\mathbb{E}}\left\|\Delta_{t_{j}}\right\|_{g} \\
\leq & \left.\left.4 \frac{k}{m}\left\|_{j=\frac{m}{100 k}}^{m}\right\|_{g} \leq 400 \frac{k}{m} \sum_{j=\frac{m}{100 k}}^{m}\left\|\Delta_{t_{j}}\right\|_{g} \right\rvert\, \forall j \leq \frac{m}{100 k}, t_{j} \notin H\right]
\end{aligned}
$$

### 20.3.4 Proof of Claim 20.2.4

Proof. For simplicity, we omit $i$ in all the subscripts in this proof.

$$
\begin{aligned}
&\left|\left\{t_{j} \mid t_{j} \in T,\left\|\Delta_{t_{j}}\right\|_{g} \geq \frac{20}{m} \sum_{j^{\prime}=\frac{m}{100 k}}^{m}\left\|\Delta_{t_{j^{\prime}}}\right\|_{g}\right\}\right| \\
& \leq\left|\left\{t_{j} \mid t_{j} \in T, j \leq \frac{m}{100 k}\right\}\right|+\left|\left\{t_{j} \mid t_{j} \in T, j>\frac{m}{100 k},\left\|\Delta_{t_{j}}\right\|_{g} \geq \frac{20}{m} \sum_{j^{\prime}=\frac{m}{100 k}}^{m}\left\|\Delta_{t_{j^{\prime}}}\right\|_{g}\right\}\right| \\
& \leq\left|\left\{t_{j} \mid t_{j} \in T, j \leq \frac{m}{100 k}\right\}\right|+\left|\left\{t_{j} \mid t_{j} \in T, j>\frac{m}{100 k},\left\|\Delta_{t_{j}}\right\|_{g} \geq \frac{10}{m-\frac{m}{100 k}} \sum_{j^{\prime}=\frac{m}{100 k}}^{m}\left\|\Delta_{t_{j^{\prime}}}\right\|_{g}\right\}\right| \\
& \leq \frac{m}{100 k}+\frac{1}{10}\left(m-\frac{m}{100 k}\right) \\
& \leq \frac{1}{5} m
\end{aligned}
$$

where the second step follows since $\frac{20}{m} \geq \frac{10}{m-m / 100 k}$

### 20.3.5 Proof of Lemma 20.2.5

Proof. For simplicity, we omit $i$ in all the subscirptis in this proof.
Let $M=A_{T}, M^{*}=A_{T}^{*}$ and $N=\Delta_{T}$. Then we can apply Lemma 20.2.2 and part (II) of Lemma 20.2.1:
$\operatorname{Pr}_{H \sim\binom{T}{2 k}}\left[\left|\left\{j \in T \mid \min _{x \in \mathbb{R}^{|H|}}\left\|A_{H} x-A_{j}\right\|_{g} \leq \operatorname{ati}_{g,|H|+1} \cdot \operatorname{mon}_{g} \cdot\left(\left\|\Delta_{j}\right\|_{g}+\sum_{j^{\prime}=1}^{|H|}\left\|\left(\Delta_{H}\right)_{j^{\prime}}\right\|_{g}\right)\right\}\right| \geq \frac{m}{4}\right] \geq \frac{1}{4}$

By Claim 20.2.3, we have

$$
\begin{equation*}
\operatorname{Pr}_{H \sim\binom{T}{2 k}}\left[\sum_{j=1}^{|H|}\left\|\left(\Delta_{H}\right)_{j}\right\|_{g} \leq 400 \frac{k}{m} \sum_{j=\frac{m}{100 k}}^{m}\left\|\Delta_{t_{j}}\right\|_{g}\right] \geq \frac{19}{20} \tag{20.4}
\end{equation*}
$$

Due to Claim 20.2.4,

$$
\left|\left\{t_{j} \mid t_{j} \in T,\left\|\Delta_{t_{j}}\right\|_{g} \geq \frac{20}{m} \sum_{j^{\prime}=\frac{m}{100 k}}^{m}\left\|\Delta_{t_{j^{\prime}}}\right\|_{g}\right\}\right| \leq \frac{1}{5} m
$$

Combining the above equation with the pigeonhole principle, for any $I \subseteq T$ with $|I| \geq m / 4$, we have

$$
\begin{equation*}
\left|\left\{t_{j} \mid t_{j} \in I,\left\|\Delta_{t_{j}}\right\|_{g}<\frac{20}{m} \sum_{j^{\prime}=\frac{m}{100 k}}^{m}\left\|\Delta_{t_{j^{\prime}}}\right\|_{g}\right\}\right| \geq \frac{1}{4} m-\frac{1}{5} m=\frac{1}{20} m \tag{20.5}
\end{equation*}
$$

Consider the quantity $\left\|\Delta_{j}\right\|_{g}+\sum_{j=1}^{|H|}\left\|\left(\Delta_{H}\right)_{j^{\prime}}\right\|_{g}$ in Eq. (20.3). We use Eq. (20.4) and Eq. (20.5) to provide an upper bound,

$$
\left\|\Delta_{j}\right\|_{g}+\sum_{j^{\prime}=1}^{|H|}\left\|\left(\Delta_{H}\right)_{j^{\prime}}\right\|_{g} \leq\left(\frac{20}{m}+\frac{400 k}{m}\right) \sum_{j^{\prime}=\frac{m}{100 k}}^{m}\left\|\Delta_{t_{j^{\prime}}}\right\|_{g} .
$$

Eq. (20.4) will decrease the final probability by $(1-19 / 20)$ (from $1 / 4$ to $1 / 4-1 / 20)$.
Eq. (20.5) will decrease the size of this set of $j$ by $\frac{1}{5} m$ (from $\frac{1}{4} m$ to $\frac{1}{4} m-\frac{1}{5} m$ ).
Putting it all together, we can update Eq. (20.3) in the following sense,
$\operatorname{Pr}_{H \sim\binom{T}{2 k}}\left[\left|\left\{j \mid j \in T, \min _{x \in \mathbb{R}^{|H|}}\left\|A_{H} x-A_{j}\right\|_{g} \leq C \cdot \frac{1}{m} \cdot \sum_{j^{\prime}=\frac{m}{100 k}}^{m}\left\|\Delta_{t_{j^{\prime}}}\right\|_{g}\right\}\right| \geq\left(\frac{1}{4}-\frac{1}{5}\right) m\right] \geq \frac{1}{4}-\frac{1}{20}$.
where

$$
C=(400+20) \cdot k \cdot \operatorname{ati}_{g,|H|+1} \cdot \operatorname{mon}_{g} .
$$

### 20.3.6 Proof of Theorem 20.1.2

Proof. The running time is discussed at the beginning of Section 20.2. In the remaining of the proof, we will focus on the correctness of Algorithm 20.1.

Firstly, let us consider the size of the output $S$. For $i \in\{0\} \cup[r]$, let $m_{i}=\left|T_{i}\right|$. We set number of rounds $r$ to be the smallest value such that $m_{r}<1000 k$. By the algorithm, we have $m_{i}=m_{i-1}-2 k-\left(m_{i-1}-2 k\right) / 20 \leq 19 / 20 \cdot m_{i-1}$. Thus, $r=O(\log n)$. In each round $i$, the size of $S_{i}$ is $2 k$. Then $|S|=\left|T_{r}\right|+\sum_{i=1}^{r}\left|S_{i}\right| \leq 1000 k+r \cdot 2 k \leq O(k \log n)$.

Next, let us consider the quality of $S$. Since each regression call has $1-1 / \operatorname{poly}(n)$ success probability, all the regression calls succeed with probability at least $1-1 / \operatorname{poly}(n)$. In the remaining of the proof, we condition on that all the regression calls succeed.

Let us fix $i \in[r], j \in[\log n]$. Recall that $T_{i}=\left\{t_{i, 1}, t_{i, 2}, \cdots, t_{i, m_{i}}\right\}$ and $\left\|\Delta_{t_{i, 1}}\right\|_{g} \geq$ $\left\|\Delta_{t, 2}\right\|_{g} \geq \cdots \geq\left\|\Delta_{t_{i, m_{i}}}\right\|_{g}$. By regression property and Lemma 20.2.5, with probability at
least $1 / 4$,

$$
\begin{aligned}
& \sum_{q \in R_{i}^{(j)}} \min _{x \in \mathbb{R}^{2 k}}\left\|A_{S_{i}^{(j)}} x-A_{q}\right\|_{g} \\
\leq & \sum_{q \in R_{i}^{(j)}} v_{i, q}^{(j)} \\
\leq & \operatorname{reg}_{g, 2 k} \cdot\left(m_{i}-2 k\right) \cdot 500 \cdot k \cdot \operatorname{ati}_{g, 2 k+1} \cdot \operatorname{mon}_{g} / m_{i} \cdot \sum_{j^{\prime}=\frac{m_{i}}{100 k}}^{m_{i}}\left\|\Delta_{t_{i, j^{\prime}}}\right\|_{g} \\
\leq & \operatorname{reg}_{g, 2 k+1} \cdot \operatorname{ati}_{g, 2 k+1} \cdot \operatorname{mon}_{g} \cdot O(k) \cdot \sum_{j^{\prime}=\frac{m_{i}}{100 k}}^{m_{i}}\left\|\Delta_{t_{i, j^{\prime}}}\right\|_{g}
\end{aligned}
$$

For each $i \in[r]$, since we repeat $\log (n)$ times, the success probability can be boosted to at least $1-1 / \operatorname{poly}(r)$, i.e., with probability at least $1-1 / \operatorname{poly}(r)$, we have

$$
\begin{equation*}
\sum_{q \in R_{i}} \min _{x \in \mathbb{R}^{2 k}}\left\|A_{S_{i}} x-A_{q}\right\|_{g} \leq \operatorname{reg}_{g, 2 k+1} \cdot \operatorname{ati}_{g, 2 k+1} \cdot \operatorname{mon}_{g} \cdot O(k) \cdot \sum_{j^{\prime}=\frac{m_{i}}{100 k}}^{m_{i}}\left\|\Delta_{t_{i, j^{\prime}}}\right\|_{g} \tag{20.6}
\end{equation*}
$$

In the remaining of the proof, we condition on above inequality for every $i \in[r]$. Without
loss of generality, we suppose $\left\|\Delta_{1}\right\|_{g} \geq\left\|\Delta_{2}\right\|_{g} \geq \cdots \geq\left\|\Delta_{n}\right\|_{g}$. We have

$$
\begin{aligned}
& \sum_{q=1}^{n} \min _{x \in \mathbb{R}^{|S|}}\left\|A_{S} x-A_{q}\right\|_{g} \\
\leq & \left(\sum_{q \in T_{r}} \min _{x \in \mathbb{R}^{\left|T_{r}\right|}}\left\|A_{T_{r}} x-A_{q}\right\|_{g}\right)+\left(\sum_{i=1}^{r} \sum_{q \in T_{i-1} \backslash T_{i}} \min _{x \in \mathbb{R}^{2 k}}\left\|A_{S_{i}} x-A_{q}\right\|_{g}\right) \\
= & \sum_{i=1}^{r} \sum_{q \in T_{i-1} \backslash T_{i}} \min _{x \in \mathbb{R}^{2 k}}\left\|A_{S_{i}} x-A_{q}\right\|_{g} \\
\leq & \sum_{i=1}^{r} \operatorname{reg}_{g, 2 k+1} \cdot \operatorname{ati}_{g, 2 k+1} \cdot \operatorname{mon}_{g} \cdot O(k) \cdot \sum_{j^{\prime}=\frac{m_{i}}{10{ }_{l}}}^{m_{i}}\left\|\Delta_{t_{i, j^{\prime}}}\right\|_{g} \\
\leq & \sum_{i=1}^{r} \operatorname{reg}_{g, 2 k+1} \cdot \operatorname{ati}_{g, 2 k+1} \cdot \operatorname{mon}_{g} \cdot O(k) \cdot \sum_{j^{\prime}=\frac{m_{i}}{100 k}}^{m_{i}}\left\|\Delta_{j^{\prime}}\right\|_{g} \\
= & \operatorname{reg}_{g, 2 k+1} \cdot \operatorname{ati}_{g, 2 k+1} \cdot \operatorname{mon}_{g} \cdot O(k) \cdot \sum_{j^{\prime}=1}^{n}\left\|\Delta_{j^{\prime}}\right\|_{g}\left(\underset{i \in[r]}{\operatorname{argmin}}\left\{m_{i}<j^{\prime}\right\}-\underset{i \in[r]}{\operatorname{argmin}}\left\{\frac{m_{i}}{100 k}<j^{\prime}\right\}+O(1)\right) \\
\leq & \operatorname{reg}_{g, 2 k+1} \cdot \operatorname{ati}_{g, 2 k+1} \cdot \operatorname{mon}_{g} \cdot O(k \log k) \cdot\|\Delta\|_{g},
\end{aligned}
$$

where the third step follows from $T_{i-1} \backslash T_{i}=S_{i} \cup R_{i}$ and Equation (20.6), the forth step follows from $\left\|\Delta_{j^{\prime}}\right\|_{g} \geq\left\|\Delta_{t_{i, j^{\prime}}}\right\|_{g}$, and the last step follows from

$$
\left(\underset{i \in[r]}{\operatorname{argmin}}\left\{m_{i}<j^{\prime}\right\}-\underset{i \in[r]}{\operatorname{argmin}}\left\{\frac{m_{i}}{100 k}<j^{\prime}\right\}+O(1)\right) \leq O(\log k) .
$$

### 20.4 Necessity of the Properties of $g$

We note that an approximate triangle inequality is necessary to obtain a column subset selection algorithm. An example function not satisfying this is the "jumping function": $g_{\tau}(x)=|x|$ if $|x| \geq \tau$, and $g_{\tau}(x)=0$ otherwise. For the identity matrix $I$ and any $k=$ $\Omega(\log n)$, the Johnson-Lindenstrauss lemma implies one can find a rank-k matrix $B$ for which $\|I-B\|_{\infty}<1 / 2$, that is, all entries of $I-B$ are at most $1 / 2$. If we set $\tau=1 / 2$, then $\|I-B\|_{g_{\tau}}=0$, but for any subset $I_{S}$ of columns of the identity matrix we choose, necessarily $\left\|I-I_{S} X\right\|_{\infty} \geq 1$, so $\|I-B\|_{g_{\tau}}>0$. Consequently, there is no subset of a small number of columns which obtains a poly $(k \log n)$-approximation with the jumping function loss measure.

While the jumping function does not satisfy the Approximate triangle inequality, it does satisfy our only other required structural property, the Monotone property.

There are interesting examples of functions $g$ which are only approximately monotone in the above sense, such as the quantile function $\rho_{\tau}(x)$, studied in [YMM14] in the context of regression, where for a given parameter $\tau, \rho_{\tau}(x)=\tau x$ if $x \geq 0$, and $\rho_{\tau}(x)=(\tau-1) x$ if $x<0$. Only when $\tau=1 / 2$ is this a monotone function with $\operatorname{mon}_{g}=1$ in the above definition, in which case it coincides with the absolute value function up to a factor of $1 / 2$. For other constant $\tau \in(0,1), \operatorname{mon}_{g}$ is a constant. The loss function $\rho_{\tau}(x)$ is also sometimes called the scalene loss, and studied in the context of low rank approximation in [UHZ $\left.{ }^{+} 16\right]$.

When $\tau=1$ this is the so-called Rectified Linear Unit (ReLU) function in machine learning, i.e., $\rho_{1}(x)=x$ if $x \geq 0$ and $\rho_{1}(x)=0$ if $x<0$. In this case $\operatorname{mon}_{g}=\infty$. and the optimal rank- $k$ approximation for any matrix $A$ is 0 , since $\left\|A-\lambda \mathbf{1 1}^{\top}\right\|_{\rho_{1}}=0$ if one sets $\lambda$ to
be a large enough positive number, thereby making all entries of $A-\lambda \mathbf{1 1}^{\top}$ negative and their corresponding cost equal to 0 . Notice though, that there are no good column subset selection algorithms for some matrices $A$, such as the $n \times n$ identity matrix. Indeed, for the identity, if we choose any subset $A_{S}$ of at most $n-1$ columns of $A$, then for any matrix $X$ there will be an entry of $A-A_{S} X$ which is positive, causing the cost to be positive. Since we will restrict ourselves to column subset selection, being approximately monotone with a small value of $\mathrm{mon}_{g}$ in the above definition is in fact necessary to obtain a good approximation with a small number of columns, as the ReLU function illustrates (see also related functions such as the leaky ReLU and squared ReLU [ZSJ+17, BHL18, GKM18]).

Note that the ReLU function is an example which satisfies the triangle inequality, showing that our additional assumption of approximate monotonicity is required.

Thus, if either property fails to hold, there need not be a small subset of columns spanning a relative error approximation. These examples are stated in more detail below.

### 20.4.1 Functions without Approximate Triangle Inequality

In this section, we show how to construct a function $f$ such that it is not possible to obtain a good entrywise- $f$ low rank approximation by selecting a small subset of columns. Furthermore, $f$ is monotone but does not have the approximate triangle inequality. Theorem 20.4.3 shows this result.

First, we show that a small subset of columns cannot give a good low rank approximation in $\ell_{\infty}$ norm. Then we reduce the $\ell_{\infty}$ column subset selection problem to the entrywise- $f$ column subset selection problem.

The following is the Johnson-Lindenstrauss lemma.

Lemma 20.4.1 (JL Lemma). For any $n \geq 1, \epsilon \in(1 / \sqrt{n}, 1 / 2)$, there exists $U \in \mathbb{R}^{n \times k}$ with $k=O\left(\epsilon^{-2} \log (n)\right)$ such that $\left\|U U^{\top}-I_{n}\right\|_{\infty} \leq O(\epsilon)$, where $I_{n} \in \mathbb{R}^{n \times n}$ is an identity matrix.

Theorem 20.4.2. For $n \geq 1$, there is a matrix $A \in \mathbb{R}^{n \times n}$ with the following properties. Let $k=\Theta\left(\epsilon^{-2} \log (n)\right)$ for an arbitrary $\epsilon \in(1 / \sqrt{n}, 1 / 2)$. Let $D \in \mathbb{R}^{n \times n}$ denote a diagonal matrix with $n-1$ nonzeros on the diagonal. We have

$$
\min _{X \in \mathbb{R}^{n \times n}}\|X D A-A\|_{\infty} \geq 1
$$

and

$$
\min _{\operatorname{rank}-k}\left\|A^{\prime}-A\right\|_{\infty}<O(\epsilon) .
$$

Proof. We choose $A$ to be the identity matrix. By Lemma 20.4.1, we can find a rank- $k$ matrix $B$ for which

$$
\|A-B\|_{\infty} \leq O(\epsilon)
$$

Since $A$ is an $n \times n$ identity matrix, even if we can use $n-1$ columns to fit the other columns, the cost is still at least 1.

In the following, we state the construction of our function $f$.

Definition 20.4.1. We define function $f(x)$ to be $f(x)=c$ if $|x|>\tau$ and $f(x)=0$ if $|x| \leq \tau$. Given matrix $A$, we define $\|A\|_{f}=\sum_{i=1}^{n} \sum_{j=1}^{n} f\left(A_{i, j}\right)$.

Theorem 20.4.3 (No good subset of columns). For any $n \geq 1$, there is a matrix $A \in \mathbb{R}^{n \times n}$ with the following property. Let $k \geq c \log n$ for a sufficiently large constant $c>0$. Let
$D \in \mathbb{R}^{n \times n}$ denote an arbitrary diagonal matrix with $n-1$ nonzeros on the diagonal. For $f$ with parameter $\tau=1 / 4$, we have

$$
\min _{X \in \mathbb{R}^{n \times n}}\|X D A-A\|_{f}>0
$$

and

$$
\min _{\operatorname{rank}-k}\left\|A^{\prime}-A\right\|_{f}=0
$$

Proof. We can set $A$ to be the identity matrix. Due to Theorem 20.4.2, there exists $A^{\prime}$ for which $\min _{\text {rank }-k A^{\prime}}\left\|A^{\prime}-A\right\|_{\infty}<1 / 4$, which implies that $\min _{\text {rank }-k A^{\prime}}\left\|A^{\prime}-A\right\|_{f}=0$. Also due to Theorem 20.4.2, we have $\min _{X \in \mathbb{R}^{n \times n}}\|X D A-A\|_{\infty}=1$, and thus, $\min _{X \in \mathbb{R}^{n \times n}} \| X D A-$ $A \|_{f}>0$.

### 20.4.2 ReLU Function Low Rank Approximation

In this section, we discuss a function which has the approximate triangle inequality but is not symmetric. The specific function we discuss in this section is ReLU. The definition of ReLU is defined in Definition 20.4.2. First, we show that ReLU low rank approximation has a trivial best rank- $k$ approximation. Second, we show that for some matrices, there is no small subset of columns which can give a good low rank approximation.

Definition 20.4.2. We define function $\operatorname{ReLU}(x)$ to be $\operatorname{ReLU}(x)=\max (0, x)$. Given matrix $A$, we define $\|A\|_{\mathrm{ReLU}}=\sum_{i=1}^{n} \sum_{j=1}^{n} \operatorname{ReLU}\left(A_{i, j}\right)$.

In the rank- $k$ approximation problem, given an input matrix $A$, the goal is to find a rank- $k$ matrix $B$ for which $\|A-B\|_{\mathrm{ReLU}}$ is minimized. A simple observation is that if we set
$B$ to be a matrix with each entry of value $\|A\|_{\infty}$, then the value of each entry of $A-B$ is at most 0 . Thus, $\|A-B\|_{\mathrm{ReLU}}=0$. Furthermore, the rank of $B$ is 1 .

Now, consider the column subset selection problem, let input matrix $A \in \mathbb{R}^{n \times n}$ be an identity matrix. Then even if we can choose $n-1$ columns, they can never fit the remaining column. Thus, the cost is at least 1 . But as discussed, the best rank- $k$ cost is always 0 . This implies that any subset of columns cannot give a good rank- $k$ approximation.

### 20.5 Regression Solvers

In this section, we discuss several regression solvers.

### 20.5.1 Regression for Convex $g$

Notice that when the function $g$ is convex, the regression problem $\min _{X \in \mathbb{R}^{d \times m}} \| A X-$ $B \|_{g}$ for any given matrices $A \in \mathbb{R}^{n \times d}, B \in \mathbb{R}^{n \times m}$ is a convex optimization problem. Thus, it can be solved exactly by convex optimization algorithms.

Fact 20.5.1. Let $g$ be a convex function. Given $A \in \mathbb{R}^{n \times d}, B \in \mathbb{R}^{n \times m}$, the regression problem $\min _{X \in \mathbb{R}^{d \times m}}\|A X-B\|_{g}$ can be solved exactly by convex optimization in $\operatorname{poly}(n, d, m)$ time.

If a function $g$ has additional properties, i.e. $g$ is symmetric, monotone and grows subquadratically, then there is a better running time constant approximation algorithm shown in [CW15b]. Here "grows quadratically" means that there is an $\alpha \in[1,2]$ and $c_{g}>0$ so that for $a, a^{\prime}$ with $|a|>\left|a^{\prime}\right|>0$,

$$
\left|\frac{a}{a^{\prime}}\right|^{\alpha} \geq \frac{g(a)}{g\left(a^{\prime}\right)} \geq c_{g}\left|\frac{a}{a^{\prime}}\right| .
$$

This kind of function $g$ is also called a "sketchable" function. Notice that the Huber function satisfies the above properties.

Theorem 20.5.2 (Modified version of Theorem 3.1 of [CW15b]). Function $g$ is symmetric, monotone and grows subquadratically (g is a G-function defined by [CW15b]). Given a matrix $A \in \mathbb{R}^{n \times d}$ and a matrix $B \in \mathbb{R}^{n \times m}$, there is an algorithm which can output a matrix $\widehat{X} \in \mathbb{R}^{d \times m}$ and a fitting cost vector $y \in \mathbb{R}^{m}$ such that with probability at least $1-1 / \operatorname{poly}(n m)$, $\forall i \in[m],\left\|A \widehat{X}_{i}-B_{i}\right\|_{g} \leq O(1) \cdot \min _{x \in \mathbb{R}^{d}}\left\|A x-B_{i}\right\|_{g}$, and $y_{i}=\Theta\left(\left\|A \widehat{X}_{i}-B_{i}\right\|_{g}\right)$. Furthermore, the running time is at most $\widetilde{O}(\operatorname{nnz}(A)+\operatorname{nnz}(B)+m \cdot \operatorname{poly}(d \log n))$.

Proof. We run $O(\log (n m))$ repetitions of the single column regression algorithm shown in Theorem 3.1 of [CW15b] for all columns $B_{i}$ for $i \in[m]$. For each regression problem $\| A x-$ $B_{i} \|_{g}$, we take the solution whose estimated cost is the median among these $O(\log (n m)$ repetitions as $\widehat{X}_{i}$. Then by the Chernoff bound, we can boost the success probability of each column to $1-1 / \operatorname{poly}(n m)$. By taking a union bound over all columns, we complete the proof.

### 20.5.2 $\ell_{p}$ Regression

One of the most important cases in regression and low rank approximation problems is when the error measure is $\ell_{p}$. For $\ell_{p}$ regression, though it can be solved by convex optimization/linear programming exactly, we can get a much faster running time if we allow some approximation ratios. In the following theorem, we show that there is an algorithm which can be used to solve $\ell_{p}$ regression for any $p \geq 1$.

Theorem 20.5.3 (Modified version of [WZ13]). Let $p \geq 1, \epsilon \in(0,1)$. Given a matrix $A \in \mathbb{R}^{n \times d}$ and a matrix $B \in \mathbb{R}^{n \times m}$, there is an algorithm which can output a matrix $\widehat{X} \in$ $\mathbb{R}^{d \times m}$ and a fitting cost vector $y \in \mathbb{R}^{m}$ such that with probability at least $1-1 / \operatorname{poly}(n m)$, $\forall i \in[m],\left\|A \widehat{X}_{i}-B_{i}\right\|_{p}^{p} \leq(1+\epsilon) \cdot \min _{x \in \mathbb{R}^{d}}\left\|A x-B_{i}\right\|_{p}^{p}$, and $y_{i}=\Theta\left(\left\|A \widehat{X}_{i}-B_{i}\right\|_{p}^{p}\right)$. Furthermore, the running time is at most $\widetilde{O}\left(\mathrm{nnz}(A)+\mathrm{nnz}(B)+m n^{\max (1-2 / p, 0)} \cdot \operatorname{poly}(d)\right)$.

Proof. As in the proof of Theorem 20.5.2, we only need to run $O(\log (n m))$ repetitions of the single column regression algorithm shown in [WZ13].

### 20.5.3 $\ell_{0}$ Regression

Definition 20.5.1 (Regular partition). Given a matrix $A \in \mathbb{R}^{n \times k}$, we say $\left\{S_{1}, S_{2}, \cdots, S_{h}\right\}$ is a regular partition for $[n]$ with respect to the matrix $A$ if, for each $i \in[h]$,

$$
\operatorname{rank}\left(A^{S_{i}}\right)=\left|S_{i}\right|, \text { and } \operatorname{rowspan}\left(A^{S_{i}}\right)=\operatorname{rowspan}\left(A^{\cup_{j=i}^{h} S_{j}}\right),
$$

where $A^{S_{i}} \in \mathbb{R}^{\left|S_{i}\right| \times k}$ denotes the matrix that selects a subset $S_{i}$ of rows of the matrix $A$.

```
Algorithm \(20.2 \ell_{0}\) regression [APY09]
    procedure L0REGRESSION \((A, b, n, k, c) \quad \triangleright\) Theorem 20.5.4
        \(x^{\prime} \leftarrow 0^{k}\)
        \(\left\{S_{1}, S_{2}, \cdots, S_{h}\right\} \leftarrow \operatorname{GenerateRegularPartition}(A, n, k)\)
        \(x^{\prime} \leftarrow 0^{k}\)
        for \(i=1 \rightarrow h\) do
            Find a \(\widetilde{x}\) such that \(A^{S_{i}} \widetilde{x}=b_{S_{i}}\)
            if \(\|A \widetilde{x}-b\|_{0}<\left\|A x^{\prime}-b\right\|_{0}\) then
                \(x^{\prime} \leftarrow \widetilde{x}\)
            end if
        end for
        return \(x^{\prime}\)
    end procedure
```

[APY09] studied the Nearest Codeword problem over finite fields $\mathbb{F}_{2}$. Their proof can be extended to the real field and generalized to Theorem 20.5.4. For completeness, we still provide the proof of the following result.

Theorem 20.5.4 (Generalization of [APY09]). Given matrix $A \in \mathbb{R}^{n \times k}$ and vector $\mathbb{R}^{n}$, for any $c \in[1, k]$, there is an algorithm (Algorithm 20.2) that runs in $n^{O(1)}$ time and outputs a vector $x^{\prime} \in \mathbb{R}^{k}$ such that

$$
\left\|A x^{\prime}-b\right\|_{0} \leq k \min _{x \in \mathbb{R}^{k}}\|A x-b\|_{0}
$$

Proof. Let $x^{*} \in \mathbb{R}^{k}$ denote the optimal solution to $\min _{x \in \mathbb{R}^{k}}\|A x-b\|_{0}$. We define set $E$ as follows

$$
E=\left\{i \in[n] \mid\left(A x^{*}\right)_{i} \neq b_{i}\right\} .
$$

We create a regular partition $\left\{S_{1}, S_{2}, \cdots, S_{h}\right\}$ for $[n]$ with respect to $A$.
Let $i$ denote the smallest index such that $\left|S_{i} \cap E\right|=0$, i.e.,

$$
i=\min \left\{j| | S_{j} \cap E \mid=0\right\} .
$$

The linear equation we want to solve is $A^{S_{i}} x=b_{S_{i}}$. Let $\widetilde{x} \in \mathbb{R}^{k}$ denote a solution to $A^{S_{i}} \widetilde{x}=A^{S_{i}} x^{*}$ (Note that, by our choice of $i, b_{S_{i}}=A^{S_{i}} x^{*}$. Then we can rewrite $\|A \widetilde{x}-b\|_{0}$ in the following sense,

$$
\begin{equation*}
\|A \widetilde{x}-b\|_{0}=\sum_{j=1}^{i-1}\left\|A^{S_{j}} \widetilde{x}-b_{S_{j}}\right\|_{0}+\sum_{j=i}^{h}\left\|A^{S_{j}} \widetilde{x}-b_{S_{j}}\right\|_{0} \tag{20.7}
\end{equation*}
$$

For each $j \in\{1,2, \cdots, i-1\}$, we have

$$
\begin{align*}
\left\|A^{S_{j}} \widetilde{x}-b_{S_{j}}\right\|_{0} & \leq k \\
& \leq\left\|A^{S_{j}} x^{*}-b_{S_{j}}\right\|_{0} \cdot\lceil k\rceil \tag{20.8}
\end{align*}
$$

where the first step follows from $\left|S_{0}\right| \leq k$, and the last step follows from $\left\|A^{S_{j}} x^{*}-b_{S_{j}}\right\|_{0} \geq 1$, $\forall j \in[i-1]$.

Note that, by our choice of $i$, we have $A^{S_{i}} \widetilde{x}=A^{S_{i}} x^{*}$. Then for each $j \in\{i, i+$ $1, \cdots, n\}$, using the regular partition property, there always exists a matrix $P_{(j)}$ such that $A^{S_{j}}=P_{(j)} A^{S_{i}}$. Then we have

$$
\begin{equation*}
A^{S_{j}} \widetilde{x}=P_{(j)} A^{S_{i}} \widetilde{x}=P_{(j)} A^{S_{i}} x^{*}=A^{S_{j}} x^{*} \tag{20.9}
\end{equation*}
$$

Plugging Eq. (20.8) and (20.9) into Eq. (20.7), we have

$$
\begin{aligned}
\|A \widetilde{x}-b\|_{0} & =\sum_{j=1}^{i-1}\left\|A^{S_{j}} \widetilde{x}-b_{S_{j}}\right\|_{0}+\sum_{j=i}^{h}\left\|A^{S_{j}} \widetilde{x}-b_{S_{j}}\right\|_{0} \\
& \leq k \sum_{j=1}^{i-1}\left\|A^{S_{j}} x^{*}-b_{S_{j}}\right\|_{0}+\sum_{j=i}^{h}\left\|A^{S_{j}} x^{*}-b_{S_{j}}\right\|_{0} \\
& \leq k\left\|A x^{*}-b\right\|_{0} .
\end{aligned}
$$

This completes the proof.

### 20.6 Hardness

### 20.6.1 Column Subset Selection for the Huber Function

The rough idea here is to define $k=\Omega(\sqrt{\log n})$ groups of columns, where we carefully choose the $i$-th group to have $n^{1-2 i \epsilon}$ columns, $\epsilon=.2 /(1.5 k)$, and in the $i$-th group each column has the form

$$
n^{1.5 i \epsilon} \cdot 1^{n}+\left[ \pm n^{-.2+i \epsilon}, \ldots, \pm n^{-.2+i \epsilon}, \pm n^{.5+2 i \epsilon}, \ldots, \pm n^{.5+2 i \epsilon}\right]
$$

where there are $n-n^{1}$ coordinates where the perturbation is randomly either $+n^{-.2+i \epsilon}$ or $-n^{-.2+i \epsilon}$, and the remaining $n^{.1}$ coordinates are randomly either $+n^{.5+2 i \epsilon}$ or $-n^{.5+2 i \epsilon}$. We call the former type of coordinates "small noise", and the latter "large noise". All remaining columns in the matrix are set to 0 . Because of the random signs, it is very hard to fit the noise in one column to that of another column. One can show, that to approximate a column in the $j$-th group by a column in the $i$-th group, $i<j$, one needs to scale by roughly $n^{1.5(j-i) \epsilon}$, just to cancel out the "mean" $n^{1.5 j \epsilon} \cdot 1^{n}$. But when doing so, since the Huber function is quadratic for small values, the scaled small noise is now magnified more than linearly compared to what it was before, and this causes a column in the $i$-th group not to be a good approximation of a column in the $j$-th group. On the other hand, if you want to approximate a column in the $j$-th group by a column in the $i$-th group, $i>j$, one again needs to scale by roughly $n^{1.5(j-i) \epsilon}$ just to cancel out the "mean", but now one can show the large noise from the column in the $i$-th group is too large and remains in the linear regime, causing a poor approximation. The details of this construction are given in the following theorem.

Theorem 20.6.1. Let $H(x)$ denote the modified Huber function with $\tau=1$, i.e.,

$$
H(x)= \begin{cases}x^{2} / \tau, & \text { if }|x|<\tau \\ |x|, & \text { if }|x| \geq \tau\end{cases}
$$

For any $n \geq 1$, there is a matrix $A \in \mathbb{R}^{n \times n}$ such that, if we select $o(\sqrt{\log n})$ columns to fit the entire matrix, there is no $O(1)$-approximation, i.e., for any subset $S \subseteq[n]$ with $|S|=o(\sqrt{\log n})$,

$$
\min _{X \in \mathbb{R} \mid \text { | } \mid \times n}\left\|A_{S} X-A\right\|_{H} \geq \omega(1) \cdot \min _{\operatorname{rank}-1 A^{\prime}}\left\|A^{\prime}-A\right\|_{H} .
$$

Proof. Suppose there is an algorithm which only finds a subset with size $k / 2=o(\sqrt{\log n})$. We want to prove a lower bound on its approximation ratio.

Let $\epsilon=0.2 /(1.5 k)$. Let $A$ denote a matrix with $k+1$ groups of columns.
For each group $i \in[k], I_{i}$ has $n^{1-2 i \epsilon}$ columns which are

$$
\left[\begin{array}{c}
n^{1.5 i \epsilon} \\
n^{1.5 i \epsilon} \\
n^{1.5 i \epsilon} \\
\vdots \\
n^{1.5 i \epsilon} \\
n^{1.5 i \epsilon} \\
n^{1.5 i \epsilon} \\
n^{1.5 i \epsilon}
\end{array}\right]+\left[\begin{array}{c} 
\pm n^{-0.2+i \epsilon} \\
\pm n^{-0.2+i \epsilon} \\
\vdots \\
\pm n^{-0.2+i \epsilon} \\
\pm n^{0.5+2 i \epsilon} \\
\pm n^{0.5+2 i \epsilon} \\
\vdots \\
\pm n^{0.5+2 i \epsilon}
\end{array}\right] \in \mathbb{R}^{n},
$$

where $\pm$ indicates i.i.d. random signs. For the error column, the first $n-n^{0.1}$ rows are $n^{-0.2+i \epsilon}$, and the last $n^{0.1}$ rows are $n^{0.5+2 i \epsilon}$.

The last group of $n-\sum_{i=1}^{k} n^{1-2 i \epsilon}$ columns are

$$
\left[\begin{array}{c}
0 \\
0 \\
\vdots \\
0 \\
0
\end{array}\right] \in \mathbb{R}^{n} .
$$

The optimal cost is at most

$$
\begin{aligned}
\sum_{i=1}^{k} n^{1-2 i \epsilon} \cdot\left(\left(n-n^{0.1}\right) H\left(n^{-0.2+i \epsilon}\right)+n^{0.1} H\left(n^{0.5+2 i \epsilon}\right)\right) & \leq \sum_{i=1}^{k} n^{1-2 i \epsilon}\left(n \cdot n^{-0.4+2 i \epsilon}+n^{0.6+2 i \epsilon}\right) \\
& \leq O\left(k n^{1.6}\right)
\end{aligned}
$$

where the second step follows since $n^{-0.2+i \epsilon}<1$ and $n^{0.5+2 i \epsilon} \geq 1$. Thus, it implies

$$
\min _{\operatorname{rank}-1}\left\|A^{\prime}-A\right\|_{H} \leq O\left(k n^{1.6}\right)
$$

Now let us consider the lower bound for using a subset of columns to fit the matrix. First, we fix a set $S=\left\{j_{1}, j_{2}, \cdots, j_{k / 2}\right\}$ of $k / 2$ columns. Since there are $k$ groups, and $|S| \leq k / 2$, the number of groups $I_{i}$ for $i \in[k]$ with $S \cap I_{i}=\emptyset$ is at least $k / 2$. It means that there are at least $k / 2$ groups for which $S$ does not have any column from them. Notice that the optimal cost is at most $O\left(k n^{1.6}\right)$, so it suffices to prove that $\forall i \in[k]$ with $I_{i} \cap S=\emptyset$, each column $j \in I_{i}$ will contribute a cost of $\omega\left(n^{0.6+2 i \epsilon}\right)$.

For notation, we $\operatorname{use} \operatorname{group}(j)$ to denote the index of the group which contains the column $j$. For each column $j$, we use $\Delta_{j}$ to denote the noise part, and use $A_{j}^{*}$ to denote the rank-1 "ground truth" part. Notice that $A_{j}=A_{j}^{*}+\Delta_{j}$.

Claim 20.6.2 (Noise cannot be used to fit other vectors). Let $x_{1}, x_{2}, \cdots, x_{s} \in \mathbb{R}^{m}$ be $s$ random sign vectors. Then with probability at least $1-2^{s} \cdot 2^{-\Theta\left(m / 2^{s}\right)}, \forall \alpha_{1}, \alpha_{2}, \cdots, \alpha_{s} \in \mathbb{R}$,
the size of $Z_{\alpha_{1}, \alpha_{2}, \cdots, \alpha_{s}}=\left\{i \in[m] \mid \operatorname{sign}\left(\left(\alpha_{1} x_{1}\right)_{i}\right)=\operatorname{sign}\left(\left(\alpha_{2} x_{2}\right)_{i}\right)=\cdots=\operatorname{sign}\left(\left(\alpha_{s} x_{s}\right)_{i}\right)=\right.$ $\operatorname{sign}(+1)\}$ is at least $\Omega\left(m / 2^{s}\right)$.

Proof. For a set of fixed $\alpha_{1}, \alpha_{2}, \cdots, \alpha_{s}$, the claim follows from the Chernoff bound. Since there are $2^{s}$ different possibilities of signs of $\alpha_{1}, \cdots, \alpha_{s}$, taking a union bound over them completes the proof.

Now we consider a specific column $j \in I_{i}$ for some $i \in[k]$, where $I_{i} \cap S=\emptyset$. Suppose the fitting coefficients are $\alpha_{1}, \alpha_{2}, \cdots, \alpha_{k / 2}$. Consider the following term

$$
\begin{aligned}
& \left(\alpha_{1} A_{j_{1}}+\alpha_{2} A_{j_{2}}+\cdots+\alpha_{k / 2} A_{k / 2}\right)-A_{j} \\
= & \left(\alpha_{1} A_{j_{1}}^{*}+\alpha_{2} A_{j_{2}}^{*}+\cdots+\alpha_{k / 2} A_{k / 2}^{*}-A_{j}^{*}\right)+\left(\alpha_{1} \Delta_{j_{1}}+\alpha_{2} \Delta_{j_{2}}+\cdots+\alpha_{k / 2} \Delta_{k / 2}-\Delta_{j}\right)
\end{aligned}
$$

Let $u^{*}=\left(\alpha_{1} A_{j_{1}}^{*}+\alpha_{2} A_{j_{2}}^{*}+\cdots+\alpha_{k / 2} A_{k / 2}^{*}-A_{j}^{*}\right)$. By Claim 20.6.2, with probability at least $1-(k / 2) \cdot 2^{-\Theta\left(n / 2^{k / 2}\right)}$,

$$
\left|\left\{t \in[n] \mid \operatorname{sign}\left(u_{t}^{*}\right)=\operatorname{sign}\left(\alpha_{1} \Delta_{j_{1}, t}\right)=\cdots=\operatorname{sign}\left(\alpha_{k / 2} \Delta_{j_{k / 2}, t}\right)=\operatorname{sign}\left(-\Delta_{j, t}\right)\right\}\right|=\Omega\left(n / 2^{k / 2}\right)
$$

Observe that all the coordinates of $u^{*}$ are the same, and the absolute value of each entry of $u^{*}$ should be at most $O\left(n^{1.5 i \epsilon}\right)$. Otherwise the column already has $\omega\left(n / 2^{k / 2} \cdot n^{1.5 i \epsilon}\right)=\omega\left(n^{0.6+2 i \epsilon}\right)$ cost. Thus, the magnitude of each entry of $\alpha_{1} A_{j_{1}}^{*}+\alpha_{2} A_{j_{2}}^{*}+\cdots+\alpha_{k / 2} A_{k / 2}^{*}$ is $\Theta\left(n^{1.5 i \epsilon}\right)$. Thus, there exists $t \in[k / 2]$, such that the absolute value of each entry of $\alpha_{t} A_{j_{t}}^{*}$ is at least $\Omega\left(n^{1.5 i \epsilon} / k\right)$. Then there are two cases.

The first case is $\operatorname{group}(t)<\operatorname{group}(j)$. Let $\operatorname{group}(t)=i^{\prime}$. Then $\left|\alpha_{t}\right|=\Omega\left(n^{1.5\left(i-i^{\prime}\right) \epsilon} / k\right)$. By Claim 20.6.2 again, with probability at least $1-(k / 2) \cdot 2^{-\Theta\left(n / 2^{k / 2}\right)}$, the size of

$$
\left\{z \in\left[n-n^{0.1}\right] \mid \operatorname{sign}\left(u_{z}^{*}\right)=\operatorname{sign}\left(\alpha_{1} \Delta_{j_{1}, z}\right)=\cdots=\operatorname{sign}\left(\alpha_{k / 2} \Delta_{j_{k / 2}, z}\right)=\operatorname{sign}\left(-\Delta_{j, z}\right)\right\}
$$



Figure 20.1: The blue curve is the Huber function which combines an $\ell_{2}$-like measure for small $x$ with an $\ell_{1}$-like measure for large $x$. The red curve is the "reverse" Huber function which combines an $\ell_{1}$-like measure for small $x$ with an $\ell_{2}$-like measure for large $x$.
is at least $\Omega\left(n / 2^{k / 2}\right)$. Thus, the cost to fit is at least $\Omega\left(n / 2^{k / 2}\right) \cdot\left(n^{-0.2+i^{\prime} \epsilon} n^{1.5\left(i-i^{\prime}\right) \epsilon} / k\right)^{2}=$ $\omega\left(n^{0.6+2 i \epsilon}\right)$.

The second case is $\operatorname{group}(t)>\operatorname{group}(j)$. Let $\operatorname{group}(t)=i^{\prime}$. Then $\left|\alpha_{t}\right|=\Omega\left(n^{1.5\left(i-i^{\prime}\right) \epsilon} / k\right)$. By Claim 20.6.2 again, with probability at least $1-(k / 2) \cdot 2^{-\Theta\left(n^{0.1} / 2^{k / 2}\right)}$, the size of $\left\{z \in\left\{n-n^{0.1}+1, \cdots, n\right\} \mid \operatorname{sign}\left(u_{z}^{*}\right)=\operatorname{sign}\left(\alpha_{1} \Delta_{j_{1}, z}\right)=\cdots=\operatorname{sign}\left(\alpha_{k / 2} \Delta_{j_{k / 2}, z}\right)=\operatorname{sign}\left(-\Delta_{j, z}\right)\right\}$ is at least $\Omega\left(n^{0.1} / 2^{k / 2}\right)$. Thus, the fitting cost is at least $\Omega\left(n^{0.1} / 2^{k / 2}\right) \cdot\left(n^{0.5+2 i^{\prime} \epsilon} n^{1.5\left(i-i^{\prime}\right) \epsilon} / k\right)=$ $\omega\left(n^{0.6+2 i \epsilon}\right)$.

By taking a union bound over all columns $j$, we have with probability at least $1-2^{n^{\ominus(1)}}$, the total cost to fit by a column subset $S$ is at least $\omega\left(k n^{1.6}\right)$.

Then, by taking a union bound over all the $\binom{n}{k}$ number of sets $S$, we complete the proof.

### 20.6.2 Column Subset Selection for the Reverse Huber Function

In this section, we consider a "reverse Huber function": $g(x)=x^{2}$ if $x \geq 1$ and $g(x)=|x|$ for $x \leq 1$.

Theorem 20.6.3. Let $g(x)$ denote the "reverse Huber function" with $\tau=1$, i.e.,

$$
H(x)= \begin{cases}x^{2} / \tau, & \text { if }|x|>\tau \\ |x|, & \text { if }|x| \leq \tau\end{cases}
$$

For any $n \geq 1$, there is a matrix $A \in \mathbb{R}^{n \times n}$ such that, if we select only 1 column to fit the entire matrix, there is no $n^{o(1)}$-approximation to the best rank-1 approximation, i.e., for any subset $S \subseteq[n]$ with $|S|=1$,

$$
\min _{X \in \mathbb{R}^{|S| \times n}}\left\|A_{S} X-A\right\|_{g} \geq n^{\Omega(1)} \cdot \min _{\operatorname{rank}-1}\left\|A^{\prime}-A\right\|_{g}
$$

Proof. Let $A \in \mathbb{R}^{n \times n}$ have one column that is $a=\left(n^{1 / 2}, 0, \ldots, 0\right)^{\top}$ and $n-1$ columns that are each equal to $b=(0,1 / n, 1 / n, \ldots, 1 / n)^{\top}$. If we choose one column which has the form as $a$ to fit the other columns, the cost is at least $(n-1)^{2} / n=\Theta(n)$. If we choose one column which has the form as $b$ to fit the other columns, the cost is at least $\left(n^{1 / 2}\right)^{2}=\Theta(n)$.

Now we consider using a vector $c=\left(1 / n^{1 / 4}, 1 / n, 1 / n, \ldots, 1 / n\right)^{\top}$ to fit all the columns. One can use $c$ to approximate $a$ with cost at most $(n-1) \cdot n^{3 / 4} / n=\Theta\left(n^{3 / 4}\right)$ by matching the first coordinate, while one can use $c$ to approximate $b$ with cost at most $1 / n^{1 / 4}$ by matching the last $n-1$ coordinates, and since there are $n-1$ columns equal to $b$, the overall total cost of using $c$ to approximate matrix $A$ is $\Theta\left(n^{3 / 4}\right)$.

## Chapter 21

## Tensor Low Rank Approximation

We consider relative error low rank approximation of tensors with respect to the Frobenius norm. Namely, given an order- $q$ tensor $A \in \mathbb{R}^{\Pi_{i=1}^{q} n_{i}}$, output a rank- $k$ tensor $B$ for which $\|A-B\|_{F}^{2} \leq(1+\epsilon)$ OPT, where OPT $=\inf _{\text {rank-k } A^{\prime}}\left\|A-A^{\prime}\right\|_{F}^{2}$. Despite much success on obtaining relative error low rank approximations for matrices, no such results were known for tensors for arbitrary $(1+\epsilon)$-approximations. One structural issue is that there may be no rank-k tensor $A_{k}$ achieving the above infinum. Another, computational issue, is that an efficient relative error low rank approximation algorithm for tensors would allow one to compute the rank of a tensor, which is NP-hard. We bypass these two issues via (1) bicriteria and (2) parameterized complexity solutions:

1. We give an algorithm which outputs a rank $k^{\prime}=O\left((k / \epsilon)^{q-1}\right)$ tensor $B$ for which $\|A-B\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}$ in $\operatorname{nnz}(A)+n \cdot \operatorname{poly}(k / \epsilon)$ time in the real RAM model, whenever either $A_{k}$ exists or OPT $>0$. Here $\mathrm{nnz}(A)$ denotes the number of nonzero entries in $A$. If both $A_{k}$ does not exist and $\mathrm{OPT}=0$, then $B$ instead satisfies $\|A-B\|_{F}^{2}<\gamma$, where $\gamma$ is any positive, arbitrarily small function of $n$.
2. We give an algorithm for any $\delta>0$ which outputs a rank $k$ tensor $B$ for which $\|A-B\|_{F}^{2} \leq(1+\epsilon)$ OPT and runs in $\left(\operatorname{nnz}(A)+n \operatorname{poly}(k / \epsilon)+\exp \left(k^{2} / \epsilon\right)\right) \cdot n^{\delta}$ time in the unit cost RAM model, whenever OPT $>2^{-O\left(n^{\delta}\right)}$ and there is a rank- $k$ tensor $B=$
$\sum_{i=1}^{k} u_{i} \otimes v_{i} \otimes w_{i}$ for which $\|A-B\|_{F}^{2} \leq(1+\epsilon / 2) \mathrm{OPT}$ and $\left\|u_{i}\right\|_{2},\left\|v_{i}\right\|_{2},\left\|w_{i}\right\|_{2} \leq 2^{O\left(n^{\delta}\right)}$. If OPT $\leq 2^{-\Omega\left(n^{\delta}\right)}$, then $B$ instead satisfies $\|A-B\|_{F}^{2} \leq 2^{-\Omega\left(n^{\delta}\right)}$.

Our first result is polynomial time, and in fact input sparsity time, in $n, k$, and $1 / \epsilon$, for any $k \geq 1$ and any $0<\epsilon<1$, while our second result is fixed parameter tractable in $k$ and $1 / \epsilon$. For outputting a rank- $k$ tensor, or even a bicriteria solution with rank- $C k$ for a certain constant $C>1$, we show a $2^{\Omega\left(k^{1-o(1)}\right)}$ time lower bound under the Exponential Time Hypothesis.

Our results are based on an "iterative existential argument", and also give the first relative error low rank approximations for tensors for a large number of error measures for which nothing was known. In particular, we give the first relative error approximation algorithms on tensors for: column row and tube subset selection, entrywise $\ell_{p}$-low rank approximation for $1 \leq p<2$, low rank approximation with respect to sum of Euclidean norms of faces or tubes, weighted low rank approximation, and low rank approximation in distributed and streaming models. We also obtain several new results for matrices, such as nnz $(A)$-time CUR decompositions, improving the previous $n n z(A) \log n$-time CUR decompositions, which may be of independent interest.

### 21.1 Introduction

Low rank approximation of matrices is one of the most well-studied problems in randomized numerical linear algebra. Given an $n \times d$ matrix $A$ with real-valued entries, we want to output a rank- $k$ matrix $B$ for which $\|A-B\|$ is small, under a given norm. While this problem can be solved exactly using the singular value decomposition for some norms like the spectral and Frobenius norms, the time complexity is still $\min \left(n d^{\omega-1}, d n^{\omega-1}\right)$, where $\omega \approx 2.373$ is the exponent of matrix multiplication [Str69, CW87, Wil12, LG14]. This time complexity is prohibitive when $n$ and $d$ are large. By now there are a number of approximation algorithms for this problem, with the Frobenius norm ${ }^{1}$ being one of the most common error measures. Initial solutions [FKV04, AM07] to this problem were based on sampling and achieved additive error in terms of $\epsilon\|A\|_{F}$, where $\epsilon>0$ is an approximation parameter, which can be arbitrarily larger than the optimal cost $\mathrm{OPT}=\min _{\text {rank-k }} B \| A-$ $B \|_{F}^{2}$. Since then a number of solutions based on the technique of oblivious sketching [Sar06, CW13, MM13, NN13a] as well as sampling based on non-uniform distributions [DMM06c, DMM06b, DMM08, DMIMW12], have been proposed which achieve the stronger notion of relative error, namely, which output a rank- $k$ matrix $B$ for which $\|A-B\|_{F}^{2} \leq(1+\epsilon)$ OPT with high probability. It is now known how to output a factorization of such a $B=U \cdot V$, where $U$ is $n \times k$ and $V$ is $k \times d$, in $\operatorname{nnz}(A)+(n+d)$ poly $(k / \epsilon)$ time [CW13, MM13, NN13a]. Such an algorithm is optimal, up to the poly $(k / \epsilon)$ factor, as any algorithm achieving relative error must read almost all of the entries.

Tensors are often more useful than matrices for capturing higher order relations

[^66]in data. Computing low rank factorizations of approximations of tensors is the primary task of interest in a number of applications, such as in psychology[Kro83], chemometrics [Paa00, SBG04], neuroscience [AAB $\left.{ }^{+} 07, \mathrm{~KB} 09, \mathrm{CLK}^{+} 15\right]$, computational biology [CV15, SC15], natural language processing [CYYM14, LZBJ14, LZMB15, BNR $\left.{ }^{+} 15\right]$, computer vision [VT02, WA03, SH05, HPS05, HD08, AFdLGTL09, PLY10, LFC ${ }^{+}$16, CLZ17], computer graphics [VT04, WWS ${ }^{+} 05$, Vas09], security [AÇKY05, ACY06, KB06], cryptography [FS99, Sch12, KYFD15, SHW ${ }^{+}$16] data mining [KS08, RST10, KABO10, Mør11], machine learning applications such as learning hidden Markov models, reinforcement learning, community detection, multi-armed bandit, ranking models, neural network, Gaussian mixture models and Latent Dirichlet allocation [MR05, AFH ${ }^{+}$12, HK13, ALB13, ABSV14, AGH ${ }^{+}$14, AGHK14, BCV14, JO14a, GHK15, PBLJ15, JSA15, ALA16, AGMR16, ZSJ+ 17], programming languages [RTP16], signal processing [Wes94, DLDM98, Com09, CMDL+ 15], and other applications [YCS11, LMWY13, OS14, ZCZJ14, STLS14, YCS16, RNSS16].

Despite the success for matrices, the situation for order- $q$ tensors for $q>2$ is much less understood. There are a number of works based on alternating minimization [CC70, Har70, FMPS13, FT15, ZG01, BS15] gradient descent or Newton methods [ES09, ZG01], methods based on the Higher-order SVD (HOSVD) [LMV00a] which provably incur $\Omega(\sqrt{n})$ inapproximability for Frobenius norm error [LMV00b], the power method or orthogonal iteration method [LMV00b], additive error guarantees in terms of the flattened (unfolded) tensor rather than the original tensor [MMD08], tensor trains [Ose11], the tree Tucker decomposition [OT09], or methods specialized to orthogonal tensors [KM11, $\mathrm{AGH}^{+} 14$, MHG15, WTSA15, WA16, SWZ16]. There are also a number of works on the problem of tensor completion, that is, recovering a low rank tensor from missing entries [WM01, AKDM10,

TSHK11, LMWY13, MHWG14, JO14b, BM16]. There is also another line of work using the sum of squares (SOS) technique to study tensor problems [BKS15a, GM15, HSS15, HSSS16, MSS16, PS17, SS17], other recent work on tensor PCA [All12b, All12a, RM14, JMZ15, ADGM16, ZX17], and work applying smoothed analysis to tensor decomposition [BCMV14]. Several previous works also consider more robust norms than the Frobenius norm for tensors, e.g., the $R_{1}$ norm ( $\ell_{1}-\ell_{2}-\ell_{2}$ norm in our work) [HD08], $\ell_{1}$-PCA [PLY10], entry-wise $\ell_{1}$ regularization [GGH14], M-estimator loss [YFS16], weighted approximation [Paa97, TK11, LRHG13], tensor-CUR [OST08, MMD08, CC10, FMMN11, FT15], or robust tensor PCA [GQ14, LFC ${ }^{+}$16, CLZ17].

Some of the above works, such as ones based on the tensor power method or alternating minimization, require incoherence or orthogonality assumptions. Others, such as those based on the simultaneous SVD, require an assumption on the minimum singular value. See the monograph of Moitra [Moi14a] for further discussion. Unlike the situation for matrices, there is no work for tensors that is able to achieve the following natural relative error guarantee: given a $q$-th order tensor $A \in \mathbb{R}^{n^{\otimes q}}$ and an arbitrary accuracy parameter $\epsilon>0$, output a rank- $k$ tensor $B$ for which

$$
\begin{equation*}
\|A-B\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT} \tag{21.1}
\end{equation*}
$$

where $\mathrm{OPT}=\inf _{\text {rank-k }} B^{\prime}\left\|A-B^{\prime}\right\|_{F}^{2}$, and where recall the rank of a tensor $B$ is the minimal integer $k$ for which $B$ can be expressed as $\sum_{i=1}^{k} U_{1}^{i} \otimes U_{2}^{i} \otimes \cdots \otimes U_{q}^{i}$. General speaking, a $q$-th order tensor can have rank at most $n^{q-1}$. A third order tensor, for example, has rank which is an integer in $\left\{0,1,2, \ldots, n^{2}\right\}$. We note that [BCV14] is able to achieve a relative error 5 -approximation for third order tensors, and an $O(q)$-approximation for $q$-th order tensors,
though it cannot achieve a $(1+\epsilon)$-approximation. We compare our work to [BCV14] in Section 21.1.4 below.

For notational simplicity, we will start by assuming third order tensors with all dimensions of equal size, but we extend all of our main theorems below to tensors of any constant order $q>3$ and dimensions of different sizes.

The first caveat regarding (21.1) for tensors is that an optimal rank- $k$ solution may not even exist! This is a well-known problem for tensors (see, e.g., [KHL89, Paa00, KDS08, Ste06, Ste08] and more details in section 4 of [DSL08]), for which for any rank- $k$ tensor $B$, there always exists another rank- $k$ tensor $B^{\prime}$ for which $\left\|A-B^{\prime}\right\|_{F}^{2}<\|A-B\|_{F}^{2}$. If OPT $=0$, then in this case for any rank- $k$ tensor $B$, necessarily $\|A-B\|_{F}^{2}>0$, and so (21.1) cannot be satisfied. This fact was known to algebraic geometers as early as the 19th century, which they refer to as the fact that the locus of $r$-th secant planes to a Segre variety may not define a (closed) algebraic variety [DSL08, Lan12]. It is also known as the phenomenon underlying the concept of border $\operatorname{rank}^{2}$ [Bin80, Bin86, BCS97, Knu98, Lan06]. In this case it is natural to allow the algorithm to output an arbitrarily small $\gamma>0$ amount of additive error. Note that unlike several additive error algorithms for matrices, the additive error here can in fact be an arbitrarily small positive function of $n$. If, however, OPT $>0$, then for any $\epsilon>0$, there exists a rank- $k$ tensor $B$ for which $\|A-B\|_{F}^{2} \leq(1+\epsilon)$ OPT, and in this case we should still require the algorithm to output a relative-error solution. If an optimal rank- $k$ solution $B$ exists, then as for matrices, it is natural to require the algorithm to output a relative-error solution.

[^67]Besides the above definitional issue, a central reason that (21.1) has not been achieved is that computing the rank of a third order tensor is well-known to be NP-hard [Hås90, HL13]. Thus, if one had such a polynomial time procedure for solving the problem above, one could determine the rank of $A$ by running the procedure on each $k \in\left\{0,1,2, \ldots, n^{2}\right\}$, and check for the first value of $k$ for which $\|A-B\|_{F}^{2}=0$, thus determining the rank of $A$. However, it is unclear if approximating the tensor rank is hard. This question will also be answered in this work.

The main question which we address is how to define a meaningful notion of (21.1) for the case of tensors and whether it is possible to obtain provably efficient algorithms which achieve this guarantee, without any assumptions on the tensor itself. Besides (21.1), there are many other notions of relative error for low rank approximation of matrices for which provable guarantees for tensors are unknown, such as tensor CURT, $R_{1}$ norm, and the weighted and $\ell_{1}$ norms mentioned above. Our goal is to provide a general technique to obtain algorithms for many of these variants as well.

### 21.1.1 Our Results

To state our results, we first consider the case when a rank- $k$ solution $A_{k}$ exists, that is, there exists a rank- $k$ tensor $A_{k}$ for which $\left\|A-A_{k}\right\|_{F}^{2}=\mathrm{OPT}$.

We first give a $\operatorname{poly}(n, k, 1 / \epsilon)$-time $(1+\epsilon)$-relative error approximation algorithm for any $0<\epsilon<1$ and any $k \geq 1$, but allow the output tensor $B$ to be of rank $O\left((k / \epsilon)^{2}\right)$ (for general $q$-order tensors, the output rank is $O\left((k / \epsilon)^{q-1}\right.$ ), whereas we measure the cost of $B$ with respect to rank- $k$ tensors. Formally, $\|A-B\|_{F}^{2} \leq(1+\epsilon)\left\|A-A_{k}\right\|_{F}^{2}$. In fact, our algorithm can be implemented in $\mathrm{nnz}(A)+n \cdot \operatorname{poly}(k / \epsilon)$ time in the real-RAM model, where $\operatorname{nnz}(A)$ is the number of non-zero entries of $A$. Since it is necessary to use $\operatorname{nnz}(A)$ time to avoid missing any (large) entry, such an algorithm is optimal (for running time) for any relative error algorithm, even bicriteria ones.

If $A_{k}$ does not exist, then our output $B$ instead satisfies $\|A-B\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}+\gamma$, where $\gamma$ is an arbitrarily small additive error. Since $\gamma$ is arbitrarily small, $(1+\epsilon)$ OPT $+\gamma$ is still a relative error whenever OPT $>0$. Our theorem is as follows.

Theorem 21.1.1 (A Version of Theorem 21.4.9, bicriteria). Given a 3rd order tensor $A \in$ $\mathbb{R}^{n \times n \times n}$, if $A_{k}$ exists then there is a randomized algorithm running in $\mathrm{nnz}(A)+n \cdot \operatorname{poly}(k / \epsilon)$ time which outputs a (factorization of a) rank- $O\left(k^{2} / \epsilon^{2}\right)$ tensor $B$ for which $\|A-B\|_{F}^{2} \leq$ $(1+\epsilon)\left\|A-A_{k}\right\|_{F}^{2}$. If $A_{k}$ does not exist, then the algorithm outputs a rank- $O\left(k^{2} / \epsilon^{2}\right)$ tensor $B$ for which $\|A-B\|_{F}^{2} \leq(1+\epsilon)$ OPT $+\gamma$, where $\gamma>0$ is an arbitrarily small positive function of $n$. In both cases, the success probability is at least $2 / 3$.

One of the main applications of matrix low rank approximation is parameter reduction, as one can store the matrix using fewer parameters in factored form or more quickly
multiply by the matrix if given in factored form, as well as remove directions that correspond to noise. In such applications, it is not essential that the low rank approximation have rank exactly $k$, since one still has a significant parameter reduction with a matrix of slightly larger rank. This same motivation applies to tensor low rank approximation; we obtain both space and time savings by representing a tensor in factored form, and in such applications bicriteria applications suffice. Moreover, the extremely efficient nnz $(A)+n \cdot \operatorname{poly}(k / \epsilon)$ time algorithm we obtain may outweigh the need for outputting a tensor of rank exactly $k$. Bicriteria algorithms are common for coping with hardness; see e.g., results on robust low rank approximation of matrices [DV07, FFSS07, CW15a], sparse recovery [CKPS16], clustering [MMSW15, HT16], and approximation algorithms more generally.

We note that there are other applications, such as unique tensor decomposition in the method of moments, see, e.g., [BCV14], where one may have a hard rank constraint of $k$ for the output. However, in such applications the so-called Tucker decomposition is still a useful dimensionality-reduction analogue of the SVD and our techniques for proving

We next consider the case when the rank parameter $k$ is small, and we try to obtain rank- $k$ solutions which are efficient for small values of $k$. As before, we first suppose that $A_{k}$ exists.

If $A_{k}=\sum_{i=1}^{k} u_{i} \otimes v_{i} \otimes w_{i}$ and the norms $\left\|u_{i}\right\|_{2},\left\|v_{i}\right\|_{2}$, and $\left\|w_{i}\right\|_{2}$ are bounded by $2^{\text {poly }(n)}$, we can return a rank- $k$ solution $B$ for which $\|A-B\|_{F}^{2} \leq(1+\epsilon)\left\|A-A_{k}\right\|_{F}^{2}+2^{-\operatorname{poly}(n)}$, in $f(k, 1 / \epsilon) \cdot \operatorname{poly}(n)$ time in the standard unit cost RAM model with words of size $O(\log n)$ bits. Thus, our algorithm is fixed parameter tractable in $k$ and $1 / \epsilon$, and in fact remains polynomial time for any values of $k$ and $1 / \epsilon$ for which $k^{2} / \epsilon=O(\log n)$. This is motivated by a number of low rank approximation applications in which $k$ is typically small. The additive error of
$2^{-\operatorname{poly}(n)}$ is only needed in order to write down our solution $B$ in the unit cost RAM model, since in general the entries of $B$ may be irrational, even if the entries of $A$ are specified by $\operatorname{poly}(n)$ bits. If instead we only want to output an approximation to the value $\left\|A-A_{k}\right\|_{F}^{2}$, then we can output a number $Z$ for which OPT $\leq Z \leq(1+\epsilon)$ OPT, that is, we do not incur additive error.

When $A_{k}$ does not exist, there still exists a rank- $k$ tensor $\widetilde{A}$ for which $\|A-\widetilde{A}\|_{F}^{2} \leq$ OPT $+\gamma$, for arbitrary $\gamma>0$. We require there exists such a $\widetilde{A}$ for which if $\widetilde{A}=\sum_{i=1}^{k} u_{i} \otimes$ $v_{i} \otimes w_{i}$, then the norms $\left\|u_{i}\right\|_{2},\left\|v_{i}\right\|_{2}$, and $\left\|w_{i}\right\|_{2}$ are bounded by $2^{\text {poly }(n)}$.

The assumption in the previous two paragraphs that the factors of $A_{k}$ and of $\widetilde{A}$ have norm bounded by $2^{\text {poly(n) }}$ is necessary in certain cases, e.g., if OPT $=0$ and we are to write down the factors in poly $(n)$ time. An abridged version of our theorem is as follows.

Theorem 21.1.2 (Combination of Theorem 21.4.1 and 21.4.2, rank-k). Given a 3rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $\delta>0$, if $A_{k}=\sum_{i=1}^{k} u_{i} \otimes v_{i} \otimes w_{i}$ exists and each of $\left\|u_{i}\right\|_{2},\left\|v_{i}\right\|_{2}$, and $\left\|w_{i}\right\|_{2}$ is bounded by $2^{O\left(n^{\delta}\right)}$, then there is a randomized algorithm running in $O(\mathrm{nnz}(A)+$ $\left.n \operatorname{poly}(k, 1 / \epsilon)+2^{O\left(k^{2} / \epsilon\right)}\right) \cdot n^{\delta}$ time in the unit cost RAM model with words of size $O(\log n)$ bits $^{3}$, which outputs a (factorization of a) rank-k tensor $B$ for which $\|A-B\|_{F}^{2} \leq(1+\epsilon) \| A-$ $A_{k} \|_{F}^{2}+2^{-O\left(n^{\delta}\right)}$. Further, we can output a number $Z$ for which $\mathrm{OPT} \leq Z \leq(1+\epsilon) \mathrm{OPT}$ in the same amount of time. When $A_{k}$ does not exist, if there exists a rank-k tensor $\widetilde{A}$ for which $\|A-\widetilde{A}\|_{F}^{2} \leq \mathrm{OPT}+2^{-O\left(n^{\delta}\right)}$ and $\widetilde{A}=\sum_{i=1}^{k} u_{i} \otimes v_{i} \otimes w_{i}$ is such that the norms $\left\|u_{i}\right\|_{2}$, $\left\|v_{i}\right\|_{2}$, and $\left\|w_{i}\right\|_{2}$ are bounded by $2^{O\left(n^{\delta}\right)}$, then we can output a (factorization of a) rank-k tensor $\widetilde{A}$ for which $\|A-\widetilde{A}\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}+2^{-O\left(n^{\delta}\right)}$.

[^68]Our techniques for proving Theorem 21.1.1 and Theorem 21.1.2 open up avenues for many other problems in linear algebra on tensors. We now define the problems and state our results for them.

There is a long line of research on matrix column subset selection and CUR decomposition [DMM08, BMD09, DR10, BDM11, FEGK13, BW14, WS15, ABF ${ }^{+}$16, SWZ17] under operator, Frobenius, and entry-wise $\ell_{1}$ norm. It is natural to consider tensor column subset selection or tensor-CURT ${ }^{4}$, however most previous works either give error bounds in terms of the tensor flattenings [DMM08], assume the original tensor has certain properties [OST08, FT15, TM17], consider the exact case which assumes the tensor has low rank [CC10], or only fit a high dimensional cross-shape to the tensor rather than to all of its entries [FMMN11]. Such works are not able to provide a $(1+\epsilon)$-approximation guarantee as in the matrix case without assumptions. We consider tensor column, row, and tube subset selection, with the goal being to find three matrices: a subset $C \in \mathbb{R}^{n \times c}$ of columns of $A$, a subset $R \in \mathbb{R}^{n \times r}$ of rows of $A$, and a subset $T \in \mathbb{R}^{n \times t}$ of tubes of $A$, such that there exists a tensor $U \in \mathbb{R}^{c \times r \times t}$ for which

$$
\begin{equation*}
\|U(C, R, T)-A\|_{\xi} \leq \alpha\left\|A_{k}-A\right\|_{\xi}+\gamma \tag{21.2}
\end{equation*}
$$

where $\gamma=0$ if $A_{k}$ exists and $\gamma=2^{-\operatorname{poly}(n)}$ otherwise, $\alpha>1$ is the approximation ratio, $\xi$ is either Frobenius norm or Entry-wise $\ell_{1}$ norm, and $U(C, R, T)=\sum_{i=1}^{c} \sum_{j=1}^{r} \sum_{l=1}^{t} U_{i, j, l} \cdot C_{i} \otimes$ $R_{j} \otimes T_{l}$. In tensor CURT decomposition, we also want to output $U$.

We provide a (nearly) input sparsity time algorithm for this, together with an alternative input sparsity time algorithm which chooses slightly larger factors $C, R$, and $T$.

[^69]To do this, we combine Theorem 21.1.1 with the following theorem which, given a factorization of a rank- $k$ tensor $B$, obtains $C, U, R$, and $T$ in terms of it:

Theorem 21.1.3 (Combination of Theorem 21.4.39 and 21.4.40, $\left\|\left\|\|_{F}\right.\right.$-norm, CURT decomposition). Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, let $k \geq 1$, and let $U_{B}, V_{B}, W_{B} \in \mathbb{R}^{n \times k}$ be given. There is an algorithm running in $O(\mathrm{nnz}(A) \log n)+\widetilde{O}\left(n^{2}\right)$ poly $(k, 1 / \epsilon)$ time (respectively, $O(\mathrm{nnz}(A))+n \operatorname{poly}(k, 1 / \epsilon)$ time) which outputs a subset $C \in \mathbb{R}^{n \times c}$ of columns of $A$, a subset $R \in \mathbb{R}^{n \times r}$ of rows of $A$, a subset $T \in \mathbb{R}^{n \times t}$ of tubes of $A$, together with a tensor $U \in \mathbb{R}^{c \times r \times t}$ with $\operatorname{rank}(U)=k$ such that $c=r=t=O(k / \epsilon)$ (respectively, $c=r=t=O(k \log k+k / \epsilon))$, and $\|U(C, R, T)-A\|_{F}^{2} \leq(1+\epsilon)\left\|U_{B} \otimes V_{B} \otimes W_{B}-A\right\|_{F}^{2}$ holds with probability at least 9/10.

Combining Theorems 21.1.2 and 21.1.3 (with $B$ being a $(1+O(\epsilon)$ )-approximation to A) we achieve Equation (21.2) with $\alpha=(1+\epsilon)$ and $\xi=F$ with the optimal number of columns, rows, tubes, and rank of $U$ (we mention our matching lower bound later), though the running time has an $2^{O\left(k^{2} / \epsilon\right)}$ term in it. We note that instead combining Theorem 21.1.1 and Theorem 21.1.3 gives a bicriteria result for CURT without a $2^{O\left(k^{2} / \epsilon\right)}$ term in the running time, though it is suboptimal in the number of columns, rows, tubes, and rank of $U$.

As a side result worth stating, our analysis improves the best matrix CUR decomposition algorithm under Frobenius norm [BW14], providing the first optimal nnz $(A)$-time algorithm:

Theorem 21.1.4 (Informal Version of Theorem 21.6.1, Matrix CUR decomposition). There is an algorithm, which given a matrix $A \in \mathbb{R}^{n \times d}$ and an integer $k \geq 1$, runs in $O(\mathrm{nnz}(A))+$ $(n+d) \operatorname{poly}(k, 1 / \epsilon)$ time and outputs three matrices: $C \in \mathbb{R}^{n \times c}$ containing c columns of $A$,
$R \in \mathbb{R}^{r \times d}$ containing $r$ rows of $A$, and $U \in \mathbb{R}^{c \times r}$ with $\operatorname{rank}(U)=k$ for which $r=c=O(k / \epsilon)$ and $\|C U R-A\|_{F}^{2} \leq(1+\epsilon) \min _{\text {rank-k }} A_{k}\left\|A_{k}-A\right\|_{F}^{2}$, holds with probability at least $9 / 10$.

### 21.1.2 Our Techniques

Many of our proofs, in particular those for Theorem 21.1.1 and Theorem 21.1.2, are based on what we call an "iterative existential proof", which we then turn into an algorithm in two different ways depending if we are proving Theorem 21.1.1 or Theorem 21.1.2.

Henceforth, we assume $A_{k}$ exists; otherwise replace $A_{k}$ with a suitably good tensor $\widetilde{A}$ in what follows. Since $A_{k}=\sum_{i=1}^{k} U_{i}^{*} \otimes V_{i}^{*} \otimes W_{i}^{* 5}$, we can create three $n \times k$ matrices $U^{*}$, $V^{*}$, and $W^{*}$ whose columns are the vectors $U_{i}^{*}, V_{i}^{*}$, and $W_{i}^{*}$, respectively. Now we consider the three different flattenings (or unfoldings) of $A_{k}$, which express $A_{k}$ as an $n \times n^{2}$ matrix. Namely, by thinking of $A_{k}$ as the sum of outer products, we can write the three flattenings of $A_{k}$ as $U^{*} \cdot Z_{1}, V^{*} \cdot Z_{2}$, and $W^{*} \cdot Z_{3}$, where the rows of $Z_{1}$ are $\operatorname{vec}\left(V_{i}^{*} \otimes W_{i}^{*}\right)^{6}$ ( For simplicity, we write $\left.Z_{1}=\left(V^{* \top} \odot W^{* \top}\right) .^{7}\right)$, the rows of $Z_{2}$ are $\operatorname{vec}\left(U_{i}^{*} \otimes W_{i}^{*}\right)$, and the rows of $Z_{3}$ are $\operatorname{vec}\left(U_{i}^{*} \otimes V_{i}^{*}\right)$, for $i \in[k] \stackrel{\text { def }}{=}\{1,2, \ldots, k\}$. Letting the three corresponding flattenings of the input tensor $A$ be $A_{1}, A_{2}$, and $A_{3}$, by the symmetry of the Frobenius norm, we have $\|A-B\|_{F}^{2}=\left\|A_{1}-U^{*} Z_{1}\right\|_{F}^{2}=\left\|A_{2}-V^{*} Z_{2}\right\|_{F}^{2}=\left\|A_{3}-W^{*} Z_{3}\right\|_{F}^{2}$.

Let us consider the hypothetical regression problem $\min _{U}\left\|A_{1}-U Z_{1}\right\|_{F}^{2}$. Note that we do not know $Z_{1}$, but we will not need to. Let $r=O(k / \epsilon)$, and suppose $S_{1}$ is an $n^{2} \times r$ matrix of i.i.d. normal random variables with mean 0 and variance $1 / r$, denoted $N(0,1 / r)$. Then by standard results for regression (see, e.g., [Woo14b] for a survey), if $\widehat{U}$ is the minimizer to

[^70]the smaller regression problem $\widehat{U}=\operatorname{argmin}_{U}\left\|U Z_{1} S_{1}-A_{1} S_{1}\right\|_{F}^{2}$, then
\[

$$
\begin{equation*}
\left\|A_{1}-\widehat{U} Z_{1}\right\|_{F}^{2} \leq(1+\epsilon) \min _{U}\left\|A_{1}-U Z_{1}\right\|_{F}^{2} \tag{21.3}
\end{equation*}
$$

\]

Moreover, $\widehat{U}=A_{1} S_{1}\left(Z_{1} S_{1}\right)^{\dagger}$. Although we do not know know $Z_{1}$, this implies $\widehat{U}$ is in the column span of $A_{1} S_{1}$, which we do know, since we can flatten $A$ to compute $A_{1}$ and then compute $A_{1} S_{1}$. Thus, this hypothetical regression argument gives us an existential statement - there exists a good rank- $k$ matrix $\widehat{U}$ in the column span of $A_{1} S_{1}$. We could similarly define $\widehat{V}=A_{2} S_{2}\left(Z_{2} S_{2}\right)^{\dagger}$ and $\widehat{W}=A_{3} S_{3}\left(Z_{3} S_{3}\right)^{\dagger}$ as solutions to the analogous regression problems for the other two flattenings of $A$, which are in the column spans of $A_{2} S_{2}$ and $A_{3} S_{3}$, respectively. Given $A_{1} S_{1}, A_{2} S_{2}$, and $A_{3} S_{3}$, which we know, we could hope there is a good rank- $k$ tensor in the span of the rank- 1 tensors

$$
\begin{equation*}
\left\{\left(A_{1} S_{1}\right)_{a} \otimes\left(A_{2} S_{2}\right)_{b} \otimes\left(A_{3} S_{3}\right)_{c}\right\}_{a, b, c \in[r]} \tag{21.4}
\end{equation*}
$$

However, an immediate issue arises. First, note that our hypothetical regression problem guarantees that $\left\|A_{1}-\widehat{U} Z_{1}\right\|_{F}^{2} \leq(1+\epsilon)\left\|A-A_{k}\right\|_{F}^{2}$, and therefore since the rows of $Z_{1}$ are of the special form $\operatorname{vec}\left(V_{i}^{*} \otimes W_{i}^{*}\right)$, we can perform a "retensorization" to create a rank- $k$ tensor $B=\sum_{i} \widehat{U}_{i} \otimes V_{i}^{*} \otimes W_{i}^{*}$ from the matrix $\widehat{U} Z_{1}$ for which $\|A-B\|_{F}^{2} \leq(1+\epsilon)\left\|A-A_{k}\right\|_{F}^{2}$. While we do not know $\widehat{U}$, since it is in the column span of $A_{1} S_{1}$, it implies that $B$ is in the span of the rank-1 tensors $\left\{\left(A_{1} S_{1}\right)_{a} \otimes V_{b}^{*} \otimes W_{c}^{*}\right\}_{a \in[r], b, c \in[k]}$. Analogously, we have that there is a good rank- $k$ tensor $B$ in the span of the rank- 1 tensors $\left\{U_{a}^{*} \otimes\left(A_{2} S_{2}\right)_{b} \otimes W_{c}^{*}\right\}_{a, c \in[k], b \in[r]}$, and a good rank- $k$ tensor $B$ in the span of the rank-1 tensors $\left\{U_{a}^{*} \otimes V_{b}^{*} \otimes\left(A_{3} S_{3}\right)_{c}\right\}_{a, b \in[k], c \in[r]}$. However, we do not know $U^{*}$ or $V^{*}$, and it is not clear there is a rank- $k$ tensor $B$ for which simultaneously its first factors are in the column span of $A_{1} S_{1}$, its second factors are in the
column span of $A_{2} S_{2}$, and its third factors are in the column span of $A_{3} S_{3}$, i.e., whether there is a good rank- $k$ tensor B in the span of rank-1 tensors in (21.4).

We fix this by an iterative argument. Namely, we first compute $A_{1} S_{1}$, and write $\widehat{U}=A_{1} S_{1}\left(Z_{1} S_{1}\right)^{\dagger}$. We now redefine $Z_{2}$ with respect to $\widehat{U}$, so the rows of $Z_{2} \operatorname{are} \operatorname{vec}\left(\widehat{U}_{i} \otimes W_{i}^{*}\right)$ for $i \in[k]$, and consider the regression problem $\min _{V}\left\|A_{2}-V Z_{2}\right\|_{F}^{2}$. While we do not know $Z_{2}$, if $S_{2}$ is an $n^{2} \times r$ matrix of i.i.d. Gaussians, we again have the statement that $\widehat{V}=A_{2} S_{2}\left(Z_{2} S_{2}\right)^{\dagger}$ satisfies

$$
\begin{aligned}
\left\|A_{2}-\widehat{V} Z_{2}\right\|_{F}^{2} & \leq(1+\epsilon) \min _{V}\left\|A_{2}-V Z_{2}\right\|_{F}^{2} \text { by the regression guarantee with Gaussians } \\
& \leq(1+\epsilon)\left\|A_{2}-V^{*} Z_{2}\right\|_{F}^{2} \text { since } V^{*} \text { is no better than the minimizer } V \\
& =(1+\epsilon)\left\|A_{1}-\widehat{U} Z_{1}\right\|_{F}^{2} \text { by retensorizing and flattening along a different dimension } \\
& \leq(1+\epsilon)^{2} \min _{U}\left\|A_{1}-U Z_{1}\right\|_{F}^{2} \text { by }(21.3) \\
& =(1+\epsilon)^{2}\left\|A-A_{k}\right\|_{F}^{2} \text { by definition of } Z_{1}
\end{aligned}
$$

Now we can retensorize $\widehat{V} Z_{2}$ to obtain a rank- $k$ tensor $B$ for which $\|A-B\|_{F}^{2}=\| A_{2}-$ $\widehat{V} Z_{2}\left\|_{F}^{2} \leq(1+\epsilon)^{2}\right\| A-A_{k} \|_{F}^{2}$. Note that since the columns of $\widehat{V}$ are in the span of $A_{2} S_{2}$, and the rows of $Z_{2}$ are $\operatorname{vec}\left(\widehat{U}_{i} \otimes W_{i}^{*}\right)$ for $i \in[k]$, where the columns of $\widehat{U}$ are in the span of $A_{1} S_{1}$, it follows that $B$ is in the span of rank-1 tensors $\left\{\left(A_{1} S_{1}\right)_{a} \otimes\left(A_{2} S_{2}\right)_{b} \otimes \widehat{V}_{c}\right\}_{a, b \in[r], c \in[k]}$.

Suppose we now redefine $Z_{3}$ so that it is now an $r^{2} \times n^{2}$ matrix with rows vec $\left(\left(A_{1} S_{1}\right)_{a} \otimes\right.$ $\left.\left(A_{2} S_{2}\right)_{b}\right)$ for all pairs $a, b \in[r]$, and consider the regression problem $\min _{W}\left\|A_{3}-W Z_{3}\right\|_{F}^{2}$. Now observe that since we know $Z_{3}$, and since we can form $A_{3}$ by flattening $A$, we can solve for $W \in \mathbb{R}^{n \times r^{2}}$ in polynomial time by solving a regression problem. Retensorizing $W Z_{3}$ to a tensor $B$, it follows that we have found a rank- $r^{2}=O\left(k^{2} / \epsilon^{2}\right)$ tensor $B$ for which
$\|A-B\|_{F}^{2} \leq(1+\epsilon)^{2}\left\|A-A_{k}\right\|_{F}^{2}=(1+O(\epsilon))\left\|A-A_{k}\right\|_{F}^{2}$, and the result follows by adjusting $\epsilon$ by a constant factor.

To obtain the $n n z(A)+n \operatorname{poly}(k / \epsilon)$ running time guarantee of Theorem 21.1.1, while we can replace $S_{1}$ and $S_{2}$ with compositions of a sparse CountSketch matrix and a Gaussian matrix (see chapter 2 of [Woo14b] for a survey), enabling us to compute $A_{1} S_{1}$ and $A_{2} S_{2}$ in $\operatorname{nnz}(A)+n \operatorname{poly}(k / \epsilon)$ time, we still need to solve the regression problem $\min _{W}\left\|A_{3}-W Z_{3}\right\|_{F}^{2}$ quickly, and note that we cannot even write down $Z_{3}$ without spending $r^{2} n^{2}$ time. Here we use a different random matrix $S_{3}$ called TensorSketch, which was introduced in [Pag13, PP13], but for which we will need the stronger properties of a subspace embedding and approximate matrix product shown to hold for it in [ANW14]. Given the latter properties, we can instead solve the regression problem $\min _{W}\left\|A_{3} S_{3}-W Z_{3} S_{3}\right\|_{F}^{2}$, and importantly $A_{3} S_{3}$ and $Z_{3} S_{3}$ can be computed in $n n z(A)+n \operatorname{poly}(k / \epsilon)$ time. Finally, this small problem can be solved in $n \operatorname{poly}(k / \epsilon)$ time.

If we want to output a rank- $k$ solution as in Theorem 21.1.2, then we need to introduce indeterminates at several places in the preceding argument and run a generic polynomial optimization procedure which runs in time exponential in the number of indeterminates. Namely, we write $\widehat{U}$ as $A_{1} S_{1} X_{1}$, where $X_{1}$ is an $r \times k$ matrix of indeterminates, we write $\widehat{V}$ as $A_{2} S_{2} X_{2}$, where $X_{2}$ is an $r \times k$ matrix of indeterminates, and we write $\widehat{W}$ as $A_{3} S_{3} X_{3}$, where $X_{3}$ is an $r \times k$ matrix of indeterminates. When executing the above iterative argument, we let the rows of $Z_{1}$ be the vectors $\operatorname{vec}\left(V_{i}^{*} \otimes W_{i}^{*}\right)$, the rows of $Z_{2}$ be the vectors $\operatorname{vec}\left(\widehat{U}_{i} \otimes W_{i}^{*}\right)$, and the rows of $Z_{3}$ be the vectors $\operatorname{vec}\left(\widehat{U}_{i} \otimes V_{i}\right)$. Then $\widehat{U}$ is a $(1+\epsilon)$-approximate minimizer to $\min _{U}\left\|A_{1}-U Z_{1}\right\|_{F}$, while $\widehat{V}$ is a $(1+\epsilon)$-approximate minimizer to $\min _{V}\left\|A_{2}-V Z_{2}\right\|_{F}$, while $\widehat{W}$ is a $(1+\epsilon)$-approximate minimizer to $\min _{W}\left\|A_{3}-W Z_{3}\right\|_{F}$. Note that by assigning
$X_{1}=\left(Z_{1} S_{1}\right)^{\dagger}, X_{2}=\left(Z_{2} S_{2}\right)^{\dagger}$, and $X_{3}=\left(Z_{3} S_{3}\right)^{\dagger}$, it follows that the rank- $k$ tensor $B=$ $\sum_{i=1}^{k}\left(A_{1} S_{1} X_{1}\right)_{i} \otimes\left(A_{2} S_{2} X_{2}\right)_{i} \otimes\left(A_{3} S_{3} X_{3}\right)_{i}$ satisfies $\|A-B\|_{F}^{2} \leq(1+\epsilon)^{3}\left\|A-A_{k}\right\|_{F}^{2}$, as desired. Note that here the rows of $Z_{2}$ are a function of $X_{1}$, while the rows of $Z_{3}$ are a function of both $X_{1}$ and $X_{2}$. What is important for us though is that it suffices to minimize the degree- 6 polynomial $\sum_{a, b, c \in[n]}\left(\sum_{i=1}^{k}\left(A_{1} S_{1} X_{1}\right)_{a, i} \cdot\left(A_{2} S_{2} X_{2}\right)_{b, i} \cdot\left(A_{3} S_{3} X_{3}\right)_{c, i}-A_{a, b, c}\right)^{2}$, over the $3 r k=$ $O\left(k^{2} / \epsilon\right)$ indeterminates $X_{1}, X_{2}, X_{3}$, since we know there exists an assignment to $X_{1}, X_{2}$, and $X_{3}$ providing a $\left(1+O(\epsilon)\right.$ )-approximate solution, and any solution $X_{1}, X_{2}$, and $X_{3}$ found by minimizing the above polynomial will be no worse than that solution. This polynomial can be minimized up to additive $2^{-\operatorname{poly}(n)}$ additive error in poly $(n)$ time [Ren92a, BPR96] assuming the entries of $U^{*}, V^{*}$, and $W^{*}$ are bounded by $2^{\text {poly }(n)}$, as assumed in Theorem 21.1.2. Similar arguments can be made for obtaining a relative error approximation to the OPT as well as handling the case when $A_{k}$ does not exist.

To optimize the running time to $n n z(A)$, we can choose CountSketch matrices $T_{1}, T_{2}, T_{3}$ of $t=\operatorname{poly}(k, 1 / \epsilon) \times n$ dimensions and reapply the above iterative argument. Then it suffices to minimize this small size degree-6 polynomial $\sum_{a, b, c \in[t]}\left(\sum_{i=1}^{k}\left(T_{1} A_{1} S_{1} X_{1}\right)_{a, i} \cdot\left(T_{2} A_{2} S_{2} X_{2}\right)_{b, i}\right.$. $\left.\left(T_{3} A_{3} S_{3} X_{3}\right)_{c, i}-\left(A\left(T_{1}, T_{2}, T_{3}\right)\right)_{a, b, c}\right)^{2}$, over the $3 r k=O\left(k^{2} / \epsilon\right)$ indeterminates $X_{1}, X_{2}, X_{3}$. Outputting $A_{1} S_{1} X_{1}, A_{2} S_{2} X_{2}, A_{3} S_{3} X_{3}$ then provides a $(1+\epsilon)$-approximate solution.

Our iterative existential argument provides a general framework for obtaining low rank approximation results for tensors for many other error measures as well.

### 21.1.3 Other Low Rank Approximation Algorithms Following Our Framework.

 Column, row, tube subset selection, and CURT decomposition. In tensor column, row, tube subset selection, the goal is to find three matrices: a subset $C$ of columns of $A$, a subset $R$ of rows of $A$, and a subset $T$ of tubes of $A$, such that there exists a small tensor $U$ for which $\|U(C, R, T)-A\|_{F}^{2} \leq(1+\epsilon)$ OPT. We first choose two Gaussian matrices $S_{1}$ and $S_{2}$ with $s_{1}=s_{2}=O(k / \epsilon)$ columns, and form a matrix $Z_{3}^{\prime} \in \mathbb{R}^{\left(s_{1} s_{2}\right) \times n^{2}}$ with $(i, j)$-th row equal to the vectorization of $\left(A_{1} S_{1}\right)_{i} \otimes\left(A_{2} S_{2}\right)_{j}$. Motivated by the regression problem $\min _{W}\left\|A_{3}-W Z_{3}^{\prime}\right\|_{F}$, we sample $d_{3}=O\left(s_{1} s_{2} / \epsilon\right)$ columns from $A_{3}$ and let $D_{3}$ denote this selection matrix. There are a few ways to do the sampling depending on the tradeoff between the number of columns and running time, which we describe below. Proceeding iteratively, we write down $Z_{2}^{\prime}$ by setting its $(i, j)$-th row to the vectorization of $\left(A_{1} S_{1}\right)_{i} \otimes\left(A_{3} D_{3}\right)_{j}$. We then sample $d_{2}=O\left(s_{1} d_{3} / \epsilon\right)$ columns from $A_{2}$ and let $D_{2}$ denote that selection matrix. Finally, we define $Z_{1}^{\prime}$ by setting its $(i, j)$-th row to be the vectorization of $\left(A_{2} D_{2}\right)_{i} \otimes\left(A_{3} D_{3}\right)_{j}$. We obtain $C=A_{1} D_{1}, R=A_{2} D_{2}$ and $T=A_{3} D_{3}$. For the sampling steps, we can use a generalized matrix column subset selection technique, which extends a column subset selection technique of [BW14] in the context of CUR decompositions to the case when $C$ is not necessarily a subset of the input. This gives $O(\mathrm{nnz}(A) \log n)+\widetilde{O}\left(n^{2}\right) \operatorname{poly}(k, 1 / \epsilon)$ time. Alternatively, we can use a technique we develop called tensor leverage score sampling described below, yielding $O(\operatorname{nnz}(A))+n \operatorname{poly}(k, 1 / \epsilon)$ time.A body of work in the matrix case has focused on finding the best possible number of columns and rows of a CUR decomposition, and we can ask the same question for tensors. It turns out that if one is given the factorization $\sum_{i=1}^{k}\left(U_{B}\right)_{i} \otimes\left(V_{B}\right)_{i} \otimes\left(W_{B}\right)_{i}$ of a rank- $k$ tensor $B \in \mathbb{R}^{n \times n \times n}$ with $U_{B}, V_{B}, W_{B} \in \mathbb{R}^{n \times k}$, then one can find a set $C$ of $O(k / \epsilon)$ columns,
a set $R$ of $O(k / \epsilon)$ rows, and a set $T$ of $O(k / \epsilon)$ tubes of $A$, together with a rank- $k$ tensor $U$ for which $\|U(C, R, T)-A\|_{F}^{2} \leq(1+\epsilon)\|A-B\|_{F}^{2}$. This is based on an iterative argument, where the initial sampling (which needs to be our generalized matrix column subset selection rather than tensor leverage score sampling to achieve optimal bounds) is done with respect to $V_{B}^{\top} \odot W_{B}^{\top}$, and then an iterative argument is carried out. Since we show a matching lower bound on the number of columns, rows, tubes and rank of $U$, these parameters are tight. The algorithm is efficient if one is given a rank- $k$ tensor $B$ which is a $(1+O(\epsilon))$-approximation to $A$; if not then one can use Theorem 21.1.2 and and this step will be exponential time in $k$. If one just wants $O(k \log k+k / \epsilon)$ columns, rows, and tubes, then one can achieve $O(\mathrm{nnz}(A))+n \operatorname{poly}(k, 1 / \epsilon)$ time, if one is given $B$.

Column-row, row-tube, tube-column face subset selection, and CURT decomposition. In tensor column-row, row-tube, tube-column face subset selection, the goal is to find three tensors: a subset $C \in \mathbb{R}^{c \times n \times n}$ of row-tube faces of $A$, a subset $R \in \mathbb{R}^{n \times r \times n}$ of tubecolumn faces of $A$, and a subset $T \in \mathbb{R}^{n \times n \times t}$ of column-row faces of $A$, such that there exists a tensor $U \in \mathbb{R}^{t n \times c n \times r n}$ with small rank for which $\left\|U\left(T_{1}, C_{2}, R_{3}\right)-A\right\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}$, where $T_{1} \in \mathbb{R}^{n \times t n}$ denotes the matrix obtained by flattening the tensor $T$ along the first dimension, $C_{2} \in \mathbb{R}^{n \times c n}$ denotes the matrix obtained by flattening the tensor $C$ along the second dimension, and $R_{3} \in \mathbb{R}^{n \times r n}$ denotes the matrix obtained by flattening the tensor $T$ along the third dimension.

We solve this problem by first choosing two Gaussian matrices $S_{1}$ and $S_{2}$ with $s_{1}=$ $s_{2}=O(k / \epsilon)$ columns, and then forming matrix $U_{3} \in \mathbb{R}^{n \times s_{1} s_{2}}$ with $(i, j)$-th column equal to $\left(A_{1} S_{1}\right)_{i}$, as well as matrix $V_{3} \in \mathbb{R}^{n \times s_{1} s_{2}}$ with $(i, j)$-th column equal to $\left(A_{2} S_{2}\right)_{j}$. Inspired by
the regression problem $\min _{W \in \mathbb{R}^{n \times s_{1} s_{2}}}\left\|V_{3} \cdot\left(W^{\top} \odot U_{3}^{\top}\right)-A_{2}\right\|_{F}$, we sample $d_{3}=O\left(s_{1} s_{2} / \epsilon\right)$ rows from $A_{2}$ and let $D_{3} \in \mathbb{R}^{n \times n}$ denote this selection matrix. In other words, $D_{3}$ selects $d_{3}$ tube-column faces from the original tensor $A$. Thus, we obtain a small regression problem: $\min _{W}\left\|D_{3} V_{3} \cdot\left(W^{\top} \odot U_{3}^{\top}\right)-D_{3} A_{2}\right\|_{F}$. By retensorizing the objective function, we obtain the problem $\min _{W}\left\|U_{3} \otimes\left(D_{3} V_{3}\right) \otimes W-A\left(I, D_{3}, I\right)\right\|_{F}$. Flattening the objective function along the third dimension, we obtain $\min _{W}\left\|W \cdot\left(U_{3}^{\top} \odot\left(D_{3} V_{3}\right)^{\top}\right)-\left(A\left(I, D_{3}, I\right)\right)_{3}\right\|_{F}$ which has optimal solution $\left(A\left(I, D_{3}, I\right)\right)_{3}\left(U_{3}^{\top} \odot\left(D_{3} V_{3}\right)^{\top}\right)^{\dagger}$. Let $W^{\prime}$ denote $\left.A\left(I, D_{3}, I\right)\right)_{3}$. In the next step, we fix $W_{2}=W^{\prime}\left(U_{3}^{\top} \odot\left(D_{3} V_{3}\right)^{\top}\right)^{\dagger}$ and $U_{2}=U_{3}$, and consider the objective function $\min _{V}\left\|U_{2} \cdot\left(V^{\top} \odot W_{2}^{\top}\right)-A_{1}\right\|_{F}$. Applying a similar argument, we obtain $V^{\prime}=\left(A\left(D_{2}, I, I\right)\right)_{2}$ and $U^{\prime}=\left(A\left(I, I, D_{1}\right)_{1}\right)$. Let $C$ denote $A\left(D_{2}, I, I\right), R$ denote $A\left(I, D_{3}, I\right)$, and $T$ denote $A\left(I, I, D_{1}\right)$. Overall, this algorithm selects poly $(k, 1 / \epsilon)$ faces from each dimension.

Similar to our column-based CURT decomposition, our face-based CURT decomposition has the property that if one is given the factorization $\sum_{i=1}^{k}\left(U_{B}\right)_{i} \otimes\left(V_{B}\right)_{i} \otimes\left(W_{B}\right)_{i}$ of a rank- $k$ tensor $B \in \mathbb{R}^{n \times n \times n}$ with $U_{B}, V_{B}, W_{B} \in \mathbb{R}^{n \times k}$ which is a $(1+O(\epsilon))$-approximation to $A$, then one can find a set $C$ of $O(k / \epsilon)$ row-tube faces, a set $R$ of $O(k / \epsilon)$ tube-column faces, and a set $T$ of $O(k / \epsilon)$ column-row faces of $A$, together with a rank- $k$ tensor $U$ for which $\left\|U\left(T_{1}, C_{2}, R_{3}\right)-A\right\|_{F}^{2} \leq(1+\epsilon)$ OPT.

Tensor multiple regression and tensor leverage score sampling. In the above we need to consider standard problems for matrices in the context of tensors. Suppose we are given a matrix $A \in \mathbb{R}^{n_{1} \times n_{2} n_{3}}$ and a matrix $B=\left(V^{\top} \odot W^{\top}\right) \in \mathbb{R}^{k \times n_{2} n_{3}}$ with rows $\left(V_{i} \otimes W_{i}\right)$ for an $n_{2} \times k$ matrix $V$ and $n_{3} \times k$ matrix $W$. Using TensorSketch [Pag13, PP13, ANW14] one can solve multiple regression $\min _{U}\|U B-A\|_{F}$ without forming $B$ in $O\left(n_{2}+n_{3}\right) \operatorname{poly}(k, 1 / \epsilon)$
time, rather than the naïve $O\left(n_{2} n_{3}\right)$ poly $(k, 1 / \epsilon)$ time. However, this does not immediately help us if we would like to sample columns of such a matrix $B$ proportional to its leverage scores. Even if we apply TensorSketch to compute a $k \times k$ change of basis matrix $R$ in $O\left(n_{2}+n_{3}\right) \operatorname{poly}\left(k, \log \left(n_{2} n_{3}\right)\right)$ time, for which the leverage scores of $B$ are (up to a constant factor) the squared column norms of $R^{-1} B$, there are still $n_{2} n_{3}$ leverage scores and we cannot write them all down! Nevertheless, we show we can still sample by them by using that the matrix of interest is formed via a tensor product, which can be rewritten as a matrix multiplication which we never need to explicily materialize. In more detail, for the $i$-th row $e_{i} R^{-1}$ of $R^{-1}$, we create a matrix $V^{\prime i}$ by scaling each of the columns of $V^{\top}$ entrywise by the entries of $z$. The squared norms of $e_{i} R^{-1} B$ are exactly the squared entries of $\left(V^{\prime}\right) W^{\top}$. We cannot compute this matrix product, but we can first sample a column of it proportional to its squared norm and then sample an entry in that column proportional to its square. To sample a column, we compute $G\left(V^{\prime} i\right) W^{\top}$ for a Gaussian matrix $G$ with $O\left(\log n_{3}\right)$ rows by computing $G \cdot V^{\prime i}$, then computing $\left(G \cdot V^{\prime i}\right) \cdot W^{\top}$, which is $O\left(n_{2}+n_{3}\right) \operatorname{poly}\left(k, \log \left(n_{2} n_{3}\right)\right)$ total time. After sampling a column, we compute the column exactly and sample a squared entry. We do this for each $i \in[k]$, first sampling an $i$ proportional to $\left\|G V^{\prime} i W^{\top}\right\|_{F}^{2}$, then running the above scheme on that $i$. The poly $(\log n)$ factor in the running time can be replaced by $\operatorname{poly}(k)$ if one wants to avoid a poly $(\log n)$ dependence in the running time.

Entry-wise $\ell_{1}$ low-rank approximation. We consider the problem of entrywise $\ell_{1}$-low rank approximation of an $n \times n \times n$ tensor $A$, namely, the problem of finding a rank- $k$ tensor $B$ for which $\|A-B\|_{1} \leq \operatorname{poly}(k, \log n)$ OPT, where OPT $=\inf _{\text {rank-k } B}\|A-B\|_{1}$, and where for a tensor $A,\|A\|_{1}=\sum_{i, j, k}\left|A_{i, j, k}\right|$. Our iterative existential argument can be applied in
much the same way as for the Frobenius norm. We iteratively flatten $A$ along each of its three dimensions, obtaining $A_{1}, A_{2}$, and $A_{3}$ as above, and iteratively build a good rank- $k$ solution $B$ of the form $\left(A_{1} S_{1} X_{1}\right) \otimes\left(A_{2} S_{2} X_{2}\right) \otimes\left(A_{3} S_{3} X_{3}\right)$, where now the $S_{i}$ are matrices of i.i.d. Cauchy random variables or sparse matrices of Cauchy random variables and the $X_{i}$ are $O(k \log k) \times k$ matrices of indeterminates. For a matrix $C$ and a matrix $S$ of i.i.d. Cauchy random variables with $k$ columns, it is known [SWZ17] that the column span of $C S$ contains a poly $(k \log n)$-approximate rank- $k$ space with respect to the entrywise $\ell_{1}$-norm for $C$. In the case of tensors, we must perform an iterative flattening and retensorizing argument to guarantee there exists a tensor $B$ of the form above. Also, if we insist on outputting a rank- $k$ solution as opposed to a bicriteria solution, $\left\|\left(A_{1} S_{1} X_{1}\right) \otimes\left(A_{2} S_{2} X_{2}\right) \otimes\left(A_{3} S_{3} X_{3}\right)-A\right\|_{1}$ is not a polynomial of the $X_{i}$, and if we introduce sign variables for the $n^{3}$ absolute values, the running time of the polynomial solver will be $2^{\# \text { of variables }}=2^{\Omega\left(n^{3}\right)}$. We perform additional dimensionality reduction by Lewis weight sampling [CP15] from the flattenings to reduce the problem size to $\operatorname{poly}(k)$. This small problem still has $\widetilde{O}\left(k^{3}\right) \operatorname{sign}$ variables, and to obtain a $2^{\widetilde{O}\left(k^{2}\right)}$ running time we relax the reduced problem to a Frobenius norm problem, mildly increasing the approximation factor by another $\operatorname{poly}(k)$ factor.

Combining the iterative existential argument with techniques in [SWZ17], we also obtain an $\ell_{1}$ CURT decomposition algorithm (which is similar to the Frobenius norm result in Theorem 21.1.3), which can find $\widetilde{O}(k)$ columns, $\widetilde{O}(k)$ rows, $\widetilde{O}(k)$ tubes, and a tensor $U$. Our algorithm starts from a given factorization of a rank- $k$ tensor $B=U_{B} \otimes V_{B} \otimes W_{B}$ found above. We compute a sampling and rescaling diagonal matrix $D_{1}$ according to the Lewis weights of matrix $B_{1}=\left(V_{B}^{\top} \odot W_{B}^{\top}\right)$, where $D_{1}$ has $\widetilde{O}(k)$ nonzero entries. Then we iteratively construct $B_{2}, D_{2}, B_{3}$ and $D_{3}$. Finally we have $C=A_{1} D_{1}$ (selecting $\widetilde{O}(k)$ columns from

```
Algorithm 21.1 Main Meta-Algorithm
    procedure TensorLowRankApproxBicriteria \((A, n, k, \epsilon) \quad \triangleright\) Theorem 21.1.1
            Choose sketching matrices \(S_{2}, S_{3}\) (Composition of Gaussian and CountSketch.)
            Choose sketching matrices \(T_{2}, T_{3}\) (CountSketch.)
            Compute \(T_{2} A_{2} S_{2}, T_{3} A_{3} S_{3}\).
            Construct \(\widehat{V}\) by setting \((i, j)\)-th column to be \(\left(A_{2} S_{2}\right)_{i}\).
            Construct \(\widehat{W}\) by setting \((i, j)\)-th column to be \(\left(A_{3} S_{3}\right)_{j}\).
            Construct matrix \(B\) by setting \((i, j)\)-th row of \(B\) is vectorization of \(\left(T_{2} A_{2} S_{2}\right)_{i} \otimes\)
    \(\left(T_{3} A_{3} S_{3}\right)_{j}\).
            Solve \(\min _{U}\left\|U B-\left(A\left(I, T_{2}, T_{3}\right)\right)_{1}\right\|_{F}^{2}\).
            return \(\widehat{U}, \widehat{V}\), and \(\widehat{W}\).
    end procedure
    procedure TensorLowRankApprox \((A, n, k, \epsilon) \quad \triangleright\) Theorem 21.1.2
            Choose sketching matrices \(S_{1}, S_{2}, S_{3}\) (Composition of Gaussian and CountSketch.)
            Choose sketching matrices \(T_{1}, T_{2}, T_{3}\) (CountSketch.)
            Compute \(T_{1} A_{1} S_{1}, T_{2} A_{2} S_{2}, T_{3} A_{3} S_{3}\).
            Solve \(\min _{X_{1}, X_{2}, X_{3}}\left\|\left(T_{1} A_{1} S_{1} X_{1}\right) \otimes\left(T_{2} A_{2} S_{2} X_{2}\right) \otimes\left(T_{3} A_{3} S_{3} X_{3}\right)-A\left(T_{1}, T_{2}, T_{3}\right)\right\|_{F}^{2}\).
            return \(A_{1} S_{1} X_{1}, A_{2} S_{2} X_{2}\), and \(A_{3} S_{3} X_{3}\).
    end procedure
```

$A), R=A_{2} D_{2}$ (selecting $\widetilde{O}(k)$ rows from $A$ ), $T=A_{3} D_{3}$ (selecting $\widetilde{O}(k)$ tubes from $A$ ) and tensor $U=\left(\left(B_{1} D_{1}\right)^{\dagger}\right) \otimes\left(\left(B_{2} D_{2}\right)^{\dagger}\right) \otimes\left(\left(B_{3} D_{3}\right)^{\dagger}\right)$.

Optimal matrix CUR decomposition. We also improve the $n n z(A) \log n+(n+d) \operatorname{poly}(\log n, k$, $1 / \epsilon)$ running time of [BW14] for CUR decomposition of $A \in \mathbb{R}^{n \times d}$ to $n n z(A)+(n+$ d) $\operatorname{poly}(k, 1 / \epsilon)$, while selecting the optimal number of columns, rows, and a rank- $k$ matrix $U$. Using [CW13, MM13, NN13a], we find a matrix $\widehat{U}$ with $k$ orthonormal columns in $\mathrm{nnz}(A)+n \operatorname{poly}(k / \epsilon)$ time for which $\min _{V}\|\widehat{U} V-A\|_{F}^{2} \leq(1+\epsilon)\left\|A-A_{k}\right\|_{F}^{2}$. Let $s_{1}=\widetilde{O}\left(k / \epsilon^{2}\right)$ and $S_{1} \in \mathbb{R}^{s_{1} \times n}$ be a sampling/rescaling matrix by the leverage scores of $\widehat{U}$. By strengthening the affine embedding analysis of [CW13] to leverage score sampling (the analysis of
[CW13] gives a weaker analysis for affine embeddings using leverage scores which does not allow approximation in the sketch space to translate to approximation in the original space), with probability at least 0.99 , for all $X^{\prime}$ which satisfy $\left\|S_{1} \widehat{U} X^{\prime}-S_{1} A\right\|_{F}^{2} \leq$ $\left(1+\epsilon^{\prime}\right) \min _{X}\left\|S_{1} \widehat{U} X-S_{1} A\right\|_{F}^{2}$, we have $\left\|\widehat{U} X^{\prime}-A\right\|_{F}^{2} \leq(1+\epsilon) \min _{X}\|\widehat{U} X-A\|_{F}^{2}$, where $\epsilon^{\prime}=0.0001 \epsilon$. Applying our generalized row subset selection procedure, we can find $Y, R$ for which $\left\|S_{1} \widehat{U} Y R-S_{1} A\right\|_{F}^{2} \leq\left(1+\epsilon^{\prime}\right) \min _{X}\left\|S_{1} \widehat{U} X-S_{1} A\right\|_{F}^{2}$, where $R$ contains $O\left(k / \epsilon^{\prime}\right)=O(k / \epsilon)$ rescaled rows of $S_{1} A$. A key point is that rescaled rows of $S_{1} A$ are also rescaled rows of $A$. Then, $\|\widehat{U} Y R-A\|_{F}^{2} \leq(1+\epsilon) \min _{X}\|\widehat{U} X-A\|_{F}^{2}$. Finding $Y, R$ can be done in $d \operatorname{poly}\left(s_{1} / \epsilon\right)=$ $d \operatorname{poly}(k / \epsilon)$ time. Now set $\widehat{V}=Y R$. We can choose $S_{2}$ to be a sampling/rescaling matrix, and then find $C, Z$ for which $\left\|C Z \widehat{V} S_{2}-A S_{2}\right\|_{F}^{2} \leq\left(1+\epsilon^{\prime}\right) \min _{X}\left\|X \widehat{V} S_{2}-A S_{2}\right\|_{F}^{2}$ in a similar way, where $C$ contains $O(k / \epsilon)$ rescaled columns of $A S_{2}$, and thus also of $A$. We thus have $\|C Z Y R-A\|_{F}^{2} \leq(1+O(\epsilon))\left\|A-A_{k}\right\|_{F}^{2}$.

### 21.1.4 Comparison to [BCV14]

[BCV14] shows for a third order $n_{1} \times n_{2} \times n_{3}$ tensor $A$ how to find a rank- $k$ tensor $B$ for which $\|A-B\|_{F}^{2} \leq 5$ OPT in $\operatorname{poly}\left(n_{1} n_{2} n_{3}\right) \exp (\operatorname{poly}(k))$ time. They generalize this to $q$-th order tensors to find a rank- $k$ tensor $B$ for which $\|A-B\|_{F}^{2}=O(q)$ OPT in $\operatorname{poly}\left(n_{1} n_{2} \cdots n_{q}\right) \exp (\operatorname{poly}(q k))$ time.

In contrast, we obtain a rank- $k$ tensor $B$ for which $\|A-B\|_{F}^{2} \leq(1+\epsilon)$ OPT in $\mathrm{nnz}(A)+n \cdot \operatorname{poly}(k / \epsilon)+\exp \left(\left(k^{2} / \epsilon\right) \operatorname{poly}(q)\right)$ time for every order $q$. Thus, we obtain a $(1+\epsilon)$ instead of an $O(q)$ approximation. The $O(q)$ approximation in [BCV14] seems inherent since the authors apply triangle inequality $q$ times, each time losing a constant factor. This seems necessary since their argument is based on the span of the top $k$ principal components in the SVD in each flattening separately containing a good space to project onto for a given mode. In contrast, our iterative existential argument chooses the space to project onto in successive modes adaptively as a function of spaces chosen for previous modes, and thus we obtain a $(1+\epsilon)^{O(q)}=(1+O(\epsilon q))$-approximation, which becomes a $(1+\epsilon)$-approximation after replacing $\epsilon$ with $\epsilon / q$. Also, importantly, our algorithm runs in $n n z(A)+n \cdot \operatorname{poly}(k / \epsilon)+$ $\exp \left(\left(k^{2} / \epsilon\right)\right.$ poly $\left.(q)\right)$ time and there are multiple hurdles we overcome to achieve this, as described in Section 21.1.2 above.

### 21.1.5 A Roadmap

Roadmap Section 21.2 introduces notation and definitions. Section 21.3 includes several tools. Our Frobenius norm low rank approximation algorithms are in Section 21.4. Section 21.5 extends our results to general $q$-th order tensors.


Figure 21.1: A 3 rd order tensor with size $8 \times 8 \times 8$.

### 21.2 Notation

For an $n \in \mathbb{N}_{+}$, let $[n]$ denote the set $\{1,2, \cdots, n\}$.
For any function $f$, we define $\widetilde{O}(f)$ to be $f \cdot \log ^{O(1)}(f)$. In addition to $O(\cdot)$ notation, for two functions $f, g$, we use the shorthand $f \lesssim g$ (resp. $\gtrsim$ ) to indicate that $f \leq C g$ (resp. $\geq)$ for an absolute constant $C$. We use $f \approx g$ to mean $c f \leq g \leq C f$ for constants $c, C$.

For a matrix $A$, we use $\|A\|_{2}$ to denote the spectral norm of $A$. For a tensor $A$, let $\|A\|$ and $\|A\|_{2}$ (which we sometimes use interchangeably) denote the spectral norm of tensor A,

$$
\|A\|=\sup _{x, y, z \neq 0} \frac{|A(x, y, z)|}{\|x\| \cdot\|y\| \cdot\|z\|}
$$

Let $\|A\|_{F}$ denote the Frobenius norm of a matrix/tensor $A$, i.e., $\|A\|_{F}$ is the square root of
sum of squares of all the entries of $A$. For $1 \leq p<2$, we use $\|A\|_{p}$ to denote the entry-wise $\ell_{p}$-norm of a matrix/tensor $A$, i.e., $\|A\|_{p}$ is the $p$-th root of the sum of $p$-th powers of the absolute values of the entries of $A .\|A\|_{1}$ will be an important special case of $\|A\|_{p}$, which corresponds to the sum of absolute values of all of the entries.

Let $n n z(A)$ denote the number of nonzero entries of $A$. Let $\operatorname{det}(A)$ denote the determinant of a square matrix $A$. Let $A^{\top}$ denote the transpose of $A$. Let $A^{\dagger}$ denote the Moore-Penrose pseudoinverse of $A$. Let $A^{-1}$ denote the inverse of a full rank square matrix.

For a 3rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, its $x$-mode fibers are called column fibers $(x=1)$, row fibers $(x=2)$ and tube fibers $(x=3)$. For tensor $A$, we use $A_{*, j, l}$ to denote its $(j, l)$-th column, we use $A_{i, *, l}$ to denote its $(i, l)$-th row, and we use $A_{i, j, *}$ to denote its $(i, j)$-th tube.

A tensor $A$ is symmetric if and only if for any $i, j, k, A_{i, j, k}=A_{i, k, j}=A_{j, i, k}=A_{j, k, i}=$ $A_{k, i, j}=A_{k, j, i}$.

For a tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$, we use $T$ to denote rotation (3 dimensional transpose) so that $A^{\top} \in \mathbb{R}^{n_{3} \times n_{1} \times n_{2}}$. For a tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$ and matrix $B \in \mathbb{R}^{n_{3} \times k}$, we define the tensor-matrix dot product to be $A \cdot B \in \mathbb{R}^{n_{1} \times n_{2} \times k}$.

We use $\otimes$ to denote outer product, $\circ$ to denote entrywise product, and $\cdot$ to denote dot product. Given two column vectors $u, v \in \mathbb{R}^{n}$, let $u \otimes v \in \mathbb{R}^{n \times n}$ and $(u \otimes v)_{i, j}=u_{i} \cdot v_{j}$, $u^{\top} v=\sum_{i=1}^{n} u_{i} v_{i} \in \mathbb{R}$ and $(u \circ v)_{i}=u_{i} v_{i}$.

Definition 21.2.1 ( $\otimes$ product for vectors). Given $q$ vectors $u_{1} \in \mathbb{R}^{n_{1}}, u_{2} \in \mathbb{R}^{n_{2}}, \cdots$, $u_{q} \in \mathbb{R}^{n_{q}}$, we use $u_{1} \otimes u_{2} \otimes \cdots \otimes u_{q}$ to denote an $n_{1} \times n_{2} \times \cdots \times n_{q}$ tensor such that, for each $\left(j_{1}, j_{2}, \cdots, j_{q}\right) \in\left[n_{1}\right] \times\left[n_{2}\right] \times \cdots \times\left[n_{q}\right]$,

$$
\left(u_{1} \otimes u_{2} \otimes \cdots \otimes u_{q}\right)_{j_{1}, j_{2}, \cdots, j_{q}}=\left(u_{1}\right)_{j_{1}}\left(u_{2}\right)_{j_{2}} \cdots\left(u_{q}\right)_{j_{q}},
$$



Figure 21.2: Flattening. We flatten a third order $4 \times 4 \times 4$ tensor along the 1st dimension to obtain a $4 \times 16$ matrix. The red blocks correspond to a column in the original third order tensor, the blue blocks correspond to a row in the original third order tensor, and the green blocks correspond to a tube in the original third order tensor.
where $\left(u_{i}\right)_{j_{i}}$ denotes the $j_{i}$-th entry of vector $u_{i}$.
Definition 21.2.2 ( $\operatorname{vec}()$, convert tensor into a vector). Given a tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times \cdots \times n_{q}}$, let $\operatorname{vec}(A) \in \mathbb{R}^{1 \times \prod_{i=1}^{q} n_{i}}$ be a row vector, such that the $t$-th entry of $\operatorname{vec}(A)$ is $A_{j_{1}, j_{2}, \cdots, j_{q}}$ where $t=\left(j_{1}-1\right) \prod_{i=2}^{q} n_{i}+\left(j_{2}-1\right) \prod_{i=3}^{q} n_{i}+\cdots+\left(j_{q-1}-1\right) n_{q}+j_{q}$.

For example if $u=\left[\begin{array}{l}1 \\ 2\end{array}\right], v=\left[\begin{array}{l}3 \\ 4 \\ 5\end{array}\right]$ then $\operatorname{vec}(u \otimes v)=\left[\begin{array}{llllll}3 & 4 & 5 & 6 & 8 & 10\end{array}\right]$.
Definition 21.2.3 ( $\otimes$ product for matrices). Given $q$ matrices $U_{1} \in \mathbb{R}^{n_{1} \times k}, U_{2} \in \mathbb{R}^{n_{2} \times k}$, $\cdots, U_{q} \in \mathbb{R}^{n_{q} \times k}$, we use $U_{1} \otimes U_{2} \otimes \cdots \otimes U_{q}$ to denote an $n_{1} \times n_{2} \times \cdots \times n_{q}$ tensor which can be written as,

$$
U_{1} \otimes U_{2} \otimes \cdots \otimes U_{q}=\sum_{i=1}^{k}\left(U_{1}\right)_{i} \otimes\left(U_{2}\right)_{i} \otimes \cdots \otimes\left(U_{q}\right)_{i} \in \mathbb{R}^{n_{1} \times n_{2} \times \cdots \times n_{q}}
$$

where $\left(U_{j}\right)_{i}$ denotes the $i$-th column of matrix $U_{j} \in \mathbb{R}^{n_{j} \times k}$.
Definition 21.2.4 ( $\odot$ product for matrices). Given $q$ matrices $U_{1} \in \mathbb{R}^{k \times n_{1}}, U_{2} \in \mathbb{R}^{k \times n_{2}}$, $\cdots, U_{q} \in \mathbb{R}^{k \times n_{q}}$, we use $U_{1} \odot U_{2} \odot \cdots \odot U_{q}$ to denote a $k \times \prod_{j=1}^{q} n_{j}$ matrix where the $i$-th
row of $U_{1} \odot U_{2} \odot \cdots \odot U_{q}$ is the vectorization of $\left(U_{1}\right)^{i} \otimes\left(U_{2}\right)^{i} \otimes \cdots \otimes\left(U_{q}\right)^{i}$, i.e.,

$$
U_{1} \odot U_{2} \odot \cdots \odot U_{q}=\left[\begin{array}{c}
\operatorname{vec}\left(\left(U_{1}\right)^{1} \otimes\left(U_{2}\right)^{1} \otimes \cdots \otimes\left(U_{q}\right)^{1}\right) \\
\operatorname{vec}\left(\left(U_{1}\right)^{2} \otimes\left(U_{2}\right)^{2} \otimes \cdots \otimes\left(U_{q}\right)^{2}\right) \\
\cdots \\
\operatorname{vec}\left(\left(U_{1}\right)^{k} \otimes\left(U_{2}\right)^{k} \otimes \cdots \otimes\left(U_{q}\right)^{k}\right)
\end{array}\right] \in \mathbb{R}^{k \times \prod_{j=1}^{q} n_{j}} .
$$

where $\left(U_{j}\right)^{i} \in \mathbb{R}^{n_{j}}$ denotes the $i$-th row of matrix $U_{j} \in \mathbb{R}^{k \times n_{j}}$.

Definition 21.2.5 (Flattening vs unflattening/retensorizing). Suppose we are given three matrices $U \in \mathbb{R}^{n_{1} \times k}, V \in \mathbb{R}_{n_{2} \times k}, W \in \mathbb{R}^{n_{3} \times k}$. Let tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$ denote $U \otimes V \otimes W$. Let $A_{1} \in \mathbb{R}^{n_{1} \times n_{2} n_{3}}$ denote a matrix obtained by flattening tensor $A$ along the 1 st dimension. Then $A_{1}=U \cdot B$, where $B=V^{\top} \odot W^{\top} \in \mathbb{R}^{k \times n_{2} n_{3}}$ denotes the matrix for which the $i$-th row is $\operatorname{vec}\left(V_{i} \otimes W_{i}\right), \forall i \in[k]$. We let the "flattening" be the operation that obtains $A_{1}$ by $A$. Given $A_{1}=U \cdot B$, we can obtain tensor $A$ by unflattening/retensorizing $A_{1}$. We let "retensorization" be the operation that obtains $A$ from $A_{1}$. Similarly, let $A_{2} \in \mathbb{R}^{n_{2} \times n_{1} n_{3}}$ denote a matrix obtained by flattening tensor $A$ along the 2nd dimension, so $A_{2}=V \cdot C$, where $C=W^{\top} \odot U^{\top} \in \mathbb{R}^{k \times n_{1} n_{3}}$ denotes the matrix for which the $i$-th row is $\operatorname{vec}\left(W_{i} \otimes U_{i}\right), \forall i \in$ $[k]$. Let $A_{3} \in \mathbb{R}^{n_{3} \times n_{1} n_{2}}$ denote a matrix obtained by flattening tensor $A$ along the 3 rd dimension. Then, $A_{3}=W \cdot D$, where $D=U^{\top} \odot V^{\top} \in \mathbb{R}^{k \times n_{1} n_{2}}$ denotes the matrix for which the $i$-th row is $\operatorname{vec}\left(U_{i} \otimes V_{i}\right), \forall i \in[k]$.

Definition 21.2.6 ( $(\cdot, \cdot, \cdot)$ operator for tensors and matrices). Given tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$ and three matrices $B_{1} \in \mathbb{R}^{n_{1} \times d_{1}}, B_{2} \in \mathbb{R}^{n_{2} \times d_{2}}, B_{3} \in \mathbb{R}^{n_{3} \times d_{3}}$, we define tensors $A\left(B_{1}, I, I\right) \in$ $\mathbb{R}^{d_{1} \times n_{2} \times n_{3}}, A\left(I, B_{2}, I\right) \in \mathbb{R}^{n_{1} \times d_{2} \times n_{3}}, A\left(I, I, B_{3}\right) \in \mathbb{R}^{n_{1} \times n_{2} \times d_{3}}, A\left(B_{1}, B_{2}, I\right) \in \mathbb{R}^{d_{1} \times d_{2} \times n_{3}}$,
$A\left(B_{1}, B_{2}, B_{3}\right) \in \mathbb{R}^{d_{1} \times d_{2} \times d_{3}}$ as follows,

$$
\begin{array}{rlrl}
A\left(B_{1}, I, I\right)_{i, j, l} & =\sum_{i^{\prime}=1}^{n_{1}} A_{i^{\prime}, j, l}\left(B_{1}\right)_{i^{\prime}, i}, & \forall(i, j, l) \in\left[d_{1}\right] \times\left[n_{2}\right] \times\left[n_{3}\right] \\
A\left(I, B_{2}, I\right)_{i, j, l} & =\sum_{j^{\prime}=1}^{n_{2}} A_{i, j^{\prime}, l}\left(B_{2}\right)_{j^{\prime}, j}, & \forall(i, j, l) \in\left[n_{1}\right] \times\left[d_{2}\right] \times\left[n_{3}\right] \\
A\left(I, I, B_{3}\right)_{i, j, l} & =\sum_{l^{\prime}=1}^{n_{3}} A_{i, j, l^{\prime}}\left(B_{3}\right)_{l^{\prime}, l}, & & \forall(i, j, l) \in\left[n_{1}\right] \times\left[n_{2}\right] \times\left[d_{3}\right] \\
A\left(B_{1}, B_{2}, I\right)_{i, j, l} & =\sum_{i^{\prime}=1}^{n_{1}} \sum_{j^{\prime}=1}^{n_{2}} A_{i^{\prime}, j^{\prime}, l}\left(B_{1}\right)_{i^{\prime}, i}\left(B_{2}\right)_{j^{\prime}, j}, & \forall(i, j, l) \in\left[d_{1}\right] \times\left[d_{2}\right] \times\left[n_{3}\right] \\
A\left(B_{1}, B_{2}, B_{3}\right)_{i, j, l} & =\sum_{i^{\prime}=1}^{n_{1}} \sum_{j^{\prime}=1}^{n_{2}} \sum_{l^{\prime}=1}^{n_{3}} A_{i^{\prime}, j^{\prime}, l^{\prime}}\left(B_{1}\right)_{i^{\prime}, i}\left(B_{2}\right)_{j^{\prime}, j}\left(B_{3}\right)_{l^{\prime}, l}, & \forall(i, j, l) \in\left[d_{1}\right] \times\left[d_{2}\right] \times\left[d_{3}\right]
\end{array}
$$

Note that $B_{1}^{\top} A=A\left(B_{1}, I, I\right), A B_{3}=A\left(I, I, B_{3}\right)$ and $B_{1}^{\top} A B_{3}=A\left(B_{1}, I, B_{3}\right)$. In our paper, if $\forall i \in[3], B_{i}$ is either a rectangular matrix or a symmetric matrix, then we sometimes use $A\left(B_{1}, B_{2}, B_{3}\right)$ to denote $A\left(B_{1}^{\top}, B_{2}^{\top}, B_{3}^{\top}\right)$ for simplicity. Similar to the $(\cdot, \cdot, \cdot)$ operator on 3rd order tensors, we can define the $(\cdot, \cdot, \cdots, \cdot)$ operator on higher order tensors.

For the matrix case, $\min _{\operatorname{rank}-k}\left\|A-A^{\prime}\right\|_{F}^{2}$ always exists. However, this is not true for tensors [DSL08]. For convenience, we redefine the notation of OPT and min.

Definition 21.2.7. Given tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}, k>0$, if $\min _{\operatorname{rank}-k}\left\|A-A^{\prime}\right\|_{F}^{2}$ does not exist, then we define $\mathrm{OPT}=\inf _{\operatorname{rank}-k}{ }_{A^{\prime}}\left\|A-A^{\prime}\right\|_{F}^{2}+\gamma$ for sufficiently small $\gamma>0$, which can be an arbitrarily small positive function of $n$. We let $\min _{\operatorname{rank}-k}{ }_{A^{\prime}}\left\|A-A^{\prime}\right\|_{F}^{2}$ be the value of OPT, and we let $\underset{\text { rank }-k A^{\prime}}{\arg \min }\left\|A-A^{\prime}\right\|_{F}^{2}$ be a rank $-k$ tensor $A_{k} \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$ which satisfies $\left\|A-A_{k}\right\|_{F}^{2}=$ OPT.


Tensor


Column


Row


Tube

Figure 21.3: A 3 rd order tensor contains $n^{2}$ columns, $n^{2}$ rows, and $n^{2}$ tubes.

### 21.3 Preliminaries

Section 21.3 .1 provides the definitions for Subspace Embeddings and Approximate Matrix Product. We introduce the definition for Tensor-CURT decomposition in Section 21.3.2. Section 21.3.3 presents a tool which we call a "polynomial system verifier". Section 21.3.4 introduces a tool which is able to determine the minimum nonzero value of the absolute value of a polynomial evaluated on a set, provided the polynomial is never equal to 0 on that set. Section 21.3.5 shows how to relax an $\ell_{p}$ problem to an $\ell_{2}$ problem. We provide definitions for CountSketch and Gaussian transforms in Section 21.3.6. We present Cauchy and $p$-stable transforms in Section 21.3.7. We introduce leverage scores and Lewis weights in Section 21.3.8 and Section 21.3.9. Finally, we explain an extension of CountSketch, which is called TensorSketch in Section 21.3.10.

### 21.3.1 Subspace Embeddings and Approximate Matrix Product

Definition 21.3.1 (Subspace Embedding). A ( $1 \pm \epsilon$ ) $\ell_{2}$-subspace embedding for the column space of an $n \times d$ matrix $A$ is a matrix $S$ for which for all $x \in \mathbb{R}^{d},\|S A x\|_{2}^{2}=(1 \pm \epsilon)\|A x\|_{2}^{2}$.

Definition 21.3.2 (Approximate Matrix Product). Let $0<\epsilon<1$ be a given approximation parameter. Given matrices $A$ and $B$, where $A$ and $B$ each have $n$ rows, the goal is to output a matrix $C$ so that $\left\|A^{\top} B-C\right\|_{F} \leq \epsilon\|A\|_{F}\|B\|_{F}$. Typically $C$ has the form $A^{\top} S^{\top} S B$, for a random matrix $S$ with a small number of rows. See, e.g., Lemma 32 of [CW13] for a number of example matrices $S$ with $O\left(\epsilon^{-2}\right)$ rows for which this property holds.


Figure 21.4: A third order tensor has three types of faces: the column-row faces, the columntube faces, and the row-tube faces

### 21.3.2 Tensor CURT decomposition

We first review matrix CUR decompositions:
Definition 21.3.3 (Matrix CUR, exact). Given a matrix $A \in \mathbb{R}^{n \times d}$, we choose $C \in \mathbb{R}^{n \times c}$ to be a subset of columns of $A$ and $R \in \mathbb{R}^{r \times n}$ to be a subset of rows of $A$. If there exists a matrix $U \in \mathbb{R}^{c \times r}$ such that $A$ can be written as,

$$
C U R=A
$$

then we say $C, U, R$ is matrix $A$ 's CUR decomposition.
Definition 21.3.4 (Matrix CUR, approximate). Given a matrix $A \in \mathbb{R}^{n \times d}$, a parameter $k \geq 1$, an approximation ratio $\alpha>1$, and a norm $\left\|\|_{\xi}\right.$, we choose $C \in \mathbb{R}^{n \times c}$ to be a subset of columns of $A$ and $R \in \mathbb{R}^{r \times n}$ to be a subset of rows of $A$. Then if there exists a matrix $U \in \mathbb{R}^{c \times r}$ such that,

$$
\|C U R-A\|_{\xi} \leq \alpha \min _{\text {rank }-k}\left\|A_{k}-A\right\|_{\xi}
$$

where $\left\|\|_{\xi}\right.$ can be operator norm, Frobenius norm or Entry-wise $\ell_{1}$ norm, we say that $C, U, R$ is matrix $A$ 's approximate CUR decomposition, and sometimes just refer to this as a CUR decomposition.

Definition 21.3.5 ([Bou11]). Given matrix $A \in \mathbb{R}^{m \times n}$, integer $k$, and matrix $C \in \mathbb{R}^{m \times r}$ with $r>k$, we define the matrix $\Pi_{C, k}^{\xi}(A) \in \mathbb{R}^{m \times n}$ to be the best approximation to $A$ (under the $\xi$-norm) within the column space of $C$ of rank at most $k$; so, $\Pi_{C, k}^{\xi}(A) \in \mathbb{R}^{m \times n}$ minimizes the residual $\|A-\widehat{A}\|_{\xi}$, over all $\widehat{A} \in \mathbb{R}^{m \times n}$ in the column space of $C$ of rank at most $k$.

We define the following notion of tensor-CURT decomposition.

Definition 21.3.6 (Tensor CURT, exact). Given a tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$, we choose three sets of pair of coordinates $S_{1} \subseteq\left[n_{2}\right] \times\left[n_{3}\right], S_{2} \subseteq\left[n_{1}\right] \times\left[n_{3}\right], S_{3} \subseteq\left[n_{1}\right] \times\left[n_{2}\right]$. We define $c=\left|S_{1}\right|, r=\left|S_{2}\right|$ and $t=\left|S_{3}\right|$. Let $C \in \mathbb{R}^{n_{1} \times c}$ denote a subset of columns of $A, R \in \mathbb{R}^{n_{2} \times r}$ denote a subset of rows of $A$, and $T \in \mathbb{R}^{n_{3} \times t}$ denote a subset of tubes of $A$. If there exists a tensor $U \in \mathbb{R}^{c \times r \times t}$ such that $A$ can be written as

$$
\left(\left(\left(U \cdot T^{\top}\right)^{\top} \cdot R^{\top}\right)^{\top} \cdot C^{\top}\right)^{\top}=A
$$

or equivalently,

$$
U(C, R, T)=A
$$

or equivalently,

$$
\forall(i, j, l) \in\left[n_{1}\right] \times\left[n_{2}\right] \times\left[n_{3}\right], A_{i, j, l}=\sum_{u_{1}=1}^{c} \sum_{u_{2}=1}^{r} \sum_{u_{3}=1}^{t} U_{u_{1}, u_{2}, u_{3}} C_{i, u_{1}} R_{j, u_{2}} T_{l, u_{3}},
$$

then we say $C, U, R, T$ is tensor $A$ 's CURT decomposition.
Definition 21.3.7 (Tensor CURT, approximate). Given a tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$, for some $k \geq 1$, for some approximation $\alpha>1$, for some norm $\left\|\|_{\xi}\right.$, we choose three sets of pair of coordinates $S_{1} \subseteq\left[n_{2}\right] \times\left[n_{3}\right], S_{2} \subseteq\left[n_{1}\right] \times\left[n_{3}\right], S_{3} \subseteq\left[n_{1}\right] \times\left[n_{2}\right]$. We define $c=\left|S_{1}\right|, r=\left|S_{2}\right|$ and $t=\left|S_{3}\right|$. Let $C \in \mathbb{R}^{n_{1} \times c}$ denote a subset of columns of $A, R \in \mathbb{R}^{n_{2} \times r}$ denote a subset of rows of $A$, and $T \in \mathbb{R}^{n_{3} \times t}$ denote a subset of tubes of $A$. If there exists a tensor $U \in \mathbb{R}^{c \times r \times t}$ such that

$$
\|U(C, R, T)-A\|_{\xi} \leq \alpha \min _{\text {rank }-k A_{k}}\left\|A_{k}-A\right\|_{\xi}
$$

where $\left\|\left\|\|_{\xi}\right.\right.$ is operator norm, Frobenius norm or Entry-wise $\ell_{1}$ norm, then we refer to $C, U, R, T$ as an approximate CUR decomposition of $A$, and sometimes just refer to this as a CURT decomposition of $A$.


Figure 21.5: Column subset selection, row subset selection and tube subset selection.

Recently, [TM17] studied a very different face-based tensor-CUR decomposition, which selects faces from tensors rather than columns. To achieve their results, [TM17] need to make several incoherence assumptions on the original tensor. Their sample complexity depends on $\log n$, and they only sample two of the three dimensions. We will provide more general face-based tensor CURT decompositions.

Definition 21.3.8 (Tensor (face-based) CURT, exact). Given a tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$, we choose three sets of coordinates $S_{1} \subseteq\left[n_{1}\right], S_{2} \subseteq\left[n_{2}\right], S_{3} \subseteq\left[n_{3}\right]$. We define $c=\left|S_{1}\right|, r=\left|S_{2}\right|$ and $t=\left|S_{3}\right|$. Let $C \in \mathbb{R}^{c \times n_{2} \times n_{3}}$ denote a subset of row-tube faces of $A, R \in \mathbb{R}^{n_{1} \times r \times n_{3}}$ denote a subset of column-tube faces of $A$, and $T \in \mathbb{R}^{n_{1} \times n_{2} \times t}$ denote a subset of column-row faces of $A$. Let $C_{2} \in \mathbb{R}^{n_{2} \times c n_{3}}$ denote the matrix obtained by flattening the tensor $C$ along the second dimension. Let $R_{3} \in \mathbb{R}^{n_{3} \times r n_{1}}$ denote the matrix obtained by flattening the tensor $R$ along the third dimension. Let $T_{1} \in \mathbb{R}^{n_{1} \times t n_{2}}$ denote the matrix obtained by flattening the tensor $T$ along the first dimension. If there exists a tensor $U \in \mathbb{R}^{t n_{2} \times c n_{3} \times r n_{1}}$ such that $A$ can be written as

$$
\begin{gathered}
\sum_{i=1}^{t n_{2}} \sum_{j=1}^{c n_{3}} \sum_{l=1}^{r n_{1}} U_{i, j, l}\left(T_{1}\right)_{l} \otimes\left(C_{2}\right)_{i} \otimes\left(R_{3}\right)_{j}=A, \\
U\left(T_{1}, C_{2}, R_{3}\right)=A,
\end{gathered}
$$

or equivalently,

$$
\forall\left(i^{\prime}, j^{\prime}, l^{\prime}\right) \in\left[n_{1}\right] \times\left[n_{2}\right] \times\left[n_{3}\right], A_{i, j, l}=\sum_{i=1}^{t n_{1}} \sum_{j=1}^{c n_{3}} \sum_{l=1}^{r n_{2}} U_{i, j, l}\left(T_{1}\right)_{i^{\prime}, i}\left(C_{2}\right)_{j^{\prime}, j}\left(R_{3}\right)_{l^{\prime}, l}
$$

then we say $C, U, R, T$ is tensor $A$ 's (face-based) CURT decomposition.

Definition 21.3 .9 (Tensor (face-based) CURT, approximate). Given a tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$, for some $k \geq 1$, for some approximation $\alpha>1$, for some norm $\left\|\|_{\xi}\right.$, we choose three sets of coordinates $S_{1} \subseteq\left[n_{1}\right], S_{2} \subseteq\left[n_{2}\right], S_{3} \subseteq\left[n_{3}\right]$. We define $c=\left|S_{1}\right|, r=\left|S_{2}\right|$ and $t=\left|S_{3}\right|$. Let $C \in \mathbb{R}^{c \times n_{2} \times n_{3}}$ denote a subset of row-tube faces of $A, R \in \mathbb{R}^{n_{1} \times r \times n_{3}}$ denote a subset of column-tube faces of $A$, and $T \in \mathbb{R}^{n_{1} \times n_{2} \times t}$ denote a subset of column-row faces of $A$. Let $C_{2} \in \mathbb{R}^{n_{2} \times c n_{3}}$ denote the matrix obtained by flattening the tensor $C$ along the second dimension. Let $R_{3} \in \mathbb{R}^{n_{3} \times r n_{1}}$ denote the matrix obtained by flattening the tensor $R$ along the third dimension. Let $T_{1} \in \mathbb{R}^{n_{1} \times t n_{2}}$ denote the matrix obtained by flattening the tensor $T$ along the first dimension. If there exists a tensor $U \in \mathbb{R}^{t n_{2} \times c n_{3} \times r n_{1}}$ such that

$$
\left\|U\left(T_{1}, C_{2}, R_{3}\right)-A\right\|_{\xi} \leq \alpha \min _{\text {rank }-k}\left\|A_{k}-A\right\|_{\xi}
$$

where $\left\|\left\|\|_{\xi}\right.\right.$ is operator norm, Frobenius norm or Entry-wise $\ell_{1}$ norm, then we refer to $C, U, R, T$ as an approximate CUR decomposition of $A$, and sometimes just refer to this as a (facebased) CURT decomposition of $A$.


Figure 21.6: An example tensor CURT decomposition.

### 21.3.3 Polynomial system verifier

We use the polynomial system verifiers independently developed by Renegar [Ren92a, Ren92b] and Basu et al. [BPR96].

Theorem 21.3.1 (Decision Problem [Ren92a, Ren92b, BPR96]). Given a real polynomial system $P\left(x_{1}, x_{2}, \cdots, x_{v}\right)$ having $v$ variables and $m$ polynomial constraints $f_{i}\left(x_{1}, x_{2}, \cdots, x_{v}\right) \Delta_{i} 0, \forall i \in$ $[m]$, where $\Delta_{i}$ is any of the "standard relations": $\{>, \geq,=, \neq, \leq,<\}$, let d denote the maximum degree of all the polynomial constraints and let $H$ denote the maximum bitsize of the coefficients of all the polynomial constraints. Then in

$$
(m d)^{O(v)} \operatorname{poly}(H),
$$

time one can determine if there exists a solution to the polynomial system $P$.

Recently, this technique has been used to solve a number of low-rank approximation and matrix factorization problems [AGKM12, Mad13, CW15a, BDL16, RSW16, SWZ17].

### 21.3.4 Lower bound on the cost of a polynomial system

An important result we use is the following lower bound on the minimum value attained by a polynomial restricted to a compact connected component of a basic closed semi-algebraic subset of $\mathbb{R}^{v}$.

Theorem 21.3.2 ([JPT13]). Let $T=\left\{x \in \mathbb{R}^{v} \mid f_{1}(x) \geq 0, \cdots, f_{\ell}(x) \geq 0, f_{\ell+1}(x)=\right.$ $\left.0, \cdots, f_{m}(x)=0\right\}$ be defined by polynomials $f_{1}, \cdots, f_{m} \in \mathbb{Z}\left[x_{1}, \cdots, x_{v}\right]$ with $n \geq 2$, degrees bounded by an even integer $d$, and coefficients of absolute value at most $H$, and let $C$ be a compact connected (in the topological sense) component of $T$. Let $g \in \mathbb{Z}\left[x_{1}, \cdots, x_{v}\right]$ be a polynomial of degree at most $d$ and coefficients of absolute value bounded by $H$. Then, the minimum value that $g$ takes over $C$ satisfies that if it is not zero, then its absolute value is greater than or equal to

$$
\left(2^{4-v / 2} \widetilde{H} d^{v}\right)^{-v 2^{v} d^{v}},
$$

where $\widetilde{H}=\max \{H, 2 v+2 m\}$.

While the above theorem involves notions from topology, we shall apply it in an elementary way. Namely, in our setting $T$ will be bounded and so every connected component, which is by definition closed, will also be bounded and therefore compact. As the connected components partition $T$ the theorem will just be applied to give a global minimum value of $g$ on $T$ provided that it is non-zero.

### 21.3.5 Frobenius norm and $\ell_{2}$ relaxation

Theorem 21.3.3 (Generalized rank-constrained matrix approximations, Theorem 2 in [FT07]). Given matrices $A \in \mathbb{R}^{n \times d}, B \in \mathbb{R}^{n \times p}$, and $C \in \mathbb{R}^{q \times d}$, let the $S V D$ of $B$ be $B=U_{B} \Sigma_{B} V_{B}^{\top}$ and the $S V D$ of $C$ be $C=U_{C} \Sigma_{C} V_{C}^{\top}$. Then,

$$
B^{\dagger}\left(U_{B} U_{B}^{\top} A V_{C} C_{C}^{\top}\right)_{k} C^{\dagger}=\underset{\operatorname{rank}-k}{\arg \min } \min _{X \in \mathbb{R}^{p \times q}}\|A-B X C\|_{F}
$$

where $\left(U_{B} U_{B}^{\top} A V_{C} V_{C}^{\top}\right)_{k} \in \mathbb{R}^{p \times q}$ is of rank at most $k$ and denotes the best rank-k approximation to $U_{B} U_{B}^{\top} A V_{C} V_{C}^{\top} \in \mathbb{R}^{p \times d}$ in Frobenius norm.

Claim 21.3.4 ( $\ell_{2}$ relaxation of $\ell_{p}$-regression). Let $p \in[1,2)$. For any $A \in \mathbb{R}^{n \times d}$ and $b \in \mathbb{R}^{n}$, define $x^{*}=\underset{x \in \mathbb{R}^{d}}{\arg \min }\|A x-b\|_{p}$ and $x^{\prime}=\underset{x \in \mathbb{R}^{d}}{\arg \min }\|A x-b\|_{2}$. Then,

$$
\left\|A x^{*}-b\right\|_{p} \leq\left\|A x^{\prime}-b\right\|_{p} \leq n^{1 / p-1 / 2} \cdot\left\|A x^{*}-b\right\|_{p}
$$

Claim 21.3.5 ((Matrix) Frobenius norm relaxation of $\ell_{p}$-low rank approximation). Let $p \in[1,2)$ and for any matrix $A \in \mathbb{R}^{n \times d}$, define $A^{*}=\underset{\operatorname{rank}-k}{\arg \min } B \in \mathbb{R}^{n \times d} \mid\|B-A\|_{p}$ and $A^{\prime}=$ $\underset{B \in \mathbb{R}^{n \times d}}{\arg \min }\|B-A\|_{F}$. Then

$$
\left\|A^{*}-A\right\|_{p} \leq\left\|A^{\prime}-A\right\|_{p} \leq(n d)^{1 / p-1 / 2}\left\|A^{*}-A\right\|_{p}
$$

Claim 21.3.6 ((Tensor) Frobenius norm relaxation of $\ell_{p}$-low rank approximation). Let $p \in$ $[1,2)$ and for any matrix $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$, define

$$
A^{*}=\underset{\operatorname{rank}-k}{\arg \min } \min _{\mathbb{R}^{n} \times n_{2} \times n_{3}}\|B-A\|_{p}
$$

and

$$
A^{\prime}=\underset{\operatorname{rank}-k}{\arg \min } \underset{B \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}}{ }\|B-A\|_{F} .
$$

Then

$$
\left\|A^{*}-A\right\|_{p} \leq\left\|A^{\prime}-A\right\|_{p} \leq\left(n_{1} n_{2} n_{3}\right)^{1 / p-1 / 2}\left\|A^{*}-A\right\|_{p} .
$$

### 21.3.6 CountSketch and Gaussian transforms

Definition 21.3.10 (Sparse embedding matrix or CountSketch transform). A CountSketch transform is defined to be $\Pi=\sigma \cdot \Phi D \in \mathbb{R}^{m \times n}$. Here, $\sigma$ is a scalar, $D$ is an $n \times n$ random diagonal matrix with each diagonal entry independently chosen to be +1 or -1 with equal probability, and $\Phi \in\{0,1\}^{m \times n}$ is an $m \times n$ binary matrix with $\Phi_{h(i), i}=1$ and all remaining entries 0 , where $h:[n] \rightarrow[m]$ is a random map such that for each $i \in[n], h(i)=j$ with probability $1 / m$ for each $j \in[m]$. For any matrix $A \in \mathbb{R}^{n \times d}, \Pi A$ can be computed in $O(\operatorname{nnz}(A))$ time. For any tensor $A \in \mathbb{R}^{n \times d_{1} \times d_{2}}, \Pi A$ can be computed in $O(\mathrm{nnz}(A))$ time. Let $\Pi_{1}, \Pi_{2}, \Pi_{3}$ denote three CountSktech transforms. For any tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$, $A\left(\Pi_{1}, \Pi_{2}, \Pi_{3}\right)$ can be computed in $O(\mathrm{nnz}(A))$ time.

If the above scalar $\sigma$ is not specified in the context, we assume the scalar $\sigma$ to be 1 . Definition 21.3.11 (Gaussian matrix or Gaussian transform). Let $S=\sigma \cdot G \in \mathbb{R}^{m \times n}$ where $\sigma$ is a scalar, and each entry of $G \in \mathbb{R}^{m \times n}$ is chosen independently from the standard Gaussian distribution. For any matrix $A \in \mathbb{R}^{n \times d}, S A$ can be computed in $O(m \cdot n n z(A))$ time. For any tensor $A \in \mathbb{R}^{n \times d_{1} \times d_{2}}, S A$ can be computed in $O(m \cdot \operatorname{nnz}(A))$ time.

If the above scalar $\sigma$ is not specified in the context, we assume the scalar $\sigma$ to be $1 / \sqrt{m}$. In most places, we can combine CountSketch and Gaussian transforms to achieve the following:

Definition 21.3.12 (CountSketch + Gaussian transform). Let $S^{\prime}=S \Pi$, where $\Pi \in \mathbb{R}^{t \times n}$ is the CountSketch transform (defined in Definition 21.3.10) and $S \in \mathbb{R}^{m \times t}$ is the Gaussian transform (defined in Definition 21.3.11). For any matrix $A \in \mathbb{R}^{n \times d}, S^{\prime} A$ can be computed in $O\left(\mathrm{nnz}(A)+d t m^{\omega-2}\right)$ time, where $\omega$ is the matrix multiplication exponent.

Lemma 21.3.7 (Affine Embedding - Theorem 39 in [CW13]). Given matrices $A \in \mathbb{R}^{n \times r}, B \in$ $\mathbb{R}^{n \times d}$, and $\operatorname{rank}(A)=k$, let $m=\operatorname{poly}(k / \epsilon), S \in \mathbb{R}^{m \times n}$ be a sparse embedding matrix (Definition 21.3.10) with scalar $\sigma=1$. Then with probability at least $0.999, \forall X \in \mathbb{R}^{r \times d}$, we have

$$
(1-\epsilon) \cdot\|A X-B\|_{F}^{2} \leq\|S(A X-B)\|_{F}^{2} \leq(1+\epsilon)\|A X-B\|_{F}^{2}
$$

Lemma 21.3 .8 (see, e.g., Lemma 10 in version 1 of $[B W Z 16]^{8}$ ). Let $m=\Omega(k / \epsilon), S=\frac{1}{\sqrt{m}} \cdot G$, where $G \in \mathbb{R}^{m \times n}$ is a random matrix where each entry is an i.i.d Gaussian $N(0,1)$. Then with probability at least 0.998 , S satisfies $(1 \pm 1 / 8)$ Subspace Embedding (Definition 21.3.1) for any fixed matrix $C \in \mathbb{R}^{n \times k}$, and it also satisfies $O(\sqrt{\epsilon / k})$ Approximate Matrix Product (Definition 21.3.2) for any fixed matrix $A$ and $B$ which has the same number of rows.

Lemma 21.3.9 (see, e.g., Lemma 11 in version 1 of [BWZ16]). Let $m=\Omega\left(k^{2}+k / \epsilon\right)$, $\Pi \in \mathbb{R}^{m \times n}$, where $\Pi$ is a sparse embedding matrix (Definition 21.3.10) with scalar $\sigma=1$, then with probability at least 0.998 , S satisfies $(1 \pm 1 / 8)$ Subspace Embedding (Definition 21.3.1) for any fixed matrix $C \in \mathbb{R}^{n \times k}$, and it also satisfies $O(\sqrt{\epsilon / k})$ Approximate Matrix Product (Definition 21.3.2) for any fixed matrix $A$ and $B$ which has the same number of rows.

Lemma 21.3.10 (see, e.g., Lemma 12 in version 1 of [BWZ16]). Let $m_{2}=\Omega\left(k^{2}+k / \epsilon\right)$, $\Pi \in \mathbb{R}^{m_{2} \times n}$, where $\Pi$ is a sparse embedding matrix (Definition 21.3.10) with scalar $\sigma=1$. Let $m_{1}=\Omega(k / \epsilon), S=\frac{1}{\sqrt{m_{1}}} \cdot G$, where $G \in \mathbb{R}^{m_{1} \times m_{2}}$ is a random matrix where each entry is an i.i.d Gaussian $N(0,1)$. Let $S^{\prime}=S \Pi$. Then with probability at least $0.99, S^{\prime}$ is a $(1 \pm 1 / 3)$ Subspace Embedding (Definition 21.3.1) for any fixed matrix $C \in \mathbb{R}^{n \times k}$, and it also satisfies

[^71]$O(\sqrt{\epsilon / k})$ Approximate Matrix Product (Definition 21.3.2) for any fixed matrix $A$ and $B$ which have the same number of rows.

Theorem 21.3.11 (Theorem 36 in [CW13]). Given $A \in \mathbb{R}^{n \times k}, B \in \mathbb{R}^{n \times d}$, suppose $S \in \mathbb{R}^{m \times n}$ is such that $S$ is a $\left(1 \pm \frac{1}{\sqrt{2}}\right)$ Subspace Embedding for A, and satisfies $O(\sqrt{\epsilon / k})$ Approximate Matrix Product for matrices $A$ and $C$ where $C$ with $n$ rows, where $C$ depends on $A$ and $B$. If

$$
\widehat{X}=\arg \min _{X \in \mathbb{R}^{k \times d}}\|S A X-S B\|_{F}^{2}
$$

then

$$
\|A \widehat{X}-B\|_{F}^{2} \leq(1+\epsilon) \min _{X \in \mathbb{R}^{k \times d}}\|A X-B\|_{F}^{2}
$$

### 21.3.7 Cauchy and $p$-stable transforms

Definition 21.3.13 (Dense Cauchy transform). Let $S=\sigma \cdot C \in \mathbb{R}^{m \times n}$ where $\sigma$ is a scalar, and each entry of $C \in \mathbb{R}^{m \times n}$ is chosen independently from the standard Cauchy distribution. For any matrix $A \in \mathbb{R}^{n \times d}, S A$ can be computed in $O(m \cdot \mathrm{nnz}(A))$ time.

Definition 21.3.14 (Sparse Cauchy transform). Let $\Pi=\sigma \cdot S C \in \mathbb{R}^{m \times n}$, where $\sigma$ is a scalar, $S \in \mathbb{R}^{m \times n}$ has each column chosen independently and uniformly from the $m$ standard basis vectors of $\mathbb{R}^{m}$, and $C \in \mathbb{R}^{n \times n}$ is a diagonal matrix with diagonals chosen independently from the standard Cauchy distribution. For any matrix $A \in \mathbb{R}^{n \times d}, \Pi A$ can be computed in $O(\operatorname{nnz}(A))$ time. For any tensor $A \in \mathbb{R}^{n \times d_{1} \times d_{2}}, \Pi A$ can be computed in $O(\mathrm{nnz}(A))$ time. Let $\Pi_{1} \in \mathbb{R}^{m_{1} \times n_{1}}, \Pi_{2} \in \mathbb{R}^{m_{2} \times n_{2}}, \Pi_{3} \in \mathbb{R}^{m_{3} \times n_{3}}$ denote three sparse Cauchy transforms. For any tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}, A\left(\Pi_{1}, \Pi_{2}, \Pi_{3}\right) \in \mathbb{R}^{m_{1} \times m_{2} \times m_{3}}$ can be computed in $O(\mathrm{nnz}(A))$ time.

Definition 21.3.15 (Dense $p$-stable transform). Let $p \in(1,2)$. Let $S=\sigma \cdot C \in \mathbb{R}^{m \times n}$, where $\sigma$ is a scalar, and each entry of $C \in \mathbb{R}^{m \times n}$ is chosen independently from the standard $p$-stable distribution. For any matrix $A \in \mathbb{R}^{n \times d}, S A$ can be computed in $O(m \mathrm{nnz}(A))$ time.

Definition 21.3.16 (Sparse $p$-stable transform). Let $p \in(1,2)$. Let $\Pi=\sigma \cdot S C \in \mathbb{R}^{m \times n}$, where $\sigma$ is a scalar, $S \in \mathbb{R}^{m \times n}$ has each column chosen independently and uniformly from the $m$ standard basis vectors of $\mathbb{R}^{m}$, and $C \in \mathbb{R}^{n \times n}$ is a diagonal matrix with diagonals chosen independently from the standard $p$-stable distribution. For any matrix $A \in \mathbb{R}^{n \times d}, \Pi A$ can be computed in $O(\mathrm{nnz}(A))$ time. For any tensor $A \in \mathbb{R}^{n \times d_{1} \times d_{2}}, \Pi A$ can be computed in $O(\mathrm{nnz}(A))$ time. Let $\Pi_{1} \in \mathbb{R}^{m_{1} \times n_{1}}, \Pi_{2} \in \mathbb{R}^{m_{2} \times n_{2}}, \Pi_{3} \in \mathbb{R}^{m_{3} \times n_{3}}$ denote three sparse $p$-stable transforms. For any tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}, A\left(\Pi_{1}, \Pi_{2}, \Pi_{3}\right) \in \mathbb{R}^{m_{1} \times m_{2} \times m_{3}}$ can be computed in $O(\mathrm{nnz}(A))$ time.

### 21.3.8 Leverage scores

Definition 21.3.17 (Leverage scores). Let $U \in \mathbb{R}^{n \times k}$ have orthonormal columns, and let $p_{i}=u_{i}^{2} / k$, where $u_{i}^{2}=\left\|e_{i}^{\top} U\right\|_{2}^{2}$ is the $i$-th leverage score of $U$.

Definition 21.3.18 (Leverage score sampling). Given $A \in \mathbb{R}^{n \times d}$ with rank $k$, let $U \in \mathbb{R}^{n \times k}$ be an orthonormal basis of the column space of $A$, and for each $i$ let $p_{i}$ be the squared row norm of the $i$-th row of $U$, i.e., the $i$-th leverage score. Let $k \cdot p_{i}$ denote the $i$-th leverage score of $U$ scaled by $k$. Let $\beta>0$ be a constant and $q=\left(q_{1}, \cdots, q_{n}\right)$ denote a distribution such that, for each $i \in[n], q_{i} \geq \beta p_{i}$. Let $s$ be a parameter. Construct an $n \times s$ sampling matrix $B$ and an $s \times s$ rescaling matrix $D$ as follows. Initially, $B=0^{n \times s}$ and $D=0^{s \times s}$. For each column $j$ of $B, D$, independently, and with replacement, pick a row index $i \in[n]$ with probability $q_{i}$, and set $B_{i, j}=1$ and $D_{j, j}=1 / \sqrt{q_{i} s}$. We denote this procedure LEVERAGE SCORE SAMPLING according to the matrix $A$.

### 21.3.9 Lewis weights

We follow the exposition of Lewis weights from [CP15].
Definition 21.3.19. For a matrix $A$, let $a_{i}$ denote the $i^{\text {th }}$ row of $A$, where $a_{i}\left(=\left(A^{i}\right)^{\top}\right)$ is a column vector. The statistical leverage score of a row $a_{i}$ is

$$
\tau_{i}(A) \stackrel{\text { def }}{=} a_{i}^{\top}\left(A^{\top} A\right)^{-1} a_{i}=\left\|\left(A^{\top} A\right)^{-1 / 2} a_{i}\right\|_{2}^{2} .
$$

For a matrix $A$ and norm $p$, the $\ell_{p}$ Lewis weights $w$ are the unique weights such that for each row $i$ we have

$$
w_{i}=\tau_{i}\left(W^{1 / 2-1 / p} A\right)
$$

or equivalently,

$$
a_{i}^{\top}\left(A^{\top} W^{1-2 / p} A\right)^{-1} a_{i}=w_{i}^{2 / p}
$$

Lemma 21.3.12 (Lemma 2.4 of [CP15] and Lemma 7 of $\left.\left[\mathrm{CLM}^{+} 15\right]\right)$. Given a matrix $A \in$ $\mathbb{R}^{n \times d}, n \geq d$, for any constant $C>0,4>p \geq 1$, there is an algorithm which can compute $C$-approximate $\ell_{p}$ Lewis weights for every row $i$ of $A$ in $O\left(\left(n n z(A)+d^{\omega} \log d\right) \log n\right)$ time, where $\omega<2.373$ is the matrix multiplication exponent/Str69, CW87, Wil12].

Lemma 21.3.13 (Theorem 7.1 of [CP15]). Given matrix $A \in \mathbb{R}^{n \times d}(n \geq d)$ with $\ell_{p}(4>$ $p \geq 1)$ Lewis weights $w$, for any set of sampling probabilities $p_{i}, \sum_{i} p_{i}=N$,

$$
p_{i} \geq f(d, p) w_{i}
$$

if $S \in \mathbb{R}^{N \times n}$ has each row chosen independently as the $i^{\text {th }}$ standard basis vector, multiplied by $1 / p_{i}^{1 / p}$, with probability $p_{i} / N$. Then, overall with probability at least 0.999,

$$
\forall x \in \mathbb{R}^{d}, \frac{1}{2}\|A x\|_{p}^{p} \leq\|S A x\|_{p}^{p} \leq 2\|A x\|_{p}^{p}
$$

Furthermore, if $p=1, N=O(d \log d)$. If $1<p<2, N=O(d \log d \log \log d)$. If $2 \leq p<4$, $N=O\left(d^{p / 2} \log d\right)$.

Lemma 21.3.14. Given matrix $A \in \mathbb{R}^{n \times d}(n \geq d)$, there is an algorithm to compute $a$ diagonal matrix $D=S S_{1}$ with $N$ nonzero entries in $O(n$ poly $(d))$ time such that, with probability at least 0.999 , for all $x \in \mathbb{R}^{d}$

$$
\frac{1}{10}\|D A x\|_{p}^{p} \leq\|A x\|_{p}^{p} \leq 10\|D A x\|_{p}^{p}
$$

where $S, S_{1}$ are two sampling/rescaling matrices. Furthermore, if $p=1$, then $N=O(d \log d)$. If $1<p<2$, then $N=O(d \log d \log \log d)$. If $2 \leq p<4$, then $N=O\left(d^{p / 2} \log d\right)$.

Given a matrix $A \in \mathbb{R}^{n \times d}(n \geq d)$, by Lemma 21.3.13 and Lemma 21.3.12, we can compute a sampling/rescaling matrix $S$ in $O\left(\left(n n z(A)+d^{\omega} \log d\right) \log n\right)$ time with $\widetilde{O}(d)$ nonzero entries such that

$$
\forall x \in \mathbb{R}^{d}, \frac{1}{2}\|A x\|_{p}^{p} \leq\|S A x\|_{p}^{p} \leq 2\|A x\|_{p}^{p}
$$

Sometimes, $\operatorname{poly}(d)$ is much smaller than $\log n$. In this case, we are also able to compute such a sampling/rescaling matrix $S$ in $n$ poly $(d)$ time in an alternative way.

To do so, we run one of the input sparsity $\ell_{p}$ embedding algorithms (see e.g., [MM13]) to compute a well conditioned basis $U$ of the column span of $A$ in $n \operatorname{poly}(d / \epsilon)$ time. By sampling according to the well conditioned basis (see e.g. [Cla05, $\mathrm{DDH}^{+} 09$, Woo14b]), we can compute a sampling/rescaling matrix $S_{1}$ such that $(1-\epsilon)\|A x\|_{p}^{p} \leq\left\|S_{1} A x\right\|_{p}^{p} \leq(1+\epsilon)\|A x\|_{p}^{p}$ where $\epsilon \in(0,1)$ is an arbitrary constant. Notice that $S_{1}$ has poly $(d / \epsilon)$ nonzero entries, and thus $S_{1} A$ has size poly $(d / \epsilon)$. Next, we apply Lewis weight sampling according to $S_{1} A$, and
we obtain a sampling/rescaling matrix $S$ for which

$$
\forall x \in \mathbb{R}^{d},\left(1-\frac{1}{3}\right)\left\|S_{1} A x\right\|_{p}^{p} \leq\left\|S S_{1} A x\right\|_{p}^{p} \leq\left(1+\frac{1}{3}\right)\left\|S_{1} A x\right\|_{p}^{p}
$$

This implies that

$$
\forall x \in \mathbb{R}^{d}, \frac{1}{2}\|A x\|_{p}^{p} \leq\left\|S S_{1} A x\right\|_{p}^{p} \leq 2\|A x\|_{p}^{p}
$$

Note that $S S_{1}$ is still a sampling/rescaling matrix according to $A$, and the number of non-zero entries is $\widetilde{O}(d)$. The total running time is thus $n \operatorname{poly}(d / \epsilon)$, as desired.

### 21.3.10 TensorSketch

Let $\phi\left(v_{1}, v_{2}, \cdots, v_{q}\right)$ denote the function that maps $q \operatorname{vectors}\left(u_{i} \in \mathbb{R}^{n_{i}}\right)$ to the $\prod_{i=1}^{q} n_{i^{-}}$ dimensional vector formed by $v_{1} \otimes v_{2} \otimes \cdots \otimes u_{q}$.

We first give the definition of TensorSketch. Similar definitions can be found in previous work [Pag13, PP13, ANW14, WTSA15].

Definition 21.3.20 (TensorSketch [Pag13]). Given $q$ points $v_{1}, v_{2}, \cdots, v_{q}$ where for each $i \in[q], v_{i} \in \mathbb{R}^{n_{i}}$, let $m$ be the target dimension. The TensorSketch transform is specified using $q 3$-wise independent hash functions, $h_{1}, \cdots, h_{q}$, where for each $i \in[q], h_{i}:\left[n_{i}\right] \rightarrow[m]$, as well as $q 4$-wise independent sign functions $s_{1}, \cdots, s_{q}$, where for each $i \in[q], s_{i}:\left[n_{i}\right] \rightarrow$ $\{-1,+1\}$.

TensorSketch applied to $v_{1}, \cdots, v_{q}$ is then CountSketch applied to $\phi\left(v_{1}, \cdots, v_{q}\right)$ with hash function $H:\left[\prod_{i=1}^{q} n_{i}\right] \rightarrow[m]$ and sign functions $S:\left[\prod_{i=1}^{q} n_{i}\right] \rightarrow\{-1,+1\}$ defined as follows:

$$
H\left(i_{1}, \cdots, i_{q}\right)=h_{1}\left(i_{1}\right)+h_{2}\left(s_{2}\right)+\cdots+h_{q}\left(i_{q}\right) \quad(\bmod m),
$$

and

$$
S\left(i_{1}, \cdots, i_{q}\right)=s_{1}\left(i_{1}\right) \cdot s_{2}\left(i_{2}\right) \cdots s_{q}\left(i_{q}\right)
$$

Using the Fast Fourier Transform, TensorSinetch $\left(v_{1}, \cdots, v_{q}\right)$ can be computed in $O\left(\sum_{i=1}^{q}\left(\mathrm{nnz}\left(v_{i}\right)+\right.\right.$ $m \log m)$ ) time.

Note that Theorem 1 in [ANW14] only defines $\phi(v)=v \otimes v \otimes \cdots \otimes v$. Here we state a stronger version of Theorem 1 than in [ANW14], though the proofs are identical; a formal derivation can be found in [DW17].

Theorem 21.3.15 (Generalized version of Theorem 1 in [ANW14]). Let $S$ be the $\left(\prod_{i=1}^{q} n_{i}\right) \times$ $m$ matrix such that TensorSketch $\left(v_{1}, v_{2}, \cdots, v_{q}\right)$ is $\phi\left(v_{1}, v_{2}, \cdots, v_{q}\right) S$ for a randomly selected TensorSketch. The matrix $S$ satisfies the following two properties.

Property I (Approximate Matrix Product). Let $A$ and $B$ be matrices with $\prod_{i=1}^{q} n_{i}$ rows. For $m \geq\left(2+3^{q}\right) /\left(\epsilon^{2} \delta\right)$, we have

$$
\operatorname{Pr}\left[\left\|A^{\top} S S^{\top} B-A^{\top} B\right\|_{F}^{2} \leq \epsilon^{2}\|A\|_{F}^{2}\|B\|_{F}^{2}\right] \geq 1-\delta
$$

Property II (Subspace Embedding). Consider a fixed $k$-dimensional subspace V. If $m \geq k^{2}\left(2+3^{q}\right) /\left(\epsilon^{2} \delta\right)$, then with probability at least $1-\delta,\|x S\|_{2}=(1 \pm \epsilon)\|x\|_{2}$ simultaneously for all $x \in V$.

### 21.4 Frobenius Norm for Arbitrary Tensors

Section 21.4.1 presents a Frobenius norm tensor low-rank approximation algorithm with $(1+\epsilon)$-approximation ratio. Section 21.4.2 introduces a tool which is able to reduce the size of the objective function from $n^{3}$ to poly $(k, 1 / \epsilon)$. Section 21.4.3 introduces a new problem called tensor multiple regression. Section 21.4.4 presents several bicriteria algorithms. Section 21.4.5 introduces a powerful tool which we call generalized matrix row subset selection. Section 21.4.6 presents an algorithm that is able to select a batch of columns, rows and tubes from a given tensor, and those samples are also able to form a low-rank solution. Section 21.4.7 presents several useful tools for tensor problems, and also two $(1+\epsilon)$ approximation CURT decomposition algorithms: one has the optimal sample complexity, and the other has the optimal running time. Section 21.4.9 shows how to solve the problem if the size of the objective function is small. Section 21.5 extends several techniques from 3rd order tensors to general $q$-th order tensors, for any $q \geq 3$. Finally, in Section 21.6 we also provide a new matrix CUR decomposition algorithm, which is faster than [BW14].

For simplicity of presentation, we assume $A_{k}$ exists in theorems (e.g., Theorem 21.4.1) which concern outputting a rank- $k$ solution, as well as the theorems (e.g., Theorem 21.4.7, Theorem 21.4.8, Theorem 21.4.13) which concern outputting a bicriteria solution (the output rank is larger than $k$ ). For each of the bicriteria theorems, we can obtain a more detailed version when $A_{k}$ does not exist, like Theorem 21.1.1 in Section 21.1 (by instead considering a tensor sufficiently close to $A_{k}$ in objective function value). Note that the theorems for column, row, tube subset selection Theorem 21.4.20 and Theorem 21.4.21 also belong to this first category. In the second category, for each of the rank- $k$ theorems we can obtain a more detailed version handling all cases, even when $A_{k}$ does not exist, like Theorem 21.1.2
in Section 21.1 (by instead considering a tensor sufficiently close to $A_{k}$ in objective function value).

Several other tensor results or tools (e.g., Theorem 21.4.4, Lemma 21.4.3, Theorem 21.4.39, Theorem 21.4.40, Theorem 21.4.14, Theorem 21.5.1) that we build in this section do not belong to the above two categories. It means those results do not depend on whether $A_{k}$ exists or not and whether OPT is zero or not.

### 21.4.1 $(1+\epsilon)$-approximate low-rank approximation

```
Algorithm 21.2 Frobenius Norm Low-rank Approximation
    procedure \(\operatorname{FLowRankApprox}(A, n, k, \epsilon) \quad \triangleright\) Theorem 21.4.1
        \(s_{1} \leftarrow s_{2} \leftarrow s_{3} \leftarrow O(k / \epsilon)\).
        Choose sketching matrices \(S_{1} \in \mathbb{R}^{n^{2} \times s_{1}}, S_{2} \in \mathbb{R}^{n^{2} \times s_{2}}, S_{3} \in \mathbb{R}^{n^{2} \times s_{3}}\). \(\quad\) Definition
    21.3.12
        Compute \(A_{i} S_{i}, \forall i \in[3]\).
        \(Y_{1}, Y_{2}, Y_{3}, C \leftarrow\) FInputSparsityReduction \(\left(A, A_{1} S_{1}, A_{2} S_{2}, A_{3} S_{3}, n, s_{1}, s_{2}, s_{3}, k, \epsilon\right)\).
                                    \(\triangleright\) Algorithm 21.3
        Create variables for \(X_{i} \in \mathbb{R}^{s_{i} \times k}, \forall i \in[3]\).
        Run polynomial system verifier for \(\left\|\left(Y_{1} X_{1}\right) \otimes\left(Y_{2} X_{2}\right) \otimes\left(Y_{3} X_{3}\right)-C\right\|_{F}^{2}\).
        return \(A_{1} S_{1} X_{1}, A_{2} S_{2} X_{2}\), and \(A_{3} S_{3} X_{3}\).
    end procedure
```

Theorem 21.4.1. Given a 3rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1, \epsilon \in(0,1)$, there exists an algorithm which takes $O(\mathrm{nnz}(A))+n \operatorname{poly}(k, 1 / \epsilon)+2^{O\left(k^{2} / \epsilon\right)}$ time and outputs three matrices $U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{n \times k}, W \in \mathbb{R}^{n \times k}$ such that

$$
\left\|\sum_{i=1}^{k} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{F}^{2} \leq(1+\epsilon)_{\mathrm{rank}-k} \min _{A_{k}}\left\|A_{k}-A\right\|_{F}^{2}
$$

holds with probability 9/10.

Proof. Given any tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$, we define three matrices $A_{1} \in \mathbb{R}^{n_{1} \times n_{2} n_{3}}, A_{2} \in$ $\mathbb{R}^{n_{2} \times n_{3} n_{1}}, A_{3} \in \mathbb{R}^{n_{3} \times n_{1} n_{2}}$ such that, for any $i \in\left[n_{1}\right], j \in\left[n_{2}\right], l \in\left[n_{3}\right]$,

$$
A_{i, j, l}=\left(A_{1}\right)_{i,(j-1) \cdot n_{3}+l}=\left(A_{2}\right)_{j,(l-1) \cdot n_{1}+i}=\left(A_{3}\right)_{l,(i-1) \cdot n_{2}+j} .
$$

We define OPT as

$$
\mathrm{OPT}=\min _{\operatorname{rank}-k A^{\prime}}\left\|A^{\prime}-A\right\|_{F}^{2} .
$$

Suppose the optimal $A_{k}=U^{*} \otimes V^{*} \otimes W^{*}$. We fix $V^{*} \in \mathbb{R}^{n \times k}$ and $W^{*} \in \mathbb{R}^{n \times k}$. We use $V_{1}^{*}, V_{2}^{*}, \cdots, V_{k}^{*}$ to denote the columns of $V^{*}$ and $W_{1}^{*}, W_{2}^{*}, \cdots, W_{k}^{*}$ to denote the columns of $W^{*}$.

We consider the following optimization problem,

$$
\min _{U_{1}, \cdots, U_{k} \in \mathbb{R}^{n}}\left\|\sum_{i=1}^{k} U_{i} \otimes V_{i}^{*} \otimes W_{i}^{*}-A\right\|_{F}^{2}
$$

which is equivalent to

$$
\min _{U_{1}, \cdots, U_{k} \in \mathbb{R}^{n}}\left\|\left[\begin{array}{llll}
U_{1} & U_{2} & \cdots & U_{k}
\end{array}\right]\left[\begin{array}{c}
V_{1}^{*} \otimes W_{1}^{*} \\
V_{2}^{*} \otimes W_{2}^{*} \\
\cdots \\
V_{k}^{*} \otimes W_{k}^{*}
\end{array}\right]-A\right\|_{F}^{2} .
$$

We use matrix $Z_{1}$ to denote $\left[\begin{array}{c}\operatorname{vec}\left(V_{1}^{*} \otimes W_{1}^{*}\right) \\ \operatorname{vec}\left(V_{2}^{*} \otimes W_{2}^{*}\right) \\ \ldots \\ \operatorname{vec}\left(V_{k}^{*} \otimes W_{k}^{*}\right)\end{array}\right] \in \mathbb{R}^{k \times n^{2}}$ and matrix $U$ to denote $\left[\begin{array}{llll}U_{1} & U_{2} & \cdots & U_{k}\end{array}\right]$. Then we can obtain the following equivalent objective function,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1}-A_{1}\right\|_{F}^{2}
$$

Notice that $\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1}-A_{1}\right\|_{F}^{2}=\mathrm{OPT}$, since $A_{k}=U^{*} Z_{1}$.
Let $S_{1}^{\top} \in \mathbb{R}^{s_{1} \times n^{2}}$ be a sketching matrix defined in Definition 21.3.12, where $s_{1}=$ $O(k / \epsilon)$. We obtain the following optimization problem,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1} S_{1}-A_{1} S_{1}\right\|_{F}^{2} .
$$

Let $\widehat{U} \in \mathbb{R}^{n \times k}$ denote the optimal solution to the above optimization problem. Then $\widehat{U}=$ $A_{1} S_{1}\left(Z_{1} S_{1}\right)^{\dagger}$. By Lemma 21.3.10 and Theorem 21.3.11, we have

$$
\left\|\widehat{U} Z_{1}-A_{1}\right\|_{F}^{2} \leq(1+\epsilon) \min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1}-A_{1}\right\|_{F}^{2}=(1+\epsilon) \mathrm{OPT},
$$

which implies

$$
\left\|\sum_{i=1}^{k} \widehat{U}_{i} \otimes V_{i}^{*} \otimes W_{i}^{*}-A\right\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT} .
$$

To write down $\widehat{U}_{1}, \cdots, \widehat{U}_{k}$, we use the given matrix $A_{1}$, and we create $s_{1} \times k$ variables for matrix $\left(Z_{1} S_{1}\right)^{\dagger}$.

As our second step, we fix $\widehat{U} \in \mathbb{R}^{n \times k}$ and $W^{*} \in \mathbb{R}^{n \times k}$, and we convert tensor $A$ into matrix $A_{2}$. Let matrix $Z_{2}$ denote $\left[\begin{array}{c}\operatorname{vec}\left(\widehat{U}_{1} \otimes W_{1}^{*}\right) \\ \operatorname{vec}\left(\widehat{U}_{2} \otimes W_{2}^{*}\right) \\ \ldots \\ \operatorname{vec}\left(\widehat{U}_{k} \otimes W_{k}^{*}\right)\end{array}\right]$. We consider the following objective function,

$$
\min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2}-A_{2}\right\|_{F}^{2}
$$

for which the optimal cost is at most $(1+\epsilon)$ OPT.
Let $S_{2}^{\top} \in \mathbb{R}^{s_{2} \times n^{2}}$ be a sketching matrix defined in Definition 21.3.12, where $s_{2}=$ $O(k / \epsilon)$. We sketch $S_{2}$ on the right of the objective function to obtain the new objective function,

$$
\min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2} S_{2}-A_{2} S_{2}\right\|_{F}^{2} .
$$

Let $\widehat{V} \in \mathbb{R}^{n \times k}$ denote the optimal solution of the above problem. Then $\widehat{V}=A_{2} S_{2}\left(Z_{2} S_{2}\right)^{\dagger}$. By Lemma 21.3.10 and Theorem 21.3.11, we have,

$$
\left\|\widehat{V} Z_{2}-A_{2}\right\|_{F}^{2} \leq(1+\epsilon) \min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2}-A_{2}\right\|_{F}^{2} \leq(1+\epsilon)^{2} \mathrm{OPT}
$$

which implies

$$
\left\|\sum_{i=1}^{k} \widehat{U}_{i} \otimes \widehat{V}_{i} \otimes W_{i}^{*}-A\right\|_{F}^{2} \leq(1+\epsilon)^{2} \mathrm{OPT}
$$

To write down $\widehat{V}_{1}, \cdots, \widehat{V}_{k}$, we need to use the given matrix $A_{2} \in \mathbb{R}^{n^{2} \times n}$, and we need to create $s_{2} \times k$ variables for matrix $\left(Z_{2} S_{2}\right)^{\dagger}$.

As our third step, we fix the matrices $\widehat{U} \in \mathbb{R}^{n \times k}$ and $\widehat{V} \in \mathbb{R}^{n \times k}$. We convert tensor $A \in \mathbb{R}^{n \times n \times n}$ into matrix $A_{3} \in \mathbb{R}^{n^{2 \times n}}$. Let matrix $Z_{3}$ denote $\left[\begin{array}{c}\operatorname{vec}\left(\widehat{U}_{1} \otimes \widehat{V}_{1}\right) \\ \operatorname{vec}\left(\widehat{U}_{2} \otimes \widehat{V}_{2}\right) \\ \cdots \\ \operatorname{vec}\left(\widehat{U}_{k} \otimes \widehat{V}_{k}\right)\end{array}\right]$. We consider the following objective function,

$$
\min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3}-A_{3}\right\|_{F}^{2},
$$

which has optimal cost at most $(1+\epsilon)^{2}$ OPT.
Let $S_{3}^{\top} \in \mathbb{R}^{s_{3} \times n^{2}}$ be a sketching matrix defined in Definition 21.3.12, where $s_{3}=$ $O(k / \epsilon)$. We sketch $S_{3}$ on the right of the objective function to obtain a new objective function,

$$
\min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3} S_{3}-A_{3} S_{3}\right\|_{F}^{2} .
$$

Let $\widehat{W} \in \mathbb{R}^{n \times k}$ denote the optimal solution of the above problem. Then $\widehat{W}=A_{3} S_{3}\left(Z_{3} S_{3}\right)^{\dagger}$. By Lemma 21.3.10 and Theorem 21.3.11, we have,

$$
\left\|\widehat{W} Z_{3}-A_{3}\right\|_{F}^{2} \leq(1+\epsilon) \min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3}-A_{3}\right\|_{F}^{2} \leq(1+\epsilon)^{3} \mathrm{OPT} .
$$

Thus, we have

$$
\min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k}\left(A_{1} S_{1} X_{1}\right)_{i} \otimes\left(A_{2} S_{2} X_{2}\right)_{i} \otimes\left(A_{3} S_{3} X_{3}\right)_{i}-A\right\|_{F}^{2} \leq(1+\epsilon)^{3} \mathrm{OPT} .
$$

Let $V_{1}=A_{1} S_{1}, V_{2}=A_{2} S_{2}, V_{3}=A_{3} S_{3}$, we then apply Lemma 21.4.3, and we obtain $\widehat{V}_{1}, \widehat{V}_{2}, \widehat{V}_{3}, C$. We then apply Theorem 21.4.44. Correctness follows by rescaling $\epsilon$ by a constant factor.

Running time. Due to Definition 21.3.12, the running time of line 4 is $O(\mathrm{nnz}(A))+$ $n$ poly $(k)$. The running time of line 5 is shown by Lemma 21.4.3, and the running time of line 7 is shown by Theorem 21.4.44.

Theorem 21.4.2. Suppose we are given a 3 rd order $n \times n \times n$ tensor $A$ such that each entry can be written using $n^{\delta}$ bits, where $\delta>0$ is a given, value which can be arbitrarily small (e.g., we could have $n^{\delta}$ being $O(\log n)$ ). Define OPT $=\inf _{\text {rank }-k A_{k}}\left\|A_{k}-A\right\|_{F}^{2}$. For any $k \geq 1$, and for any $0<\epsilon<1$, define $n^{\delta^{\prime}}=O\left(n^{\delta} 2^{O\left(k^{2} / \epsilon\right)}\right)$. (I) If OPT $>0$, and there exists a rank-k $A_{k}=U^{*} \otimes V^{*} \otimes W^{*}$ tensor, with size $n \times n \times n$, such that $\left\|A_{k}-A\right\|_{F}^{2}=$ OPT, and $\max \left(\left\|U^{*}\right\|_{F},\left\|V^{*}\right\|_{F},\left\|W^{*}\right\|_{F}\right) \leq 2^{O\left(n^{\delta^{\prime}}\right)}$, then there exists an algorithm that takes $\left(\mathrm{nnz}(A)+n \operatorname{poly}(k, 1 / \epsilon)+2^{O\left(k^{2} / \epsilon\right)}\right) n^{\delta}$ time in the unit cost RAM model with word size $O(\log n)$ bits ${ }^{9}$ and outputs three $n \times k$ matrices $U, V, W$ such that

$$
\begin{equation*}
\|U \otimes V \otimes W-A\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT} \tag{21.5}
\end{equation*}
$$

holds with probability 9/10, and each entry of each of $U, V, W$ fits in $n^{\delta^{\prime}}$ bits.
(II) If $\mathrm{OPT}>0$, and $A_{k}$ does not exist, and there exist three $n \times k$ matrices $U^{\prime}, V^{\prime}, W^{\prime}$ for which $\max \left(\left\|U^{\prime}\right\|_{F},\left\|V^{\prime}\right\|_{F},\left\|W^{\prime}\right\|_{F}\right) \leq 2^{O\left(n^{\delta^{\prime}}\right)}$ and $\left\|U^{\prime} \otimes V^{\prime} \otimes W^{\prime}-A\right\|_{F}^{2} \leq(1+\epsilon / 2)$ OPT, then we can find $U, V, W$ such that (21.5) holds.

[^72](III) If $\mathrm{OPT}=0$ and $A_{k}$ does exist, and there exists a solution $U^{*}, V^{*}, W^{*}$ such that each entry can be written by $n^{\delta^{\prime}}$ bits, then we can obtain (21.5).
(IV) If $\mathrm{OPT}=0$, and there exist three $n \times k$ matrices $U, V$, $W$ such that
$$
\max \left(\|U\|_{F},\|V\|_{F},\|W\|_{F}\right) \leq 2^{O\left(n^{\delta^{\prime}}\right)}
$$
and
\[

$$
\begin{equation*}
\|U \otimes V \otimes W-A\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}+2^{-\Omega\left(n^{\delta^{\prime}}\right)}=2^{-\Omega\left(n^{\delta^{\prime}}\right)} \tag{21.6}
\end{equation*}
$$

\]

then we can output $U, V, W$ such that (21.6) holds.
Further if $A_{k}$ exists, we can output a number $Z$ for which $\mathrm{OPT} \leq Z \leq(1+\epsilon) \mathrm{OPT}$. For all the cases above, the algorithm runs in the same time as (I) and succeeds with probability at least 9/10.

Proof. This follows by the discussion in Section 21.1, Theorem 21.4.1 and Theorem 21.4.44 in Section 21.4.9.

Part (I) Suppose $\delta>0$ and $A_{k}=U^{*} \otimes V^{*} \otimes W^{*}$ exists and each of $\left\|U^{*}\right\|_{F},\left\|V^{*}\right\|_{F}$, and $\left\|W^{*}\right\|_{F}$ is bounded by $2^{O\left(n^{\delta^{\prime}}\right)}$. We assume the computation model is the unit cost RAM model with words of size $O(\log n)$ bits, and allow each number of the input tensor $A$ to be written using $n^{\delta}$ bits. For the case when OPT is nonzero, using the proof of Theorem 21.4.1 and Theorems 21.4.44, 21.3.2, there exists a lower bound on the cost OPT, which is at least
 using $O\left(n^{\delta^{\prime}}\right)$ bits to obtain $U^{\prime}, V^{\prime}, W^{\prime}$. Using the triangle inequality and our lower bound on OPT, it follows that $U^{\prime}, V^{\prime}, W^{\prime}$ provide a $(1+\epsilon)$-approximation.

Thus, applying Theorem 21.4.1 by fixing $U^{\prime}, V^{\prime}, W^{\prime}$ and using Theorem 21.4.44 at the end, we can output three matrices $U, V, W$, where each entry can be written using $n^{\delta^{\prime}}$ bits, so that we satisfy $\|U \otimes V \otimes W-A\|_{F}^{2} \leq(1+\epsilon)$ OPT.

For the running time, since each entry of the input is bounded by $n^{\delta}$ bits, due to Theorem 21.4.1, we need $(\operatorname{nnz}(A)+n \operatorname{poly}(k / \epsilon)) \cdot n^{\delta}$ time to reduce the size of the problem to $\operatorname{poly}(k / \epsilon)$ size (with each number represented using $O\left(n^{\delta}\right)$ bits). According to Theorem 21.4.44, the running time of using a polynomial system verifier to get the solution is $2^{O\left(k^{2} / \epsilon\right)} n^{O\left(\delta^{\prime}\right)}=2^{O\left(k^{2} / \epsilon\right)} n^{O(\delta)}$ time. Thus the total running time is $(\mathrm{nnz}(A)+n \operatorname{poly}(k / \epsilon)) n^{\delta}+$ $2^{O\left(k^{2} / \epsilon\right)} \cdot n^{O(\delta)}$.

Part (II) is similar to Part (I). Part (III) is trivial to prove since there exists a solution which can be written down in the bit model, so we obtain a $(1+\epsilon)$-approximation. Part (IV) is also very similar to Part (II).

### 21.4.2 Input sparsity reduction

```
Algorithm 21.3 Reducing the Size of the Objective Function from poly \((n)\) to poly \((k)\)
    procedure FInputSparsityReduction \(\left(A, V_{1}, V_{2}, V_{3}, n, b_{1}, b_{2}, b_{3}, k, \epsilon\right) \quad \triangleright\) Lemma
    21.4.3
        \(c_{1} \leftarrow c_{2} \leftarrow c_{3} \leftarrow \operatorname{poly}(k, 1 / \epsilon)\).
        Choose sparse embedding matrices \(T_{1} \in \mathbb{R}^{c_{1} \times n}, T_{2} \in \mathbb{R}^{c_{2} \times n}, T_{3} \in \mathbb{R}^{c_{3} \times n}\). \(\triangleright\)
    Definition 21.3.10
        \(\widehat{V}_{i} \leftarrow T_{i} V_{i} \in \mathbb{R}^{c_{i} \times b_{i}}, \forall i \in[3]\).
        \(C \leftarrow A\left(T_{1}, T_{2}, T_{3}\right) \in \mathbb{R}^{c_{1} \times c_{2} \times c_{3}}\).
        return \(\widehat{V}_{1}, \widehat{V}_{2}, \widehat{V}_{3}\) and \(C\).
    end procedure
```

Lemma 21.4.3. Let $\operatorname{poly}(k, 1 / \epsilon) \geq b_{1} b_{2} b_{3} \geq k$. Given a tensor $A \in \mathbb{R}^{n \times n \times n}$ and three matrices $V_{1} \in \mathbb{R}^{n \times b_{1}}$, $V_{2} \in \mathbb{R}^{n \times b_{2}}$, and $V_{3} \in \mathbb{R}^{n \times b_{3}}$, there exists an algorithm that takes $O\left(\mathrm{nnz}(A)+\mathrm{nnz}\left(V_{1}\right)+\mathrm{nnz}\left(V_{2}\right)+\mathrm{nnz}\left(V_{3}\right)\right)=O(\mathrm{nnz}(A)+n \operatorname{poly}(k / \epsilon))$ time and outputs a tensor $C \in \mathbb{R}^{c_{1} \times c_{2} \times c_{3}}$ and three matrices $\widehat{V}_{1} \in \mathbb{R}^{c_{1} \times b_{1}}, \widehat{V}_{2} \in \mathbb{R}^{c_{2} \times b_{2}}$ and $\widehat{V}_{3} \in \mathbb{R}^{c_{3} \times b_{3}}$ with $c_{1}=c_{2}=c_{3}=\operatorname{poly}(k, 1 / \epsilon)$, such that with probability at least 0.99 , for all $\alpha>0, X_{1}, X_{1}^{\prime} \in$ $\mathbb{R}^{b_{1} \times k}, X_{2}, X_{2}^{\prime} \in \mathbb{R}^{b_{2} \times k}, X_{3}, X_{3}^{\prime} \in \mathbb{R}^{b_{3} \times k}$ satisfy that,

$$
\left\|\sum_{i=1}^{k}\left(\widehat{V}_{1} X_{1}^{\prime}\right)_{i} \otimes\left(\widehat{V}_{2} X_{2}^{\prime}\right)_{i} \otimes\left(\widehat{V}_{3} X_{3}^{\prime}\right)_{i}-C\right\|_{F}^{2} \leq \alpha\left\|\sum_{i=1}^{k}\left(\widehat{V}_{1} X_{1}\right)_{i} \otimes\left(\widehat{V}_{2} X_{2}\right)_{i} \otimes\left(\widehat{V}_{3} X_{3}\right)_{i}-C\right\|_{F}^{2},
$$

then,

$$
\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}^{\prime}\right)_{i} \otimes\left(V_{2} X_{2}^{\prime}\right)_{i} \otimes\left(V_{3} X_{3}^{\prime}\right)_{i}-A\right\|_{F}^{2} \leq(1+\epsilon) \alpha\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}\right)_{i} \otimes\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}-A\right\|_{F}^{2}
$$

Proof. Let $X_{1} \in \mathbb{R}^{b_{1} \times k}, X_{2} \in \mathbb{R}^{b_{2} \times k}, X_{3} \in \mathbb{R}^{b_{3} \times k}$. First, we define $Z_{1}=\left(\left(V_{2} X_{2}\right)^{\top} \odot\left(V_{3} X_{3}\right)^{\top}\right) \in$ $\mathbb{R}^{k \times n^{2}}$. (Note that, for each $i \in[k]$, the $i$-th row of matrix $Z_{1}$ is $\operatorname{vec}\left(\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}\right)$.) Then, by flattening we have

$$
\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}\right)_{i} \otimes\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}-A\right\|_{F}^{2}=\left\|V_{1} X_{1} \cdot Z_{1}-A_{1}\right\|_{F}^{2} .
$$

We choose a sparse embedding matrix (Definition 21.3.10) $T_{1} \in \mathbb{R}^{c_{1} \times n}$ with $c_{1}=\operatorname{poly}(k, 1 / \epsilon)$ rows. Since $V_{1}$ has $b_{1} \leq \operatorname{poly}(k / \epsilon)$ columns, according to Lemma 21.3.7 with probability 0.999, for all $X_{1} \in \mathbb{R}^{b_{1} \times k}, Z \in \mathbb{R}^{k \times n^{2}}$,

$$
(1-\epsilon)\left\|V_{1} X_{1} Z-A_{1}\right\|_{F}^{2} \leq\left\|T_{1} V_{1} X_{1} Z-T_{1} A_{1}\right\|_{F}^{2} \leq(1+\epsilon)\left\|V_{1} X_{1} Z-A_{1}\right\|_{F}^{2}
$$

Therefore, we have

$$
\left\|T_{1} V_{1} X_{1} \cdot Z_{1}-T_{1} A_{1}\right\|_{F}^{2}=(1 \pm \epsilon)\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}\right)_{i} \otimes\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}-A\right\|_{F}^{2}
$$

Second, we unflatten matrix $T_{1} A_{1} \in \mathbb{R}^{c_{1} \times n^{2}}$ to obtain a tensor $A^{\prime} \in \mathbb{R}^{c_{1} \times n \times n}$. Then we flatten $A^{\prime}$ along the second direction to obtain $A_{2} \in \mathbb{R}^{n \times c_{1} n}$. We define $Z_{2}=\left(T_{1} V_{1} X_{1}\right)^{\top} \odot$ $\left(V_{3} X_{3}\right)^{\top} \in \mathbb{R}^{k \times c_{1} n}$. Then, by flattening,

$$
\begin{aligned}
\left\|V_{2} X_{2} \cdot Z_{2}-A_{2}\right\|_{F}^{2} & =\left\|T_{1} V_{1} X_{1} \cdot Z_{1}-T_{1} A_{1}\right\|_{F}^{2} \\
& =(1 \pm \epsilon)\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}\right)_{i} \otimes\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}-A\right\|_{F}^{2}
\end{aligned}
$$

We choose a sparse embedding matrix (Definition 21.3.10) $T_{2} \in \mathbb{R}^{c_{2} \times n}$ with $c_{2}=\operatorname{poly}(k, 1 / \epsilon)$ rows. Then according to Lemma 21.3 .7 with probability 0.999 , for all $X_{2} \in \mathbb{R}^{b_{2} \times k}, Z \in$ $\mathbb{R}^{k \times c_{1} n}$,

$$
(1-\epsilon)\left\|V_{2} X_{2} Z-A_{2}\right\|_{F}^{2} \leq\left\|T_{2} V_{2} X_{2} Z-T_{2} A_{2}\right\|_{F}^{2} \leq(1+\epsilon)\left\|V_{2} X_{2} Z-A_{2}\right\|_{F}^{2}
$$

Therefore, we have

$$
\begin{aligned}
\left\|T_{2} V_{2} X_{2} \cdot Z_{2}-T_{2} A_{2}\right\|_{F}^{2} & =(1 \pm \epsilon)\left\|V_{2} X_{2} \cdot Z_{2}-A_{2}\right\|_{F}^{2} \\
& =(1 \pm \epsilon)^{2}\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}\right)_{i} \otimes\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}-A\right\|_{F}^{2} .
\end{aligned}
$$

Third, we unflatten matrix $T_{2} A_{2} \in \mathbb{R}^{c_{2} \times c_{1} n}$ to obtain a tensor $A^{\prime \prime}\left(=A\left(T_{1}, T_{2}, I\right)\right) \in$ $\mathbb{R}^{c_{1} \times c_{2} \times n}$. Then we flatten tensor $A^{\prime \prime}$ along the last direction (the third direction) to obtain matrix $A_{3} \in \mathbb{R}^{n \times c_{1} c_{2}}$. We define $Z_{3}=\left(T_{1} V_{1} X_{1}\right)^{\top} \odot\left(T_{2} V_{2} X_{2}\right)^{\top} \in \mathbb{R}^{k \times c_{1} c_{2}}$. Then, by flattening, we have

$$
\begin{aligned}
\left\|V_{3} X_{3} \cdot Z_{3}-A_{3}\right\|_{F}^{2} & =\left\|T_{2} V_{2} X_{2} \cdot Z_{2}-T_{2} A_{2}\right\|_{F}^{2} \\
& =(1 \pm \epsilon)^{2}\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}\right)_{i} \otimes\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}-A\right\|_{F}^{2}
\end{aligned}
$$

We choose a sparse embedding matrix (Definition 21.3.10) $T_{3} \in \mathbb{R}^{c_{3} \times n}$ with $c_{3}=\operatorname{poly}(k, 1 / \epsilon)$ rows. Then according to Lemma 21.3 .7 with probability 0.999 , for all $X_{3} \in \mathbb{R}^{b_{3} \times k}, Z \in$ $\mathbb{R}^{k \times c_{1} c_{2}}$,

$$
(1-\epsilon)\left\|V_{3} X_{3} Z-A_{3}\right\|_{F}^{2} \leq\left\|T_{3} V_{3} X_{3} Z-T_{3} A_{3}\right\|_{F}^{2} \leq(1+\epsilon)\left\|V_{3} X_{3} Z-A_{3}\right\|_{F}^{2}
$$

Therefore, we have

$$
\left\|T_{3} V_{3} X_{3} \cdot Z_{3}-T_{3} A_{3}\right\|_{F}^{2}=(1 \pm \epsilon)^{3}\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}\right)_{i} \otimes\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}-A\right\|_{F}^{2} .
$$

Note that

$$
\left\|T_{3} V_{3} X_{3} \cdot Z_{3}-T_{3} A_{3}\right\|_{F}^{2}=\left\|\sum_{i=1}^{k}\left(T_{1} V_{1} X_{1}\right)_{i} \otimes\left(T_{2} V_{2} X_{2}\right)_{i} \otimes\left(T_{3} V_{3} X_{3}\right)_{i}-A\left(T_{1}, T_{2}, T_{3}\right)\right\|_{F}^{2}
$$

and thus, we have $\forall X_{1} \in \mathbb{R}^{b_{1} \times k}, X_{2} \in \mathbb{R}^{b_{2} \times k}, X_{3} \in \mathbb{R}^{b_{3} \times k}$

$$
\begin{aligned}
& \left\|\sum_{i=1}^{k}\left(T_{1} V_{1} X_{1}\right)_{i} \otimes\left(T_{2} V_{2} X_{2}\right)_{i} \otimes\left(T_{3} V_{3} X_{3}\right)_{i}-A\left(T_{1}, T_{2}, T_{3}\right)\right\|_{F}^{2} \\
= & (1 \pm \epsilon)^{3}\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}\right)_{i} \otimes\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}-A\right\|_{F}^{2}
\end{aligned}
$$

Let $\widehat{V}_{i}$ denote $T_{i} V_{i}$, for each $i \in[3]$. Let $C \in \mathbb{R}^{c_{1} \times c_{2} \times c_{3}}$ denote $A\left(T_{1}, T_{2}, T_{3}\right)$. For $\alpha>1$, if

$$
\left\|\sum_{i=1}^{k}\left(\widehat{V}_{1} X_{1}^{\prime}\right)_{i} \otimes\left(\widehat{V}_{2} X_{2}^{\prime}\right)_{i} \otimes\left(\widehat{V}_{3} X_{3}^{\prime}\right)_{i}-C\right\|_{F}^{2} \leq \alpha\left\|\sum_{i=1}^{k}\left(\widehat{V}_{1} X_{1}\right)_{i} \otimes\left(\widehat{V}_{2} X_{2}\right)_{i} \otimes\left(\widehat{V}_{3} X_{3}\right)_{i}-C\right\|_{F}^{2}
$$

then

$$
\begin{aligned}
& \left\|\sum_{i=1}^{k}\left(V_{1} X_{1}^{\prime}\right)_{i} \otimes\left(V_{2} X_{2}^{\prime}\right)_{i} \otimes\left(V_{3} X_{3}^{\prime}\right)_{i}-C\right\|_{F}^{2} \\
\leq & \frac{1}{(1-\epsilon)^{3}}\left\|\sum_{i=1}^{k}\left(\widehat{V}_{1} X_{1}^{\prime}\right)_{i} \otimes\left(\widehat{V}_{2} X_{2}^{\prime}\right)_{i} \otimes\left(\widehat{V}_{3} X_{3}^{\prime}\right)_{i}-C\right\|_{F}^{2} \\
\leq & \frac{1}{(1-\epsilon)^{3}} \alpha\left\|\sum_{i=1}^{k}\left(\widehat{V}_{1} X_{1}\right)_{i} \otimes\left(\widehat{V}_{2} X_{2}\right)_{i} \otimes\left(\widehat{V}_{3} X_{3}\right)_{i}-C\right\|_{F}^{2} \\
\leq & \frac{(1+\epsilon)^{3}}{(1-\epsilon)^{3}} \alpha\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}\right)_{i} \otimes\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}-C\right\|_{F}^{2}
\end{aligned}
$$

By rescaling $\epsilon$ by a constant, we complete the proof of correctness.

Running time. According to Section 21.3.6, for each $i \in[3], T_{i} V_{i}$ can be computed in $O\left(\mathrm{nnz}\left(V_{i}\right)\right)$ time, and $A\left(T_{1}, T_{2}, T_{3}\right)$ can be computed in $O(\mathrm{nnz}(A))$ time.

By the analysis above, the proof is complete.

### 21.4.3 Tensor multiple regression

```
Algorithm 21.4 Frobenius Norm Tensor Multiple Regression
    procedure FTensorMultipleRegression \((A, U, V, d, n, k) \quad \triangleright\) Theorem 21.4.4
        \(s \leftarrow O\left(k^{2}+k / \epsilon\right)\).
        Choose \(S \in \mathbb{R}^{n^{2} \times s}\) to be a TensorSketch. \(\triangleright\) Definition 21.3.20
        Compute \(A \cdot S\).
        Compute \(B \cdot S . \quad \triangleright B=U^{\top} \odot V^{\top}\)
        \(W \leftarrow(A S)(B S)^{\dagger}\)
        return \(W\).
    end procedure
```

Theorem 21.4.4. Given matrices $A \in \mathbb{R}^{d \times n^{2}}, U, V \in \mathbb{R}^{n \times k}$, let $B \in \mathbb{R}^{k \times n^{2}}$ denote $U^{\top} \odot V^{\top}$. There exists an algorithm that takes $O(\mathrm{nnz}(A)+\mathrm{nnz}(U)+\mathrm{nnz}(V)+d$ poly $(k, 1 / \epsilon))$ time and outputs a matrix $W^{\prime} \in \mathbb{R}^{d \times k}$ such that,

$$
\left\|W^{\prime} B-A\right\|_{F}^{2} \leq(1+\epsilon) \min _{W \in \mathbb{R}^{d \times k}}\|W B-A\|_{F}^{2}
$$

Proof. We choose a TensorSketch (Definition 21.3.20) $S \in \mathbb{R}^{n^{2} \times s}$ to reduce the problem to a smaller problem,

$$
\min _{W \in \mathbb{R}^{d \times k}}\|W B S-A S\|_{F}^{2}
$$

Let $W^{\prime}$ denote the optimal solution to the above problem. Following a similar proof to that in Section 21.4.7.3, if $S$ is a $(1 \pm 1 / 2)$-subspace embedding and satisfies $\sqrt{\epsilon / k}$-approximate matrix product, then $W^{\prime}$ provides a $(1+\epsilon)$-approximation to the original problem. By Theorem 21.3.15, we have $s=O\left(k^{2}+k / \epsilon\right)$.

Running time. According to Definition 21.3.20, $B S$ can be computed in $O(\mathrm{nnz}(U)+$ $n n z(V))+\operatorname{poly}(k / \epsilon)$ time. Notice that each row of $S$ has exactly 1 nonzero entry, thus $A S$ can be computed in $O(\mathrm{nnz}(A))$ time. Since $B S \in \mathbb{R}^{k \times s}$ and $A S \in \mathbb{R}^{d \times s}, \min _{W \in \mathbb{R}^{d \times k}} \| W B S-$ $A S \|_{F}^{2}$ can be solved in $d \operatorname{poly}(s k)=d \operatorname{poly}(k / \epsilon)$ time.

### 21.4.4 Bicriteria algorithms

### 21.4.4.1 Solving a small regression problem

Lemma 21.4.5. Given tensor $A \in \mathbb{R}^{n \times n \times n}$ and three matrices $U \in \mathbb{R}^{n \times s_{1}}, V \in \mathbb{R}^{n \times s_{2}}$ and $W \in \mathbb{R}^{n \times s_{3}}$, there exists an algorithm that takes $O\left(\mathrm{nnz}(A)+n \operatorname{poly}\left(s_{1}, s_{2}, s_{3}, 1 / \epsilon\right)\right)$ time and outputs $\alpha^{\prime} \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}$ such that

$$
\begin{aligned}
& \left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} \alpha_{i, j, l}^{\prime} \cdot U_{i} \otimes V_{j} \otimes W_{l}-A\right\|_{F}^{2} \\
\leq & (1+\epsilon) \min _{\alpha \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} \alpha_{i, j, l} \cdot U_{i} \otimes V_{j} \otimes W_{l}-A\right\|_{F}^{2} .
\end{aligned}
$$

holds with probability at least . 99 .

Proof. We define $\widetilde{b} \in \mathbb{R}^{n^{3}}$ to be the vector where the $i+(j-1) n+(l-1) n^{2}$-th entry of $\widetilde{b}$ is $A_{i, j, l}$. We define $\widetilde{A} \in \mathbb{R}^{n^{3} \times s_{1} s_{2} s_{3}}$ to be the matrix where the $\left(i+(j-1) n+(l-1) n^{2}, i^{\prime}+\left(j^{\prime}-\right.\right.$ 1) $\left.s_{2}+\left(l^{\prime}-1\right) s_{2} s_{3}\right)$ entry is $U_{i^{\prime}, i} \cdot V_{j^{\prime}, j} \cdot W_{l^{\prime}, l}$. This problem is equivalent to a linear regression problem,

$$
\min _{x \in \mathbb{R}_{s_{1} s_{2} s_{3}}}\|\widetilde{A} x-\widetilde{b}\|_{2}^{2}
$$

where $\widetilde{A} \in \mathbb{R}^{n^{3} \times s_{1} s_{2} s_{3}}, \widetilde{b} \in \mathbb{R}^{n^{3}}$. Thus, it can be solved fairly quickly using recent work [CW13, MM13, NN13a]. However, the running time of this naïvely is $\Omega\left(n^{3}\right)$, since we have to write down each entry of $\widetilde{A}$. In the next few paragraphs, we show how to improve the running time to $\operatorname{nnz}(A)+n \operatorname{poly}\left(s_{1}, s_{2}, s_{3}\right)$.

Since $\alpha \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}, \alpha$ can be always written as $\alpha=X_{1} \otimes X_{2} \otimes X_{3}$, where $X_{1} \in$
$\mathbb{R}^{s_{1} \times s_{1} s_{2} s_{3}}, X_{2} \in \mathbb{R}^{s_{2} \times s_{1} s_{2} s_{3}}, X_{3} \in \mathbb{R}^{s_{3} \times s_{1} s_{2} s_{3}}$, we have

$$
\begin{aligned}
& \min _{\alpha \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} \alpha_{i, j, l} \cdot U_{i} \otimes V_{j} \otimes W_{l}-A\right\|_{F}^{2} \\
& =\min _{\substack{X_{1} \in \mathbb{R}^{s_{1} \times s_{1} s_{2} s_{3}} \\
X_{2} \in \in^{s} s_{2} s_{1} s_{2} s_{3} \\
X_{3} \in \mathbb{R}^{s} s_{3} \times s_{1} s_{2} s_{3}}}\left\|\left(U X_{1}\right) \otimes\left(V X_{2}\right) \otimes\left(W X_{3}\right)-A\right\|_{F}^{2} .
\end{aligned}
$$

By Lemma 21.4.3, we can reduce the problem size $n \times n \times n$ to a smaller problem that has size $t_{1} \times t_{2} \times t_{3}$,

$$
\min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{s_{1} s_{2} s_{3}}\left(T_{1} U X_{1}\right)_{i} \otimes\left(T_{2} V X_{2}\right)_{i} \otimes\left(T_{3} W X_{3}\right)_{i}-A\left(T_{1}, T_{2}, T_{3}\right)\right\|_{F}^{2}
$$

where $T_{1} \in \mathbb{R}^{t_{1} \times n}, T_{2} \in \mathbb{R}^{t_{2} \times n}, T_{3} \in \mathbb{R}^{t_{3} \times n}, t_{1}=t_{2}=t_{3}=\operatorname{poly}\left(s_{1} s_{2} s_{3} / \epsilon\right)$. Notice that

$$
\begin{aligned}
& \min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{s_{1} s_{2} s_{3}}\left(T_{1} U X_{1}\right)_{i} \otimes\left(T_{2} V X_{2}\right)_{i} \otimes\left(T_{3} W X_{3}\right)_{i}-A\left(T_{1}, T_{2}, T_{3}\right)\right\|_{F}^{2} \\
= & \min _{\alpha \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} \alpha_{i, j, l} \cdot\left(T_{1} U\right)_{i} \otimes\left(T_{2} V\right)_{j} \otimes\left(T_{3} W\right)_{l}-A\left(T_{1}, T_{2}, T_{3}\right)\right\|_{F}^{2} .
\end{aligned}
$$

Let

$$
\alpha^{\prime}=\underset{\alpha \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}}{\arg \min }\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} \alpha_{i, j, l} \cdot\left(T_{1} U\right)_{i} \otimes\left(T_{2} V\right)_{j} \otimes\left(T_{3} W\right)_{l}-A\left(T_{1}, T_{2}, T_{3}\right)\right\|_{F}^{2}
$$

then we have

$$
\begin{aligned}
& \left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} \alpha_{i, j, l}^{\prime} \cdot U_{i} \otimes V_{j} \otimes W_{l}-A\right\|_{F}^{2} \\
\leq & (1+\epsilon) \min _{\alpha \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} \alpha_{i, j, l} \cdot U_{i} \otimes V_{j} \otimes W_{l}-A\right\|_{F}^{2} .
\end{aligned}
$$

Again, according to Lemma 21.4.3, the total running time is then $O\left(\mathrm{nnz}(A)+n\right.$ poly $\left.\left(s_{1}, s_{2}, s_{3}, 1 / \epsilon\right)\right)$.

Lemma 21.4.6. Given tensor $A \in \mathbb{R}^{n \times n \times n}$, and two matrices $U \in \mathbb{R}^{n \times s}, V \in \mathbb{R}^{n \times s}$ with $\operatorname{rank}(U)=r_{1}, \operatorname{rank}(V)=r_{2}$, let $T_{1} \in \mathbb{R}^{t_{1} \times n}, T_{2} \in \mathbb{R}^{t_{2} \times n}$ be two sparse embedding matrices (Definition 21.3.10) with $t_{1}=\operatorname{poly}\left(r_{1} / \epsilon\right), t_{2}=\operatorname{poly}\left(r_{2} / \epsilon\right)$. Then with probability at least 0.99, $\forall X \in \mathbb{R}^{n \times s}$,
$(1-\epsilon)\|U \otimes V \otimes X-A\|_{F}^{2} \leq\left\|T_{1} U \otimes T_{2} V \otimes X-A\left(T_{1}, T_{2}, I\right)\right\|_{F}^{2} \leq(1+\epsilon)\|U \otimes V \otimes X-A\|_{F}^{2}$.

Proof. Let $X \in \mathbb{R}^{n \times s}$. We define $Z_{1}=\left(V^{\top} \odot X^{\top}\right) \in \mathbb{R}^{s \times n^{2}}$. We choose a sparse embedding matrix (Definition 21.3.10) $T_{1} \in \mathbb{R}^{t_{1} \times n}$ with $t_{1}=\operatorname{poly}\left(r_{1} / \epsilon\right)$ rows. According to Lemma 21.3.7 with probability 0.999 , for all $Z \in \mathbb{R}^{s \times n^{2}}$,

$$
(1-\epsilon)\left\|U Z-A_{1}\right\|_{F}^{2} \leq\left\|T_{1} U Z-T_{1} A_{1}\right\|_{F}^{2} \leq(1+\epsilon)\left\|T_{1} U Z-A_{1}\right\|_{F}^{2}
$$

It means that

$$
(1-\epsilon)\left\|U Z_{1}-A_{1}\right\|_{F}^{2} \leq\left\|T_{1} U Z_{1}-T_{1} A_{1}\right\|_{F}^{2} \leq(1+\epsilon)\left\|T_{1} U Z_{1}-A_{1}\right\|_{F}^{2} .
$$

Second, we unflatten matrix $T_{1} A_{1} \in \mathbb{R}^{t_{1} \times n^{2}}$ to obtain a tensor $A^{\prime} \in \mathbb{R}^{t_{1} \times n \times n}$. Then we flatten $A^{\prime}$ along the second direction to obtain $A_{2}^{\prime} \in \mathbb{R}^{n \times t_{1} n}$. We define $Z_{2}=\left(\left(T_{1} U\right)^{\top} \odot X^{\top}\right) \in \mathbb{R}^{s \times t_{1} n}$. Then, by flattening,

$$
\left\|V \cdot Z_{2}-A_{2}^{\prime}\right\|_{F}^{2}=\left\|T_{1} U \cdot Z_{1}-T_{1} A_{1}\right\|_{F}^{2}=(1 \pm \epsilon)\|U \otimes V \otimes X-A\|_{F}^{2}
$$

We choose a sparse embedding matrix (Definition 21.3.10) $T_{2} \in \mathbb{R}^{t_{2} \times n}$ with $t_{2}=\operatorname{poly}\left(r_{2} / \epsilon\right)$ rows. Then according to Lemma 21.3 .7 with probability 0.999 , for all $Z \in \mathbb{R}^{s \times t_{1} n}$,

$$
(1-\epsilon)\left\|V Z-A_{2}^{\prime}\right\|_{F}^{2} \leq\left\|T_{2} V Z-T_{2} A_{2}^{\prime}\right\|_{F}^{2} \leq(1+\epsilon)\left\|V Z-A_{2}^{\prime}\right\|_{F}^{2} .
$$

Thus,

$$
\left\|T_{2} V \cdot Z_{2}-T_{2} A_{2}^{\prime}\right\|_{F}^{2}=(1 \pm \epsilon)^{2}\|U \otimes V \otimes X-A\|_{F}^{2}
$$

After rescaling $\epsilon$ by a constant, with probability at least $0.99, \forall X \in \mathbb{R}^{n \times s}$,

$$
(1-\epsilon)\|U \otimes V \otimes X-A\|_{F}^{2} \leq\left\|T_{1} U \otimes T_{2} V \otimes X-A\left(T_{1}, T_{2}, I\right)\right\|_{F}^{2} \leq(1+\epsilon)\|U \otimes V \otimes X-A\|_{F}^{2}
$$

### 21.4.4.2 Algorithm I

We start with a slightly unoptimized bicriteria low rank approximation algorithm.

```
Algorithm 21.5 Frobenius Norm Bicriteria Low Rank Approximation Algorithm, rank-
\(O\left(k^{3} / \epsilon^{3}\right)\)
    procedure FTEnsorLowRankBicriteriaCubicRank \((A, n, k) \quad \triangleright\) Theorem 21.4.7
        \(s_{1} \leftarrow s_{2} \leftarrow s_{3} \leftarrow O(k / \epsilon)\).
        \(t_{1} \leftarrow t_{2} \leftarrow t_{3} \leftarrow \operatorname{poly}(k / \epsilon)\).
        Choose \(S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}\) to be a Sketching matrix, \(\forall i \in[3]\). \(\triangleright\) Definition 21.3.12
        Choose \(T_{i} \in \mathbb{R}^{t_{i} \times n}\) to be a Sketching matrix, \(\forall i \in[3]\). \(\triangleright\) Definition 21.3.10
        Compute \(U \leftarrow T_{1} \cdot\left(A_{1} \cdot S_{1}\right), V \leftarrow T_{2} \cdot\left(A_{2} \cdot S_{2}\right)\), \(W \leftarrow T_{3} \cdot\left(A_{3} \cdot S_{3}\right)\).
        Compute \(C \leftarrow A\left(T_{1}, T_{2}, T_{3}\right)\).
        \(X \leftarrow\) FTensorRegression \(\left(C, U, V, W, t_{1}, s_{1}, t_{2}, s_{2}, t_{3}, s_{3}\right)\). \(\quad\) Linear regression
        return \(X\left(A_{1} S_{1}, A_{2} S_{2}, A_{3} S_{3}\right)\).
    end procedure
```

Theorem 21.4.7. Given a 3rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1, \epsilon \in(0,1)$, let $r=O\left(k^{3} / \epsilon^{3}\right)$. There exists an algorithm that takes $O(\operatorname{nnz}(A)+n \operatorname{poly}(k, 1 / \epsilon))$ time and outputs three matrices $U \in \mathbb{R}^{n \times r}, V \in \mathbb{R}^{n \times r}, W \in \mathbb{R}^{n \times r}$ such that

$$
\left\|\sum_{i=1}^{r} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{F}^{2} \leq(1+\epsilon)_{\operatorname{rank}-k} \min _{A_{k}}\left\|A_{k}-A\right\|_{F}^{2}
$$

holds with probability 9/10.

Proof. At the end of Theorem 21.4.1, we need to run a polynomial system verifier. This is why we obtain exponential in $k$ running time. Instead of running the polynomial system verifier, we can use Lemma 21.4.5. This reduces the running time to be polynomial in all parameters: $n, k, 1 / \epsilon$. However, the output tensor has rank $(k / \epsilon)^{3}$ (Here we mean that we do not obtain a better decomposition than $(k / \epsilon)^{3}$ components). According to Section 21.3.6, for each $i, A_{i} S_{i}$ can be computed in $O(\operatorname{nnz}(A))+n \operatorname{poly}(k / \epsilon)$ time. Then $T_{i}\left(A_{i} S_{i}\right)$ can be computed in $n \operatorname{poly}(k, 1 / \epsilon)$ time and $A\left(T_{1}, T_{2}, T_{3}\right)$ also can be computed in $O(\mathrm{nnz}(A))$ time. The running time for the regression is $\operatorname{poly}(k / \epsilon)$.

Now we present an optimized bicriteria algorithm.
Theorem 21.4.8. Given a 3rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1, \epsilon \in(0,1)$, let $r=O\left(k^{2} / \epsilon^{2}\right)$. There exists an algorithm that takes $O(\operatorname{nnz}(A)+n \operatorname{poly}(k, 1 / \epsilon))$ time and outputs three matrices $U \in \mathbb{R}^{n \times r}, V \in \mathbb{R}^{n \times r}, W \in \mathbb{R}^{n \times r}$ such that

$$
\left\|\sum_{i=1}^{r} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{F}^{2} \leq(1+\epsilon)_{\operatorname{rank}-k} \min _{A_{k}}\left\|A_{k}-A\right\|_{F}^{2}
$$

holds with probability 9/10.

Note that there are two different ways to implement algorithm FTEnsorLowRAnkBiCRITERIAQUADRATICRANK. We present the proofs for both of them here.

Approach I.

Proof. Let OPT $=\min _{\operatorname{rank}-k A_{k}}\left\|A_{k}-A\right\|_{F}^{2}$. According to Theorem 21.4.1, we know that there exists a sketching matrix $S_{3} \in \mathbb{R}^{n^{2} \times s_{3}}$ where $s_{3}=O(k / \epsilon)$, such that

$$
\min _{X_{1} \in \mathbb{R}^{s_{1} \times k}, X_{2} \in \mathbb{R}^{s_{2} \times k}, X_{3} \in \mathbb{R}^{s_{3} \times k}}\left\|\sum_{l=1}^{k}\left(A_{1} S_{1} X_{1}\right)_{l} \otimes\left(A_{2} S_{2} X_{2}\right)_{l} \otimes\left(A_{3} S_{3} X_{3}\right)_{l}-A\right\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}
$$

```
Algorithm 21.6 Frobenius Norm Low Rank Approximation Algorithm, rank- \(O\left(k^{2} / \epsilon^{2}\right)\)
    procedure FTEnsorLowRankBicriteriaQuadraticRank \((A, n, k) \quad \triangleright\)
    Theorem 21.4.8
        \(s_{1} \leftarrow s_{2} \leftarrow O(k / \epsilon)\).
        Choose \(S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}\) to be a sketching matrix, \(\forall i \in[3]\). \(\triangleright\) Definition 21.3.12
        Compute \(A_{1} \cdot S_{1}, A_{2} \cdot S_{2}\).
        Form \(\widehat{U}\) by using \(A_{1} S_{1}\) according to Equation (21.9).
        Form \(\widehat{V}\) by using \(A_{2} S_{2}\) according to Equation (21.10).
        \(\widehat{W} \leftarrow\) FTEnSorMultipleRegression \(\left(A, \widehat{U}, \widehat{V}, n, n, s_{1} s_{2}\right) . \quad \triangleright\) Algorithm 21.4
        return \(\widehat{U}, \widehat{V}, \widehat{W}\).
    end procedure
    procedure FTensorLowRankBicriteriaQuadraticRank \((A, n, k) \quad \triangleright\)
    Theorem 21.4.8
        \(s_{1} \leftarrow s_{2} \leftarrow O(k / \epsilon)\).
        \(t_{1} \leftarrow t_{2} \leftarrow \operatorname{poly}(k / \epsilon)\).
        Choose \(S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}\) to be a Sketching matrix, \(\forall i \in[2]\). \(\triangleright\) Definition 21.3.12
        Choose \(T_{i} \in \mathbb{R}^{t_{i} \times n}\) to be a Sketching matrix, \(\forall i \in[2]\). \(\triangleright\) Definition 21.3.10
        Form \(\widehat{U}\) by using \(A_{1} S_{1}\) according to Equation (21.9).
        Form \(\widehat{V}\) by using \(A_{2} S_{2}\) according to Equation (21.10).
        Compute \(C \leftarrow A\left(T_{1}, T_{2}, I\right)\). \(\quad \triangleright C \in \mathbb{R}^{t_{1} \times t_{2} \times n}\)
        Compute \(B \leftarrow\left(T_{1} \widehat{U}\right)^{\top} \odot\left(T_{2} \widehat{V}\right)^{\top}\).
        \(\widehat{W} \leftarrow \underset{X \in \mathbb{R}^{n \times s_{1} s_{2}}}{\arg \min }\left\|X B-C_{3}\right\|_{F}^{2}\).
        return \(\widehat{U}, \widehat{V}, \widehat{W}\).
    end procedure
```

Now we fix an $l$ and we have:

$$
\begin{aligned}
& \left(A_{1} S_{1} X_{1}\right)_{l} \otimes\left(A_{2} S_{2} X_{2}\right)_{l} \otimes\left(A_{3} S_{3} X_{3}\right)_{l} \\
= & \left(\sum_{i=1}^{s_{1}}\left(A_{1} S_{1}\right)_{i}\left(X_{1}\right)_{i, l}\right) \otimes\left(\sum_{j=1}^{s_{2}}\left(A_{2} S_{2}\right)_{j}\left(X_{2}\right)_{j, l}\right) \otimes\left(A_{3} S_{3} X_{3}\right)_{l} \\
= & \sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}}\left(A_{1} S_{1}\right)_{i} \otimes\left(A_{2} S_{2}\right)_{j} \otimes\left(A_{3} S_{3} X_{3}\right)_{l}\left(X_{1}\right)_{i, l}\left(X_{2}\right)_{j, l}
\end{aligned}
$$

Thus, we have

$$
\begin{equation*}
\min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}}\left(A_{1} S_{1}\right)_{i} \otimes\left(A_{2} S_{2}\right)_{j} \otimes\left(\sum_{l=1}^{k}\left(A_{3} S_{3} X_{3}\right)_{l}\left(X_{1}\right)_{i, l}\left(X_{2}\right)_{j, l}\right)-A\right\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT} \tag{21.7}
\end{equation*}
$$

We use matrices $A_{1} S_{1} \in \mathbb{R}^{n \times s_{1}}$ and $A_{2} S_{2} \in \mathbb{R}^{n \times s_{2}}$ to construct a matrix $B \in \mathbb{R}^{s_{1} s_{2} \times n^{2}}$ in the following way: each row of $B$ is the vector corresponding to the matrix generated by the $\otimes$ product between one column vector in $A_{1} S_{1}$ and the other column vector in $A_{2} S_{2}$, i.e.,

$$
\begin{equation*}
B^{i+(j-1) s_{1}}=\operatorname{vec}\left(\left(A_{1} S_{1}\right)_{i} \otimes\left(A_{2} S_{2}\right)_{j}\right), \forall i \in\left[s_{1}\right], j \in\left[s_{2}\right], \tag{21.8}
\end{equation*}
$$

where $\left(A_{1} S_{1}\right)_{i}$ denotes the $i$-th column of $A_{1} S_{1}$ and $\left(A_{2} S_{2}\right)_{j}$ denote the $j$-th column of $A_{2} S_{2}$.
We create matrix $\widehat{U} \in \mathbb{R}^{n \times s_{1} s_{2}}$ by copying matrix $A_{1} S_{1} s_{2}$ times, i.e.,

$$
\widehat{U}=\left[\begin{array}{llll}
A_{1} S_{1} & A_{1} S_{1} & \cdots & A_{1} S_{1} \tag{21.9}
\end{array}\right] .
$$

We create matrix $\widehat{V} \in \mathbb{R}^{n \times s_{1} s_{2}}$ by copying the $i$-th column of $A_{2} S_{2}$ a total of $s_{1}$ times, into columns $(i-1) s_{1}, \cdots, i s_{1}$ of $\widehat{V}$, for each $i \in\left[s_{2}\right]$, i.e.,

$$
\widehat{V}=\left[\begin{array}{llllllll}
\left(A_{2} S_{2}\right)_{1} & \cdots & \left(A_{2} S_{2}\right)_{1} & \left(A_{2} S_{2}\right)_{2} & \cdots & \left(A_{2} S_{2}\right)_{2} & \cdots & \left(A_{2} S_{2}\right)_{s_{2}} \tag{21.10}
\end{array} \cdots \quad\left(A_{2} S_{2}\right)_{s_{2}}\right] .
$$

Thus, we can use $\widehat{U}$ and $\widehat{V}$ to represent $B$,

$$
B=\left(\widehat{U}^{\top} \odot \widehat{V}^{\top}\right) \in \mathbb{R}^{s_{1} s_{2} \times n^{2}}
$$

According to Equation (21.7), we have:

$$
\min _{W \in \mathbb{R}^{n \times s_{1} s_{2}}}\left\|W B-A_{3}\right\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT} .
$$

Next, we want to find matrix $W \in \mathbb{R}^{n \times s_{1} s_{2}}$ by solving the following optimization problem,

$$
\min _{W \in \mathbb{R}^{n \times s_{1} s_{2}}}\left\|W B-A_{3}\right\|_{F}^{2}
$$

Note that $B$ has size $s_{1} s_{2} \times n^{2}$. Naïvely writing down $B$ already requires $\Omega\left(n^{2}\right)$ time. In order to achieve nearly linear time in $n$, we cannot write down $B$. We choose $S_{3} \in \mathbb{R}^{n_{1} n_{2} \times s_{3}}$ to be a TensorSketch (Definition 21.3.20). In order to solve multiple regression, we need to set $s_{3}=O\left(\left(s_{1} s_{2}\right)^{2}+\left(s_{1} s_{2}\right) / \epsilon\right)$. Let $\widehat{W}$ denote the optimal solution to $\left\|W B S_{3}-A_{3} S_{3}\right\|_{F}^{2}$. Then $\widehat{W}=\left(A_{3} S_{3}\right)\left(B S_{3}\right)^{\dagger}$. Since each row of $S_{3}$ has exactly 1 nonzero entry, $A_{3} S_{3}$ can be computed in $O(\operatorname{nnz}(A))$ time. Since $B=\left(\widehat{U}^{\top} \odot \widehat{V}^{\top}\right)$, according to Definition 21.3.20, $B S_{3}$ can be computed in $n \operatorname{poly}\left(s_{1} s_{2} / \epsilon\right)=n \operatorname{poly}(k / \epsilon)$ time. By Theorem 21.4.4, we have

$$
\left\|\widehat{W} B-A_{3}\right\|_{F}^{2} \leq(1+\epsilon) \min _{W \in \mathbb{R}^{n \times s_{1} s_{2}}}\left\|W B-A_{3}\right\|_{F}^{2}
$$

Thus, we have

$$
\|\widehat{U} \otimes \widehat{V} \otimes \widehat{W}-A\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}
$$

According to Definition 21.3.12, $A_{1} S_{1}, A_{2} S_{2}$ can be computed in $O(\mathrm{nnz}(A)+\operatorname{poly}(k / \epsilon))$ time. Te total running time is thus $O(\mathrm{nnz}(A)+\operatorname{poly}(k / \epsilon))$.

## Approach II.

Proof. Let OPT $=\min _{\operatorname{rank}-k A_{k}}\left\|A_{k}-A\right\|_{F}^{2}$. Choose sketching matrices (Definition 21.3.12) $S_{1} \in$ $\mathbb{R}^{n^{2} \times s_{1}}, S_{2} \in \mathbb{R}^{n^{2} \times s_{2}}, S_{3} \in \mathbb{R}^{n^{2} \times s_{3}}$, and sketching matrices (Definition 21.3.10) $T_{1} \in \mathbb{R}^{t_{1} \times n}$ and
$T_{2} \in \mathbb{R}^{t_{2} \times n}$ with $s_{1}=s_{2}=s_{3}=O(k / \epsilon), t_{1}=t_{2}=\operatorname{poly}(k / \epsilon)$. We create matrix $\widehat{U} \in \mathbb{R}^{n \times s_{1} s_{2}}$ by copying matrix $A_{1} S_{1} s_{2}$ times, i.e.,

$$
\widehat{U}=\left[\begin{array}{llll}
A_{1} S_{1} & A_{1} S_{1} & \cdots & A_{1} S_{1}
\end{array}\right] .
$$

We create matrix $\widehat{V} \in \mathbb{R}^{n \times s_{1} s_{2}}$ by copying the $i$-th column of $A_{2} S_{2}$ a total of $s_{1}$ times, into columns $(i-1) s_{1}, \cdots, i s_{1}$ of $\widehat{V}$, for each $i \in\left[s_{2}\right]$, i.e.,

$$
\left.\widehat{V}=\left[\begin{array}{llllllll}
\left(A_{2} S_{2}\right)_{1} & \cdots & \left(A_{2} S_{2}\right)_{1} & \left(A_{2} S_{2}\right)_{2} & \cdots & \left(A_{2} S_{2}\right)_{2} & \cdots & \left(A_{2} S_{2}\right)_{s_{2}}
\end{array} \cdots{ }^{\cdots} A_{2} S_{2}\right)_{s_{2}}\right] .
$$

As we proved in Approach I, we have

$$
\min _{X \in \mathbb{R}^{n \times s_{1} s_{2}}}\|\widehat{U} \otimes \widehat{V} \otimes X-A\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}
$$

Let $B=\left(\left(T_{1} \widehat{U}\right)^{\top} \odot\left(T_{2} \widehat{V}\right)^{\top}\right) \in \mathbb{R}^{s_{1} s_{2} \times t_{1} t_{2}}$, and flatten $A\left(T_{1}, T_{2}, I\right)$ along the third direction to obtain $C_{3} \in \mathbb{R}^{n \times t_{1} t_{2}}$. Let

$$
\widehat{W}=\underset{X \in \mathbb{R}^{n \times s_{1} s_{2}}}{\arg \min }\left\|T_{1} \widehat{U} \otimes T_{2} \widehat{V} \otimes X-A\left(T_{1}, T_{2}, I\right)\right\|_{F}^{2}=\underset{X \in \mathbb{R}^{n \times s_{1} s_{2}}}{\arg \min }\left\|X B-C_{3}\right\|_{F}^{2}
$$

Let

$$
W^{*}=\underset{X \in \mathbb{R}^{n \times s_{1} s_{2}}}{\arg \min }\|\widehat{U} \otimes \widehat{V} \otimes X-A\|_{F}^{2}
$$

According to Lemma 21.4.6,

$$
\begin{aligned}
& \|\widehat{U} \otimes \widehat{V} \otimes \widehat{W}-A\|_{F}^{2} \\
\leq & \frac{1}{1-\epsilon}\left\|T_{1} \widehat{U} \otimes T_{2} \widehat{V} \otimes \widehat{W}-A\left(T_{1}, T_{2}, I\right)\right\|_{F}^{2} \\
\leq & \frac{1}{1-\epsilon}\left\|T_{1} \widehat{U} \otimes T_{2} \widehat{V} \otimes W^{*}-A\left(T_{1}, T_{2}, I\right)\right\|_{F}^{2} \\
\leq & \frac{1+\epsilon}{1-\epsilon}\left\|\widehat{U} \otimes \widehat{V} \otimes W^{*}-A\right\|_{F}^{2} \\
\leq & \frac{(1+\epsilon)^{2}}{1-\epsilon} \mathrm{OPT} .
\end{aligned}
$$

According to Definition 21.3.12, $A_{1} S_{1}, A_{2} S_{2}$ can be computed in $O(\mathrm{nnz}(A)+\operatorname{poly}(k / \epsilon))$ time. The total running time is thus $O(\mathrm{nnz}(A)+\operatorname{poly}(k / \epsilon))$. Since $T_{1}, T_{2}$ are sparse embedding matrices, $T_{1} \widehat{U}, T_{2} \widehat{V}$ can be computed in $O(\operatorname{nnz}(A)+\operatorname{poly}(k / \epsilon))$ time. The total running time is in $O(\mathrm{nnz}(A)+\operatorname{poly}(k / \epsilon))$.

Theorem 21.4.9. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$ and any $0<\epsilon<1$, if $A_{k}$ exists then there is a randomized algorithm running in $\operatorname{nnz}(A)+n \cdot \operatorname{poly}(k / \epsilon)$ time which outputs a rank- $O\left(k^{2} / \epsilon^{2}\right)$ tensor $B$ for which $\|A-B\|_{F}^{2} \leq(1+\epsilon)\left\|A-A_{k}\right\|_{F}^{2}$. If $A_{k}$ does not exist, then the algorithm outputs a rank- $O\left(k^{2} / \epsilon^{2}\right)$ tensor $B$ for which $\|A-B\|_{F}^{2} \leq$ $(1+\epsilon) \mathrm{OPT}+\gamma$, where $\gamma$ is an arbitrarily small positive function of $n$. In both cases, the algorithm succeeds with probability at least 9/10.

Proof. If $A_{k}$ exists, then the proof directly follows the proof of Theorem 21.4.1 and Theorem 21.4.8. If $A_{k}$ does not exist, then for any $\gamma>0$, there exist $U^{*} \in \mathbb{R}^{n \times k}, V^{*} \in \mathbb{R}^{n \times k}, W^{*} \in$ $\mathbb{R}^{n \times k}$ such that

$$
\left\|U^{*} \otimes V^{*} \otimes W^{*}-A\right\|_{F}^{2} \leq \inf _{\operatorname{rank}-k A^{\prime}}\left\|A-A^{\prime}\right\|_{F}^{2}+\frac{1}{10} \gamma
$$

Then we just regard $U^{*} \otimes V^{*} \otimes W^{*}$ as the "best" rank $k$ approximation to $A$, and follow the same argument as in the proof of Theorem 21.4.1 and the proof of Theorem 21.4.8. We can finally output a tensor $B \in \mathbb{R}^{n \times n \times n}$ with rank- $O\left(k^{2} / \epsilon^{2}\right)$ such that

$$
\begin{aligned}
\|B-A\|_{F}^{2} & \leq(1+\epsilon)\left\|U^{*} \otimes V^{*} \otimes W^{*}-A\right\|_{F}^{2} \\
& \leq(1+\epsilon)\left(\inf _{\operatorname{rank}-k}\left\|A-A^{\prime}\right\|_{F}^{2}+\frac{1}{10} \gamma\right) \\
& \leq(1+\epsilon) \inf _{\operatorname{rank}-k A^{\prime}}\left\|A-A^{\prime}\right\|_{F}^{2}+\gamma
\end{aligned}
$$

where the first inequality follows by the proof of Theorem 21.4.1 and the proof of theorem 21.4.8. The second inequality follows by our choice of $U^{*}, V^{*}, W^{*}$. The third inequality follows since $1+\epsilon<2$ and $\gamma>0$.

### 21.4.4.3 $\operatorname{poly}(k)$-approximation to multiple regression

Lemma 21.4.10 ((1.4) and (1.9) in [RV09]). Let $s \geq k$. Let $U \in \mathbb{R}^{n \times k}$ denote a matrix that has orthonormal columns, and $S \in \mathbb{R}^{s \times n}$ denote an i.i.d. $N(0,1 / s)$ Gaussian matrix. Then $S U$ is also an $s \times k$ i.i.d. Gaussian matrix with each entry draw from $N(0,1 / s)$, and furthermore, we have with arbitrarily large constant probability,

$$
\sigma_{\max }(S U)=O(1) \text { and } \sigma_{\min }(S U)=\Omega(1 / \sqrt{s})
$$

Proof. Note that $\sqrt{s}-\sqrt{k-1}=\frac{s-k-1}{\sqrt{s}+\sqrt{k-1}}=\Omega(1 / \sqrt{s})$.

Lemma 21.4.11. Given matrices $A \in \mathbb{R}^{n \times k}, B \in \mathbb{R}^{n \times d}$, let $S \in \mathbb{R}^{s \times n}$ denote a standard Gaussian $N(0,1)$ matrix with $s=k$. Let $X^{*}=\min _{X \in \mathbb{R}^{k \times d}}\|A X-B\|_{F}$. Let $X^{\prime}=\min _{X \in \mathbb{R}^{k \times d}} \| S A X-$ $S B \|_{F}$. Then, we have that

$$
\left\|A X^{\prime}-B\right\|_{F} \leq O(\sqrt{k})\left\|A X^{*}-B\right\|_{F},
$$

holds with probability at least 0.99.

Proof. Let $X^{*} \in \mathbb{R}^{k \times d}$ denote the optimal solution such that

$$
\left\|A X^{*}-B\right\|_{F}=\min _{X \in \mathbb{R}^{k \times d}}\|A X-B\|_{F} .
$$

Consider a standard Gaussian matrix $S \in \mathbb{R}^{k \times n}$ scaled by $1 / \sqrt{k}$ with exactly $k$ rows. Then for any $X \in \mathbb{R}^{k \times d}$, by the triangle inequality, we have

$$
\|S A X-S B\|_{F} \leq\left\|S A X-S A X^{*}\right\|_{F}+\left\|S A X^{*}-S B\right\|_{F}
$$

and

$$
\|S A X-S B\|_{F} \geq\left\|S A X-S A X^{*}\right\|_{F}-\left\|S A X^{*}-S B\right\|_{F}
$$

We first show how to bound $\left\|S A X-S A X^{*}\right\|_{F}$, and then show how to bound $\left\|S A X^{*}-S B\right\|_{F}$.
Note that Lemma 21.4.10 implies the following result,
Claim 21.4.12. For any $X \in \mathbb{R}^{k \times d}$, with probability 0.999 , we have

$$
\frac{1}{\sqrt{k}}\left\|A X-A X^{*}\right\|_{F} \lesssim\left\|S A X-S A X^{*}\right\|_{F} \lesssim\left\|A X-A X^{*}\right\|_{F}
$$

Proof. First, we can write $A=U R \in \mathbb{R}^{n \times k}$ where $U \in \mathbb{R}^{n \times k}$ has orthonormal columns and $R \in \mathbb{R}^{k \times k}$. It gives,

$$
\left\|S A X-S A X^{*}\right\|_{F}=\left\|S U\left(R X-R X^{*}\right)\right\|_{F}
$$

Second, applying Lemma 21.4.10 to $S U \in \mathbb{R}^{s \times k}$ completes the proof.

Using Markov's inequality, for any fixed matrix $A X^{*}-B$, choosing a Gaussian matrix $S$, we have that

$$
\left\|S A X^{*}-S B\right\|_{F}^{2}=O\left(\left\|A X^{*}-B\right\|_{F}^{2}\right)
$$

holds with probability at least 0.999 . This is equivalent to

$$
\begin{equation*}
\left\|S A X^{*}-S B\right\|_{F}=O\left(\left\|A X^{*}-B\right\|_{F}\right) \tag{21.11}
\end{equation*}
$$

holding with probability at least 0.999.
Let $X^{\prime}=\underset{X \in \mathbb{R}^{k \times d}}{\arg \min }\|S A X-S B\|_{F}$. Putting it all together, we have

$$
\begin{array}{rr} 
& \left\|A X^{\prime}-B\right\|_{F} \\
\leq & \left\|A X^{\prime}-A X^{*}\right\|_{F}+\left\|A X^{*}-B\right\|_{F} \\
\leq O(\sqrt{k})\left\|S A X^{\prime}-S A X^{*}\right\|_{F}+\left\|A X^{*}-B\right\|_{F} & \text { by triangle inequality } \\
\leq O(\sqrt{k})\left\|S A X^{\prime}-S B\right\|_{F}+O(\sqrt{k})\left\|S A X^{*}-S B\right\|_{F}+\left\|A X^{*}-B\right\|_{F} & \text { by Claim 21.4.12 } \\
\leq O(\sqrt{k})\left\|S A X^{*}-S B\right\|_{F}+O(\sqrt{k})\left\|S A X^{*}-S B\right\|_{F}+\left\|A X^{*}-B\right\|_{F} & \text { by definition of } X^{\prime} \\
\leq O(\sqrt{k})\left\|A X^{*}-B\right\|_{F} . & \text { by Equation (21.11) } \tag{21.11}
\end{array}
$$

by Claim 21.4.12

### 21.4.4.4 Algorithm II

Theorem 21.4.13. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, let $r=k^{2}$. There exists an algorithm which takes $O(\mathrm{nnz}(A) k)+n \operatorname{poly}(k)$ time and outputs three matrices $U, V, W \in \mathbb{R}^{n \times r}$ such that,

$$
\left\|\sum_{i=1}^{r} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{F} \leq \operatorname{poly}(k) \min _{\operatorname{rank}-k A^{\prime}}\left\|A^{\prime}-A\right\|_{F}
$$

holds with probability 9/10.

Proof. Let OPT $=\min _{\operatorname{rank}-k}\left\|A^{\prime}-A\right\|_{F}$, we fix $V^{*} \in \mathbb{R}^{n \times k}, W^{*} \in \mathbb{R}^{n \times k}$ to be the optimal solution of the original problem. We use $Z_{1}=\left(V^{* \top} \odot W^{* \top}\right) \in \mathbb{R}^{k \times n^{2}}$ to denote the matrix where the $i$-th row is the vectorization of $V_{i}^{*} \otimes W_{i}^{*}$. Let $A_{1} \in \mathbb{R}^{n \times n^{2}}$ denote the matrix
obtained by flattening tensor $A \in \mathbb{R}^{n \times n \times n}$ along the first direction. Then, we have

$$
\min _{U}\left\|U Z_{1}-A_{1}\right\|_{F} \leq \mathrm{OPT}
$$

Choosing an $N(0,1 / k)$ Gaussian sketching matrix $S_{1} \in \mathbb{R}^{n^{2} \times s_{1}}$ with $s_{1}=k$, we can obtain the smaller problem,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1} S_{1}-A_{1} S_{1}\right\|_{F}
$$

Define $\widehat{U}=A_{1} S_{1}\left(Z_{1} S_{1}\right)^{\dagger}$. Define $\alpha=O(\sqrt{k})$. By Lemma 21.4.11, we have

$$
\left\|\widehat{U} Z_{1}-A_{1}\right\|_{F} \leq \alpha \mathrm{OPT}
$$

Second, we fix $\widehat{U}$ and $W^{*}$. Define $Z_{2}, A_{2}$ similarly as above. Choosing an $N(0,1 / k)$ Gaussian sketching matrix $S_{2} \in \mathbb{R}^{n^{2} \times s_{2}}$ with $s_{2}=k$, we can obtain another smaller problem,

$$
\min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2} S_{2}-A_{2} S_{2}\right\|_{F}
$$

Define $\widehat{V}=A_{2} S_{2}\left(Z_{2} S_{2}\right)^{\dagger}$. By Lemma 21.4.11 again, we have

$$
\left\|\widehat{V} Z_{2}-A_{2}\right\|_{F} \leq \alpha^{2} \mathrm{OPT}
$$

Thus, we now have

$$
\min _{X_{1}, X_{2}, W}\left\|A_{1} S_{1} X_{1} \otimes A_{2} S_{2} X_{2} \otimes W-A\right\|_{F} \leq \alpha^{2} \mathrm{OPT}
$$

We use a similar idea as in the proof of Theorem 21.4.8. We create matrix $\widetilde{U} \in \mathbb{R}^{n \times s_{1} s_{2}}$ by copying matrix $A_{1} S_{1} s_{2}$ times, i.e.,

$$
\widetilde{U}=\left[\begin{array}{llll}
A_{1} S_{1} & A_{1} S_{1} & \cdots & A_{1} S_{1}
\end{array}\right]
$$

We create matrix $\widetilde{V} \in \mathbb{R}^{n \times s_{1} s_{2}}$ by copying the $i$-th column of $A_{2} S_{2}$ a total of $s_{1}$ times, into columns $(i-1) s_{1}, \cdots, i s_{1}$ of $\widetilde{V}$, for each $i \in\left[s_{2}\right]$, i.e.,

$$
\widetilde{V}=\left[\begin{array}{lllllllll}
\left(A_{2} S_{2}\right)_{1} & \cdots & \left(A_{2} S_{2}\right)_{1} & \left(A_{2} S_{2}\right)_{2} & \cdots & \left(A_{2} S_{2}\right)_{2} & \cdots & \left(A_{2} S_{2}\right)_{s_{2}} & \cdots
\end{array}\left(A_{2} S_{2}\right)_{s_{2}}\right] .
$$

We have

$$
\min _{X \in \mathbb{R}^{n \times s_{1} s_{2}}}\|\widetilde{U} \otimes \widetilde{V} \otimes X-A\|_{F} \leq \alpha^{2} \mathrm{OPT}
$$

Choose $T_{i} \in \mathbb{R}^{t_{i} \times n}$ to be a sparse embedding matrix (Definition 21.3.10) with $t_{i}=$ $\operatorname{poly}(k / \epsilon)$, for each $i \in[2]$. By applying Lemma 21.4.6, we have, if $W^{\prime}$ satisfies,

$$
\left\|T_{1} \widetilde{U} \otimes T_{2} \widetilde{V} \otimes W^{\prime}-A\left(T_{1}, T_{2}, I\right)\right\|_{F}=\min _{X \in \mathbb{R}^{n \times s_{1} s_{2}}}\left\|T_{1} \widetilde{U} \otimes T_{2} \widetilde{V} \otimes X-A\left(T_{1}, T_{2}, I\right)\right\|_{F}
$$

then,

$$
\left\|\widetilde{U} \otimes \widetilde{V} \otimes W^{\prime}-A\right\|_{F} \leq(1+\epsilon) \min _{X \in \mathbb{R}^{n \times s_{1} s_{2}}}\|\widetilde{U} \otimes \widetilde{V} \otimes X-A\|_{F} \leq(1+\epsilon) \alpha^{2} \mathrm{OPT}
$$

Thus, we only need to solve

$$
\min _{X \in \mathbb{R}^{n \times s_{1} s_{2}}}\left\|T_{1} \widetilde{U} \otimes T_{2} \widetilde{V} \otimes X-A\left(T_{1}, T_{2}, I\right)\right\|_{F}
$$

which is similar to the proof of Theorem 21.4.8. Therefore, we complete the proof of correctness. For the running time, $A_{1} S_{1}, A_{2} S_{2}$ can be computed in $O(\mathrm{nnz}(A) k)$ time, $T_{1} \widetilde{U}, T_{2} \widetilde{V}$ can be computed in $n$ poly $(k)$ time. The final regression problem can be computed in $n$ poly $(k)$ running time.

### 21.4.5 Generalized matrix row subset selection

Note that in this section, the notation $\Pi_{C, k}^{\xi}$ is given in Definition 21.3.5.

```
Algorithm 21.7 Generalized Matrix Row Subset Selection: Constructing \(R\) with \(r=O(k+\)
\(k / \epsilon)\) Rows and a rank- \(k U \in \mathbb{R}^{k \times r}\)
    procedure GeneralizedMatrixRowSubsetSelection \((A, C, n, m, k, \epsilon) \quad \triangleright\)
    Theorem 21.4.14
        \(Y, \Phi, \Delta \leftarrow \operatorname{ApproxSubSpaceSVD}(A, C, k) . \quad\) Claim 21.4.16 and Lemma 3.12 in
    [BW14]
        \(B \leftarrow Y \Delta\).
        \(Z_{2}, D \leftarrow \operatorname{QR}(B) . \quad \triangleright Z_{2} \in \mathbb{R}^{m \times k}, Z_{2}^{\top} Z_{2}=I_{k}, D \in \mathbb{R}^{k \times k}\)
        \(h_{2} \leftarrow 8 k \ln (20 k)\).
        \(\Omega_{2}, D_{2} \leftarrow\) RandSampling \(\left(Z_{2}, h_{2}, 1\right) \quad \triangleright\) Definition 3.6 in [BW14]
        \(M_{2} \leftarrow Z_{2}^{\top} \Omega_{2} D_{2} \in \mathbb{R}^{k \times h_{2}}\).
        \(U_{M_{2}}, \Sigma_{M_{2}}, V_{M_{2}}^{\top} \leftarrow \operatorname{SVD}\left(M_{2}\right) . \quad \triangleright \operatorname{rank}\left(M_{2}\right)=k\) and \(V_{M_{2}} \in \mathbb{R}^{h_{2} \times k}\)
        \(r_{1} \leftarrow 4 k\).
        \(S_{2} \leftarrow \operatorname{BSSSAmPLINGSPARSE}\left(V_{M_{2}},\left(\left(A^{\top}-A^{\top} Z_{2} Z_{2}^{\top}\right) \Omega_{2} D_{2}\right)^{\top}, r_{1}, 0.5\right) \triangleright\) Lemma 4.3 in
    [BW14]
        \(R_{1} \leftarrow\left(A^{\top} \Omega_{2} D_{2} S_{2}\right)^{\top} \in \mathbb{R}^{r_{1 \times n}}\) containing rescaled rows from \(A\).
        \(r_{2} \leftarrow 4820 k / \epsilon\).
        \(R_{2} \leftarrow \operatorname{AdAPtiveRowsSparse}\left(A, Z_{2}, R_{1}, r_{2}\right) \quad \triangleright\) Lemma 4.5 in [BW14]
        \(R \leftarrow\left[R_{1}^{\top}, R_{2}^{\top}\right]^{\top} . \quad \triangleright R \in \mathbb{R}^{\left(r_{1}+r_{2}\right) \times n}\) containing \(r=4 k+4820 k / \epsilon\) rescaled rows of \(A\).
        Choose \(W \in \mathbb{R}^{\xi \times m}\) to be a randomly chosen sparse subspace embedding with \(\xi=\)
    \(\Omega\left(k^{2} \epsilon^{-2}\right)\).
        \(U \leftarrow \Phi^{-1} \Delta D^{-1}\left(W C \Phi^{-1} \Delta D^{-1}\right)^{\dagger} W A R^{\dagger}=\Phi^{-1} \Delta \Delta^{\top}(W C)^{\dagger} W A R^{\dagger}\).
        return \(R, U\).
    end procedure
```

Theorem 21.4.14. Given matrices $A \in \mathbb{R}^{n \times m}$ and $C \in \mathbb{R}^{n \times k}$, there exists an algorithm which takes $O(\mathrm{nnz}(A) \log n)+(m+n)$ poly $(k, 1 / \epsilon)$ time and outputs a diagonal matrix $D \in$ $\mathbb{R}^{n \times n}$ with $d=O(k / \epsilon)$ nonzeros (or equivalently a matrix $R$ that contains $d=O(k / \epsilon)$ rescaled
rows of $A$ ) and a matrix $U \in \mathbb{R}^{k \times d}$ such that

$$
\|C U D A-A\|_{F}^{2} \leq(1+\epsilon) \min _{X \in \mathbb{R}^{k \times m}}\|C X-A\|_{F}^{2}
$$

holds with probability . 99 .

Proof. This follows by combining Lemma 21.4.17 and 21.4.18. Let $U, R$ denote the output of procedure GeneralizedMatrixRowSubsetSelection,

$$
\begin{aligned}
\|A-C U R\|_{F}^{2} & \leq(1+\epsilon)\left\|A-Z_{2} Z_{2}^{\top} A R^{\dagger} R\right\|_{F}^{2} \\
& \leq(1+\epsilon)(1+60 \epsilon)\left\|A-\Pi_{C, k}^{F}(A)\right\|_{F}^{2} \\
& \leq(1+130 \epsilon)\left\|A-\Pi_{C, k}^{F}(A)\right\|_{F}^{2}
\end{aligned}
$$

Because $R$ is a subset of rows of $A$ and $R$ has size $O(k / \epsilon) \times m$, there must exist a diagonal matrix $D \in \mathbb{R}^{n \times n}$ with $O(k / \epsilon)$ nonzeros such that $R=D A$. This completes the proof.

Corollary 21.4.15 (A slightly different version of Theorem 21.4.14, faster running time, and small input matrix). Given matrices $A \in \mathbb{R}^{n \times m}$ and $C \in \mathbb{R}^{n \times k}$, if $\min (m, n)=\operatorname{poly}(k, 1 / \epsilon)$, then there exists an algorithm which takes $O(\mathrm{nnz}(A))+(m+n) \operatorname{poly}(k, 1 / \epsilon)$ time and outputs a diagonal matrix $D \in \mathbb{R}^{n \times n}$ with $d=O(k / \epsilon)$ nonzeros (or equivalently a matrix $R$ that contains $d=O(k / \epsilon)$ rescaled rows of $A)$ and a matrix $U \in \mathbb{R}^{k \times d}$ such that

$$
\|C U D A-A\|_{F}^{2} \leq(1+\epsilon) \min _{X \in \mathbb{R}^{k \times m}}\|C X-A\|_{F}^{2}
$$

holds with probability . 99 .

Proof. The $\log n$ factor comes from the adaptive sampling where we need to choose a Gaussian matrix with $O(\log n)$ rows and compute $S A$. If $A$ has poly $(k, 1 / \epsilon)$ columns, it is sufficient
to choose $S$ to be a CountSketch matrix with $\operatorname{poly}(k, 1 / \epsilon)$ rows. Then, we do not need a $\log n$ factor in the running time. If $S$ has $\operatorname{poly}(k, 1 / \epsilon)$ rows, then we no longer need the matrix $S$.

Claim 21.4.16. Given matrices $A \in \mathbb{R}^{m \times n}$ and $C \in \mathbb{R}^{m \times c}$, let $Y \in \mathbb{R}^{m \times c}, \Phi \in \mathbb{R}^{c \times c}$ and $\Delta \in \mathbb{R}^{c \times k}$ denote the output of procedure $\operatorname{ApproxSubspaceSVD}(A, C, k, \epsilon)$. Then with probability .99, we have,

$$
\left\|A-Y \Delta \Delta^{\top} Y^{\top} A\right\|_{F}^{2} \leq(1+30 \epsilon)\left\|A-\Pi_{C, k}^{F}(A)\right\|_{F}^{2}
$$

Proof. This follows by Lemma 3.12 in [BW14].
Lemma 21.4.17. The matrices $R$ and $Z_{2}$ in procedure GeneralizedMatrixRowSubSETSELECTION (Algorithm 21.7) satisfy with probability at least $0.17-2 / n$,

$$
\left\|A-Z_{2} Z_{2}^{\top} A R^{\dagger} R\right\|_{F}^{2} \leq\left\|A-\Pi_{C, k}^{F}(A)\right\|_{F}^{2}+60 \epsilon\left\|A-\Pi_{C, k}^{F}(A)\right\|_{F}^{2}
$$

Proof. We can show,

$$
\begin{aligned}
& \left\|A-Z_{2} Z_{2}^{\top} A\right\|_{F}^{2}+\frac{30 \epsilon}{4820}\left\|A-A R_{1}^{\dagger} R_{1}\right\|_{F}^{2} \\
= & \left\|A-B B^{\dagger} A\right\|_{F}^{2}+\frac{30 \epsilon}{4820}\left\|A-A R_{1}^{\dagger} R_{1}\right\|_{F}^{2} \\
\leq & \left\|A-B B^{\dagger} A\right\|_{F}^{2}+30 \epsilon\left\|A-A_{k}\right\|_{F}^{2} \\
\leq & \left\|A-Y \Delta \Delta^{\top} Y A\right\|_{F}^{2}+30 \epsilon\left\|A-\Pi_{C, k}^{F}(A)\right\|_{F}^{2} \\
\leq & (1+30 \epsilon)\left\|A-\Pi_{C, k}^{F}(A)\right\|_{F}^{2}+30 \epsilon\left\|A-\Pi_{C, k}^{F}(A)\right\|_{F}^{2},
\end{aligned}
$$

where the first step follows by the fact that $Z_{2} Z_{2}^{\top}=Z_{2} D D^{-1} Z_{2}^{\top}=\left(Z_{2} D\right)\left(Z_{2} D\right)^{\dagger}=B B^{\dagger}$, the second step follows by $\left\|A-A R_{1}^{\dagger} R_{1}\right\|_{F}^{2} \leq 4820\left\|A-A_{k}\right\|_{F}^{2}$, the third step follows by $B=Y \Delta$ and $B^{\dagger}=(Y \Delta)^{\dagger}=\Delta^{\dagger} Y^{\dagger}=\Delta^{\top} Y^{\top}$, and the last step follows by Claim 21.4.16.

Lemma 21.4.18. The matrices $C, U$ and $R$ in procedure GeneralizedMatrixRowSubSETSELECTION (Algorithm 21.7) satisfy that

$$
\|A-C U R\|_{F}^{2} \leq(1+\epsilon)\left\|A-Z_{2} Z_{2}^{\top} A R^{\dagger} R\right\|_{F}^{2}
$$

with probability at least .99.

Proof. Let $U_{R}, \Sigma_{R}, V_{R}$ denote the SVD of $R$. Then $V_{R} V_{R}^{\top}=R^{\dagger} R$.
We define $Y^{*}$ to be the optimal solution of

$$
\min _{X \in \mathbb{R}^{k \times r}}\left\|W A V_{R} V_{R}^{\top}-W C \Phi^{-1} \Delta D^{-1} Y R\right\|_{F}^{2}
$$

We define $\widehat{X}^{*}$ to be $Y^{*} R \in \mathbb{R}^{k \times n}$, which is also equivalent to defining $\widehat{X}^{*}$ to be the optimal solution of

$$
\min _{X \in \mathbb{R}^{k \times n}}\left\|W A V_{R} V_{R}^{\top}-W C \Phi^{-1} \Delta D^{-1} X\right\|_{F}^{2}
$$

Furthermore, it implies $\widehat{X}^{*}=\left(W C \Phi^{-1} \Delta D^{-1}\right)^{\dagger} W A V_{R} V_{R}^{\dagger}$.
We also define $X^{*}$ to be the optimal solution of

$$
\min _{X \in \mathbb{R}^{k \times n}}\left\|A V_{R} V_{R}^{\dagger}-C \Phi^{-1} \Delta D^{-1} X\right\|_{F}^{2}
$$

which implies that,

$$
X^{*}=\left(C \Phi^{-1} \Delta D^{-1}\right)^{\dagger} A V_{R} V_{R}^{\top}=Z_{2}^{\top} A V_{R} V_{R}^{\top}
$$

Now, we start to prove an upper bound on $\|A-C U R\|_{F}^{2}$,

$$
\begin{array}{rlr}
\|A-C U R\|_{F}^{2} & =\left\|A-C \Phi^{-1} \Delta D^{-1} Y^{*} R\right\|_{F}^{2} & \text { by definition of } U \\
& =\left\|A-C \Phi^{-1} \Delta D^{-1} \widehat{X}^{*}\right\|_{F}^{2} & \text { by } \widehat{X}^{*}=Y^{*} R \\
& =\left\|A V_{R} V_{R}^{\top}-C \Phi^{-1} \Delta D^{-1} \widehat{X}^{*}+A-A V_{R} V_{R}^{\top}\right\|_{F}^{2} & \\
& =\underbrace{\left\|A V_{R} V_{R}^{\top}-C \Phi^{-1} \Delta D^{-1} \widehat{X}^{*}\right\|_{F}^{2}}_{\alpha}+\underbrace{\left\|A-A V_{R} V_{R}^{\top}\right\|_{F}^{2}}_{\beta}, \tag{21.12}
\end{array}
$$

where the last step follows by $\widehat{X}^{*}=M V_{R}^{\top}, A-A V_{R} V_{R}^{\top}=A\left(I-V_{R} V_{R}^{\top}\right)$ and the Pythagorean theorem. We show how to upper bound the term $\alpha$,

$$
\begin{align*}
\alpha & \leq(1+\epsilon)\left\|A V_{R} V_{R}^{\top}-C \Phi^{-1} \Delta D^{-1} X^{*}\right\|_{F}^{2} \\
& =\epsilon\left\|A V_{R} V_{R}^{\top}-C \Phi^{-1} \Delta D^{-1} X^{*}\right\|_{F}^{2}+\left\|A V_{R} V_{R}^{\top}-C \Phi^{-1} \Delta D^{-1} X^{*}\right\|_{F}^{2} \\
& =\epsilon\left\|A V_{R} V_{R}^{\top}-C \Phi^{-1} \Delta D^{-1} X^{*}\right\|_{F}^{2}+\left\|A V_{R} V_{R}^{\top}-C \Phi^{-1} \Delta D^{-1}\left(Z_{2}^{\top} A R^{\dagger} R\right)\right\|_{F}^{2} . \tag{21.13}
\end{align*}
$$

by Lemma 21.4.19

By the Pythagorean theorem and the definition of $Z_{2}$ (which means $Z_{2}=C \Phi^{-1} \Delta D^{-1}$ ), we have,

$$
\begin{align*}
& \left\|A V_{R} V_{R}^{\top}-C \Phi^{-1} \Delta D^{-1}\left(Z_{2}^{\top} A R^{\dagger} R\right)\right\|_{F}^{2}+\beta \\
= & \left\|A V_{R} V_{R}^{\top}-C \Phi^{-1} \Delta D^{-1}\left(Z_{2}^{\top} A R^{\dagger} R\right)\right\|_{F}^{2}+\left\|A-A V_{R} V_{R}^{\top}\right\|_{F}^{2} \\
= & \left\|A-C \Phi^{-1} \Delta D^{-1}\left(Z_{2}^{\top} A R^{\dagger} R\right)\right\|_{F}^{2} \\
= & \left\|A-Z_{2} Z_{2}^{\top} A R^{\dagger} R\right\|_{F}^{2} . \tag{21.14}
\end{align*}
$$

Combining Equations (21.12), (21.13) and (21.14) together, we obtain,

$$
\|A-C U R\|_{F}^{2} \leq \epsilon\left\|A V_{R} V_{R}^{\top}-C \Phi^{-1} \Delta D^{-1} X^{*}\right\|_{F}^{2}+\left\|A-Z_{2} Z_{2}^{\top} A R^{\dagger} R\right\|_{F}^{2} .
$$

We want to show $\left\|A V_{R} V_{R}^{\top}-C \Phi^{-1} \Delta D^{-1} X^{*}\right\|_{F}^{2} \leq\left\|A-Z_{2} Z_{2}^{\top} A R^{\dagger} R\right\|_{F}^{2}$,

$$
\begin{aligned}
& \left\|A V_{R} V_{R}^{\top}-C \Phi^{-1} \Delta D^{-1} X^{*}\right\|_{F}^{2} \\
= & \left\|A V_{R} V_{R}^{\top}-C \Phi^{-1} \Delta D^{-1} Z_{2}^{\top} A V_{R} V_{R}^{\top}\right\|_{F}^{2} \\
\leq\left\|A-C \Phi^{-1} \Delta D^{-1} Z_{2}^{\top} A\right\|_{F}^{2} & \text { by } X^{*}=Z_{2}^{\top} A V_{R} V_{R}^{\top} \\
\leq & \left\|A-C \Phi^{-1} \Delta D^{-1} Z_{2}^{\top} A R^{\dagger} R\right\|_{F}^{2} \\
= & \text { by properties of projections } \\
= & \text { by properties of projections } \\
Z_{2} Z_{2}^{\top} A R^{\dagger} R \|_{F}^{2} . & \text { by } Z_{2}=C \Phi^{-1} \Delta D^{-1}
\end{aligned}
$$

This completes the proof.
Lemma 21.4.19 ([CW13]). Let $A \in \mathbb{R}^{n \times d}$ have rank $\rho$ and $B \in \mathbb{R}^{n \times r}$. Let $W \in \mathbb{R}^{r \times n}$ be a randomly chosen sparse subspace embedding with $r=\Omega\left(\rho^{2} \epsilon^{-2}\right)$. Let $\widehat{X}^{*}=\underset{X \in \mathbb{R}^{d \times r}}{\arg \min } \| W A X-$ $W B \|_{F}^{2}$ and let $X^{*}=\underset{X \in \mathbb{R}^{d \times r}}{\arg \min }\|A X-B\|_{F}^{2}$. Then with probability at least .99,

$$
\left\|A \widetilde{X}^{*}-B\right\|_{F}^{2} \leq(1+\epsilon)\left\|A X^{*}-B\right\|_{F}^{2}
$$

### 21.4.6 Column, row, and tube subset selection, ( $1+\epsilon$ )-approximation

```
Algorithm 21.8 Frobenius Norm Tensor Column, Row and Tube Subset Selection, Poly-
nomial Time
    procedure \(\operatorname{FCRTSELECTION}(A, n, k, \epsilon) \quad \triangleright\) Theorem 21.4.20
        \(s_{1} \leftarrow s_{2} \leftarrow O(k / \epsilon)\).
        Choose a Gaussian matrix \(S_{1}\) with \(s_{1}\) columns. \(\triangleright\) Definition 21.3.12
        Choose a Gaussian matrix \(S_{2}\) with \(s_{2}\) columns. \(\triangleright\) Definition 21.3.12
        Form matrix \(Z_{3}^{\prime}\) by setting the \((i, j)\)-th row to be the vectorization of \(\left(A_{1} S_{1}\right)_{i} \otimes\)
    \(\left(A_{2} S_{2}\right)_{j}\).
    \(D_{3} \leftarrow\) GEnERALIZEDMATRIXROWSUBSETSELECTION \(\left(A_{3}^{\top},\left(Z_{3}^{\prime}\right)^{\top}, n^{2}, n, s_{1} s_{2}, \epsilon\right) . \quad \triangleright\)
    Algorithm 21.7
    Let \(d_{3}\) denote the number of nonzero entries in \(D_{3}\). \(\triangleright d_{3}=O\left(s_{1} s_{2} / \epsilon\right)\)
    Form matrix \(Z_{2}^{\prime}\) by setting the \((i, j)\)-th row to be the vectorization of \(\left(A_{1} S_{1}\right)_{i} \otimes\)
    \(\left(A_{3} S_{3}^{\prime}\right)_{j}\).
    \(D_{2} \leftarrow\) GeneralizedMatrixRowSubsetSelection \(\left(A_{2}^{\top},\left(Z_{2}^{\prime}\right)^{\top}, n^{2}, n, s_{1} d_{3}, \epsilon\right)\).
    Let \(d_{2}\) denote the number of nonzero entries in \(D_{2}\). \(\quad \triangleright d_{2}=O\left(s_{1} d_{3} / \epsilon\right)\)
    Form matrix \(Z_{1}^{\prime}\) by setting the \((i, j)\)-th row to be the vectorization of \(\left(A_{2} D_{2}\right)_{i} \otimes\)
    \(\left(A_{3} D_{3}\right)_{j}\).
    \(D_{1} \leftarrow\) GeneralizedMatrixRowSubsetSelection \(\left(A_{1}^{\top},\left(Z_{1}^{\prime}\right)^{\top}, n^{2}, n, d_{2} d_{3}, \epsilon\right)\).
    Let \(d_{1}\) denote the number of nonzero entries in \(D_{1}\). \(\quad \triangleright d_{1}=O\left(d_{2} d_{3} / \epsilon\right)\)
    \(C \leftarrow A_{1} D_{1}, R \leftarrow A_{2} D_{2}\) and \(T \leftarrow A_{3} D_{3}\).
    return \(C, R\) and \(T\).
    end procedure
```

We provide two bicriteria CURT results in this Section. We first present a warm-up result. That result (Theorem 21.4.20) does not output tensor $U$ and only guarantees that there is a rank-poly $(k / \epsilon)$ tensor $U$. Then we show the second result (Theorem 21.4.21), our second result is able to output tensor $U$. The $U$ has rank poly $(k / \epsilon)$, but not $k$.

Theorem 21.4.20. Given a 3rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, there exists an algorithm which takes $O(\mathrm{nnz}(A))+n$ poly $(k, 1 / \epsilon)$ time and outputs three matrices: $C \in \mathbb{R}^{n \times c}$, a subset of columns of $A, R \in \mathbb{R}^{n \times r}$ a subset of rows of $A$, and $T \in \mathbb{R}^{n \times t}$, a subset of tubes
of $A$ where $c=r=t=\operatorname{poly}(k, 1 / \epsilon)$, and there exists a tensor $U \in \mathbb{R}^{c \times r \times t}$ such that

$$
\left\|\left(\left(\left(U \cdot T^{\top}\right)^{\top} \cdot R^{\top}\right)^{\top} \cdot C^{\top}\right)^{\top}-A\right\|_{F}^{2} \leq(1+\epsilon) \min _{\text {rank }-k A_{k}}\left\|A_{k}-A\right\|_{F}^{2}
$$

or equivalently,

$$
\left\|\sum_{i=1}^{c} \sum_{j=1}^{r} \sum_{l=1}^{t} U_{i, j, l} \cdot C_{i} \otimes R_{j} \otimes T_{l}-A\right\|_{F}^{2} \leq(1+\epsilon)_{\operatorname{rank}-k A_{k}}\left\|A_{k}-A\right\|_{F}^{2}
$$

holds with probability 9/10.

Proof. We mainly analyze Algorithm 21.8 and it is easy to extend to Algorithm 21.9.
We fix $V^{*} \in \mathbb{R}^{n \times k}$ and $W^{*} \in \mathbb{R}^{n \times k}$. We define $Z_{1} \in \mathbb{R}^{k \times n^{2}}$ where the $i$-th row of $Z_{1}$ is the vector $V_{i} \otimes W_{i}$. Choose sketching (Gaussian) matrix $S_{1} \in \mathbb{R}^{n^{2} \times s_{1}}$ (Definition 21.3.12), and let $\widehat{U}=A_{1} S_{1}\left(Z_{1} S_{1}\right)^{\dagger} \in \mathbb{R}^{n \times k}$. Following a similar argument as in the previous theorem, we have

$$
\left\|\widehat{U} Z_{1}-A_{1}\right\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}
$$

We fix $\widehat{U}$ and $W^{*}$. We define $Z_{2} \in \mathbb{R}^{k \times n^{2}}$ where the $i$-th row of $Z_{2}$ is the vector $\widehat{U}_{i} \otimes$ $W_{i}^{*}$. Choose sketching (Gaussian) matrix $S_{2} \in \mathbb{R}^{n^{2} \times s_{2}}$ (Definition 21.3.12), and let $\widehat{V}=$ $A_{2} S_{2}\left(Z_{2} S_{2}\right)^{\dagger} \in \mathbb{R}^{n \times k}$. Following a similar argument as in the previous theorem, we have

$$
\left\|\widehat{V} Z_{2}-A_{2}\right\|_{F}^{2} \leq(1+\epsilon)^{2} \mathrm{OPT}
$$

We fix $\widehat{U}$ and $\widehat{V}$. Note that $\widehat{U}=A_{1} S_{1}\left(Z_{1} S_{1}\right)^{\dagger}$ and $\widehat{V}=A_{2} S_{2}\left(Z_{2} S_{2}\right)^{\dagger}$. We define $Z_{3} \in \mathbb{R}^{k \times n^{2}}$ such that the $i$-th row of $Z_{3}$ is the vector $\widehat{U}_{i} \otimes \widehat{V}_{i}$. Let $z_{3}=s_{1} \cdot s_{2}$. We define
$Z_{3}^{\prime} \in \mathbb{R}^{z_{3} \times n^{2}}$ such that, $\forall i \in\left[s_{1}\right], \forall j \in\left[s_{2}\right]$, the $i+(j-1) s_{1}$-th row of $Z_{3}^{\prime}$ is the vector $\left(A_{1} S_{1}\right)_{i} \otimes\left(A_{2} S_{2}\right)_{j}$. We consider the following objective function,

$$
\min _{W \in \mathbb{R}^{n \times k}, X \in \mathbb{R}^{k \times z_{3}}}\left\|W X Z_{3}^{\prime}-A_{3}\right\|_{F}^{2} \leq \min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3}-A_{3}\right\|_{F}^{2} \leq(1+\epsilon)^{2} \mathrm{OPT}
$$

Using Theorem 21.4.14, we can find a diagonal matrix $D_{3} \in \mathbb{R}^{n^{2} \times n^{2}}$ with $d_{3}=$ $O\left(z_{3} / \epsilon\right)=O\left(k^{2} / \epsilon^{3}\right)$ nonzero entries such that

$$
\min _{X \in \mathbb{R}^{d_{3} \times z_{3}}}\left\|A_{3} D_{3} X Z_{3}^{\prime}-A_{3}\right\|_{F}^{2} \leq(1+\epsilon)^{3} \mathrm{OPT}
$$

In the following, we abuse notation and let $A_{3} D_{3} \in \mathbb{R}^{n \times d_{3}}$ by deleting zero columns. Let $W^{\prime}$ denote $A_{3} D_{3} \in \mathbb{R}^{n \times d_{3}}$. Then,

$$
\min _{X \in \mathbb{R}^{d_{3} \times z_{3}}}\left\|W^{\prime} X Z_{3}^{\prime}-A_{3}\right\|_{F}^{2} \leq(1+\epsilon)^{3} \mathrm{OPT}
$$

We fix $\widehat{U}$ and $W^{\prime}$. Let $z_{2}=s_{1} \cdot d_{3}$. We define $Z_{2}^{\prime} \in \mathbb{R}^{z_{2} \times n^{2}}$ such that, $\forall i \in\left[s_{1}\right], \forall j \in\left[d_{3}\right]$, the $i+(j-1) s_{1}$-th row of $Z_{2}^{\prime}$ is the vector $\left(A_{1} S_{1}\right)_{i} \otimes\left(A_{3} D_{3}\right)_{j}$.

Using Theorem 21.4.14, we can find a diagonal matrix $D_{2} \in \mathbb{R}^{n^{2} \times n^{2}}$ with $d_{2}=$ $O\left(z_{2} / \epsilon\right)=O\left(s_{1} d_{3} / \epsilon\right)=O\left(k^{3} / \epsilon^{5}\right)$ nonzero entries such that

$$
\min _{X \in \mathbb{R}^{d_{2} \times z_{2}}}\left\|A_{2} D_{2} X Z_{2}^{\prime}-A_{2}\right\|_{F}^{2} \leq(1+\epsilon)^{4} \mathrm{OPT}
$$

Let $V^{\prime}$ denote $A_{2} D_{2}$. Then,

$$
\min _{X \in \mathbb{R}^{d_{2} \times z_{2}}}\left\|V^{\prime} X Z_{2}^{\prime}-A_{2}\right\|_{F}^{2} \leq(1+\epsilon)^{4} \mathrm{OPT}
$$

We fix $V^{\prime}$ and $W^{\prime}$. Let $z_{1}=d_{2} \cdot d_{3}$. We define $Z_{1}^{\prime} \in \mathbb{R}^{z_{1} \times n^{2}}$ such that, $\forall i \in\left[d_{2}\right], \forall j \in$ $\left[d_{3}\right]$, the $i+(j-1) s_{1}$-th row of $Z_{1}^{\prime}$ is the vector $\left(A_{2} D_{2}\right)_{i} \otimes\left(A_{3} D_{3}\right)_{j}$.

Using Theorem 21.4.14, we can find a diagonal matrix $D_{1} \in \mathbb{R}^{n^{2} \times n^{2}}$ with $d_{1}=$ $O\left(z_{1} / \epsilon\right)=O\left(d_{2} d_{3} / \epsilon\right)=O\left(k^{5} / \epsilon^{9}\right)$ nonzero entries such that

$$
\min _{X \in \mathbb{R}^{d_{1} \times z_{1}}}\left\|A_{1} D_{1} X Z_{1}^{\prime}-A_{1}\right\|_{F}^{2} \leq(1+\epsilon)^{5} \mathrm{OPT}
$$

Let $U^{\prime}$ denote $A_{1} D_{1}$. Then,

$$
\min _{X \in \mathbb{R}^{d_{1} \times z_{1}}}\left\|U^{\prime} X Z_{1}^{\prime}-A_{1}\right\|_{F}^{2} \leq(1+\epsilon)^{5} \text { OPT } .
$$

Putting $U^{\prime}, V^{\prime}, W^{\prime}$ all together, we complete the proof.
All the above analysis gives the running time $O(\mathrm{nnz}(A)) \log n+n^{2}$ poly $(\log n, k, 1 / \epsilon)$. To improve the running time, we need to use Algorithm 21.9, the similar analysis will go through, the running time will be improved to $O(\operatorname{nnz}(A)+n$ poly $(k, 1 / \epsilon))$, but the sample complexity of $c, r, k$ will be slightly worse (poly $\log$ factors).

Theorem 21.4.21. Given a 3rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, there exists an algorithm which takes $O(\mathrm{nnz}(A)+n$ poly $(k, 1 / \epsilon))$ time and outputs three matrices: $C \in \mathbb{R}^{n \times c}$, a subset of columns of $A, R \in \mathbb{R}^{n \times r}$ a subset of rows of $A$, and $T \in \mathbb{R}^{n \times t}$, a subset of tubes of $A$, together with a tensor $U \in \mathbb{R}^{c \times r \times t}$ with $\operatorname{rank}(U)=k^{\prime}$ where $c=r=t=\operatorname{poly}(k, 1 / \epsilon)$ and $k^{\prime}=\operatorname{poly}(k, 1 / \epsilon)$ such that

$$
\|U(C, R, T)-A\|_{F}^{2} \leq(1+\epsilon) \min _{\text {rank }-k A_{k}}\left\|A_{k}-A\right\|_{F}^{2},
$$

or equivalently,

$$
\left\|\sum_{i=1}^{c} \sum_{j=1}^{r} \sum_{l=1}^{t} U_{i, j, l} \cdot C_{i} \otimes R_{j} \otimes T_{l}-A\right\|_{F}^{2} \leq(1+\epsilon)_{\operatorname{rank}-k A_{k}}\left\|A_{k}-A\right\|_{F}^{2}
$$

holds with probability 9/10.

Proof. The proof follows by combining Theorem 21.1.1 and Theorem 21.1.3 directly.

```
Algorithm 21.9 Frobenius Norm Tensor Column, Row and Tube Subset Selection, Input
Sparsity Time
    procedure FCRTSELECTION \((A, n, k, \epsilon) \quad \triangleright\) Theorem 21.4.20
    \(s_{1} \leftarrow s_{2} \leftarrow O(k / \epsilon)\).
    \(\epsilon_{0} \leftarrow 0.001\).
    Choose a Gaussian matrix \(S_{1}\) with \(s_{1}\) columns. \(\triangleright\) Definition 21.3.12
    Choose a Gaussian matrix \(S_{2}\) with \(s_{2}\) columns. \(\triangleright\) Definition 21.3.12
    Form matrix \(B_{1}\) by setting \((i, j)\)-th column to be \(\left(A_{1} S_{1}\right)_{i}\).
    Form matrix \(B_{2}\) by setting \((i, j)\)-th column to be \(\left(A_{2} S_{2}\right)_{j}\). \(\quad \triangleright Z_{3}^{\prime}=B_{1}^{\top} \odot B_{2}^{\top}\)
    \(d_{3} \leftarrow O\left(s_{1} s_{2} \log \left(s_{1} s_{2}\right)+\left(s_{1} s_{2} / \epsilon\right)\right)\).
    \(D_{3} \leftarrow \operatorname{FASTTENSORLEVERAGESCOREGENERALORDER}\left(B_{1}^{\top}, B_{2}^{\top}, n, n, s_{1} s_{2}, \epsilon_{0}, d_{1}\right) . \triangleright\)
    Algorithm 21.15
        Form matrix \(B_{1}\) by setting \((i, j)\)-th column to be \(\left(A_{1} S_{1}\right)_{i}\).
        Form matrix \(B_{3}\) by setting \((i, j)\)-th column to be \(\left(A_{3} D_{3}\right)_{j}\). \(\quad \triangleright Z_{2}^{\prime}=B_{1}^{\top} \odot B_{3}^{\top}\)
        \(d_{2} \leftarrow O\left(s_{1} d_{3} \log \left(s_{1} d_{3}\right)+\left(s_{1} d_{3} / \epsilon\right)\right)\).
        \(D_{2} \leftarrow \operatorname{FastTEnsorLEVERAGESCOREGENERALORDER}\left(B_{1}^{\top}, B_{3}^{\top}, n, n, s_{1} d_{3}, \epsilon_{0}, d_{2}\right)\).
        Form matrix \(B_{2}\) by setting \((i, j)\)-th column to be \(\left(A_{2} D_{2}\right)_{i}\).
        Form matrix \(B_{3}\) by setting \((i, j)\)-th column to be \(\left(A_{3} D_{3}\right)_{j}\). \(\quad \triangleright Z_{1}^{\prime}=B_{2}^{\top} \odot B_{3}^{\top}\)
        \(d_{1} \leftarrow O\left(d_{2} d_{3} \log \left(d_{2} d_{3}\right)+\left(d_{2} d_{3} / \epsilon\right)\right)\).
        \(D_{1} \leftarrow \operatorname{FastTensorLeverageScoreGEneralOrder}\left(B_{2}^{\top}, B_{3}^{\top}, n, n, d_{2} d_{3}, \epsilon_{0}, d_{1}\right)\).
        \(C \leftarrow A_{1} D_{1}, R \leftarrow A_{2} D_{2}\) and \(T \leftarrow A_{3} D_{3}\).
        return \(C, R\) and \(T\).
    end procedure
```


### 21.4.7 CURT decomposition, $(1+\epsilon)$-approximation

### 21.4.7.1 Properties of leverage score sampling and BSS sampling

Notice that, the BSS algorithm is a deterministic procedure developed in [BSS12] for selecting rows from a matrix $A \in \mathbb{R}^{n \times d}$ (with $\|A\|_{2} \leq 1$ and $\|A\|_{F}^{2} \leq k$ ) using a selection matrix $S$ so that

$$
\left\|A^{\top} S^{\top} S A-A^{\top} A\right\|_{2} \leq \epsilon
$$

The algorithm runs in poly $(n, d, 1 / \epsilon)$ time. Using the ideas from [BW14] and [CEM $\left.{ }^{+} 15\right]$, we are able to reduce the number of nonzero entries from $O\left(\epsilon^{-2} k \log k\right)$ to $O\left(\epsilon^{-2} k\right)$, and also improve the running time to input sparsity.

Lemma 21.4.22 (Leverage score preserves subspace embedding - Theorem 2.11 in [Woo14b]). Given a rank- $k$ matrix $A \in \mathbb{R}^{n \times d}$, via leverage score sampling, we can obtain a diagonal matrix $D$ with $m$ nonzero entries such that, letting $B=D A$, if $m=O\left(\epsilon^{-2} k \log k\right)$, then, with probability at least 0.999 , for all $x \in \mathbb{R}^{d}$,

$$
(1-\epsilon)\|A x\|_{2} \leq\|B x\|_{2} \leq(1+\epsilon)\|A x\|_{2}
$$

Lemma 21.4.23. Given a rank- $k$ matrix $A \in \mathbb{R}^{n \times d}$, there exists an algorithm that runs in $O(\operatorname{nnz}(A)+n \operatorname{poly}(k, 1 / \epsilon))$ time and outputs a matrix $B$ containing $O\left(\epsilon^{-2} k \log k\right)$ re-weighted rows of $A$, such that with probability at least 0.999 , for all $x \in \mathbb{R}^{d}$,

$$
(1-\epsilon)\|A x\|_{2} \leq\|B x\|_{2} \leq(1+\epsilon)\|A x\|_{2}
$$

Proof. We choose a sparse embedding matrix (Definition 21.3.10) $\Pi \in \mathbb{R}^{d \times s}$ with $s=$ $\operatorname{poly}(k / \epsilon)$. With probability at least $0.999, \Pi^{\top}$ is a subspace embedding of $A^{\top}$. Thus, $\operatorname{rank}(A \Pi)=\operatorname{rank}(A)$. Also, the leverage scores of $A \Pi$ are the same as those of $A$. Thus, we can compute the leverage scores of $A \Pi$. The running time of computing $A \Pi$ is $O(\mathrm{nnz}(A))$. Thus the total running time is $O(\mathrm{nnz}(A)+n \operatorname{poly}(k, 1 / \epsilon))$.

Lemma 21.4.24. Let $B$ denote a matrix which contains $O\left(\epsilon^{-2} k \log k\right)$ rows of $A \in \mathbb{R}^{n \times d}$. Choosing $\Pi$ to be a sparse subspace embedding matrix of size $d \times O\left(\epsilon^{-6}(k \log k)^{2}\right)$, with probability at least 0.999,

$$
\left\|B \Pi \Pi^{\top} B^{\top}-B B^{\top}\right\|_{2} \leq \epsilon\|B\|_{2}^{2}
$$

Combining Lemma 21.4.23, 21.4.24 and the BSS algorithm, we obtain:

Lemma 21.4.25. Given a rank- $k$ matrix $A \in \mathbb{R}^{n \times d}$, there exists an algorithm that runs in $O(\mathrm{nnz}(A)+n \operatorname{poly}(k, 1 / \epsilon))$ time and outputs a sampling and rescaling diagonal matrix $S$ that selects $O\left(\epsilon^{-2} k\right)$ re-weighted rows of $A$, such that, with probability at least 0.999,

$$
\left\|A^{\top} S^{\top} S A-A^{\top} A\right\|_{2} \leq \epsilon\|A\|_{2}^{2}
$$

or equivalently, for all $x \in \mathbb{R}^{d}$,

$$
(1-\epsilon)\|A x\|_{2} \leq\|S A x\|_{2} \leq(1+\epsilon)\|A x\|_{2}
$$

Proof. Using Lemma 21.4.23, we can obtain $B$. Then we apply a sparse subspace embedding matrix $\Pi$ on the right of $B$. At the end, we run the BSS algorithm on $B \Pi$ and we are able to output $O\left(\epsilon^{-2} k\right)$ re-weighted rows of $B \Pi$. Using these rows, we are able to determine $O\left(\epsilon^{-2} k\right)$ re-weighted rows of $A$.

### 21.4.7.2 Row sampling for linear regression

Theorem 21.4.26 (Theorem 5 in [CNW15]). We are given $A \in \mathbb{R}^{n \times d}$ with $\|A\|_{2}^{2} \leq 1$ and $\|A\|_{F}^{2} \leq k$, and an $\epsilon \in(0,1)$. There exists a diagonal matrix $S$ with $O\left(k / \epsilon^{2}\right)$ nonzero entries such that

$$
\left\|(S A)^{\top} S A-A^{\top} A\right\|_{2} \leq \epsilon
$$

Corollary 21.4.27. Given a rank-k matrix $A \in \mathbb{R}^{n \times d}$, vector $b \in \mathbb{R}^{n}$, and parameter $\epsilon>0$, let $U \in \mathbb{R}^{n \times(k+1)}$ denote an orthonormal basis of $[A, b]$. Let $S \in \mathbb{R}^{n \times n}$ denote a sampling and rescaling diagonal matrix according to Leverage score sampling and sparse BSS sampling of
$U$ with $m$ nonzero entries. If $m=O(k)$, then $S$ is a $(1 \pm 1 / 2)$ subspace embedding for $U$; if $m=O(k / \epsilon)$, then $S$ satisfies $\sqrt{\epsilon}$-operator norm approximate matrix product for $U$.

Proof. This follows by Lemma 21.4.22, Lemma 21.4.24 and Theorem 21.4.26.

Lemma 21.4.28 ([NW14]). Given $A \in \mathbb{R}^{n \times d}$ and $b \in \mathbb{R}^{n}$, let $S \in \mathbb{R}^{n \times n}$ denote a sampling and rescaling diagonal matrix. Let $x^{*}$ denote $\arg \min _{x}\|A x-b\|_{2}^{2}$ and $x^{\prime}$ denote $\arg \min _{x} \| S A x-$ $S b \|_{2}^{2}$. If $S$ is a $(1 \pm 1 / 2)$ subspace embedding for the column span of $A$, and $\epsilon^{\prime}(=\sqrt{\epsilon})$-operator norm approximate matrix product for $U$ adjoined with $b-A x^{*}$, then, with probability at least .999,

$$
\left\|A x^{\prime}-b\right\|_{2}^{2} \leq(1+\epsilon)\left\|A x^{*}-b\right\|_{2}^{2} .
$$

Proof. We define OPT $=\min _{x}\|A x-b\|_{2}$. We define $x^{\prime}=\underset{x}{\arg \min }\|S A x-S b\|_{2}^{2}$ and $x^{*}=$ $\arg \min \|A x-b\|_{2}^{2}$. Let $w=b-A x^{*}$. Let $U$ denote an orthonormal basis of $A$. We can write $A x^{\prime}-A x^{*}=U \beta$. Then, we have,

$$
\begin{array}{rlr}
\left\|A x^{\prime}-b\right\|_{2}^{2} & =\left\|A x^{\prime}-A x^{*}+A A^{\dagger} b-b\right\|_{2}^{2} & \text { by } x^{*}=A^{\dagger} b \\
& =\left\|U \beta+\left(U U^{\top}-I\right) b\right\|_{2}^{2} & \\
& =\left\|A x^{*}-A x^{\prime}\right\|_{2}^{2}+\left\|A x^{*}-b\right\|_{2}^{2} \quad \text { by Pythagorean Theorem } \\
& =\|U \beta\|_{2}^{2}+\mathrm{OPT}^{2} & \\
& =\|\beta\|_{2}^{2}+\mathrm{OPT}^{2} . &
\end{array}
$$

If S is a $(1 \pm 1 / 2)$ subspace embedding for $U$, then we can show

$$
\begin{aligned}
& \|\beta\|_{2}-\left\|U^{\top} S^{\top} S U \beta\right\|_{2} \\
\leq & \left\|\beta-U^{\top} S^{\top} S U \beta\right\|_{2} \\
= & \left\|\left(I-U^{\top} S^{\top} S U\right) \beta\right\|_{2} \\
\leq & \left\|I-U^{\top} S^{\top} S U\right\|_{2} \cdot\|\beta\|_{2} \\
\leq & \frac{1}{2}\|\beta\|_{2} .
\end{aligned}
$$

$$
\leq\left\|\beta-U^{\top} S^{\top} S U \beta\right\|_{2} \quad \text { by triangle inequality }
$$

Thus, we obtain

$$
\left\|U^{\top} S^{\top} S U \beta\right\|_{2} \geq\|\beta\|_{2} / 2
$$

Next, we can show

$$
\begin{aligned}
\left\|U^{\top} S^{\top} S U \beta\right\|_{2} & =\left\|U^{\top} S^{\top} S\left(A x^{\prime}-A x^{*}\right)\right\|_{2}^{2} & & \\
& =\left\|U^{\top} S^{\top} S\left(A(S A)^{\dagger} S b-A x^{*}\right)\right\|_{2} & & \text { by } x^{\prime}=(S A)^{\dagger} S b \\
& =\left\|U^{\top} S^{\top} S\left(b-A x^{*}\right)\right\|_{2} & & \text { by } S A(S A)^{\dagger}=I \\
& =\left\|U^{\top} S^{\top} S w\right\|_{2} . & & \text { by } w=b-A x^{*}
\end{aligned}
$$

We define $U^{\prime}=\left[\begin{array}{ll}U & w /\|w\|_{2}\end{array}\right]$. We define $X$ and $y$ to satisfy $U=U^{\prime} X$ and $w=U^{\prime} y$. Then,
we have

$$
\begin{aligned}
& \left\|U^{\top} S^{\top} S w\right\|_{2} \\
= & \left\|U^{\top} S^{\top} S w-U^{\top} w\right\|_{2} \\
= & \left\|X^{\top} U^{\prime \top} S^{\top} S U^{\prime} y-X^{\top} U^{\prime \top} U^{\prime} y\right\|_{2} \\
= & \left\|X^{\top}\left(U^{\prime \top} S^{\top} S U^{\prime}-I\right) y\right\|_{2} \\
\leq & \|X\|_{2} \cdot\left\|U^{\prime \top} S^{\top} S U^{\prime}-I\right\|_{2} \cdot\|y\|_{2} \\
\leq & \epsilon^{\prime}\|X\|_{2}\|y\|_{2} \\
= & \epsilon^{\prime}\|U\|_{2}\|w\|_{2} \\
= & \epsilon^{\prime} \mathrm{OPT},
\end{aligned} \quad \text { by } U^{\top} w=0
$$

where the fifth inequality follows since $S$ satisfies $\epsilon^{\prime}$-operator norm approximate matrix product for the column span of $U$ adjoined with $w$.

Putting it all together, we have

$$
\begin{array}{rlr}
\left\|A x^{\prime}-b\right\|_{2}^{2} & =\left\|A x^{*}-b\right\|_{2}^{2}+\left\|A x^{*}-A x^{\prime}\right\|_{2}^{2} & \\
& =\mathrm{OPT}^{2}+\|\beta\|_{2}^{2} & \\
& \leq \mathrm{OPT}^{2}+4\left\|U^{\top} S^{\top} S w\right\|_{2}^{2} & \\
& \leq \mathrm{OPT}^{2}+4\left(\epsilon^{\prime} \mathrm{OPT}\right)^{2} & \text { by } \epsilon^{\prime}=\frac{1}{2} \sqrt{\epsilon} .
\end{array}
$$

Finally, note that $S$ satisfies $\epsilon^{\prime}$-operator norm approximate matrix product for $U$ adjoined with $w$ if it is a $\left(1 \pm \epsilon^{\prime}\right)$-subspace embedding for $U$ adjoined with $w$, which holds using BSS sampling by Theorem 5 of [CNW15] with $O(d / \epsilon)$ samples.

### 21.4.7.3 Leverage scores for multiple regression

Lemma 21.4.29 (see, e.g., Lemma 32 in [CW13] among other places). Given matrix $A \in$ $\mathbb{R}^{n \times d}$ with orthonormal columns, and parameter $\epsilon>0$, if $S \in \mathbb{R}^{n \times n}$ is a sampling and rescaling diagonal matrix according to the leverage scores of $A$ where the number of nonzero entries is $t=O\left(1 / \epsilon^{2}\right)$, then, for any $B \in \mathbb{R}^{n \times m}$, we have

$$
\left\|A^{\top} S^{\top} S B-A^{\top} B\right\|_{F}^{2}<\epsilon^{2}\|A\|_{F}^{2}\|B\|_{F}^{2},
$$

holds with probability at least 0.9999.

Corollary 21.4.30. Given matrix $A \in \mathbb{R}^{n \times d}$ with orthonormal columns, and parameter $\epsilon>0$, if $S \in \mathbb{R}^{n \times n}$ is a sampling and rescaling diagonal matrix according to the leverage scores of $A$ with $m$ nonzero entries, then if $m=O(d \log d)$, then $S$ is a $(1 \pm 1 / 2)$ subspace embedding for $A$. If $m=O(d / \epsilon)$, then $S$ satisfies $\sqrt{\epsilon / d}$-Frobenius norm approximate matrix product for $A$.

Proof. This follows by Lemma 21.4.22 and Lemma 21.4.29.

Lemma 21.4.31 ([NW14]). Given $A \in \mathbb{R}^{n \times d}$ and $B \in \mathbb{R}^{n \times m}$, let $S \in \mathbb{R}^{n \times n}$ denote a sampling and rescaling matrix according to $A$. Let $X^{*}$ denote $\arg \min _{X}\|A X-B\|_{F}^{2}$ and $X^{\prime}$ denote $\arg \min _{X}\|S A X-S B\|_{F}^{2}$. Let $U$ denote an orthonormal basis for $A$. If $S$ is a ( $1 \pm 1 / 2$ ) subspace embedding for $U$, and satisfies $\epsilon^{\prime}(=\sqrt{\epsilon / d})$-Frobenius norm approximate matrix product for $U$, then, we have that

$$
\left\|A X^{\prime}-B\right\|_{F}^{2} \leq(1+\epsilon)\left\|A X^{*}-B\right\|_{F}^{2}
$$

holds with probability at least 0.999.

Proof. We define OPT $=\min _{X}\|A X-B\|_{F}$. Let $A=U \Sigma V^{\top}$ denote the SVD of $A$. Since $A$ has rank $k, U$ and $V$ have $k$ columns. We can write $A\left(X^{\prime}-X^{*}\right)=U \beta$. Then, we have

$$
\begin{align*}
\left\|A X^{\prime}-B\right\|_{F}^{2} & =\left\|A X^{\prime}-A X^{*}+A A^{\dagger} B-B\right\|_{F}^{2} \quad \text { by } X^{*}=A^{\dagger} B \\
& =\left\|U \beta+\left(U U^{\top}-I\right) B\right\|_{F}^{2} \\
& =\left\|A X^{*}-A X^{\prime}\right\|_{F}^{2}+\left\|A X^{*}-B\right\|_{F}^{2} \quad \text { by Pythagorean Theorem } \\
& =\|U \beta\|_{F}^{2}+\mathrm{OPT}^{2} \\
& =\|\beta\|_{F}^{2}+\mathrm{OPT}^{2} \tag{21.15}
\end{align*}
$$

If $S$ is a $(1 \pm 1 / 2)$ subspace embedding for $U$, then we can show,

$$
\begin{aligned}
& \|\beta\|_{F}-\left\|U^{\top} S^{\top} S S U \beta\right\|_{F} \\
\leq & \left\|\beta-U^{\top} S^{\top} S U \beta\right\|_{F} \\
= & \text { by triangle inequality } \\
\leq & \left\|\left(I-U^{\top} S^{\top} S U\right) \beta\right\|_{F} \\
\leq & \frac{1}{2}\|\beta\|_{F} .
\end{aligned}
$$

Thus, we obtain

$$
\begin{equation*}
\left\|U^{\top} S^{\top} S U \beta\right\|_{F} \geq\|\beta\|_{F} / 2 \tag{21.16}
\end{equation*}
$$

Next, we can show

$$
\begin{array}{rlrl}
\left\|U^{\top} S^{\top} S U \beta\right\|_{F} & =\left\|U^{\top} S^{\top} S\left(A X^{\prime}-A X^{*}\right)\right\|_{F} & \\
& =\left\|U^{\top} S^{\top} S\left(A(S A)^{\dagger} S b-A X^{*}\right)\right\|_{F} & & \text { by } X^{\prime}=(S A)^{\dagger} S B \\
& =\left\|U^{\top} S^{\top} S\left(B-A X^{*}\right)\right\|_{F} . & & \text { by } S A(S A)^{\dagger}=I
\end{array}
$$

Then, we can show

$$
\begin{align*}
\left\|U^{\top} S^{\top} S\left(B-A X^{*}\right)\right\|_{F} & \leq \epsilon^{\prime}\left\|U^{\top}\right\|_{F}\left\|B-A X^{*}\right\|_{F} \\
& =\epsilon^{\prime} \sqrt{d} \mathrm{OPT} \tag{21.17}
\end{align*}
$$

where the first step follows from Lemma 21.4.29, the second step follows from $\|U\|_{F}=$ $\sqrt{d}$ and $\left\|B-A X^{*}\right\|_{F}=\mathrm{OPT}$.

Putting it all together, we have

$$
\begin{array}{rlr}
\left\|A X^{\prime}-B\right\|_{F}^{2} & =\left\|A X^{*}-B\right\|_{F}^{2}+\left\|A X^{*}-A X^{\prime}\right\|_{F}^{2} & \\
& =\mathrm{OPT}^{2}+\|\beta\|_{F}^{2} & \\
& \leq \mathrm{OPT}^{2}+4\left\|U^{\top} S^{\top} S w\right\|_{F}^{2} & \\
& \leq \mathrm{OPT}^{2}+4\left(\epsilon^{\prime} \sqrt{d} \mathrm{OPT}\right)^{2} & \\
& \leq(1+\epsilon) \mathrm{OPT}^{2} . & \\
\text { by Equation }(21.15) \\
& \text { bquation }(21.16) \\
& 21.17) \\
& & \text { by } \epsilon^{\prime}=\frac{1}{2} \sqrt{\epsilon / d}
\end{array}
$$

### 21.4.7.4 Sampling columns according to leverage scores implicitly, improving polynomial running time to nearly linear running time

This section explains an algorithm that is able to sample from the leverage scores from the $\odot$ product of two matrices $U, V$ without explicitly writing down $U \odot V$. To build this algorithm we combine TensorSketch, some ideas from [DMIMW12] and some ideas from [AKO11, MW10]. Finally, we are able to improve the running time of sampling columns according to leverage scores from $\Omega\left(n^{2}\right)$ to $\widetilde{O}(n)$. Given two matrices $U, V \in \mathbb{R}^{k \times n}$, we define $A \in \mathbb{R}^{k \times n_{1} n_{2}}$ to be the matrix where the $i$-th row of $A$ is the vectorization of $U^{i} \otimes V^{i}, \forall i \in[k]$.

Naïvely, in order to sample $O($ poly $(k, 1 / \epsilon))$ rows from $A^{\top}$ according to leverage scores, we need to write down $n^{2}$ leverage scores. This approach will take at least $\Omega\left(n^{2}\right)$ running time. In the rest of this section, we will explain how to do it in $O(n \cdot \operatorname{poly}(\log n, k, 1 / \epsilon))$ time. In Section 21.5.1, we will explain how to extend this idea from 3rd order tensors to general $q$-th order tensors and remove the poly $(\log n)$ factor from running time, i.e., obtain $O(n \cdot \operatorname{poly}(k, 1 / \epsilon))$ time.

Lemma 21.4.32. Given two matrices $U \in \mathbb{R}^{k \times n_{1}}$ and $V \in \mathbb{R}^{k \times n_{2}}$, there exists an algorithm that takes $O\left(\left(n_{1}+n_{2}\right) \cdot \operatorname{poly}\left(\log \left(n_{1} n_{2}\right), k\right) \cdot R_{\text {samples }}\right)$ time and samples $R_{\text {samples }}$ columns of $U \odot V \in \mathbb{R}^{k \times n_{1} n_{2}}$ according to the leverage scores of $R^{-1}(U \odot V)$, where $R$ is the $R$ of a $Q R$ factorization.

Proof. We choose $\Pi \in \mathbb{R}^{n_{1} n_{2} \times s_{1}}$ to be a TensorSketch. Then, according to Section 21.3.10, we can compute $R^{-1}$ in $n \cdot \operatorname{poly}(\log n, k, 1 / \epsilon)$ time, where $R$ is the $R$ in a QR-factorization. We want to sample columns from $U \odot V$ according to the square of the $\ell_{2}$-norms of each column of $R^{-1}(U \odot V)$. However, explicitly writing down the matrix $R^{-1}(U \odot V)$ takes $k n_{1} n_{2}$ time, and the number of columns is already $n_{1} n_{2}$. The goal is to sample columns from $R^{-1}(U \odot V)$ without explicitly computing the square of the $\ell_{2}$-norm of each column.

The first simple observation is that the following two sampling procedures are equivalent in terms of the column samples of a matrix that they take. (1) We sample a single entry from the matrix $R^{-1}(U \odot V)$ proportional to its squared value. (2) We sample a column from the matrix $R^{-1}(U \odot V)$ proportional to its squared $\ell_{2}$-norm. Let the $\left(i, j_{1}, j_{2}\right)$-th entry denote the entry in the $i$-th row and the $\left(j_{1}-1\right) n_{2}+j_{2}$-th column. We can show, for a

```
Algorithm 21.10 Fast Tensor Leverage Score Sampling
    procedure FastTensorLeverageScore \(\left(U, V, n_{1}, n_{2}, k, \epsilon, R_{\text {samples }}\right) \triangleright\) Lemma 21.4.32
        \(s_{1} \leftarrow \operatorname{poly}(k, 1 / \epsilon)\).
        \(g_{1} \leftarrow g_{2} \leftarrow g_{3} \leftarrow O\left(\epsilon^{-2} \log \left(n_{1} n_{2}\right)\right)\).
        Choose \(\Pi \in \mathbb{R}^{n_{1} n_{2} \times s_{1}}\) to be a TensorSketch. \(\triangleright\) Definition 21.3.20
        Compute \(R^{-1} \in \mathbb{R}^{k \times k}\) by using \((U \odot V) \Pi\). \(\triangleright U \in \mathbb{R}^{k \times n_{1}}, V \in \mathbb{R}^{k \times n_{2}}\)
        Choose \(G_{1} \in \mathbb{R}^{g_{1} \times k}\) to be a Gaussian sketching matrix.
        for \(i=1 \rightarrow g_{1}\) do
            \(w \leftarrow\left(G^{i} R^{-1}\right)^{\top} \quad \triangleright G^{i}\) denotes the \(i\)-th row of \(G\)
            for \(j=1 \rightarrow\left[n_{1}\right]\) do \(\quad \triangleright\) Form matrix \(U^{\prime i} \in \mathbb{R}^{k \times n_{1}}\)
                \(U_{j}^{\prime i} \leftarrow w \circ U_{j}, \forall j \in\left[n_{1}\right] . \quad \triangleright U_{j}\) denotes the \(j\)-th column of \(U \in \mathbb{R}^{k \times n_{1}}\)
            end for
        end for
        Choose \(G_{2, i} \in \mathbb{R}^{g_{2} \times n_{1}}\) to be a Gaussian sketching matrix.
        for \(i=1 \rightarrow g_{1}\) do
            \(\alpha_{i} \leftarrow\left\|\left(G_{2, i} U^{\prime \top}\right) V\right\|_{F}^{2}\).
            Choose \(G_{3, i} \in \mathbb{R}^{g_{3} \times n_{1}}\) to be a Gaussian sketching matrix.
            for \(j_{2}=1 \rightarrow n_{2}\) do
                \(\beta_{i, j} \leftarrow\left\|G_{3, i}\left(U^{\prime \top}\right) V_{j_{2}}\right\|_{2}^{2}\).
            end for
        end for
        \(\mathcal{S} \leftarrow \emptyset\).
        for \(r=1 \rightarrow R_{\text {samples }}\) do
            Sample \(i\) from \(\left[g_{1}\right]\) with probability \(\alpha_{i} / \sum_{i^{\prime}=1}^{g_{1}} \alpha_{i^{\prime}}\).
            Sample \(j_{2}\) from \(\left[n_{2}\right]\) with probability \(\beta_{i, j_{2}} / \sum_{j_{2}^{\prime}=1}^{n_{2}} \beta_{i, j_{2}^{\prime}}\).
            for \(j_{1}=1 \rightarrow n_{1}\) do
                \(\gamma_{j_{1}} \leftarrow\left(\left(U^{\prime i \top}\right)^{j_{1}} V_{j_{2}}\right)^{2}\).
            end for
            Sample \(j_{1}\) from \(\left[n_{1}\right]\) with probability \(\gamma_{j_{1}} / \sum_{j_{1}^{\prime}=1}^{n_{1}} \gamma_{j_{1}^{\prime}}\).
            \(\mathcal{S} \leftarrow \mathcal{S} \cup\left(j_{1}, j_{2}\right)\).
        end for
        Convert \(\mathcal{S}\) into a diagonal matrix \(D\) with at most \(R_{\text {samples }}\) nonzero entries.
        return \(D\).
        \(\triangleright\) Diagonal matrix \(D \in \mathbb{R}^{n_{1} n_{2} \times n_{1} n_{2}}\)
    end procedure
```

particular column $\left(j_{1}-1\right) n_{2}+j_{2}$,
$\operatorname{Pr}\left[\right.$ sample an entry from the $\left(j_{1}-1\right) n_{2}+j_{2}$ th column of a matrix]

$$
\begin{align*}
& =\sum_{i=1}^{k} \operatorname{Pr}\left[\text { sample the }\left(i, j_{1}, j_{2}\right) \text {-th entry of matrix }\right] \\
& =\sum_{i=1}^{k} \frac{\left|\left(R^{-1}(U \odot V)\right)_{i,\left(j_{1}-1\right) n_{2}+j_{2}}\right|^{2}}{\left\|R^{-1}(U \odot V)\right\|_{F}^{2}} \\
& =\frac{\left\|\left(R^{-1}(U \odot V)\right)_{\left(j_{1}-1\right) n_{2}+j_{2}}\right\|^{2}}{\left\|R^{-1}(U \odot V)\right\|_{F}^{2}} \\
& =\operatorname{Pr}\left[\text { sample the }\left(j_{1}-1\right) n_{2}+j_{2} \text { th column of matrix }\right] . \tag{21.18}
\end{align*}
$$

Thus, it is sufficient to show how to sample a single entry from matrix $R^{-1}(U \odot V)$ proportional to its squared value without writing down all of the entries of a $k \times n_{1} n_{2}$ matrix.

We choose a Gaussian matrix $G_{1} \in \mathbb{R}^{g_{1} \times k}$ with $g_{1}=O\left(\epsilon^{-2} \log \left(n_{1} n_{2}\right)\right)$. By Claim 21.4.33 we can reduce the length of each column vector of matrix $R^{-1} U \odot V$ from $k$ to $g_{1}$ while preserving the squared $\ell_{2}$-norm of all columns simultaneously. Thus, we obtain a new matrix $G R^{-1}(U \odot V) \in \mathbb{R}^{g_{1} \times n_{1} n_{2}}$, and sampling from this new matrix is equivalent to sampling from the original matrix $R^{-1}(U \odot V)$.

In the following paragraphs, we explain a sampling procedure (also described in Procedure Fast TensorLeverageScore in Algorithm 21.10) which contains three sampling steps. The first step is sampling $i$ from $\left[g_{1}\right]$, the second step is sampling $j_{2}$ from $\left[n_{2}\right]$, and the last step is sampling $j_{1}$ from $\left[n_{1}\right]$.

For each $j_{1} \in\left[n_{1}\right]$, let $U_{j_{1}}$ denote the $j_{1}$-th column of $U$. For each $i \in\left[g_{1}\right]$, let $G_{1}^{i}$ denote the $i$-th row of matrix $G_{1} \in \mathbb{R}^{g_{1} \times k}$, let $U^{\prime i} \in \mathbb{R}^{k \times n_{1}}$ denote a matrix where the $j_{1}$-th column is $\left(G^{i} R^{-1}\right)^{\top} \circ U_{j_{1}} \in \mathbb{R}^{k}, \forall j \in\left[n_{1}\right]$. Then, using Claim 21.4.37, we have that $\left(G^{i} R^{-1}\right) \cdot(U \odot V) \in$
$\mathbb{R}^{n_{1} n_{2}}$ is a row vector where the entry in the $\left(j_{1}-1\right) n_{2}+j_{2}$-th coordinate is the entry in the $j_{1^{-}}$ th row and $j_{2}$-th column of matrix $\left(U^{\prime i \top} V\right) \in \mathbb{R}^{n_{1} \times n_{2}}$. Further, the squared $\ell_{2}$-norm of vector $\left(G^{i} R^{-1}\right) \cdot(U \odot V)$ is equal to the squared Frobenius norm of matrix $\left(U^{\prime \top} T V\right)$. Thus, sampling $i$ proportional to the squared $\ell_{2}$-norm of vector $\left(G^{i} R^{-1}\right) \cdot(U \odot V)$ is equivalent to sampling $i$ proportional to the squared Frobenius norm of matrix $\left(U^{\prime i \top} V\right)$. Naïvely, computing the Frobenius norm of an $n_{1} \times n_{2}$ matrix requires $O\left(n_{1} n_{2}\right)$ time. However, we can choose a Gaussian matrix $G_{2, i} \in \mathbb{R}^{g_{2} \times n_{1}}$ to sample according to the value $\left\|\left(G_{2, i} U^{\prime i \top}\right) V\right\|_{F}^{2}$, which can be computed in $O\left(\left(n_{1}+n_{2}\right) g_{2} k\right)$ time. By claim 21.4.35, $\left\|\left(G_{2, i} U^{\prime i \top}\right) V\right\|_{F}^{2} \approx\left\|\left(U^{\prime \top}\right) V\right\|_{F}^{2}$ with high probability. So far, we have finished the first step of the sampling procedure.

For the second step of the sampling procedure, we need to sample $j_{2}$ from $\left[n_{2}\right]$. To do that, we need to compute the squared $\ell_{2}$-norm of each column of $U^{\prime i \top} V \in \mathbb{R}^{n_{1} \times n_{2}}$. This can be done by choosing another Gaussian matrix $G_{3, i} \in \mathbb{R}^{g_{3} \times n_{1}}$. For all $j_{2} \in\left[n_{2}\right]$, by Claim 21.4.36, we have $\left\|G_{3, i} U^{\prime i \top} V_{j_{2}}\right\|_{2}^{2} \approx\left\|U^{\prime i \top} V_{j_{2}}\right\|_{2}^{2}$. Also, for $j_{2} \in\left[n_{2}\right],\left\|G_{3, i} U^{\prime i \top} V_{j_{2}}\right\|_{2}^{2}$ can be computed in nearly linear in $n_{1}+n_{2}$ time.

For the third step of the sampling procedure, we need to sample $j_{1}$ from $\left[n_{1}\right]$. Since we already have $i$ and $j_{2}$ from the previous two steps, we can directly compute $\left|\left(U^{\prime i \top}\right)^{j_{1}} V_{j_{2}}\right|^{2}$, for all $j_{1}$. This only takes $O\left(n_{1} k\right)$ time.

Overall, the running time is $O\left(\left(n_{1}+n_{2}\right) \cdot \operatorname{poly}\left(\log \left(n_{1} n_{2}\right), k, 1 / \epsilon\right)\right)$. Because our estimates are accurate enough, our sampling probabilities are also good approximations to the leverage score sampling probabilities. Putting it all together, we complete the proof.

Claim 21.4.33. Given matrix $R^{-1}(U \odot V) \in \mathbb{R}^{k \times n_{1} n_{2}}$, let $G_{1} \in \mathbb{R}^{g_{1} \times k}$ denote a Gaussian matrix with $g_{1}=\left(\epsilon^{-2} \log \left(n_{1} n_{2}\right)\right)$. Then with probability at least $1-1 / \operatorname{poly}\left(n_{1} n_{2}\right)$, we have:
for all $j \in\left[n_{1} n_{2}\right]$,

$$
(1-\epsilon)\left\|R^{-1}(U \odot V)_{j}\right\|_{2}^{2} \leq\left\|G_{1} R^{-1}(U \odot V)_{j}\right\|_{2}^{2} \leq(1+\epsilon)\left\|R^{-1}(U \odot V)_{j}\right\|_{2}^{2}
$$

Proof. This follows by the Johnson-Lindenstrauss Lemma.

Claim 21.4.34. For a fixed $i \in\left[g_{1}\right]$, let $G_{2, i} \in \mathbb{R}^{g_{2} \times n_{1}}$ denote a Gaussian matrix with $g_{2}=O\left(\epsilon^{-2} \log \left(n_{1} n_{2}\right)\right)$. Then with probability at least $1-1 / \operatorname{poly}\left(n_{1} n_{2}\right)$, we have: for all $j_{2} \in\left[n_{2}\right]$,

$$
(1-\epsilon)\left\|U^{\prime \top \top} V_{j_{2}}\right\|_{2}^{2} \leq\left\|\left(G_{2, i} U^{\prime i \top}\right) V_{j_{2}}\right\|_{2} \leq(1+\epsilon)\left\|U^{\prime i \top} V_{j_{2}}\right\|_{2}^{2}
$$

By taking the union bound over all $i \in\left[g_{1}\right]$, we obtain a stronger claim,
Claim 21.4.35. With probability at least $1-1 / \operatorname{poly}\left(n_{1} n_{2}\right)$, we have : for all $i \in\left[g_{1}\right]$, for all $j_{2} \in\left[n_{2}\right]$,

$$
(1-\epsilon)\left\|U^{\prime i \top} V_{j_{2}}\right\|_{2}^{2} \leq\left\|\left(G_{2, i} U^{\prime i \top}\right) V_{j_{2}}\right\|_{2} \leq(1+\epsilon)\left\|U^{\prime i \top} V_{j_{2}}\right\|_{2}^{2}
$$

Similarly, if we choose $G_{3, i}$ to be a Gaussian matrix, we can obtain the same result as for $G_{2, i}$ :

Claim 21.4.36. With probability at least $1-1 / \operatorname{poly}\left(n_{1} n_{2}\right)$, we have : for all $i \in\left[g_{1}\right]$, for all $j_{2} \in\left[n_{2}\right]$,

$$
(1-\epsilon)\left\|U^{\prime i \top} V_{j_{2}}\right\|_{2}^{2} \leq\left\|\left(G_{3, i} U^{\prime i \top}\right) V_{j_{2}}\right\|_{2} \leq(1+\epsilon)\left\|U^{\prime i \top} V_{j_{2}}\right\|_{2}^{2}
$$

Claim 21.4.37. For any $i \in\left[g_{1}\right], j_{1} \in\left[n_{1}\right], j_{2} \in\left[n_{2}\right]$, let $G_{1}^{i}$ denote the $i$-th row of matrix $G_{1} \in \mathbb{R}^{g_{1} \times k}$. Let $(U \odot V)_{\left(j_{1}-1\right) n_{2}+j_{2}}$ denote the $\left(j_{1}-1\right) n_{2}+j_{2}$-th column of matrix $\mathbb{R}^{k \times n_{1} n_{2}}$.

Let $\left(U^{\prime i \top}\right)^{j_{1}}$ denote the $j_{1}$-th row of matrix $\left(U^{\prime i \top}\right) \in \mathbb{R}^{n_{1} \times k}$. Let $V_{j_{2}}$ denote the $j_{2}$-th column of matrix $V \in \mathbb{R}^{k \times n_{2}}$. Then, we have

$$
G_{1}^{i} R^{-1}(U \odot V)_{\left(j_{1}-1\right) n_{2}+j_{2}}=\left(U^{\prime i \top}\right)^{j_{1}} V_{j_{2}} .
$$

Proof. This follows by,

$$
G_{1}^{i} R^{-1}(U \odot V)_{\left(j_{1}-1\right) n_{2}+j_{2}}=G_{1}^{i} R^{-1}\left(U_{j_{1}} \circ V_{j_{2}}\right)=\left(G_{1}^{i} R^{-1} \circ\left(U_{j_{1}}\right)^{\top}\right) V_{j_{2}}=\left(U^{\prime i \top}\right)^{j_{1}} V_{j_{2}}
$$

Lemma 21.4.38. Given $A \in \mathbb{R}^{n \times n^{2}}, V, W \in \mathbb{R}^{k \times n}$, for any $\epsilon>0$, there exists an algorithm that runs in $O(n \cdot \operatorname{poly}(k, 1 / \epsilon))$ time and outputs a diagonal matrix $D \in \mathbb{R}^{n^{2} \times n^{2}}$ with $m=$ $O(k \log k+k / \epsilon)$ nonzero entries such that,

$$
\|\widehat{U}(V \odot W)-A\|_{F}^{2} \leq(1+\epsilon) \min _{U \in \mathbb{R}^{n \times k}}\|U(V \odot W)-A\|_{F}^{2},
$$

holds with probability at least 0.999 , where $\widehat{U}$ denotes the optimal solution to $\min _{U} \| U(V \odot$ W) $D-A D \|_{F}^{2}$.

Proof. This follows by combining Theorem 21.5.1, Corollary 21.4.30, and Lemma 21.4.31.

Remark 21.4.1. Replacing Theorem 21.5.1 (Algorithm 21.15) by Lemma 21.4.32 (Algorithm 21.10), we can obtain a slightly different version of Lemma 21.4.38 with $n$ poly $(\log n, k, 1 / \epsilon)$ running time, where the dependence on $k$ is better.

```
Algorithm 21.11 Frobenius Norm CURT Decomposition Algorithm, Input Sparsity Time
and Nearly Optimal Number of Samples
    procedure FCURTInputSparsity \(\left(A, U_{B}, V_{B}, W_{B}, n, k, \epsilon\right) \quad \triangleright\) Theorem 21.4.39
    \(d_{1} \leftarrow d_{2} \leftarrow d_{3} \leftarrow O(k \log k+k / \epsilon)\).
    \(\epsilon_{0} \leftarrow 0.01\).
    Form \(B_{1}=V_{B}^{\top} \odot W_{B}^{\top} \in \mathbb{R}^{k \times n^{2}}\).
    \(D_{1} \leftarrow \operatorname{FastTEnsorLEVERAGESCOREGENERALORDER}\left(V_{B}^{\top}, W_{B}^{\top}, n, n, k, \epsilon_{0}, d_{1}\right) . \quad \triangleright\)
    Algorithm 21.15
    Form \(\widehat{U}=A_{1} D_{1}\left(B_{1} D_{1}\right)^{\dagger} \in \mathbb{R}^{n \times k}\).
    Form \(B_{2}=\widehat{U}^{\top} \odot W_{B}^{\top} \in \mathbb{R}^{k \times n^{2}}\).
    \(D_{2} \leftarrow\) FastTensorLeveragescoreGeneralOrder \(\left(\widehat{U}^{\top}, W_{B}^{\top}, n, n, k, \epsilon_{0}, d_{2}\right)\).
    Form \(\widehat{V}=A_{2} D_{2}\left(B_{2} D_{2}\right)^{\dagger} \in \mathbb{R}^{n \times k}\).
    Form \(B_{3}=\widehat{U}^{\top} \odot \widehat{V}^{\top} \in \mathbb{R}^{k \times n^{2}}\).
    \(D_{3} \leftarrow\) FAStTENSORLEVERAGESCOREGENERALORDER \(\left(\widehat{U}^{\top}, \widehat{V}^{\top}, n, n, k, \epsilon_{0}, d_{3}\right)\).
    \(C \leftarrow A_{1} D_{1}, R \leftarrow A_{2} D_{2}, T \leftarrow A_{3} D_{3}\).
    \(U \leftarrow \sum_{i=1}^{k}\left(\left(B_{1} D_{1}\right)^{\dagger}\right)_{i} \otimes\left(\left(B_{2} D_{2}\right)^{\dagger}\right)_{i} \otimes\left(\left(B_{3} D_{3}\right)^{\dagger}\right)_{i}\).
    return \(C, R, T\) and \(U\).
    end procedure
```


### 21.4.7.5 Input sparsity time algorithm

Theorem 21.4.39. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, let $k \geq 1$, and let $U_{B}, V_{B}, W_{B} \in$ $\mathbb{R}^{n \times k}$ denote a rank-k, $\alpha$-approximation to $A$. Then there exists an algorithm which takes $O(\operatorname{nnz}(A)+n \operatorname{poly}(k, 1 / \epsilon))$ time and outputs three matrices $C \in \mathbb{R}^{n \times c}$ with columns from $A$, $R \in \mathbb{R}^{n \times r}$ with rows from $A, T \in \mathbb{R}^{n \times t}$ with tubes from $A$, and a tensor $U \in \mathbb{R}^{c \times r \times t}$ with $\operatorname{rank}(U)=k$ such that $c=r=t=O(k \log k+k / \epsilon)$, and

$$
\left\|\sum_{i=1}^{c} \sum_{j=1}^{r} \sum_{l=1}^{t} U_{i, j, l} \cdot C_{i} \otimes R_{j} \otimes T_{l}-A\right\|_{F}^{2} \leq(1+\epsilon) \alpha \min _{\operatorname{rank}-k A^{\prime}}\left\|A^{\prime}-A\right\|_{F}^{2}
$$

holds with probability 9/10.

Proof. We define

$$
\mathrm{OPT}:=\min _{\operatorname{rank}-k A^{\prime}}\left\|A^{\prime}-A\right\|_{F}^{2}
$$

We already have three matrices $U_{B} \in \mathbb{R}^{n \times k}, V_{B} \in \mathbb{R}^{n \times k}$ and $W_{B} \in \mathbb{R}^{n \times k}$ and these three matrices provide a rank- $k, \alpha$-approximation to $A$, i.e.,

$$
\begin{equation*}
\left\|\sum_{i=1}^{k}\left(U_{B}\right)_{i} \otimes\left(V_{B}\right)_{i} \otimes\left(W_{B}\right)_{i}-A\right\|_{F}^{2} \leq \alpha \mathrm{OPT} \tag{21.19}
\end{equation*}
$$

Let $B_{1}=V_{B}^{\top} \odot W_{B}^{\top} \in \mathbb{R}^{k \times n^{2}}$ denote the matrix where the $i$-th row is the vectorization of $\left(V_{B}\right)_{i} \otimes\left(W_{B}\right)_{i}$. Let $D_{1} \in \mathbb{R}^{n^{2} \times n^{2}}$ be a sampling and rescaling matrix corresponding to sampling by the leverage scores of $B_{1}^{\top}$; there are $d_{1}$ nonzero entries on the diagonal of $D_{1}$. Let $A_{i} \in \mathbb{R}^{n \times n^{2}}$ denote the matrix obtained by flattening $A$ along the $i$-th direction, for each $i \in[3]$.

Define $U^{*} \in \mathbb{R}^{n \times k}$ to be the optimal solution to $\min _{U \in \mathbb{R}^{n \times k}}\left\|U B_{1}-A_{1}\right\|_{F}^{2}, \widehat{U}=A_{1} D_{1}\left(B_{1} D_{1}\right)^{\dagger} \in$ $\mathbb{R}^{n \times k}$, and $V_{0} \in \mathbb{R}^{n \times k}$ to be the optimal solution to $\min _{V \in \mathbb{R}^{n \times k}}\left\|V \cdot\left(\widehat{U}^{\top} \odot W_{B}^{\top}\right)-A_{2}\right\|_{F}^{2}$. Due to Lemma 21.4.38, if $d_{1}=O(k \log k+k / \epsilon)$ then with constant probability, we have

$$
\begin{equation*}
\left\|\widehat{U} B_{1}-A_{1}\right\|_{F}^{2} \leq \alpha_{D_{1}}\left\|U^{*} B_{1}-A_{1}\right\|_{F}^{2} \tag{21.20}
\end{equation*}
$$

Recall that $\left(\widehat{U}^{\top} \odot W_{B}^{\top}\right) \in \mathbb{R}^{k \times n^{2}}$ denotes the matrix where the $i$-th row is the vector-
ization of $\widehat{U}_{i} \otimes\left(W_{B}\right)_{i}, \forall i \in[k]$. Now, we can show,

$$
\begin{align*}
& \left\|V_{0} \cdot\left(\widehat{U}^{\top} \odot W_{B}^{\top}\right)-A_{2}\right\|_{F}^{2} \\
& \leq\left\|\widehat{U} B_{1}-A_{1}\right\|_{F}^{2} \quad \text { by } V_{0}=\underset{V \in \mathbb{R}^{n \times k}}{\arg \min }\left\|V \cdot\left(\widehat{U}^{\top} \odot W_{B}^{\top}\right)-A_{2}\right\|_{F}^{2} \\
& \leq \alpha_{D_{1}}\left\|U^{*} B_{1}-A_{1}\right\|_{F}^{2} \quad \text { by Equation (21.20) } \\
& \leq \alpha_{D_{1}}\left\|U_{B} B_{1}-A_{1}\right\|_{F}^{2} \quad \text { by } U^{*}=\underset{U \in \mathbb{R}^{n \times k}}{\arg \min }\left\|U B_{1}-A_{1}\right\|_{F}^{2} \\
& \leq \alpha_{D_{1}} \alpha \mathrm{OPT} \text {. }  \tag{21.21}\\
& \text { by Equation (21.19) }
\end{align*}
$$

We define $B_{2}=\widehat{U}^{\top} \odot W_{B}^{\top}$. Let $D_{2} \in \mathbb{R}^{n^{2} \times n^{2}}$ be a sampling and rescaling matrix corresponding to the leverage scores of $B_{2}^{\top}$. Suppose there are $d_{2}$ nonzero entries on the diagonal of $D_{2}$.

Define $V^{*} \in \mathbb{R}^{n \times k}$ to be the optimal solution to $\min _{V \in \mathbb{R}^{n \times k}}\left\|V B_{2}-A_{2}\right\|_{F}^{2}, \widehat{V}=$ $A_{2} D_{2}\left(B_{2} D_{2}\right)^{\dagger} \in \mathbb{R}^{n \times k}, W_{0} \in \mathbb{R}^{n \times k}$ to be the optimal solution to $\min _{W \in \mathbb{R}^{n \times k}}\left\|W \cdot\left(\widehat{U}^{\top} \odot \widehat{V}^{\top}\right)-A_{3}\right\|_{F}^{2}$, and $V^{\prime}$ to be the optimal solution to $\min _{V \in \mathbb{R}^{n \times k}}\left\|V B_{2} D_{2}-A_{2} D_{2}\right\|_{F}^{2}$.

Due to Lemma 21.4.38, with constant probability, we have

$$
\begin{equation*}
\left\|\widehat{V} B_{2}-A_{2}\right\|_{F}^{2} \leq \alpha_{D_{2}}\left\|V^{*} B_{2}-A_{2}\right\|_{F}^{2} \tag{21.22}
\end{equation*}
$$

Recall that $\left(\widehat{U}^{\top} \odot \widehat{V}^{\top}\right) \in \mathbb{R}^{k \times n^{2}}$ denotes the matrix where the $i$-th row is the vector-
ization of $\widehat{U}_{i} \otimes \widehat{V}_{i}, \forall i \in[k]$. Now, we can show,

$$
\begin{array}{rr} 
& \left\|W_{0} \cdot\left(\widehat{U}^{\top} \odot \widehat{V}^{\top}\right)-A_{3}\right\|_{F}^{2} \\
\leq\left\|\widehat{V} B_{2}-A_{2}\right\|_{F}^{2} & \text { by } W_{0}=\underset{W \in \mathbb{R}^{n \times k}}{\arg \min }\left\|W \cdot\left(\widehat{U}^{\top} \odot \widehat{V}^{\top}\right)-A_{3}\right\|_{F}^{2} \\
\leq & \alpha_{D_{2}}\left\|V^{*} B_{2}-A_{2}\right\|_{F}^{2} \\
\leq & \text { by Equation (21.22) } \\
\leq \alpha_{D_{2}}\left\|V_{0} B_{2}-A_{2}\right\|_{F}^{2} & \text { by } V^{*}=\underset{V \in \mathbb{R}^{n \times k}}{\arg \min }\left\|V B_{2}-A_{2}\right\|_{F}^{2}  \tag{21.23}\\
\leq & \alpha_{D_{2}} \alpha_{D_{1}} \alpha \text { OPT. }
\end{array}
$$

We define $B_{3}=\widehat{U}^{\top} \odot \widehat{V}^{\top}$. Let $D_{3} \in \mathbb{R}^{n^{2} \times n^{2}}$ denote a sampling and rescaling matrix corresponding to sampling by the leverage scores of $B_{3}^{\top}$. Suppose there are $d_{3}$ nonzero entries on the diagonal of $D_{3}$.

Define $W^{*} \in \mathbb{R}^{n \times k}$ to be the optimal solution to $\min _{W \in \mathbb{R}^{n \times k}}\left\|W B_{3}-A_{3}\right\|_{F}^{2}, \widehat{W}=$ $A_{3} D_{3}\left(B_{3} D_{3}\right)^{\dagger} \in \mathbb{R}^{n \times k}$, and $W^{\prime}$ to be the optimal solution to $\min _{W \in \mathbb{R}^{n \times k}}\left\|W B_{3} D_{3}-A_{3} D_{3}\right\|_{F}^{2}$.

Due to Lemma 21.4.38 with constant probability, we have

$$
\begin{equation*}
\left\|\widehat{W} B_{3}-A_{3}\right\|_{F}^{2} \leq \alpha_{D_{3}}\left\|W^{*} B_{3}-A_{3}\right\|_{F}^{2} \tag{21.24}
\end{equation*}
$$

Now we can show,

$$
\begin{array}{rlr}
\left\|\widehat{W} B_{3}-A_{3}\right\|_{F}^{2} & \leq \alpha_{D_{3}}\left\|W^{*} B_{3}-A_{3}\right\|_{F}^{2}, & \text { by Equation }(21.24) \\
& \leq \alpha_{D_{3}}\left\|W_{0} B_{3}-A_{3}\right\|_{F}^{2}, & \text { by } W^{*}=\underset{W \in \mathbb{R}^{n \times k}}{\arg \min }\left\|W B_{3}-A_{3}\right\|_{F}^{2} \\
& \leq \alpha_{D_{3}} \alpha_{D_{2}} \alpha_{D_{1}} \alpha \mathrm{OPT} . & \text { by Equation }(21.23)
\end{array}
$$

This implies,

$$
\left\|\sum_{i=1}^{k} \widehat{U}_{i} \otimes \widehat{V}_{i} \otimes \widehat{W}_{i}-A\right\|_{F}^{2} \leq O(1) \alpha \mathrm{OPT}^{2}
$$

where $\widehat{U}=A_{1} D_{1}\left(B_{1} D_{1}\right)^{\dagger}, \widehat{V}=A_{2} D_{2}\left(B_{2} D_{2}\right)^{\dagger}, \widehat{W}=A_{3} D_{3}\left(B_{3} D_{3}\right)^{\dagger}$.
By Lemma 21.4.38, we need to set $d_{1}=d_{2}=d_{3}=O(k \log k+k / \epsilon)$. Note that $B_{1}=\left(V_{B}^{\top} \odot W_{B}^{\top}\right)$. Thus $D_{1}$ can be found in $n \cdot \operatorname{poly}(k, 1 / \epsilon)$ time. Because $D_{1}$ has a small number of nonzero entries on the diagonal, we can compute $B_{1} D_{1}$ quickly without explicitly writing down $B_{1}$. Also $A_{1} D_{1}$ can be computed in nnz $(A)$ time. Using $\left(A_{1} D_{1}\right)$ and $\left(B_{1} D_{1}\right)$, we can compute $\widehat{U}$ in $n \operatorname{poly}(k, 1 / \epsilon)$ time. In a similar way, we can compute $B_{2}, D_{2}, B_{3}$, and $D_{3}$. Since tensor $U$ is constructed based on three poly $(k, 1 / \epsilon)$ size matrices, $\left(B_{1} D_{1}\right)^{\dagger}$, $\left(B_{2} D_{2}\right)^{\dagger}$, and $\left(B_{3} D_{3}\right)^{\dagger}$, the overall running time is $O(\mathrm{nnz}(A)+n \operatorname{poly}(k, 1 / \epsilon))$

### 21.4.7.6 Optimal sample complexity algorithm

Theorem 21.4.40. Given a $3 r d$ order tensor $A \in \mathbb{R}^{n \times n \times n}$, let $k \geq 1$, and let $U_{B}, V_{B}, W_{B} \in$ $\mathbb{R}^{n \times k}$ denote a rank- $k$, $\alpha$-approximation to $A$. Then there exists an algorithm which takes $O\left(\operatorname{nnz}(A) \log n+n^{2} \operatorname{poly}(\log n, k, 1 / \epsilon)\right)$ time and outputs three matrices: $C \in \mathbb{R}^{n \times c}$ with columns from $A, R \in \mathbb{R}^{n \times r}$ with rows from $A, T \in \mathbb{R}^{n \times t}$ with tubes from $A$, and a tensor $U \in \mathbb{R}^{c \times r \times t}$ with $\operatorname{rank}(U)=k$ such that $c=r=t=O(k / \epsilon)$, and

$$
\left\|\sum_{i=1}^{c} \sum_{j=1}^{r} \sum_{l=1}^{t} U_{i, j, l} \cdot C_{i} \otimes R_{j} \otimes T_{l}-A\right\|_{F}^{2} \leq(1+\epsilon) \alpha \min _{\text {rank }-k A^{\prime}}\left\|A^{\prime}-A\right\|_{F}^{2}
$$

holds with probability 9/10.

Proof. The proof is almost the same as the proof of Theorem 21.4.39. The only difference is that instead of using Theorem 21.4.38, we use Theorem 21.4.14.

```
Algorithm 21.12 Frobenius Norm CURT Decomposition Algorithm, Optimal Sample Com-
plexity
    procedure FCURTOptimalSamples \(\left(A, U_{B}, V_{B}, W_{B}, n, k\right) \quad \triangleright\) Theorem 21.4.40
    \(d_{1} \leftarrow d_{2} \leftarrow d_{3} \leftarrow O(k / \epsilon)\).
    Form \(B_{1}=V_{B}^{\top} \odot W_{B}^{\top} \in \mathbb{R}^{k \times n^{2}}\).
    \(D_{1} \leftarrow \operatorname{GENERALIZEDMATRIXROWSUBSETSELECTION}\left(A_{1}^{\top}, B_{1}^{\top}, n^{2}, n, k, \epsilon\right) . \quad \triangleright\)
    Algorithm 21.7
    Let \(d_{1}\) denote the number of nonzero entries in \(D_{1}\). \(\triangleright d_{1}=O(k / \epsilon)\)
    Form \(\widehat{U}=A_{1} D_{1}\left(B_{1} D_{1}\right)^{\dagger} \in \mathbb{R}^{n \times k}\).
    Form \(B_{2}=\widehat{U}^{\top} \odot W_{B}^{\top} \in \mathbb{R}^{k \times n^{2}}\).
    \(D_{2} \leftarrow\) GeneralizedMatrixRowSubsetSelection \(\left(A_{2}^{\top}, B_{2}^{\top}, n^{2}, n, k, \epsilon\right)\). \(\triangleright\)
    Algorithm 21.7
    Let \(d_{2}\) denote the number of nonzero entries in \(D_{2}\). \(\triangleright d_{2}=O(k / \epsilon)\)
    Form \(\widehat{V}=A_{2} D_{2}\left(B_{2} D_{2}\right)^{\dagger} \in \mathbb{R}^{n \times k}\).
    Form \(B_{3}=\widehat{U}^{\top} \odot \widehat{V}^{\top} \in \mathbb{R}^{k \times n^{2}}\).
    \(D_{3} \leftarrow \operatorname{GEnERALIZEDMATRIXROWSubSETSELECTION}\left(A_{3}^{\top}, B_{3}^{\top}, n^{2}, n, k, \epsilon\right) . \quad \triangleright\)
    Algorithm 21.7
    \(d_{3}\) denote the number of nonzero entries in \(D_{3} . \quad \triangleright d_{3}=O(k / \epsilon)\)
    \(C \leftarrow A_{1} D_{1}, R \leftarrow A_{2} D_{2}, T \leftarrow A_{3} D_{3}\).
    \(U \leftarrow \sum_{i=1}^{k}\left(\left(B_{1} D_{1}\right)^{\dagger}\right)_{i} \otimes\left(\left(B_{2} D_{2}\right)^{\dagger}\right)_{i} \otimes\left(\left(B_{3} D_{3}\right)^{\dagger}\right)_{i}\).
    return \(C, R, T\) and \(U\).
    end procedure
```


### 21.4.8 Face-based selection and decomposition

Previously we provided column-based tensor CURT algorithms, which are algorithms that can select a subset of columns from each of the three dimensions. Here we provide two face-based tensor CURT decomposition algorithms. The first algorithm runs in polynomial time and is a bicriteria algorithm (the number of samples is poly $(k / \epsilon)$ ). The second algorithm needs to start with a rank- $k(1+O(\epsilon))$-approximate solution, which we then show how to combine with our previous algorithm. Both of our algorithms are able to select a subset of column-row faces, a subset of row-tube faces and a subset of column-tube faces. The second
algorithm is able to output $U$, but the first algorithm is not.

### 21.4.8.1 Column-row, column-tube, row-tube face subset selection

```
Algorithm 21.13 Frobenius Norm Tensor Column-row, Row-tube and Tube-column Face
Subset Selection
    procedure FFACECRTSELECTION \((A, n, k, \epsilon) \quad \triangleright\) Theorem 21.4.41
        \(s_{1} \leftarrow s_{2} \leftarrow O(k / \epsilon)\).
        Choose a Gaussian matrix \(S_{1}\) with \(s_{1}\) columns. \(\triangleright\) Definition 21.3.12
        Choose a Gaussian matrix \(S_{2}\) with \(s_{2}\) columns. \(\triangleright\) Definition 21.3.12
        Form matrix \(V_{3}\) by setting the \((i, j)\)-th column to be \(\left(A_{2} S_{2}\right)_{j}\).
    \(D_{3} \leftarrow \operatorname{GENERALIZEDMATRIXROWSUBSETSELECTION}\left(A_{2}, V_{3}, n, n^{2}, s_{1} s_{2}, \epsilon\right)\). \(\triangleright\)
    Algorithm 21.7
    Let \(d_{3}\) denote the number of nonzero entries in \(D_{3} . \quad \triangleright d_{3}=O\left(s_{1} s_{2} / \epsilon\right)\)
    Form matrix \(U_{2}\) by setting the \((i, j)\)-th column to be \(\left(A_{1} S_{1}\right)_{i}\).
    \(D_{2} \leftarrow\) GeneralizedMatrixRowSubsetSelection \(\left(A_{1}, U_{2}, n, n^{2}, s_{1} s_{2}, \epsilon\right)\).
    Let \(d_{2}\) denote the number of nonzero entries in \(D_{2}\). \(\triangleright d_{2}=O\left(s_{1} s_{2} / \epsilon\right)\)
    Form matrix \(W_{1}\) by setting the \((i, j)\)-th column to be \(\left(A\left(I, D_{3}, I\right)_{3}\right)_{j}\).
    \(D_{1} \leftarrow \operatorname{GeneralizedMatrixRowSubsetSelection}\left(A_{3}, W_{1}, n, n^{2}, s_{1} s_{2}, \epsilon\right)\).
        Let \(d_{1}\) denote the number of nonzero entries in \(D_{1}\). \(\triangleright d_{1}=O\left(s_{1} s_{2} / \epsilon\right)\)
        \(T \leftarrow A\left(I, I, D_{1}\right), C \leftarrow A\left(D_{2}, I, I\right)\), and \(R \leftarrow A\left(I, D_{3}, I\right)\).
        return \(C, R\) and \(T\).
    end procedure
```

Theorem 21.4.41. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, there exists an algorithm which takes $O(\mathrm{nnz}(A)) \log n+n^{2}$ poly $(\log n, k, 1 / \epsilon)$ time and outputs three tensors : a subset $C \in \mathbb{R}^{c \times n \times n}$ of row-tube faces of $A$, a subset $R \in \mathbb{R}^{n \times r \times n}$ of column-tube faces of A, and a subset $T \in \mathbb{R}^{n \times n \times t}$ of column-row faces of $A$, where $c=r=t=\operatorname{poly}(k, 1 / \epsilon)$, and for which there exists a tensor $U \in \mathbb{R}^{t n \times c n \times r n}$ for which

$$
\left\|U\left(T_{1}, C_{2}, R_{3}\right)-A\right\|_{F}^{2} \leq(1+\epsilon) \min _{\text {rank }-k}\left\|A^{\prime}-A\right\|_{F}^{2},
$$

or equivalently,

$$
\left\|\sum_{i=1}^{t n} \sum_{j=1}^{c n} \sum_{l=1}^{r n} U_{i, j, l} \cdot\left(T_{1}\right)_{i} \otimes\left(C_{2}\right)_{j} \otimes\left(R_{3}\right)_{l}-A\right\|_{F}^{2} \leq(1+\epsilon)_{\operatorname{rank}-k A^{\prime}}\left\|A^{\prime}-A\right\|_{F}^{2}
$$

Proof. We fix $V^{*} \in \mathbb{R}^{n \times k}$ and $W^{*} \in \mathbb{R}^{n \times k}$. We define $Z_{1} \in \mathbb{R}^{k \times n^{2}}$ where the $i$-th row of $Z_{1}$ is the vector $V_{i} \otimes W_{i}$. Choose a sketching (Gaussian) matrix $S_{1} \in \mathbb{R}^{n^{2} \times s_{1}}$ (Definition 21.3.12), and let $\widehat{U}=A_{1} S_{1}\left(Z_{1} S_{1}\right)^{\dagger} \in \mathbb{R}^{n \times k}$. Following a similar argument as in the previous theorem, we have

$$
\left\|\widehat{U} Z_{1}-A_{1}\right\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}
$$

We fix $\widehat{U}$ and $W^{*}$. We define $Z_{2} \in \mathbb{R}^{k \times n^{2}}$ where the $i$-th row of $Z_{2}$ is the vector $\widehat{U}_{i} \otimes$ $W_{i}^{*}$. Choose a sketching (Gaussian) matrix $S_{2} \in \mathbb{R}^{n^{2} \times s_{2}}$ (Definition 21.3.12), and let $\widehat{V}=$ $A_{2} S_{2}\left(Z_{2} S_{2}\right)^{\dagger} \in \mathbb{R}^{n \times k}$. Following a similar argument as in the previous theorem, we have

$$
\left\|\widehat{V} Z_{2}-A_{2}\right\|_{F}^{2} \leq(1+\epsilon)^{2} \mathrm{OPT}
$$

We fix $\widehat{U}$ and $\widehat{V}$. Note that $\widehat{U}=A_{1} S_{1}\left(Z_{1} S_{1}\right)^{\dagger}$ and $\widehat{V}=A_{2} S_{2}\left(Z_{2} S_{2}\right)^{\dagger}$. We define $Z_{3} \in \mathbb{R}^{k \times n^{2}}$ such that the $i$-th row of $Z_{3}$ is the vector $\widehat{U}_{i} \otimes \widehat{V}_{i}$. Let $z_{3}=s_{1} \cdot s_{2}$. We define $Z_{3}^{\prime} \in \mathbb{R}^{z_{3} \times n^{2}}$ such that, $\forall i \in\left[s_{1}\right], \forall j \in\left[s_{2}\right]$, the $i+(j-1) s_{1}$-th row of $Z_{3}^{\prime}$ is the vector $\left(A_{1} S_{1}\right)_{i} \otimes\left(A_{2} S_{2}\right)_{j}$.

We define $U_{3} \in \mathbb{R}^{n \times z_{3}}$ to be the matrix where the $i+(j-1) s_{1}$-th column is $\left(A_{1} S_{1}\right)_{i}$ and $V_{3} \in \mathbb{R}^{n \times z_{3}}$ to be the matrix where the $i+(j-1) s_{1}$-th column is $\left(A_{2} S_{2}\right)_{j}$. Then $Z_{3}^{\prime}=\left(U_{3}^{\top} \odot V_{3}^{\top}\right)$.

We first have,

$$
\min _{W \in \mathbb{R}^{n \times k}, X \in \mathbb{R}^{k \times z_{3}}}\left\|W X Z_{3}^{\prime}-A_{3}\right\|_{F}^{2} \leq \min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3}-A_{3}\right\|_{F}^{2} \leq(1+\epsilon)^{2} \mathrm{OPT} .
$$

Now consider the following objective function,

$$
\min _{W \in \mathbb{R}^{n \times z_{3}}}\left\|V_{3} \cdot\left(W^{\top} \odot U_{3}^{\top}\right)-A_{2}\right\|_{F}^{2} .
$$

Let $D_{3}$ denote a sampling and rescaling diagonal matrix according to $V_{1} \in \mathbb{R}^{n \times z_{3}}$, let $d_{3}$ denote the number of nonzero entries of $D_{3}$. Then we have

$$
\begin{aligned}
& \min _{W \in \mathbb{R}^{n \times z_{3}}}\left\|D_{3} V_{3} \cdot\left(W^{\top} \odot U_{3}^{\top}\right)-D_{3} A_{2}\right\|_{F}^{2} \\
= & \min _{W \in \mathbb{R}^{n \times z_{3}}}\left\|U_{3} \otimes\left(D_{3} V_{3}\right) \otimes W-A\left(I, D_{3}, I\right)\right\|_{F}^{2} \\
= & \min _{W \in \mathbb{R}^{n \times z_{3}}}\left\|W \cdot\left(U_{3}^{\top} \odot\left(D_{3} V_{3}\right)^{\top}\right)-\left(A\left(I, D_{3}, I\right)\right)_{3}\right\|_{F}^{2},
\end{aligned}
$$

where the first equality follows by retensorizing the objective function, and the second equality follows by flattening the tensor along the third dimension.

Let $\bar{Z}_{3}$ denote $\left(U_{3}^{\top} \odot\left(D_{3} V_{3}\right)^{\top}\right) \in \mathbb{R}^{z_{3} \times n d_{3}}$ and $W^{\prime}=\left(A\left(I, D_{3}, I\right)\right)_{3} \in \mathbb{R}^{n \times n d_{3}}$. Using Theorem 21.4.14, we can find a diagonal matrix $D_{3} \in \mathbb{R}^{n^{2} \times n^{2}}$ with $d_{3}=O\left(z_{3} / \epsilon\right)=O\left(k^{2} / \epsilon^{3}\right)$ nonzero entries such that

$$
\left\|U_{3} \otimes V_{3} \otimes\left(W^{\prime} Z_{3}^{\dagger}\right)-A\right\|_{F}^{2} \leq(1+\epsilon)^{3} \mathrm{OPT}
$$

We define $U_{2}=U_{3} \in \mathbb{R}^{n \times z_{2}}$ with $z_{2}=z_{3}$. We define $W_{2}=W^{\prime} \bar{Z}_{3}^{\dagger} \in \mathbb{R}^{n \times z_{2}}$ with $z_{2}=z_{3}$. We consider,

$$
\min _{V \in \mathbb{R}^{n \times z_{2}}}\left\|U_{2} \cdot\left(V^{\top} \odot W_{2}^{\top}\right)-A_{1}\right\|_{F}^{2}
$$

Let $D_{2}$ denote a sampling and rescaling matrix according to $U_{2}$, and let $d_{2}$ denote the number
of nonzero entries of $D_{2}$. Then, we have

$$
\begin{aligned}
& \min _{V \in \mathbb{R}^{n \times z_{2}}}\left\|D_{2} U_{2} \cdot\left(V^{\top} \odot W_{2}^{\top}\right)-D_{2} A_{1}\right\|_{F}^{2} \\
= & \min _{V \in \mathbb{R}^{n \times z_{2}}}\left\|D_{2} U_{2} \otimes V \otimes W_{2}-A\left(D_{2}, I, I\right)\right\|_{F}^{2} \\
= & \min _{V \in \mathbb{R}^{n \times z_{2}}}\left\|V \cdot\left(W_{2}^{\top} \odot\left(D_{2} U_{2}\right)^{\top}\right)-\left(A\left(D_{2}, I, I\right)\right)_{2}\right\|_{F}^{2},
\end{aligned}
$$

where the first equality follows by retensorizing the objective function, and the second equality follows by flattening the tensor along the second dimension.

Let $\bar{Z}_{2}$ denote $\left(W_{2}^{\top} \odot\left(D_{2} U_{2}\right)^{\top}\right) \in \mathbb{R}^{z_{2} \times n d_{2}}$ and $V^{\prime}=\left(A\left(D_{2}, I, I\right)\right)_{2} \in \mathbb{R}^{n \times n d_{2}}$. Using Theorem 21.4.14, we can find a diagonal matrix $D_{2} \in \mathbb{R}^{n^{2} \times n^{2}}$ with $d_{2}=O\left(z_{2} / \epsilon\right)$ nonzero entries such that

$$
\left\|U_{2} \otimes\left(V^{\prime} \bar{Z}_{2}^{\dagger}\right) \otimes W_{2}-A\right\|_{F}^{2} \leq(1+\epsilon)^{4} \mathrm{OPT}
$$

We define $W_{1}=W_{2} \in \mathbb{R}^{n \times z_{1}}$ with $z_{1}=z_{2}$, and define $V_{1}=\left(V^{\prime} \bar{Z}_{2}^{\dagger}\right) \in \mathbb{R}^{n \times z_{1}}$ with $z_{1}=z_{2}$.

Let $D_{1}$ denote a sampling and rescaling matrix according to $W_{1}$, and let $d_{1}$ denote the number of nonzero entries of $D_{1}$. Then we have

$$
\begin{aligned}
& \min _{U \in \mathbb{R}^{n \times z_{1}}}\left\|D_{1} W_{1} \cdot\left(U^{\top} \odot V_{1}^{\top}\right)-D_{1} A_{3}\right\|_{F}^{2} \\
= & \min _{U \in \mathbb{R}^{n \times z_{1}}}\left\|U \otimes V_{1} \otimes\left(D_{1} W_{1}\right)-A\left(I, I, D_{1}\right)\right\|_{F}^{2} \\
= & \min _{U \in \mathbb{R}^{n \times z_{1}}}\left\|U \cdot\left(V_{1}^{\top} \odot\left(D_{1} W_{1}\right)^{\top}\right)-A\left(I, I, D_{1}\right)_{1}\right\|_{F}^{2}
\end{aligned}
$$

where the first equality follows by unflattening the objective function, and second equality follows by flattening the tensor along the first dimension.

Let $\bar{Z}_{1}$ denote $\left(V_{1}^{\top} \odot\left(D_{1} W_{1}\right)^{\top}\right) \in \mathbb{R}^{z_{1} \times n d_{1}}$, and $U^{\prime}=A\left(I, I, D_{1}\right)_{1} \in \mathbb{R}^{n \times n d_{1}}$. Using Theorem 21.4.14, we can find a diagonal matrix $D_{1} \in \mathbb{R}^{n^{2} \times n^{2}}$ with $d_{1}=O\left(z_{1} / \epsilon\right)$ nonzero entries such that

$$
\left\|\left(U^{\prime} \bar{Z}_{1}^{\dagger}\right) \otimes\left(V_{1}\right) \otimes W_{1}-A\right\|_{F}^{2} \leq(1+\epsilon)^{5} \mathrm{OPT},
$$

which means,

$$
\left\|\left(U^{\prime} \bar{Z}_{1}^{\dagger}\right) \otimes\left(V^{\prime} \bar{Z}_{2}^{\dagger}\right) \otimes\left(W^{\prime} \bar{Z}_{3}^{\dagger}\right)-A\right\|_{F}^{2} \leq(1+\epsilon)^{5} \mathrm{OPT}
$$

Putting $U^{\prime}, V^{\prime}, W^{\prime}$ together completes the proof.

Corollary 21.4.42. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, there exists an algorithm which takes $O(\mathrm{nnz}(A))+n^{2}$ poly $(k, 1 / \epsilon)$ time and outputs three tensors : a subset $C \in \mathbb{R}^{c \times n \times n}$ of row-tube faces of $A$, a subset $R \in \mathbb{R}^{n \times r \times n}$ of column-tube faces of $A$, and $a$ subset $T \in \mathbb{R}^{n \times n \times t}$ of column-row faces of $A$, where $c=r=t=\operatorname{poly}(k, 1 / \epsilon)$, so that there exists a tensor $U \in \mathbb{R}^{t n \times c n \times r n}$ for which

$$
\left\|U\left(T_{1}, C_{2}, R_{3}\right)-A\right\|_{F}^{2} \leq(1+\epsilon) \min _{\text {rank }-k}\left\|A^{\prime}-A\right\|_{F}^{2},
$$

or equivalently,

$$
\left\|\sum_{i=1}^{t n} \sum_{j=1}^{c n} \sum_{l=1}^{r n} U_{i, j, l} \cdot\left(T_{1}\right)_{i} \otimes\left(C_{2}\right)_{j} \otimes\left(R_{3}\right)_{l}-A\right\|_{F}^{2} \leq(1+\epsilon) \min _{\text {rank }-k}\left\|A^{\prime}-A\right\|_{F}^{2}
$$

Proof. If we allow a poly $(k / \epsilon)$ factor increase in running time and a $\operatorname{poly}(k / \epsilon)$ factor increase in the number of faces selected, then instead of using generalized row subset selection, which has running time depending on $\log n$, we can use the technique in Section 21.6 to avoid the $\log n$ factor.

```
Algorithm 21.14 Frobenius Norm (Face-based) CURT Decomposition Algorithm, Optimal
Sample Complexity
    procedure FFACECURTDECOMPOSITION \(\left(A, U_{B}, V_{B}, W_{B}, n, k\right) \quad \triangleright\) Theorem 21.4.43
    \(D_{1} \leftarrow\) GeneralizedMatrixRowSubsetSelection \(\left(A_{3}, W_{B}, n, n^{2}, k, \epsilon\right) . \quad \triangleright\)
    Algorithm 21.7, the number of nonzero entries is \(d_{1}=O(k / \epsilon)\)
    Form \(Z_{1}=V_{B}^{\top} \odot\left(D_{1} W_{B}\right)^{\top}\).
    Form \(\widehat{U}=\left(A\left(I, I, D_{1}\right)\right)_{1} Z_{1}^{\dagger} \in \mathbb{R}^{n \times k}\).
    \(D_{2} \leftarrow\) GeneralizedMatrixRowSubsetSelection \(\left(A_{1}, \widehat{U}, n, n^{2}, k, \epsilon\right) . \quad \triangleright\) The
number of nonzero entries is \(d_{2}=O(k / \epsilon)\)
    Form \(Z_{2}=\left(W_{B}^{\top} \odot\left(D_{2} \widehat{U}\right)\right)\).
    Form \(\widehat{V}=\left(A\left(D_{2}, I, I\right)\right)_{2} Z_{2}^{\dagger} \in \mathbb{R}^{n \times k}\).
    \(D_{3} \leftarrow \operatorname{GenERALIZEDMATRIxROWSUBSETSELECtion}\left(A_{2}, \widehat{V}, n, n^{2}, k, \epsilon\right)\). \(\quad\) The
number of nonzero entries is \(d_{3}=O(k / \epsilon)\)
    Form \(Z_{3}=\widehat{U}^{\top} \odot\left(D_{3} \widehat{V}\right)^{\top}\).
    Form \(\widehat{W}=\left(A\left(I, D_{3}, I\right)\right)_{3}\left(Z_{3}\right)^{\dagger} \in \mathbb{R}^{n \times k}\).
    \(T \leftarrow A\left(I, I, D_{1}\right), C \leftarrow A\left(D_{2}, I, I\right), R \leftarrow A\left(I, D_{3}, I\right)\).
    \(U \leftarrow \sum_{i=1}^{k}\left(\left(Z_{1}\right)^{\dagger}\right)_{i} \otimes\left(\left(Z_{2}\right)^{\dagger}\right)_{i} \otimes\left(\left(Z_{3}\right)^{\dagger}\right)_{i}\).
    return \(C, R, T\) and \(U\).
end procedure
```


### 21.4.8.2 CURT decomposition

Theorem 21.4.43. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, let $k \geq 1$, and let $U_{B}, V_{B}, W_{B} \in$ $\mathbb{R}^{n \times k}$ denote a rank-k, $\alpha$-approximation to $A$. Then there exists an algorithm which takes $O(\mathrm{nnz}(A)) \log n+n^{2}$ poly $(\log n, k, 1 / \epsilon)$ time and outputs three tensors: $C \in \mathbb{R}^{c \times n \times n}$ with rowtube faces from $A, R \in \mathbb{R}^{n \times r \times n}$ with colum-tube faces from $A, T \in \mathbb{R}^{n \times n \times t}$ with column-row faces from $A$, and $a$ (factorization of a) tensor $U \in \mathbb{R}^{t n \times c n \times r n}$ with $\operatorname{rank}(U)=k$ for which $c=r=t=O(k / \epsilon)$ and

$$
\left\|U\left(T_{1}, C_{2}, R_{3}\right)-A\right\|_{F}^{2} \leq(1+\epsilon) \alpha \min _{\text {rank }-k}\left\|A^{\prime}-A\right\|_{F}^{2},
$$

or equivalently,

$$
\left\|\sum_{i=1}^{t n} \sum_{j=1}^{c n} \sum_{l=1}^{r n} U_{i, j, l} \cdot\left(T_{1}\right)_{i} \otimes\left(C_{2}\right)_{j} \otimes\left(R_{3}\right)_{l}-A\right\|_{F}^{2} \leq(1+\epsilon) \alpha \min _{\operatorname{rank}-k A^{\prime}}\left\|A^{\prime}-A\right\|_{F}^{2}
$$

holds with probability 9/10.

Proof. We already have three matrices $U_{B} \in \mathbb{R}^{n \times k}, V_{B} \in \mathbb{R}^{n \times k}$ and $W_{B} \in \mathbb{R}^{n \times k}$ and these three matrices provide a rank- $k, \alpha$-approximation to $A$, i.e.,

$$
\left\|U_{B} \otimes V_{B} \otimes W_{B}-A\right\|_{F}^{2} \leq \alpha \underbrace{\min _{\text {rank }-k A^{\prime}}\left\|A^{\prime}-A\right\|_{F}^{2}}_{\text {OPT }}
$$

We can consider the following problem,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|W_{B} \cdot\left(U^{\top} \odot V_{B}^{\top}\right)-A_{3}\right\|_{F}^{2}
$$

Let $D_{1}$ denote a sampling and rescaling diagonal matrix according to $W_{B}$, and let $d_{1}$ denote the number of nonzero entries of $D_{1}$. Then we have

$$
\begin{aligned}
& \min _{U \in \mathbb{R}^{n \times k}}\left\|\left(D_{1} W_{B}\right) \cdot\left(U^{\top} \odot V_{B}^{\top}\right)-D_{1} A_{3}\right\|_{F}^{2} \\
= & \min _{U \in \mathbb{R}^{n \times k}}\left\|U \otimes V_{B} \otimes D_{1} W_{B}-A\left(I, I, D_{1}\right)\right\|_{F}^{2} \\
= & \min _{U \in \mathbb{R}^{n \times k}}\left\|U \cdot\left(V_{B}^{\top} \odot\left(D_{1} W_{B}\right)^{\top}\right)-\left(A\left(I, I, D_{1}\right)\right)_{1}\right\|_{F}^{2},
\end{aligned}
$$

where the first equality follows by retensorizing the objective function, and the second equality follows by flattening the tensor along the first dimension. Let $Z_{1}$ denote $V_{B}^{\top} \odot\left(D_{1} W_{B}\right)^{\top} \in$ $\mathbb{R}^{k \times n d_{1}}$, and define $\widehat{U}=\left(A\left(I, I, D_{1}\right)\right)_{1} Z_{1}^{\dagger} \in \mathbb{R}^{n \times k}$. Then we have

$$
\left\|\widehat{U} \otimes V_{B} \otimes W_{B}-A\right\|_{F}^{2} \leq(1+\epsilon) \alpha \mathrm{OPT}
$$

In the second step, we fix $\widehat{U}$ and $W_{B}$, and consider the following objective function,

$$
\min _{V \in \mathbb{R}^{n \times k}}\left\|\widehat{U} \cdot\left(V^{\top} \odot W_{B}\right)-A_{1}\right\|_{F}^{2}
$$

Let $D_{2}$ denote a sampling and rescaling matrix according to $\widehat{U}$, and let $d_{2}$ denote the number of nonzero entries of $D_{2}$. Then we have,

$$
\begin{aligned}
& \min _{V \in \mathbb{R}^{n \times k}}\left\|\left(D_{2} \widehat{U}\right) \cdot\left(V^{\top} \odot W_{B}^{\top}\right)-D_{2} A_{1}\right\|_{F}^{2} \\
= & \min _{V \in \mathbb{R}^{n \times k}}\left\|\left(D_{2} \widehat{U}\right) \otimes V \otimes W_{B}-A\left(D_{2}, I, I\right)\right\|_{F}^{2} \\
= & \min _{V \in \mathbb{R}^{n \times k}}\left\|V \cdot\left(W_{B}^{\top} \odot\left(D_{2} \widehat{U}\right)^{\top}\right)-\left(A\left(D_{2}, I, I\right)\right)_{2}\right\|_{F}^{2},
\end{aligned}
$$

where the first equality follows by unflattening the objective function, and the second equality follows by flattening the tensor along the second dimension. Let $Z_{2}$ denote $\left(W_{B}^{\top} \odot\left(D_{2} \widehat{U}\right)^{\top}\right) \in$ $\mathbb{R}^{k \times n d_{2}}$, and define $\widehat{V}=\left(A\left(D_{2}, I, I\right)\right)_{2}\left(Z_{2}\right)^{\dagger} \in \mathbb{R}^{n \times k}$. Then we have,

$$
\left\|\widehat{U} \otimes \widehat{V} \otimes W_{B}-A\right\|_{F}^{2} \leq(1+\epsilon)^{2} \alpha \mathrm{OPT}
$$

In the third step, we fix $\widehat{U}$ and $\widehat{V}$, and consider the following objective function,

$$
\min _{W \in \mathbb{R}^{n \times k}}\left\|\widehat{V} \cdot(W \odot \widehat{U})-A_{2}\right\|_{F}^{2}
$$

Let $D_{3}$ denote a sampling and rescaling matrix according to $\widehat{V}$, and let $d_{3}$ denote the number of nonzero entries of $D_{3}$. Then we have,

$$
\begin{aligned}
& \min _{W \in \mathbb{R}^{n \times k}}\left\|\left(D_{3} \widehat{V}\right) \cdot\left(W^{\top} \odot \widehat{U}^{\top}\right)-D_{3} A_{2}\right\|_{F}^{2} \\
= & \min _{W \in \mathbb{R}^{n \times k}}\left\|\widehat{U} \otimes\left(D_{3} \widehat{V}\right) \otimes W-A\left(I, D_{3}, I\right)\right\|_{F}^{2} \\
= & \min _{W \in \mathbb{R}^{n \times k}}\left\|W \cdot\left(\widehat{U}^{\top} \odot\left(D_{3} \widehat{V}\right)^{\top}\right)-\left(A\left(I, D_{3}, I\right)\right)_{3}\right\|_{F}^{2},
\end{aligned}
$$

where the first equality follows by retensorizing the objective function, and the second equality follows by flattening the tensor along the third dimension. Let $Z_{3}$ denote $\left(\widehat{U}^{\top} \odot\left(D_{3} \widehat{V}\right)^{\top}\right) \in$ $\mathbb{R}^{k \times n d_{3}}$, and define $\widehat{W}=\left(A\left(I, D_{3}, I\right)\right)_{3}\left(Z_{3}\right)^{\dagger}$. Putting it all together, we have,

$$
\|\widehat{U} \otimes \widehat{V} \otimes \widehat{W}-A\|_{F}^{2} \leq(1+\epsilon)^{3} \alpha \mathrm{OPT}
$$

This implies

$$
\left\|\left(A\left(I, I, D_{1}\right)\right)_{1} Z_{1}^{\dagger} \otimes\left(A\left(D_{2}, I, I\right)\right)_{2} Z_{2}^{\dagger} \otimes\left(A\left(I, D_{3}, I\right)\right)_{3} Z_{3}^{\dagger}-A\right\|_{F}^{2} \leq(1+\epsilon)^{3} \alpha \mathrm{OPT}
$$

### 21.4.9 Solving small problems

Theorem 21.4.44. Let $\max _{i}\left\{t_{i}, d_{i}\right\} \leq n$. Given a $t_{1} \times t_{2} \times t_{3}$ tensor $A$ and three matrices: a $t_{1} \times d_{1}$ matrix $T_{1}$, a $t_{2} \times d_{2}$ matrix $T_{2}$, and a $t_{3} \times d_{3}$ matrix $T_{3}$, if for any $\delta>0$ there exists a solution to

$$
\min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k}\left(T_{1} X_{1}\right)_{i} \otimes\left(T_{2} X_{2}\right)_{i} \otimes\left(T_{3} X_{3}\right)_{i}-A\right\|_{F}^{2}:=\mathrm{OPT}
$$

and each entry of $X_{i}$ can be expressed using $O\left(n^{\delta}\right)$ bits, then there exists an algorithm that takes $n^{O(\delta)} \cdot 2^{O\left(d_{1} k+d_{2} k+d_{3} k\right)}$ time and outputs three matrices: $\widehat{X}_{1}, \widehat{X}_{2}$, and $\widehat{X}_{3}$ such that $\left\|\left(T_{1} \widehat{X}_{1}\right) \otimes\left(T_{2} \widehat{X}_{2}\right) \otimes\left(T_{3} \widehat{X}_{3}\right)-A\right\|_{F}^{2}=\mathrm{OPT}$.

Proof. For each $i \in[3]$, we can create $t_{i} \times d_{i}$ variables to represent matrix $X_{i}$. Let $x$ denote this list of variables. Let $B$ denote tensor $\sum_{i=1}^{k}\left(T_{1} X_{1}\right)_{i} \otimes\left(T_{2} X_{2}\right)_{i} \otimes\left(T_{3} X_{3}\right)_{i}$ and let $B_{i, j, l}(x)$ denote an entry of tensor $B$ (which can be thought of as a polynomial written in terms of $x)$. Then we can write the following objective function,

$$
\min _{x} \sum_{i=1}^{t_{1}} \sum_{j=1}^{t_{2}} \sum_{l=1}^{t_{3}}\left(B_{i, j, l}(x)-A_{i, j, l}\right)^{2} .
$$

We slightly modify the above objective function to obtain a new objective function,

$$
\begin{aligned}
& \min _{x, \sigma} \sum_{i=1}^{t_{1}} \sum_{j=1}^{t_{2}} \sum_{l=1}^{t_{3}}\left(B_{i, j, l}(x)-A_{i, j, l}\right)^{2}, \\
& \text { s.t. }\|x\|_{2}^{2} \leq 2^{O\left(n^{\delta}\right)}
\end{aligned}
$$

where the last constraint is unharmful, because there exists a solution that can be written using $O\left(n^{\delta}\right)$ bits. Note that the number of inequality constraints in the above system is
$O(1)$, the degree is $O(1)$, and the number of variables is $v=\left(d_{1} k+d_{2} k+d_{3} k\right)$. Thus by Theorem 21.3.2, the minimum nonzero cost is at least

$$
\left(2^{O\left(n^{\delta}\right)}\right)^{-2^{O(v)}} .
$$

It is clear that the upper bound on the cost is at most $2^{O\left(n^{\delta}\right)}$. Thus the number of binary search steps is at most $\log \left(2^{O\left(n^{\delta}\right)}\right) 2^{O(v)}$. In each step of the binary search, we need to choose a cost $C$ between the lower bound and the upper bound, and write down the polynomial system,

$$
\begin{aligned}
& \sum_{i=1}^{t_{1}} \sum_{j=1}^{t_{2}} \sum_{l=1}^{t_{3}}\left(B_{i, j, l}(x)-A_{i, j, l}\right)^{2} \leq C \\
& \|x\|_{2}^{2} \leq 2^{O\left(n^{\delta}\right)}
\end{aligned}
$$

Using Theorem 21.3.1, we can determine if there exists a solution to the above polynomial system. Since the number of variables is $v$, and the degree is $O(1)$, the number of inequality constraints is $O(1)$. Thus, the running time is

$$
\operatorname{poly}(\text { bitsize }) \cdot(\# \text { constraints } \cdot \text { degree })^{\# \text { variables }}=n^{O(\delta)} 2^{O(v)} .
$$

### 21.5 Extension to general $q$-th order tensors

This section provides the details for our extensions from 3rd order tensors to general $q$-th order tensors. In most practical applications, the order $q$ is a constant. Thus, to simplify the analysis, we use $O_{q}(\cdot)$ to hide dependencies on $q$.

### 21.5.1 Fast sampling of columns according to leverage scores, implicitly

This section explains an algorithm that is able to sample from the leverage scores from the $\odot$ product of $q$ matrices $U_{1}, U_{2}, \cdots, U_{q}$ without explicitly writing down $U_{1} \odot U_{2} \odot$ $\cdots U_{q}$. To build this algorithm we combine TensorSketch, some ideas from [DMIMW12], and some techniques from [AKO11, MW10]. Finally, we improve the running time for sampling columns according to the leverage scores from $\operatorname{poly}(n)$ to $\widetilde{O}(n)$. Given $q$ matrices $U_{1}, U_{2}, \cdots, U_{q}$, with each such matrix $U_{i}$ having size $k \times n_{i}$, we define $A \in \mathbb{R}^{k \times \prod_{i=1}^{q} n_{i}}$ to be the matrix where the $i$-th row of $A$ is the vectorization of $U_{1}^{i} \otimes U_{2}^{i} \otimes \cdots \otimes U_{q}^{i}, \forall i \in[k]$. Naïvely, in order to sample poly $(k, 1 / \epsilon)$ rows from $A$ according to the leverage scores, we need to write down $\prod_{i=1}^{q} n_{i}$ leverage scores. This approach will take at least $\prod_{i=1}^{q} n_{i}$ running time. In the remainder of this section, we will explain how to do it in $O_{q}(n \cdot \operatorname{poly}(k, 1 / \epsilon))$ time for any constant $p$, and $\max _{i \in[q]} n_{i} \leq n$.

Theorem 21.5.1. Given $q$ matrices $U_{1} \in \mathbb{R}^{k \times n_{1}}, U_{2} \in \mathbb{R}^{k \times n_{2}}, \cdots, U_{q} \in \mathbb{R}^{k \times n_{q}}$, let $\max _{i} n_{i} \leq$ $n$. There exists an algorithm that takes $O_{q}\left(n \cdot \operatorname{poly}(k, 1 / \epsilon) \cdot R_{\text {samples }}\right)$ time and samples $R_{\text {samples }}$ columns of $U_{1} \odot U_{2} \odot \cdots \odot U_{q} \in \mathbb{R}^{k \times \prod_{i=1}^{q} n_{i}}$ according to the leverage scores of $U_{1} \odot U_{2} \odot \cdots \odot U_{q}$.

Proof. Let $\max _{i} n_{i} \leq n$. First, choosing $\Pi_{0}$ to be a TensorSketch, we can compute $R^{-1}$ in $O_{q}(n \operatorname{poly}(k, 1 / \epsilon))$ time, where $R$ is the $R$ in a QR-factorization. We want to sample columns from $U_{1} \odot U_{2} \odot \cdots \odot U_{q}$ according to the square of the $\ell_{2}$-norm of each column of $R^{-1}\left(U_{1} \odot U_{2} \odot \cdots U_{q}\right)$. The issue is the number of columns of this matrix is already $\prod_{i=1}^{q} n_{i}$. The goal is to sample columns from $R^{-1}\left(U_{1} \odot U_{2} \odot \cdots U_{q}\right)$ without explicitly computing the square of the $\ell_{2}$-norm of each column.

Similarly as in the proof of Lemma 21.4.32, we have the observation that the following

```
Algorithm 21.15 Fast Tensor Leverage Score Sampling, for General \(q\)-th Order
    procedure Fast TensorLeverageScoreGeneralorder \(\left(\left\{U_{i}\right\}_{i \in[q]},\left\{n_{i}\right\}_{i \in[q]}, k, \epsilon, R_{\text {samples }}\right)\)
    \(\triangleright\) Theorem 21.5.1
        \(s_{1} \leftarrow \operatorname{poly}(k, 1 / \epsilon)\).
        Choose \(\Pi_{0}, \Pi_{1} \in \mathbb{R}^{n_{1} n_{2} \cdots n_{q} \times s_{1}}\) to each be a TensorSketch. \(\triangleright\) Definition 21.3.20
        Compute \(R^{-1} \in \mathbb{R}^{k \times k}\) by using \(\left(U_{1} \odot U_{2} \odot \cdots \odot U_{q}\right) \Pi_{0} . \quad \triangleright U_{i} \in \mathbb{R}^{k \times n_{i}}, \forall i \in[q]\)
        \(V_{0} \leftarrow R^{-1}, n_{0} \leftarrow k\).
        for \(i=1 \rightarrow\left[n_{0}\right]\) do
            \(\alpha_{i} \leftarrow\left\|\left(V_{0}\right)^{i}\left(\left(U_{1} \odot U_{2} \odot \cdots \odot U_{q}\right) \Pi_{1}\right)\right\|_{2}^{2}\).
        end for
        for \(r=1 \rightarrow R_{\text {samples }}\) do
            Sample \(\widehat{j}_{0}\) from \(\left[n_{0}\right.\) ] with probability \(\alpha_{i} / \sum_{i^{\prime}=1}^{n_{0}} \alpha_{i^{\prime}}\).
            for \(l=1 \rightarrow q-1\) do
                \(s_{l+1} \leftarrow O_{q}(\operatorname{poly}(k, 1 / \epsilon))\).
                Choose \(\Pi_{l+1} \in \mathbb{R}^{n_{l+1} \cdots n_{q} \times s_{l+1}}\) to be a TensorSketch.
                for \(j_{l}=1 \rightarrow\left[n_{l}\right]\) do \(\quad \triangleright\) Form \(V_{l} \in \mathbb{R}^{n_{l} \times k}\)
                    \(\left(V_{l}\right)^{j_{l}} \leftarrow\left(V_{l-1}\right)^{\hat{j}_{l-1}} \circ\left(U_{l}\right)_{j_{l}}^{\top}\).
                    end for
                    for \(i=1 \rightarrow n_{q}\) do
                    \(\beta_{i} \leftarrow\left\|\left(V_{l}\right)^{i}\left(\left(U_{l+1} \odot \cdots \odot U_{q}\right) \Pi_{l+1}\right)\right\|_{2}^{2}\).
            end for
            Sample \(\widehat{j}_{l}\) from \(\left[n_{l}\right]\) with probability \(\beta_{i} / \sum_{i^{\prime}=1}^{n_{l}} \beta_{i^{\prime}}\).
            end for
            for \(i=1 \rightarrow n_{q}\) do
                \(\beta_{i} \leftarrow\left|\left(V_{q-1}\right)^{\hat{j}_{q-1}}\left(U_{q}\right)_{i}\right|^{2}\).
            end for
            Sample \(\widehat{j}_{q}\) from \(\left[n_{q}\right]\) with probability \(\beta_{i} / \sum_{i^{\prime}=1}^{n_{q}} \beta_{i^{\prime}}\).
            \(\mathcal{S} \leftarrow \mathcal{S} \cup\left(\widehat{j}_{1}, \cdots, \widehat{j}_{q}\right)\).
        end for
        Convert \(\mathcal{S}\) into a diagonal matrix \(D\) with at most \(R_{\text {samples }}\) nonzero entries.
        return \(D\). \(\quad\) Diagonal matrix \(D \in \mathbb{R}^{n_{1} n_{2} \cdots n_{q} \times n_{1} n_{2} \cdots n_{q}}\)
    end procedure
```

two sampling procedures are equivalent in terms of sampling a column of a matrix: (1) We sample a single entry from matrix $R^{-1}\left(U_{1} \odot U_{2} \odot \cdots \odot U_{q}\right)$ proportional to its squared value,
(2) We sample a column from matrix $R^{-1}\left(U_{1} \odot U_{2} \odot \cdots \odot U_{q}\right)$ proportional to its squared $\ell_{2}$-norm. Let the $\left(i, j_{1}, j_{2}, \cdots, j_{q}\right)$-th entry denote the entry in the $i$-th row and the $j$-th column, where

$$
j=\sum_{l=1}^{q-1}\left(j_{l}-1\right) \prod_{t=l+1}^{q} n_{t}+j_{q}
$$

Similarly to Equation (21.18), we can show, for a particular column $j$,
$\operatorname{Pr}[$ we sample an entry from the $j$-th column of matrix $]=\operatorname{Pr}$ [we sample the $j$-th column of a matrix].

Thus, it is sufficient to show how to sample a single entry from matrix $R^{-1}\left(U_{1} \odot U_{2} \odot \cdots \odot U_{q}\right)$ proportional to its squared value without writing down all the entries of the $k \times \prod_{i=1}^{q} n_{i}$ matrix.

Let $V_{0}$ denote $R^{-1}$. Let $n_{0}$ denote the number of rows of $V_{0}$.
In the next few paragraphs, we describe a sampling procedure (procedure Fast TenSORLEVERAGESCOREGENERALORDER in Algorithm 21.15) which first samples $\widehat{j}_{0}$ from $\left[n_{0}\right]$, then samples $\widehat{j}_{1}$ from $\left[n_{1}\right], \cdots$, and at the end samples $\widehat{j}_{q}$ from $\left[n_{q}\right]$.

In the first step, we want to sample $\widehat{j}_{0}$ from $\left[n_{0}\right]$ proportional to the squared $\ell_{2}$-norm of that row. To do this efficiently, we choose $\Pi_{1} \in \mathbb{R}^{\Pi_{i=1}^{q} n_{i} \times s_{1}}$ to be a TENSORSKETCH to sketch on the right of $V_{0}\left(U_{1} \odot U_{2} \odot \cdots \odot U_{q}\right)$. By Section 21.3.10, as long as $s_{1}=O_{q}(\operatorname{poly}(k, 1 / \epsilon))$, then $\Pi_{1}$ is a $(1 \pm \epsilon)$-subspace embedding matrix. Thus with probability $1-1 / \Omega(q)$, for all $i \in\left[n_{0}\right]$,

$$
\left\|\left(V_{0}\right)^{i}\left(\left(U_{1} \odot U_{2} \odot \cdots \odot U_{q}\right) \Pi_{1}\right)\right\|_{2}^{2}=(1 \pm \epsilon)\left\|\left(V_{0}\right)^{i}\left(\left(U_{1} \odot U_{2} \odot \cdots \odot U_{q}\right)\right)\right\|_{2}^{2}
$$

which means we can sample $\widehat{j}_{0}$ from $\left[n_{0}\right]$ in $O_{q}(n \operatorname{poly}(k, 1 / \epsilon))$ time.

In the second step, we have already obtained $\widehat{j}_{0}$. Using that row of $V_{0}$ with $U_{1}$, we can form a new matrix $V_{1} \in \mathbb{R}^{n_{1} \times k}$ in the following sense,

$$
\left(V_{1}\right)^{i}=\left(V_{0}\right)^{\hat{j}_{0}} \circ\left(U_{1}\right)_{i}^{\top}, \forall i \in\left[n_{1}\right],
$$

where $\left(V_{1}\right)^{i}$ denotes the $i$-th row of matrix $V_{1},\left(V_{0}\right)^{\hat{j}_{0}}$ denotes the $\widehat{j}_{0}$-th row of $V_{0}$ and $\left(U_{1}\right)_{i}$ is the $i$-th column of $U_{1}$. Another important observation is, the entry in the $\left(j_{1}, j_{2}, \cdots, j_{q}\right)$-th coordinate of vector $\left(V_{0}\right)^{\hat{j}_{0}}\left(U_{1} \odot U_{2} \odot \cdots \odot U_{q}\right)$ is the same as the entry in the $j_{1}$-th row and $\left(j_{2}, \cdots, j_{q}\right)$-th column of matrix $V_{1}\left(U_{2} \odot U_{3} \odot \cdots \odot U_{q}\right)$. Thus, sampling $j_{1}$ is equivalent to sampling $j_{1}$ from the new matrix $V_{1}\left(U_{2} \odot U_{3} \odot \cdots \odot U_{q}\right)$ proportional to the squared $\ell_{2}$-norm of that row. We still have the computational issue that the length of the row vector is very long. To deal with this, we can choose $\Pi_{2} \in \mathbb{R}^{\Pi_{i=2}^{q} n_{i} \times s_{2}}$ to be a TEnsorSketch to multiply on the right of $V_{1}\left(U_{2} \odot U_{3} \odot \cdots \odot U_{q}\right)$.

By Section 21.3.10, as long as $s_{2}=O_{q}(\operatorname{poly}(k, 1 / \epsilon))$, then $\Pi_{2}$ is a $(1 \pm \epsilon)$-subspace embedding matrix. Thus with probability $1-1 / \Omega(q)$, for all $i \in\left[n_{1}\right]$,

$$
\left\|\left(V_{1}\right)^{i}\left(\left(U_{2} \odot \cdots \odot U_{q}\right) \Pi_{2}\right)\right\|_{2}^{2}=(1 \pm \epsilon)\left\|\left(V_{1}\right)^{i}\left(\left(U_{2} \odot \cdots \odot U_{q}\right)\right)\right\|_{2}^{2}
$$

which means we can sample $\widehat{j}_{1}$ from $\left[n_{1}\right]$ in $O_{q}(n \operatorname{poly}(k, 1 / \epsilon))$ time.
We repeat the above procedure until we obtain each of $\widehat{j}_{0}, \widehat{j}_{1}, \cdots, \widehat{j}_{q}$. Note that the last one, $\widehat{j}_{q}$, is easier, since the length of the vector is already small enough, and so we do not need to use TensorSketch for it.

By Section 21.3.10, the time for multiplying by TensorSketch is $O_{q}(n \operatorname{poly}(k, 1 / \epsilon))$. Setting $\epsilon$ to be a small constant, and taking a union bound over $O(q)$ events completes the proof.

Lemma 21.5.2. Given $A \in \mathbb{R}^{n_{0} \times \prod_{i=1}^{q} n_{i}}, U_{1}, U_{2}, \cdots, U_{q} \in \mathbb{R}^{k \times n}$, for any $\epsilon>0$, there exists an algorithm that runs in $O(n \cdot \operatorname{poly}(k, 1 / \epsilon))$ time and outputs a diagonal matrix $D \in \mathbb{R}^{\Pi_{i=1}^{q} n_{i} \times \prod_{i=1}^{q} n_{i}}$ with $m=O(k \log k+k / \epsilon)$ nonzero entries such that,

$$
\left\|\widehat{U}\left(U_{1} \odot U_{2} \odot \cdots \odot U_{q}\right)-A\right\|_{F}^{2} \leq(1+\epsilon) \min _{U \in \mathbb{R}^{n \times k}}\left\|U\left(U_{1} \odot U_{2} \odot \cdots \odot U_{q}\right)-A\right\|_{F}^{2},
$$

holds with probability at least 0.999, where $\widehat{U}$ denotes the optimal solution of

$$
\min _{U \in \mathbb{R}^{n} \times k}\left\|U\left(U_{1} \odot U_{2} \odot \cdots \odot U_{q}\right) D-A D\right\|_{F}^{2}
$$

Proof. This follows by combining Theorem 21.5.1, Corollary 21.4.30, and Lemma 21.4.31.

### 21.5.2 General iterative existential proof

```
Algorithm 21.16 General \(q\)-th Order Iterative Existential Proof
    procedure GeneraliterativeExistentialProof \((A, n, k, q, \epsilon) \quad \triangleright\) Section 21.5.2
        Fix \(U_{1}^{*}, U_{2}^{*}, \cdots, U_{q}^{*} \in \mathbb{R}^{n \times k}\).
        for \(i=1 \rightarrow q\) do
            Choose sketching matrix \(S_{i} \in \mathbb{R}^{q^{q-1} \times s_{i}}\) with \(s_{i}=O_{q}(k / \epsilon)\).
            Define \(Z_{i} \in \mathbb{R}^{k \times n^{q-1}}\) to be \(\underset{j<i}{\odot} \widehat{U}_{j}^{\top} \odot \underset{j^{\prime}>i}{\odot} U_{j^{\prime}}^{* \top}\).
            Let \(A_{i}\) denote the matrix obtained by flattening tensor \(A\) along the \(i\)-th dimension.
            Define \(\widehat{U}_{i}\) to be \(A_{i} S_{i}\left(Z_{i} S_{i}\right)^{\dagger}\).
        end for
        return \(\widehat{U}_{1}, \widehat{U}_{2}, \cdots, \widehat{U}_{q}\).
    end procedure
```

Given a $q$-th order tensor $A \in \mathbb{R}^{n \times n \times \cdots \times n}$, we fix $U_{1}^{*}, U_{2}^{*}, \cdots, U_{q}^{*} \in \mathbb{R}^{n \times k}$ to be the best rank- $k$ solution (if it does not exist, then we replace it by a good approximation, as discussed). We define OPT $=\left\|U_{1}^{*} \otimes U_{2}^{*} \otimes \cdots \otimes U_{q}^{*}-A\right\|_{F}^{2}$. Our iterative proof works as follows. We first obtain the objective function,

$$
\min _{U_{1} \in \mathbb{R}^{n \times k}}\left\|U_{1} \cdot Z_{1}-A_{1}\right\|_{F}^{2} \leq \mathrm{OPT}
$$

where $A_{1}$ is a matrix obtained by flattening tensor $A$ along the first dimension, $Z_{1}=\left(U_{2}^{* \top} \odot\right.$ $\left.U_{3}^{* \top} \odot \cdots \odot U_{q}^{* \top}\right)$ denotes a $k \times n^{q-1}$ matrix. Choosing $S_{1} \in \mathbb{R}^{q^{q-1} \times s_{1}}$ to be a Gaussian sketching matrix with $s_{1}=O(k / \epsilon)$, we obtain a smaller problem,

$$
\min _{U_{1} \in \mathbb{R}^{n \times k}}\left\|U_{1} \cdot Z_{1} S_{1}-A_{1} S_{1}\right\|_{F}^{2}
$$

We define $\widehat{U}_{1}$ to be $A_{1} S_{1}\left(Z_{1} S_{1}\right)^{\dagger} \in \mathbb{R}^{n \times k}$, which gives,

$$
\left\|\widehat{U}_{1} \cdot Z_{1}-A_{1}\right\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}
$$

After retensorizing the above, we have,

$$
\left\|\widehat{U}_{1} \otimes U_{2}^{*} \otimes \cdots \otimes U_{q}^{*}-A\right\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}
$$

In the second round, we fix $\widehat{U}_{1}, U_{3}^{*}, \cdots, U_{q}^{*} \in \mathbb{R}^{n \times k}$, and choose $S_{2} \in \mathbb{R}^{n^{q-1} \times s_{2}}$ to be a Gaussian sketching matrix with $s_{2}=O(k / \epsilon)$. We define $Z_{2} \in \mathbb{R}^{k \times n^{q-1}}$ to be $\left(\widehat{U}_{1}^{\top} \odot U_{3}^{* \top} \odot\right.$ $\left.\cdots \odot U_{q}^{* \top}\right)$. We define $\widehat{U}_{2}$ to be $A_{2} S_{2}\left(Z_{2} S_{2}\right)^{\dagger} \in \mathbb{R}^{n \times k}$. Then, we have

$$
\left\|\widehat{U}_{1} \otimes \widehat{U}_{2} \otimes U_{3}^{*} \otimes \cdots \otimes U_{q}^{*}-A\right\|_{F}^{2} \leq(1+\epsilon)^{2} \mathrm{OPT}
$$

We repeat the above process, where in the $i$-th round we fix $\widehat{U}_{1}, \cdots, \widehat{U}_{i-1}, U_{i+1}^{*}, \cdots, U_{q}^{*} \in$ $\mathbb{R}^{n \times k}$, and choose $S_{i} \in \mathbb{R}^{n^{q-1} \times s_{i}}$ to be a Gaussian sketching matrix with $s_{i}=O(k / \epsilon)$. We define $Z_{i} \in \mathbb{R}^{k \times n^{q-1}}$ to be $\left(\widehat{U}_{1}^{\top} \odot \cdots \odot \widehat{U}_{i-1}^{\top} \odot U_{i+1}^{* \top} \odot \cdots \odot U_{q}^{* \top}\right)$. We define $\widehat{U}_{i}$ to be $A_{i} S_{i}\left(Z_{i} S_{i}\right)^{\dagger} \in \mathbb{R}^{n \times k}$. Then, we have

$$
\left\|\widehat{U}_{1} \otimes \cdots \otimes \widehat{U}_{i-1} \otimes \widehat{U}_{i} \otimes U_{i+1}^{*} \otimes \cdots \otimes U_{q}^{*}-A\right\|_{F}^{2} \leq(1+\epsilon)^{2} \mathrm{OPT}
$$

At the end of the $q$-th round, we have

$$
\left\|\widehat{U}_{1} \otimes \cdots \otimes \widehat{U}_{q}-A\right\|_{F}^{2} \leq(1+\epsilon)^{q} \mathrm{OPT}
$$

Replacing $\epsilon=\epsilon^{\prime} /(2 q)$, we obtain

$$
\left\|\widehat{U}_{1} \otimes \cdots \otimes \widehat{U}_{q}-A\right\|_{F}^{2} \leq\left(1+\epsilon^{\prime}\right) \mathrm{OPT}
$$

where for all $i \in[q], s_{i}=O\left(k q / \epsilon^{\prime}\right)=O_{q}\left(k / \epsilon^{\prime}\right)$.

### 21.5.3 General input sparsity reduction

This section shows how to extend the input sparsity reduction from third order tensors to general $q$-th order tensors. Given a tensor $A \in \mathbb{R}^{n \times n \times \cdots \times n}$ and $q$ matrices, for each $i \in[q]$, matrix $V_{i}$ has size $V_{i} \in \mathbb{R}^{n \times b_{i}}$, with $b_{i} \leq \operatorname{poly}(k, 1 / \epsilon)$. We choose a batch of sparse embedding matrices $T_{i} \in \mathbb{R}^{t_{i} \times n}$. Define $\widehat{V}_{i}=T_{i} V_{i}$, and $C=A\left(T_{1}, T_{2}, \cdots, T_{q}\right)$. Thus we have with probability $99 / 100$, for any $\alpha \geq 0$, for all $\left\{X_{i}, X_{i}^{\prime} \in \mathbb{R}^{b_{i} \times k}\right\}_{i \in[q]}$, if

$$
\left\|\widehat{V}_{1} X_{1}^{\prime} \otimes \widehat{V}_{2} X_{2}^{\prime} \otimes \cdots \otimes \widehat{V}_{q} X_{q}^{\prime}-C\right\|_{F}^{2} \leq \alpha\left\|\widehat{V}_{1} X_{1} \otimes \widehat{V}_{2} X_{2} \otimes \cdots \otimes \widehat{V}_{q} X_{q}-C\right\|_{F}^{2},
$$

then

$$
\left\|V_{1} X_{1}^{\prime} \otimes V_{2} X_{2}^{\prime} \otimes \cdots \otimes V_{q} X_{q}^{\prime}-A\right\|_{F}^{2} \leq(1+\epsilon) \alpha\left\|V_{1} X_{1} \otimes V_{2} X_{2} \otimes \cdots \otimes V_{q} X_{q}-A\right\|_{F}^{2}
$$

where $t_{i}=O_{q}\left(\operatorname{poly}\left(b_{i}, 1 / \epsilon\right)\right)$.

```
Algorithm 21.17 General \(q\)-th Order Input Sparsity Reduction
    procedure GeneralinputSparsityReduction \(\left(A,\left\{V_{i}\right\}_{i \in[q]}, n, k, q, \epsilon\right)\)
    Section 21.5.3
        for \(i=1 \rightarrow q\) do
            Choose sketching matrix \(T_{i} \in \mathbb{R}^{t_{i} \times n}\) with \(t_{i}=\operatorname{poly}(k, q, 1 / \epsilon)\).
            \(\widehat{V}_{i} \leftarrow T_{i} V_{i}\).
        end for
        \(C \leftarrow A\left(T_{1}, T_{2}, \cdots, T_{q}\right)\).
        return \(\left\{\widehat{V}_{i}\right\}_{i \in[q]}, C\).
    end procedure
```


### 21.5.4 Bicriteria algorithm

This section explains how to extend the bicriteria algorithm from third order tensors (Section 21.4.4) to general $q$-th order tensors. Given any $q$-th order tensor $A \in \mathbb{R}^{n \times n \times \cdots \times n}$, we can output a rank- $r$ tensor (or equivalently $q$ matrices $U_{1}, U_{2}, \cdots, U_{q} \in \mathbb{R}^{n \times r}$ ) such that,

$$
\left\|U_{1} \otimes U_{2} \otimes \cdots \otimes U_{q}-A\right\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}
$$

where $r=O_{q}\left((k / \epsilon)^{q-1}\right)$ and the algorithm takes $O_{q}(\operatorname{nnz}(A)+n \cdot \operatorname{poly}(k, 1 / \epsilon))$.

```
Algorithm 21.18 General \(q\)-th Order Bicriteria Algorithm
    procedure GeneralBicriteriaAlgorithm \((A, n, k, q, \epsilon) \quad \triangleright\) Section 21.5.4
        for \(i=2 \rightarrow q\) do
            Choose sketching matrix \(S_{i} \in \mathbb{R}^{n^{q-1} \times s_{i}}\) with \(s_{i}=O(k q / \epsilon)\).
            Choose sketching matrix \(T_{i} \in \mathbb{R}^{t_{i} \times n}\) with \(t_{i}=\operatorname{poly}(k, q, 1 / \epsilon)\).
            Form matrix \(\widehat{U}_{i}\) by setting \(\left(j_{2}, j_{3}, \cdots, j_{q}\right)\)-th column to be \(\left(A_{i} S_{i}\right)_{j_{i}}\).
        end for
        Solve \(\min _{U_{1}}\left\|U_{1} B-\left(A\left(I, T_{2}, \cdots, T_{q}\right)\right)_{1}\right\|_{F}^{2}\).
        return \(\left\{\widehat{U}_{i}\right\}_{i \in[q]}\).
    end procedure
```


### 21.5.5 CURT decomposition

This section extends the tensor CURT algorithm from 3rd order tensors (Section 21.4.7) to general $q$-th order tensors. Given a $q$-th order tensor $A \in \mathbb{R}^{n \times n \times \cdots \times n}$ and a batch of matrices $U_{1}, U_{2}, \cdots, U_{q} \in \mathbb{R}^{n \times k}$, we iteratively apply the proof in Theorem 21.4.39 (or Theorem 21.4.40) $q$ times. Then for each $i \in[q]$, we are able to select $d_{i}$ columns from the $i$-th dimension of tensor $A$ (let $C_{i}$ denote those columns) and also find a tensor $U \in \mathbb{R}^{d_{1} \times d_{2} \times \cdots \times d_{q}}$ such that,

$$
\left\|U\left(C_{1}, C_{2}, \cdots, C_{q}\right)-A\right\|_{F}^{2} \leq(1+\epsilon)\left\|U_{1} \otimes U_{2} \otimes \cdots \otimes U_{q}-A\right\|_{F}^{2},
$$

where either $d_{i}=O_{q}(k \log k+k / \epsilon)$ (similar to Theorem 21.4.39) or $d_{i}=O_{q}(k / \epsilon)$ (similar to Theorem 21.4.40).

```
Algorithm 21.19 General \(q\)-th Order CURT Decomposition
    procedure GeneralcURTDecomposition \(\left(A,\left\{U_{i}\right\}_{i \in[q]}, n, k, q, \epsilon\right) \quad \triangleright\) Section 21.5.5
        for \(i=1 \rightarrow q\) do
            Form \(B_{i}=\underset{j<i}{\odot} \widehat{U}_{j}^{\top} \odot \underset{j>i}{\odot} U_{j}^{\top} \in \mathbb{R}^{k \times n^{q-1}}\).
            if fast \(=\) true then \(\quad \triangleright\) Optimal running time
            \(\epsilon_{0} \leftarrow 0.01\).
            \(d_{i} \leftarrow O_{q}(k \log k+k / \epsilon)\).
            \(D_{i} \leftarrow\)
    \(\left(\left\{\widehat{U}_{j}\right\}_{j<i},\left\{U_{j}\right\}_{j>i}, n, k, \epsilon_{0}, d_{i}\right) . \quad \triangleright\) Algorithm 21.15
                                    FastTensorLeverageScoreGeneralorder
            else \(\triangleright\) Optimal sample complexity
                    \(\epsilon_{0} \leftarrow O_{q}(\epsilon)\).
                            \(D_{i} \quad \leftarrow \quad\) GeneralizedMatrixRowSubsetSelection
    \(\left(A_{i}^{\top}, B_{i}^{\top}, n^{q-1}, n, k, \epsilon_{0}\right) . \quad \triangleright\) Algorithm 21.4.5, \(d_{i}=O_{q}(k / \epsilon)\).
            end if
            \(\widehat{U}_{i} \leftarrow A_{i} D_{i}\left(B_{i} D_{i}\right)^{\dagger}\).
            \(C_{i} \leftarrow A_{i} D_{i}\).
        end for
        \(U \leftarrow\left(B_{1} D_{1}\right)^{\dagger} \otimes\left(B_{2} D_{2}\right)^{\dagger} \otimes \cdots \otimes\left(B_{q} D_{q}\right)^{\dagger}\).
        return \(\left\{C_{i}\right\}_{i \in[q]}, U\).
    end procedure
```


### 21.6 Matrix CUR decomposition

There is a long line of research on matrix CUR decomposition under operator, Frobenius or recently, entry-wise $\ell_{1}$ norm [DMM08, BMD09, DR10, BDM11, BW14, SWZ17]. We provide the first algorithm that runs in $n n z(A)$ time, which improves the previous best matrix CUR decomposition algorithm under Frobenius norm [BW14].

### 21.6.1 Algorithm

```
Algorithm 21.20 Optimal Matrix CUR Decomposition Algorithm
    procedure OptimalMatrixCUR \((A, n, k, \epsilon) \quad \triangleright\) Theorem 21.6.1
        \(\epsilon^{\prime} \leftarrow 0.1 \epsilon . \epsilon^{\prime \prime} \leftarrow 0.001 \epsilon^{\prime}\).
        \(\widehat{U} \leftarrow \operatorname{SpARSESVD}\left(A, k, \epsilon^{\prime}\right) . \quad \triangleright \widehat{U} \in \mathbb{R}^{n \times k}\)
        Choose \(S_{1} \in \mathbb{R}^{n \times n}\) to be a sampling and rescaling diagonal matrix according to the
    leverage scores of \(\widehat{U}\) with \(s_{1}=O\left(\epsilon^{-2} k \log k\right)\) nonzero entries.
        \(R, Y \leftarrow \operatorname{GENERALIZEDMATRIXROWSUBSEtSELECTION}\left(S_{1} A, S_{1} \widehat{U}, s_{1}, n, k, \epsilon^{\prime \prime}\right) . \quad \triangleright\)
    Algorithm 21.7, \(R \in \mathbb{R}^{r \times n}, Y \in \mathbb{R}^{k \times r}\) and \(r=O(k / \epsilon)\)
        \(\widehat{V} \leftarrow Y R \in \mathbb{R}^{k \times n}\).
        Choose \(S_{2}^{\top} \in \mathbb{R}^{n \times n}\) to be a sampling and rescaling diagonal matrix according to the
    leverage scores of \(\widehat{V}^{\top} \in \mathbb{R}^{n \times k}\) with \(s_{2}=O\left(\epsilon^{-2} k \log k\right)\) nonzero entries.
        \(C^{\top}, Z^{\top} \quad \leftarrow \quad\) GeneralizedMatrixRowSubsetSelection
    \(\left(\left(A S_{2}\right)^{\top},\left(\widehat{V} S_{2}\right)^{\top}, s_{2}, n, k, \epsilon^{\prime \prime}\right) . \quad \triangleright\) Algorithm 21.7, \(C \in \mathbb{R}^{n \times c}, Z \in \mathbb{R}^{c \times k}\), and \(c=O(k / \epsilon)\)
        \(U \leftarrow Z Y . \quad \triangleright U \in \mathbb{R}^{c \times r}\) and \(\operatorname{rank}(U)=k\)
        return \(C, U, R\).
    end procedure
```

Theorem 21.6.1. Given matrix $A \in \mathbb{R}^{n \times n}$, for any $k \geq 1$ and $\epsilon \in(0,1)$, there exists an algorithm that takes $O(\mathrm{nnz}(A)+n$ poly $(k, 1 / \epsilon))$ time and outputs three matrices $C \in \mathbb{R}^{n \times c}$ with c columns from $A, R \in \mathbb{R}^{r \times n}$ with $r$ rows from $A$, and $U \in \mathbb{R}^{c \times r}$ with $\operatorname{rank}(U)=k$ such that $r=c=O(k / \epsilon)$ and,

$$
\|C U R-A\|_{F}^{2} \leq(1+\epsilon) \min _{\text {rank }-k}\left\|A_{k}-A\right\|_{F}^{2},
$$

holds with probability at least 9/10.

Proof. We define

$$
\mathrm{OPT}=\min _{\text {rank }-k}\left\|A_{k}-A\right\|_{F}^{2}
$$

We first compute $\widehat{U} \in \mathbb{R}^{n \times k}$ by using the result of [CW13], so that $\widehat{U}$ satisfies:

$$
\begin{equation*}
\min _{X \in \mathbb{R}^{k \times n}}\|\widehat{U} X-A\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT} . \tag{21.25}
\end{equation*}
$$

This step can be done in $O(n n z(A)+n \operatorname{poly}(k, 1 / \epsilon))$ time.
We choose $S_{1} \in \mathbb{R}^{n \times n}$ to be a sampling and rescaling diagonal matrix according to the leverage scores of $\widehat{U}$, where here $s_{1}=O\left(\epsilon^{-2} k \log k\right)$ is the number of samples. This step also can be done in $O(n \operatorname{poly}(k, 1 / \epsilon))$ time.

We run GeneralizedMatrixRowSubsetSelection(Algorithm 21.7) on matrices $S_{1} A$ and $S_{1} \widehat{U}$. Then we obtain two new matrices $R$ and $Y$, where $R$ contains $r=O(k / \epsilon)$ rows of $S_{1} A$ and $Y$ has size $k \times r$. According to Theorem 21.4.14 and Corollary 21.4.15, this step takes $n$ poly $(k, 1 / \epsilon)$ time.

We construct $\widehat{V}=Y R$, and choose $S_{2}^{\top}$ to be another sampling and rescaling diagonal matrix according to the leverage scores of $\widehat{V}^{\top}$ with $s_{2}=O\left(\epsilon^{-2} k \log k\right)$ nonzero entries. As in the case of constructing $S_{1}$, this step can be done in $O(n \operatorname{poly}(k, 1 / \epsilon))$ time.

We run GeneralizedMatrixRowSubsetSelection(Algorithm 21.7) on matrices $\left(A S_{2}\right)^{\top}$ and $\left(\widehat{V} S_{2}\right)^{\top}$. Then we can obtain two new matrices $C^{\top}$ and $Y^{\top}$, where $C^{\top}$ contains $c=O(k / \epsilon)$ rows of $\left(A S_{2}\right)^{\top}$ and $Z^{\top}$ has size $k \times c$. According to Theorem 21.4.14 and Corollary 21.4.15, this step takes $n$ poly $(k, 1 / \epsilon)$ time.

Thus, overall the running time is $O(\mathrm{nnz}(A)+n \operatorname{poly}(k, 1 / \epsilon))$.

Correctness. Let

$$
X^{*}=\arg \min _{X \in \mathbb{R}^{n \times k}}\|X \widehat{V}-A\|_{F}^{2}
$$

According to Corollary 21.4.15,

$$
\left\|C Z \widehat{V} S_{2}-A S_{2}\right\|_{F}^{2} \leq\left(1+\epsilon^{\prime \prime}\right) \min _{X \in \mathbb{R}^{n \times k}}\left\|X \widehat{V} S_{2}-A S_{2}\right\|_{F}^{2} \leq\left(1+\epsilon^{\prime \prime}\right)\left\|X^{*} \widehat{V} S_{2}-A S_{2}\right\|_{F}^{2}
$$

According to Theorem 21.6.5, $\epsilon^{\prime \prime}=0.001 \epsilon^{\prime}$,

$$
\begin{equation*}
\|C Z \widehat{V}-A\|_{F}^{2} \leq\left(1+\epsilon^{\prime}\right)\left\|X^{*} \widehat{V}-A\right\|_{F}^{2} \tag{21.26}
\end{equation*}
$$

Let

$$
\tilde{X}=\arg \min _{X \in \mathbb{R}^{k \times n}}\|\widehat{U} X-A\|_{F}^{2}
$$

According to Corollary 21.4.15,

$$
\left\|S_{1} \widehat{U} Y R-S_{1} A\right\|_{F}^{2} \leq\left(1+\epsilon^{\prime \prime}\right) \min _{X \in \mathbb{R}^{k \times n}}\left\|S_{1} \widehat{U} X-S_{1} A\right\|_{F}^{2} \leq\left(1+\epsilon^{\prime \prime}\right)\left\|S_{1} \widehat{U} \widetilde{X}-S_{1} A\right\|_{F}^{2}
$$

According to Theorem 21.6.5, since $\epsilon^{\prime \prime}=0.001 \epsilon^{\prime}$,

$$
\begin{equation*}
\|\widehat{U} Y R-A\|_{F}^{2} \leq\left(1+\epsilon^{\prime}\right)\|\widehat{U} \widetilde{X}-A\|_{F}^{2} \tag{21.27}
\end{equation*}
$$

Then, we can conclude

$$
\begin{aligned}
\|C U R-A\|_{F}^{2} & =\|C Z Y R-A\|_{F}^{2} \\
& =\|C Z \widehat{V}-A\|_{F}^{2} \\
& \leq\left(1+\epsilon^{\prime}\right) \min _{X \in \mathbb{R}^{n \times k}}\|X \widehat{V}-A\|_{F}^{2} \\
& \leq\left(1+\epsilon^{\prime}\right)\|\widehat{U} \widehat{V}-A\|_{F}^{2} \\
& \leq\left(1+\epsilon^{\prime}\right)^{2} \min _{X \in \mathbb{R}^{k \times n}}\|\widehat{U} X-A\|_{F}^{2} \\
& \leq\left(1+\epsilon^{\prime}\right)^{3} \mathrm{OPT} \\
& \leq(1+\epsilon) \mathrm{OPT} .
\end{aligned}
$$

The first equality follows since $U=Z Y$. The second equality follows since $Y R=\widehat{V}$. The first inequality follows by Equation (21.26). The third inequality follows by Equation (21.27). The fourth inequality follows by Equation (21.25). The last inequality follows since $\epsilon^{\prime}=0.1 \epsilon$.

Notice that $C$ has $O(k / \epsilon)$ reweighted columns of $A S_{2}$, and $A S_{2}$ is a subset of reweighted columns of $A$, so $C$ has $O(k / \epsilon)$ reweighted columns of $A$. Similarly, we can prove that $R$ has $O(k / \epsilon)$ reweighted rows of $A$. Thus, $C U R$ is a CUR decomposition of $A$.

### 21.6.2 Stronger property achieved by leverage scores

Claim 21.6.2. Given matrix $A \in \mathbb{R}^{n \times m}$, for any distribution $p=\left(p_{1}, p_{2}, \cdots, p_{n}\right)$ define random variable $X$ such that $X=\left\|A_{i}\right\|_{2}^{2} / p_{i}$ with probability $p_{i}$, where $A_{i}$ is the $i$-th row of matrix $A$. Then take $m$ independent samples $X^{1}, X^{2}, \cdots, X^{m}$, and let $Y=\frac{1}{m} \sum_{j=1}^{m} X^{j}$. We have

$$
\operatorname{Pr}\left[Y \leq 100\|A\|_{F}^{2}\right] \geq .99
$$

Proof. We can compute the expectation of $X^{j}$, for any $j \in[m]$,

$$
\mathbb{E}\left[X^{j}\right]=\sum_{i=1}^{n} \frac{\left\|A_{i}\right\|_{2}^{2}}{p_{i}} \cdot p_{i}=\|A\|_{F}^{2}
$$

Then $\mathbb{E}[Y]=\frac{1}{m} \sum_{j=1}^{m} \mathbb{E}\left[X^{j}\right]=\|A\|_{F}^{2}$. Using Markov's inequality, we have

$$
\operatorname{Pr}\left[Y \geq\|A\|_{F}^{2}\right] \leq .01
$$

Theorem 21.6.3 (The leverage score case of Theorem 39 in [CW13]). Let $A \in \mathbb{R}^{n \times k}$, $B \in \mathbb{R}^{n \times d}$. Let $S \in \mathbb{R}^{n \times n}$ denote a sampling and rescaling diagonal matrix according to the leverage scores of $A$. If the event occurs that $S$ satisfies $(\epsilon / \sqrt{k})$-Frobenius norm approximate matrix product for $A$, and also $S$ is a $(1+\epsilon)$-subspace embedding for $A$, then let $X^{*}$ be the optimal solution of $\min _{X}\|A X-B\|_{F}^{2}$, and $\widetilde{B} \equiv A X^{*}-B$. Then, for all $X \in \mathbb{R}^{k \times d}$,

$$
(1-2 \epsilon)\|A X-B\|_{F}^{2} \leq\|S(A X-B)\|_{F}^{2}+\|\widetilde{B}\|_{F}^{2}-\|S \widetilde{B}\|_{F}^{2} \leq(1+2 \epsilon)\|A X-B\|_{F}^{2}
$$

Furthermore, if $S$ has $m=O\left(\epsilon^{-2} k \log (k)\right)$ nonzero entries, the above event happens with probability at least 0.99.

Note that Theorem 39 in [CW13] is stated in a way that holds for general sketching matrices. However, we are only interested in the case when $S$ is a sampling and rescaling diagonal matrix according to the leverage scores. For completeness, we provide the full proof of the leverage score case with certain parameters.

Proof. Suppose $S$ is a sampling and rescaling diagonal matrix according to the leverage scores of $A$, and it has $m=O\left(\epsilon^{-2} k \log k\right)$ nonzero entries. Then, according to Lemma 21.4.22, $S$ is a $(1+\epsilon)$-subspace embedding for $A$ with probability at least 0.999 , and according to Lemma 21.4.29, $S$ satisfies $(\epsilon / \sqrt{k})$-Frobenius norm approximate matrix product for $A$ with probability at least 0.999.

Let $U \in \mathbb{R}^{n \times k}$ denote an orthonormal basis of the column span of $A$. Then the leverage scores of $U$ are the same as the leverage scores of $A$. Furthermore, for any $X \in \mathbb{R}^{k \times d}$, there is a matrix $Y$ such that $A X=U Y$, and vice versa, so we can now assume $A$ has $k$ orthonormal columns.

Then,

$$
\begin{align*}
& \|S(A X-B)\|_{F}^{2}-\|S \widetilde{B}\|_{F}^{2} \\
= & \left\|S A\left(X-X^{*}\right)+S\left(A X^{*}-B\right)\right\|_{F}^{2}-\|S \widetilde{B}\|_{F}^{2} \\
= & \left\|S A\left(X-X^{*}\right)\right\|_{F}^{2}+\left\|S\left(A X^{*}-B\right)\right\|_{F}^{2}+2 \operatorname{tr}\left(\left(X-X^{*}\right)^{\top} A^{\top} S^{\top} S\left(A X^{*}-B\right)\right)-\|S \widetilde{B}\|_{F}^{2} \\
= & \underbrace{\left\|S A\left(X-X^{*}\right)\right\|_{F}^{2}+2 \operatorname{tr}\left(\left(X-X^{*}\right)^{\top} A^{\top} S^{\top} S \widetilde{B}\right)}_{\alpha} . \tag{21.28}
\end{align*}
$$

The second equality follows using $\|C+D\|_{F}^{2}=\|C\|_{F}^{2}+\|D\|_{F}^{2}+2 \operatorname{tr}\left(C^{\top} D\right)$. The third equality
follows from $\widetilde{B}=A X^{*}-B$. Now, let us first upper bound the term $\alpha$ in Equation (21.28):

$$
\begin{aligned}
& \left\|S A\left(X-X^{*}\right)\right\|_{F}^{2}+2 \operatorname{tr}\left(\left(X-X^{*}\right)^{\top} A^{\top} S^{\top} S \widetilde{B}\right) \\
\leq & (1+\epsilon)\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}+2\left\|X-X^{*}\right\|_{F}\left\|A^{\top} S^{\top} S \widetilde{B}\right\|_{F} \\
\leq & (1+\epsilon)\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}+2(\epsilon / \sqrt{k}) \cdot\left\|X-X^{*}\right\|_{F}\|A\|_{F}\|\widetilde{B}\|_{F} \\
\leq & (1+\epsilon)\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}+2 \epsilon\left\|A\left(X-X^{*}\right)\right\|_{F}\|\widetilde{B}\|_{F}
\end{aligned}
$$

The first inequality follows since $S$ is a $(1+\epsilon)$ subspace embedding of $A$, and $\operatorname{tr}\left(C^{\top} D\right) \leq$ $\|C\|_{F}\|D\|_{F}$. The second inequality follows since $S$ satisfies $(\epsilon / \sqrt{k})$-Frobenius norm approximate matrix product for $A$. The last inequality follows using that $\|A\|_{F} \leq \sqrt{k}$ since $A$ only has $k$ orthonormal columns. Now, let us lower bound the term $\alpha$ in Equation (21.28):

$$
\begin{aligned}
& \left\|S A\left(X-X^{*}\right)\right\|_{F}^{2}+2 \operatorname{tr}\left(\left(X-X^{*}\right)^{\top} A^{\top} S^{\top} S \widetilde{B}\right) \\
\geq & (1-\epsilon)\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}-2\left\|X-X^{*}\right\|_{F}\left\|A^{\top} S^{\top} S \widetilde{B}\right\|_{F} \\
\geq & (1-\epsilon)\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}-2(\epsilon / \sqrt{k}) \cdot\left\|X-X^{*}\right\|_{F}\|A\|_{F}\|\widetilde{B}\|_{F} \\
\geq & (1-\epsilon)\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}-2 \epsilon\left\|A\left(X-X^{*}\right)\right\|_{F}\|\widetilde{B}\|_{F}
\end{aligned}
$$

The first inequality follows since $S$ is a $(1+\epsilon)$ subspace embedding of $A$, and $\operatorname{tr}\left(C^{\top} D\right) \geq$ $-\|C\|_{F}\|D\|_{F}$. The second inequality follows since $S$ satisfies $(\epsilon / \sqrt{k})$-Frobenius norm approximate matrix product for $A$. The last inequality follows using that $\|A\|_{F} \leq \sqrt{k}$ since $A$ only has $k$ orthonormal columns.

Therefore,

$$
\begin{equation*}
(1-\epsilon)\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}-2 \epsilon\left\|A\left(X-X^{*}\right)\right\|_{F}\|\widetilde{B}\|_{F} \leq\|S(A X-B)\|_{F}^{2}-\|S \widetilde{B}\|_{F}^{2} \tag{21.29}
\end{equation*}
$$

and

$$
\begin{equation*}
(1+\epsilon)\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}+2 \epsilon\left\|A\left(X-X^{*}\right)\right\|_{F}\|\widetilde{B}\|_{F} \geq\|S(A X-B)\|_{F}^{2}-\|S \widetilde{B}\|_{F}^{2} \tag{21.30}
\end{equation*}
$$

Notice that $\widetilde{B}=A X^{*}-B=A A^{\dagger} B-B=\left(A A^{\dagger}-I\right) B$, so according to the Pythagorean theorem, we have

$$
\|A X-B\|_{F}^{2}=\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}+\|\widetilde{B}\|_{F}^{2}
$$

which means that

$$
\begin{equation*}
\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}=\|A X-B\|_{F}^{2}-\|\widetilde{B}\|_{F}^{2} \tag{21.31}
\end{equation*}
$$

Using Equation (21.31), we can rewrite and lower bound the LHS of Equation (21.29),

$$
\begin{align*}
& (1-\epsilon)\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}-2 \epsilon\left\|A\left(X-X^{*}\right)\right\|_{F}\|\widetilde{B}\|_{F} \\
= & \left\|A\left(X-X^{*}\right)\right\|_{F}^{2}-\epsilon\left(\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}+2\left\|A\left(X-X^{*}\right)\right\|_{F}\|\widetilde{B}\|_{F}\right) \\
= & \|A X-B\|_{F}^{2}-\|\widetilde{B}\|_{F}^{2}-\epsilon\left(\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}+2\left\|A\left(X-X^{*}\right)\right\|_{F}\|\widetilde{B}\|_{F}\right) \\
\geq & \|A X-B\|_{F}^{2}-\|\widetilde{B}\|_{F}^{2}-\epsilon\left(\left\|A\left(X-X^{*}\right)\right\|_{F}+\|\widetilde{B}\|_{F}\right)^{2} \\
\geq & \|A X-B\|_{F}^{2}-\|\widetilde{B}\|_{F}^{2}-2 \epsilon\left(\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}+\|\widetilde{B}\|_{F}^{2}\right) \\
= & (1-2 \epsilon)\|A X-B\|_{F}^{2}-\|\widetilde{B}\|_{F}^{2} . \tag{21.32}
\end{align*}
$$

The second step follows by Equation (21.31). The first inequality follows using $a^{2}+2 a b<$ $(a+b)^{2}$. The second inequality follows using $(a+b)^{2} \leq 2\left(a^{2}+b^{2}\right)$. The last equality follows using $\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}+\|\widetilde{B}\|_{F}^{2}=\|A X-B\|_{F}^{2}$. Similarly, using Equation (21.31), we can rewrite and upper bound the LHS of Equation (21.30)

$$
\begin{equation*}
(1+\epsilon)\left\|A\left(X-X^{*}\right)\right\|_{F}^{2}+2 \epsilon\left\|A\left(X-X^{*}\right)\right\|_{F}\|\widetilde{B}\|_{F} \leq(1+2 \epsilon)\|A X-B\|_{F}^{2}-\|\widetilde{B}\|_{F}^{2} \tag{21.33}
\end{equation*}
$$

Combining Equations (21.29),(21.32),(21.30),(21.33), we conclude that

$$
(1-2 \epsilon)\|A X-B\|_{F}^{2}-\|\widetilde{B}\|_{F}^{2} \leq\|S(A X-B)\|_{F}^{2}-\|S \widetilde{B}\|_{F}^{2} \leq(1+2 \epsilon)\|A X-B\|_{F}^{2}-\|\widetilde{B}\|_{F}^{2}
$$

Theorem 21.6.4. Let $A \in \mathbb{R}^{n \times k}, B \in \mathbb{R}^{n \times d}$, and $\frac{1}{2}>\epsilon>0$. Let $X^{*}$ be the optimal solution to $\min _{X}\|A X-B\|_{F}^{2}$, and $\widetilde{B} \equiv A X^{*}-B$. Let $S \in \mathbb{R}^{n \times n}$ denote a sketching matrix which satisfies the following:

1. $\|S \widetilde{B}\|_{F}^{2} \leq 100 \cdot\|\widetilde{B}\|_{F}^{2}$,
2. for all $X \in \mathbb{R}^{k \times d}$,

$$
(1-\epsilon)\|A X-B\|_{F}^{2} \leq\|S(A X-B)\|_{F}^{2}+\|\widetilde{B}\|_{F}^{2}-\|S \widetilde{B}\|_{F}^{2} \leq(1+\epsilon)\|A X-B\|_{F}^{2}
$$

Then, for all $X_{1}, X_{2} \in \mathbb{R}^{k \times d}$ satisfying

$$
\left\|S A X_{1}-S B\right\|_{F}^{2} \leq\left(1+\frac{\epsilon}{100}\right) \cdot\left\|S A X_{2}-S B\right\|_{F}^{2}
$$

we have

$$
\left\|A X_{1}-B\right\|_{F}^{2} \leq(1+5 \epsilon) \cdot\left\|A X_{2}-B\right\|_{F}^{2}
$$

Proof. Let $A, B, S, \epsilon$ be the same as in the statement of the theorem, and suppose $S$ satisfies those two conditions. Let $X_{1}, X_{2} \in \mathbb{R}^{k \times d}$ satisfy

$$
\left\|S A X_{1}-S B\right\|_{F}^{2} \leq\left(1+\frac{\epsilon}{100}\right)\left\|S A X_{2}-S B\right\|_{F}^{2}
$$

We have

$$
\begin{aligned}
& \left\|A X_{1}-B\right\|_{F}^{2} \\
\leq & \frac{1}{1-\epsilon}\left(\left\|S\left(A X_{1}-B\right)\right\|_{F}^{2}+\|\widetilde{B}\|_{F}^{2}-\|S \widetilde{B}\|_{F}^{2}\right) \\
\leq & \frac{1}{1-\epsilon}\left(\left(1+\frac{\epsilon}{100}\right) \cdot\left\|S\left(A X_{2}-B\right)\right\|_{F}^{2}+\|\widetilde{B}\|_{F}^{2}-\|S \widetilde{B}\|_{F}^{2}\right) \\
= & \frac{1}{1-\epsilon}\left(\left(1+\frac{\epsilon}{100}\right) \cdot\left(\left\|S\left(A X_{2}-B\right)\right\|_{F}^{2}+\|\widetilde{B}\|_{F}^{2}-\|S \widetilde{B}\|_{F}^{2}\right)-\frac{\epsilon}{100} \cdot\left(\|\widetilde{B}\|_{F}^{2}-\|S \widetilde{B}\|_{F}^{2}\right)\right) \\
\leq & \frac{1}{1-\epsilon} \cdot\left(1+\frac{\epsilon}{100}\right) \cdot\left\|A X_{2}-B\right\|_{F}^{2}-\frac{1}{1-\epsilon} \cdot \frac{\epsilon}{100} \cdot\left(\|\widetilde{B}\|_{F}^{2}-\|S \widetilde{B}\|_{F}^{2}\right) \\
\leq & (1+3 \epsilon)\left\|A X_{2}-B\right\|_{F}^{2}+\frac{1}{1-\epsilon} \cdot \frac{\epsilon}{100}\|S \widetilde{B}\|_{F}^{2} \\
\leq & (1+3 \epsilon)\left\|A X_{2}-B\right\|_{F}^{2}+2 \epsilon\|\widetilde{B}\|_{F}^{2} \\
\leq & (1+5 \epsilon)\left\|A X_{2}-B\right\|_{F}^{2} .
\end{aligned}
$$

The first inequality follows since $S$ satisfies the second condition. The second inequality follows by the relationship between $X_{1}$ and $X_{2}$. The third inequality follows since $S$ satisfies the second condition. The fifth inequality follows using that $\epsilon<\frac{1}{2}$ and that $S$ satisfies the first condition. The last inequality follows using that $\|\widetilde{B}\|_{F}^{2}=\left\|A X^{*}-B\right\|_{F}^{2} \leq\left\|A X_{2}-B\right\|_{F}^{2}$.

Theorem 21.6.5. Let $A \in \mathbb{R}^{n \times k}, B \in \mathbb{R}^{n \times d}$, and $\frac{1}{2}>\epsilon>0$. Let $S \in \mathbb{R}^{n \times n}$ denote $a$ sampling and rescaling diagonal matrix according to the leverage scores of $A$. If $S$ has at least $m=O\left(k \log (k) / \epsilon^{2}\right)$ nonzero entries, then with probability at least 0.98 , for all $X_{1}, X_{2} \in \mathbb{R}^{k \times d}$ satisfying

$$
\left\|S A X_{1}-S B\right\|_{F}^{2} \leq\left(1+\frac{\epsilon}{500}\right) \cdot\left\|S A X_{2}-S B\right\|_{F}^{2},
$$

we have

$$
\left\|A X_{1}-B\right\|_{F}^{2} \leq(1+\epsilon) \cdot\left\|A X_{2}-B\right\|_{F}^{2}
$$

Proof. The proof directly follows by Claim 21.6.2, Theorem 21.6.3 and Theorem 21.6.4. Because of Claim 21.6.2, $S$ satisfies the first condition in the statement of Theorem 21.6.4 with probability at least 0.99 . According to Theorem 21.6.3, $S$ satisfies the second condition in the statement of Theorem 21.6.4 with probability at least 0.99 . Thus, with probability 0.98 , by Theorem 21.6.4, we complete the proof.

### 21.7 Entry-wise $\ell_{1}$ Norm for Arbitrary Tensors

In this section, we provide several different algorithms for tensor $\ell_{1}$-low rank approximation. Section 21.7.1 provides some useful facts and definitions. Section 21.7.2 presents several existence results. Section 21.7.3 describes a tool that is able to reduce the size of the objective function from $\operatorname{poly}(n)$ to poly $(k)$. Section 21.7.4 discusses the case when the problem size is small. Section 21.7.5 provides several bicriteria algorithms. Section 21.7.6 summarizes a batch of algorithms. Section 21.7 .7 provides an algorithm for $\ell_{1}$ norm CURT decomposition.

Notice that if the rank $-k$ solution does not exist, then every bicriteria algorithm in Section 21.7.5 can be stated in a form similar to Theorem 21.1.1, and every algorithm which can output a rank $-k$ solution in Section 21.7.6 can be stated in a form similar to Theorem 21.1.2. See Section 21.1 for more details.

### 21.7.1 Facts

We present a method that is able to reduce the entry-wise $\ell_{1}$-norm objective function to the Frobenius norm objective function.

Fact 21.7.1. Given a 3 rd order tensor $C \in \mathbb{R}^{c_{1} \times c_{2} \times c_{3}}$, three matrices $V_{1} \in \mathbb{R}^{c_{1} \times b_{1}}, V_{2} \in$ $\mathbb{R}^{c_{2} \times b_{2}}, V_{3} \in \mathbb{R}^{c_{3} \times b_{3}}$, for any $k \in\left[1, \min _{i} b_{i}\right]$, if $X_{1}^{\prime} \in \mathbb{R}^{b_{1} \times k}, X_{2}^{\prime} \in \mathbb{R}^{b_{2} \times k}, X_{3}^{\prime} \in \mathbb{R}^{b_{3} \times k}$ satisfies that,

$$
\left\|\left(V_{1} X_{1}^{\prime}\right) \otimes\left(V_{2} X_{2}^{\prime}\right) \otimes\left(V_{3} X_{3}^{\prime}\right)-C\right\|_{F} \leq \alpha \min _{X_{1}, X_{2}, X_{3}}\left\|\left(V_{1} X_{1}\right) \otimes\left(V_{2} X_{2}\right) \otimes\left(V_{3} X_{3}\right)-C\right\|_{F}
$$

then

$$
\left\|\left(V_{1} X_{1}^{\prime}\right) \otimes\left(V_{2} X_{2}^{\prime}\right) \otimes\left(V_{3} X_{3}^{\prime}\right)-C\right\|_{1} \leq \alpha \sqrt{c_{1} c_{2} c_{3}} \min _{X_{1}, X_{2}, X_{3}}\left\|\left(V_{1} X_{1}\right) \otimes\left(V_{2} X_{2}\right) \otimes\left(V_{3} X_{3}\right)-C\right\|_{1}
$$

We extend Lemma C. 15 in [SWZ17] to tensors:

Fact 21.7.2. Given tensor $A \in \mathbb{R}^{n \times n \times n}$, let $\mathrm{OPT}=\min _{\operatorname{rank}-k A_{k}}\left\|A-A_{k}\right\|_{1}$. For any $r \geq k$, if rank-r tensor $B \in \mathbb{R}^{n \times n \times n}$ is an $f$-approximation to $A$, i.e.,

$$
\|B-A\|_{1} \leq f \cdot \mathrm{OPT}
$$

and $U, V, W \in \mathbb{R}^{n \times k}$ is a $g$-approximation to $B$, i.e.,

$$
\|U \otimes V \otimes W-B\|_{1} \leq g \cdot \min _{\operatorname{rank}-k B_{k}}\left\|B_{k}-B\right\|_{1},
$$

then,

$$
\|U \otimes V \otimes W-A\|_{1} \lesssim g f \cdot \mathrm{OPT}
$$

Proof. We define $\widetilde{U}, \widetilde{V}, \widetilde{W} \in \mathbb{R}^{n \times k}$ to be three matrices, such that

$$
\|\widetilde{U} \otimes \widetilde{V} \otimes \widetilde{W}-B\|_{1} \leq g \min _{\text {rank }-k B_{k}}\left\|B_{k}-B\right\|_{1},
$$

and also define,
$\widehat{U}, \widehat{V}, \widehat{W}=\underset{U, V, W \in \mathbb{R}^{n \times k}}{\arg \min }\|U \otimes V \otimes W-B\|_{1}$ and $U^{*}, V^{*}, W^{*}=\underset{U, V, W \in \mathbb{R}^{n \times k}}{\arg \min }\|U \otimes V \otimes W-A\|_{1}$.
It is obvious that,

$$
\begin{equation*}
\|\widehat{U} \otimes \widehat{V} \otimes \widehat{W}-B\|_{1} \leq\left\|U^{*} \otimes V^{*} \otimes W^{*}-B\right\|_{1} \tag{21.34}
\end{equation*}
$$

Then,

$$
\begin{aligned}
& \|\widetilde{U} \otimes \widetilde{V} \otimes \widetilde{W}-A\|_{1} \\
\leq & \|\widetilde{U} \otimes \widetilde{V} \otimes \widetilde{W}-B\|_{1}+\|B-A\|_{1} \\
\leq g\|\widehat{U} \otimes \widehat{V} \otimes \widehat{W}-B\|_{1}+\|B-A\|_{1} & \text { by the triangle inequality } \\
\leq g\left\|U^{*} \otimes V^{*} \otimes W^{*}-B\right\|_{1}+\|B-A\|_{1} & \text { by definition } \\
\leq g\left\|U^{*} \otimes V^{*} \otimes W^{*}-A\right\|_{1}+g\|B-A\|_{1}+\|B-A\|_{1} & \text { by Equation (21.34) } \\
=g \mathrm{OPT}+(g+1)\|B-A\|_{1} & \text { by the triangle inequality } \\
\leq g \mathrm{OPT}+(g+1) f \cdot \mathrm{OPT} & \text { sy definition of OPT } \\
\lesssim g f \mathrm{OPT} . &
\end{aligned}
$$

This completes the proof.

Using the above fact, we are able to optimize our approximation ratio.

### 21.7.2 Existence results

Definition 21.7.1 ( $\ell_{1}$ multiple regression cost preserving sketch - Definition D. 5 in [SWZ17]). Given matrices $U \in \mathbb{R}^{n \times r}, A \in \mathbb{R}^{n \times d}$, let $S \in \mathbb{R}^{m \times n}$. If $\forall \beta \geq 1, \widehat{V} \in \mathbb{R}^{r \times d}$ which satisfy

$$
\|S U \widehat{V}-S A\|_{1} \leq \beta \cdot \min _{V \in \mathbb{R}^{r \times d}}\|S U V-S A\|_{1}
$$

it holds that

$$
\|U \widehat{V}-A\|_{1} \leq \beta \cdot c \cdot \min _{V \in \mathbb{R}^{r \times d}}\|U V-A\|_{1},
$$

then $S$ provides a $c-\ell_{1}$-multiple-regression-cost-preserving-sketch for $(U, A)$.

Theorem 21.7.3. Given a 3rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, there exist three matrices $S_{1} \in \mathbb{R}^{n^{2} \times s_{1}}, S_{2} \in \mathbb{R}^{n^{2} \times s_{2}}, S_{3} \in \mathbb{R}^{n^{2} \times s_{3}}$ such that

$$
\min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k}\left(A_{1} S_{1} X_{1}\right)_{i} \otimes\left(A_{2} S_{2} X_{2}\right)_{i} \otimes\left(A_{3} S_{3} X_{3}\right)_{i}-A\right\|_{1} \leq \alpha_{\text {rank }-k} \min _{A_{k} \in \mathbb{R}^{n \times n \times n}}\left\|A_{k}-A\right\|_{1},
$$

holds with probability 99/100.
(I). Using a dense Cauchy transform, $s_{1}=s_{2}=s_{3}=\widetilde{O}(k), \alpha=\widetilde{O}\left(k^{1.5}\right) \log ^{3} n$.
(II). Using a sparse Cauchy transform, $s_{1}=s_{2}=s_{3}=\widetilde{O}\left(k^{5}\right), \alpha=\widetilde{O}\left(k^{13.5}\right) \log ^{3} n$.
(III). Guessing Lewis weights, $s_{1}=s_{2}=s_{3}=\widetilde{O}(k), \alpha=\widetilde{O}\left(k^{1.5}\right)$.

Proof. We use OPT to denote

$$
\mathrm{OPT}:=\min _{\mathrm{rank}-k} \operatorname{Ain}_{A_{k} \in \mathbb{R}^{n \times n \times n}}\left\|A_{k}-A\right\|_{1} .
$$

Given a tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$, we define three matrices $A_{1} \in \mathbb{R}^{n_{1} \times n_{2} n_{3}}, A_{2} \in \mathbb{R}^{n_{2} \times n_{3} n_{1}}, A_{3} \in$ $\mathbb{R}^{n_{3} \times n_{1} n_{2}}$ such that, for any $i \in\left[n_{1}\right], j \in\left[n_{2}\right], l \in\left[n_{3}\right]$,

$$
A_{i, j, l}=\left(A_{1}\right)_{i,(j-1) \cdot n_{3}+l}=\left(A_{2}\right)_{j,(l-1) \cdot n_{1}+i}=\left(A_{3}\right)_{l,(i-1) \cdot n_{2}+j} .
$$

We fix $V^{*} \in \mathbb{R}^{n \times k}$ and $W^{*} \in \mathbb{R}^{n \times k}$, and use $V_{1}^{*}, V_{2}^{*}, \cdots, V_{k}^{*}$ to denote the columns of $V^{*}$ and $W_{1}^{*}, W_{2}^{*}, \cdots, W_{k}^{*}$ to denote the columns of $W^{*}$.

We consider the following optimization problem,

$$
\min _{U_{1}, \cdots, U_{k} \in \mathbb{R}^{n}}\left\|\sum_{i=1}^{k} U_{i} \otimes V_{i}^{*} \otimes W_{i}^{*}-A\right\|_{1},
$$

which is equivalent to

$$
\min _{U_{1}, \cdots, U_{k} \in \mathbb{R}^{n}}\left\|\left[\begin{array}{llll}
U_{1} & U_{2} & \cdots & U_{k}
\end{array}\right]\left[\begin{array}{c}
V_{1}^{*} \otimes W_{1}^{*} \\
V_{2}^{*} \otimes W_{2}^{*} \\
\cdots \\
V_{k}^{*} \otimes W_{k}^{*}
\end{array}\right]-A\right\|_{1} .
$$

We use matrix $Z_{1}$ to denote $V^{* \top} \odot W^{* \top} \in \mathbb{R}^{k \times n^{2}}$ and matrix $U$ to denote $\left[\begin{array}{llll}U_{1} & U_{2} & \cdots & U_{k}\end{array}\right]$. Then we can obtain the following equivalent objective function,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1}-A_{1}\right\|_{1} .
$$

Choose an $\ell_{1}$ multiple regression cost preserving sketch $S_{1} \in \mathbb{R}^{n^{2} \times s_{1}}$ for $\left(Z_{1}^{\top}, A_{1}^{\top}\right)$. We can obtain the optimization problem,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1} S_{1}-A_{1} S_{1}\right\|_{1}=\min _{U \in \mathbb{R}^{n \times k}} \sum_{i=1}^{n}\left\|U^{i} Z_{1} S_{1}-\left(A_{1} S_{1}\right)^{i}\right\|_{1}
$$

where $U^{i}$ denotes the $i$-th row of matrix $U \in \mathbb{R}^{n \times k}$ and $\left(A_{1} S_{1}\right)^{i}$ denotes the $i$-th row of matrix $A_{1} S_{1}$. Instead of solving it under the $\ell_{1}$-norm, we consider the $\ell_{2}$-norm relaxation,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1} S_{1}-A_{1} S_{1}\right\|_{F}^{2}=\min _{U \in \mathbb{R}^{n \times k}} \sum_{i=1}^{n}\left\|U^{i} Z_{1} S_{1}-\left(A_{1} S_{1}\right)^{i}\right\|_{2}^{2}
$$

Let $\widehat{U} \in \mathbb{R}^{n \times k}$ denote the optimal solution of the above optimization problem. Then, $\widehat{U}=$ $A_{1} S_{1}\left(Z_{1} S_{1}\right)^{\dagger}$. We plug $\widehat{U}$ into the objective function under the $\ell_{1}$-norm. According to Claim 21.3.4, we have,

$$
\left\|\widehat{U} Z_{1} S_{1}-A_{1} S_{1}\right\|_{1}=\sum_{i=1}^{n}\left\|\widehat{U}^{i} Z_{1} S_{1}-\left(A_{1} S_{1}\right)^{i}\right\|_{1} \leq \sqrt{s_{1}} \min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1} S_{1}-A_{1} S_{1}\right\|_{1} .
$$

Since $S_{1} \in \mathbb{R}^{n^{2} \times s_{1}}$ satisfies Definition 21.7.1, we have

$$
\left\|\widehat{U} Z_{1}-A_{1}\right\|_{1} \leq \alpha \min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1}-A_{1}\right\|_{1}=\alpha \mathrm{OPT}
$$

where $\alpha=\sqrt{s_{1}} \beta$ and $\beta$ (see Definition 21.7.1) is a parameter which depends on which kind of sketching matrix we actually choose. It implies

$$
\left\|\widehat{U} \otimes V^{*} \otimes W^{*}-A\right\|_{1} \leq \alpha \mathrm{OPT}
$$

As a second step, we fix $\widehat{U} \in \mathbb{R}^{n \times k}$ and $W^{*} \in \mathbb{R}^{n \times k}$, and convert tensor $A$ into matrix $A_{2}$. Let matrix $Z_{2}$ denote $\widehat{U}^{\top} \odot W^{* \top}$. We consider the following objective function,

$$
\min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2}-A_{2}\right\|_{1}
$$

and the optimal cost of it is at most $\alpha$ OPT.
Choose an $\ell_{1}$ multiple regression cost preserving sketch $S_{2} \in \mathbb{R}^{n^{2} \times s_{2}}$ for $\left(Z_{2}^{\top}, A_{2}^{\top}\right)$, and sketch on the right of the objective function to obtain this new objective function,

$$
\min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2} S_{2}-A_{2} S_{2}\right\|_{1}=\min _{U \in \mathbb{R}^{n \times k}} \sum_{i=1}^{n}\left\|V^{i} Z_{2} S_{2}-\left(A_{2} S_{2}\right)^{i}\right\|_{1}
$$

where $V^{i}$ denotes the $i$-th row of matrix $V$ and $\left(A_{2} S_{2}\right)^{i}$ denotes the $i$-th row of matrix $A_{2} S_{2}$. Instead of solving this under the $\ell_{1}$-norm, we consider the $\ell_{2}$-norm relaxation,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|V Z_{2} S_{2}-A_{2} S_{2}\right\|_{F}^{2}=\min _{V \in \mathbb{R}^{n \times k}}\left\|V^{i}\left(Z_{2} S_{2}\right)-\left(A_{2} S_{2}\right)^{i}\right\|_{2}^{2}
$$

Let $\widehat{V} \in \mathbb{R}^{n \times k}$ denote the optimal solution of the above problem. Then $\widehat{V}=$ $A_{2} S_{2}\left(Z_{2} S_{2}\right)^{\dagger}$. By properties of the sketching matrix $S_{2} \in \mathbb{R}^{n^{2} \times s_{2}}$, we have,

$$
\left\|\widehat{V} Z_{2}-A_{2}\right\|_{1} \leq \alpha \min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2}-A_{2}\right\|_{1} \leq \alpha^{2} \mathrm{OPT}
$$

which implies

$$
\left\|\widehat{U} \otimes \widehat{V} \otimes W^{*}-A\right\|_{1} \leq \alpha^{2} \mathrm{OPT}
$$

As a third step, we fix the matrices $\widehat{U} \in \mathbb{R}^{n \times k}$ and $\widehat{V} \in \mathbb{R}^{n \times k}$. We can convert tensor $A \in \mathbb{R}^{n \times n \times n}$ into matrix $A_{3} \in \mathbb{R}^{n^{2} \times n}$. Let matrix $Z_{3}$ denote $\widehat{U}^{\top} \odot \widehat{V}^{\top} \in \mathbb{R}^{k \times n^{2}}$. We consider the following objective function,

$$
\min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3}-A_{3}\right\|_{1},
$$

and the optimal cost of it is at most $\alpha^{2}$ OPT.
Choose an $\ell_{1}$ multiple regression cost preserving sketch $S_{3} \in \mathbb{R}^{n^{2} \times s_{3}}$ for $\left(Z_{3}^{\top}, A_{3}^{\top}\right)$ and sketch on the right of the objective function to obtain the new objective function,

$$
\min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3} S_{3}-A_{3} S_{3}\right\|_{1} .
$$

Let $\widehat{W} \in \mathbb{R}^{n \times k}$ denote the optimal solution of the above problem. Then $\widehat{W}=A_{3} S_{3}\left(Z_{3} S_{3}\right)^{\dagger}$. By properties of sketching matrix $S_{3} \in \mathbb{R}^{n^{2} \times s_{3}}$, we have,

$$
\left\|\widehat{W} Z_{3}-A_{3}\right\|_{1} \leq \alpha \min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3}-A_{3}\right\|_{1} \leq \alpha^{3} \mathrm{OPT}
$$

Thus, we obtain,

$$
\min _{X_{1} \in \mathbb{R}^{s_{1} \times k}, X_{2} \in \mathbb{R}^{s_{2} \times k}, X_{3} \in \mathbb{R}^{s_{3} \times k}}\left\|\sum_{i=1}^{k}\left(A_{1} S_{1} X_{1}\right)_{i} \otimes\left(A_{2} S_{2} X_{2}\right)_{i} \otimes\left(A_{3} S_{3} X_{3}\right)_{i}-A\right\|_{1} \leq \alpha^{3} \mathrm{OPT}
$$

Proof of (I) By Theorem C. 1 in [SWZ17], we can use dense Cauchy transforms for $S_{1}, S_{2}, S_{3}$, and then $s_{1}=s_{2}=s_{3}=O(k \log k)$ and $\alpha=O(\sqrt{k \log k} \log n)$.

Proof of (II) By Theorem C. 1 in [SWZ17], we can use sparse Cauchy transforms for $S_{1}, S_{2}, S_{3}$, and then $s_{1}=s_{2}=s_{3}=O\left(k^{5} \log ^{5} k\right)$ and $\alpha=O\left(k^{4.5} \log ^{4.5} k \log n\right)$.

Proof of (III) By Theorem C. 1 in [SWZ17], we can sample by Lewis weights. Then $S_{1}, S_{2}, S_{3} \in \mathbb{R}^{n^{2} \times n^{2}}$ are diagonal matrices, and each of them has $O(k \log k)$ nonzero rows. This gives $\alpha=O(\sqrt{k \log k})$.

### 21.7.3 Polynomial in $k$ size reduction

Definition 21.7.2 (Definition D. 1 in [SWZ17]). Given a matrix $M \in \mathbb{R}^{n \times d}$, if matrix $S \in$ $\mathbb{R}^{m \times n}$ satisfies

$$
\|S M\|_{1} \leq \beta\|M\|_{1}
$$

then $S$ has at most $\beta$ dilation on $M$.

Definition 21.7.3 (Definition D. 2 in [SWZ17]). Given a matrix $U \in \mathbb{R}^{n \times k}$, if matrix $S \in$ $\mathbb{R}^{m \times n}$ satisfies

$$
\forall x \in \mathbb{R}^{k},\|S U x\|_{1} \geq \frac{1}{\beta}\|U x\|_{1}
$$

then $S$ has at most $\beta$ contraction on $U$.

Theorem 21.7.4. Given a tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$ and three matrices $V_{1} \in \mathbb{R}^{n_{1} \times b_{1}}, V_{2} \in$ $\mathbb{R}^{n_{2} \times b_{2}}, V_{3} \in \mathbb{R}^{n_{3} \times b_{3}}$, let $X_{1}^{*} \in \mathbb{R}^{b_{1} \times k}, X_{2}^{*} \in \mathbb{R}^{b_{2} \times k}, X_{3}^{*} \in \mathbb{R}^{b_{3} \times k}$ satisfies

$$
X_{1}^{*}, X_{2}^{*}, X_{3}^{*}=\underset{X_{1} \in \mathbb{R}^{b_{1} \times k}, X_{2} \in \mathbb{R}^{b_{2} \times k}, X_{3} \in \mathbb{R}^{b_{3} \times k}}{\arg \min _{1}}\left\|V_{1} X_{1} \otimes V_{2} X_{2} \otimes V_{3} X_{3}-A\right\|_{1} .
$$

Let $S \in \mathbb{R}^{m \times n}$ have at most $\beta_{1} \geq 1$ dilation on $V_{1} X_{1}^{*} \cdot\left(\left(V_{2} X_{2}^{*}\right)^{\top} \odot\left(V_{3} X_{3}^{*}\right)^{\top}\right)-A_{1}$ and $S$ have at most $\beta_{2} \geq 1$ contraction on $V_{1}$. If $\widehat{X}_{1} \in \mathbb{R}^{b_{1} \times k}, \widehat{X}_{2} \in \mathbb{R}^{b_{2} \times k}, \widehat{X}_{3} \in \mathbb{R}^{b_{3} \times k}$ satisfies $\left\|S V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-S A\right\|_{1} \leq \beta_{X_{1} \in \mathbb{R}^{b_{1} \times k}, X_{2} \in \mathbb{R}^{b_{2} \times k}, X_{3} \in \mathbb{R}^{b_{3} \times k}}\left\|S V_{1} X_{1} \otimes V_{2} X_{2} \otimes V_{3} X_{3}-S A\right\|_{1}$, where $\beta \geq 1$, then

$$
\left\|V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-A\right\|_{1} \lesssim \beta_{1} \beta_{2} \beta \min _{X_{1}, X_{2}, X_{3}}\left\|V_{1} X_{1} \otimes V_{2} X_{2} \otimes V_{3} X_{3}-A\right\|_{1}
$$

The proof idea is similar to [SWZ17].

Proof. Let $A, V_{1}, V_{2}, V_{3}, S, X_{1}^{*}, X_{2}^{*}, X_{3}^{*}, \beta_{1}, \beta_{2}$ be the same as stated in the theorem. Let $\widehat{X}_{1} \in$ $\mathbb{R}^{b_{1} \times k}, \widehat{X}_{2} \in \mathbb{R}^{b_{2} \times k}, \widehat{X}_{3} \in \mathbb{R}^{b_{3} \times k}$ satisfy
$\left\|S V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-S A\right\|_{1} \leq \beta_{X_{1} \in \mathbb{R}^{b_{1} \times k}, X_{2} \in \mathbb{R}^{b_{2} \times k}, X_{3} \in \mathbb{R}^{b_{3} \times k}}\left\|S V_{1} X_{1} \otimes V_{2} X_{2} \otimes V_{3} X_{3}-S A\right\|_{1}$.
We have,

$$
\begin{align*}
& \left\|S V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-S A\right\|_{1} \\
\geq & \left\|S V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-S V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}\right\|_{1}-\left\|S V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-S A\right\|_{1} \\
\geq & \frac{1}{\beta_{2}}\left\|V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}\right\|_{1}-\beta_{1}\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{1} \\
\geq & \frac{1}{\beta_{2}}\left\|V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-A\right\|_{1}-\frac{1}{\beta_{2}}\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{1} \\
& -\beta_{1}\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{1} \\
= & \frac{1}{\beta_{2}}\left\|V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-A\right\|_{1}-\left(\frac{1}{\beta_{2}}+\beta_{1}\right)\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{1} . \tag{21.35}
\end{align*}
$$

The first and the third inequality follow by the triangle inequalities. The second inequality follows using that

$$
\begin{aligned}
& \left\|S V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-S V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}\right\|_{1} \\
= & \left\|S V_{1}\left(\widehat{X}_{1}-X_{1}^{*}\right) \cdot\left(\left(V_{2}\left(\widehat{X}_{2}-X_{2}^{*}\right)\right)^{\top} \odot\left(V_{3}\left(\widehat{X}_{3}-X_{3}^{*}\right)\right)^{\top}\right)\right\|_{1} \\
\geq & \frac{1}{\beta_{2}}\left\|V_{1}\left(\widehat{X}_{1}-X_{1}^{*}\right) \cdot\left(\left(V_{2}\left(\widehat{X}_{2}-X_{2}^{*}\right)\right)^{\top} \odot\left(V_{3}\left(\widehat{X}_{3}-X_{3}^{*}\right)\right)^{\top}\right)\right\|_{1} \\
\geq & \frac{1}{\beta_{2}}\left\|V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}\right\|_{1},
\end{aligned}
$$

and

$$
\begin{align*}
& \left\|S V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-S A\right\|_{1} \\
= & \left\|S\left(V_{1} X_{1}^{*} \cdot\left(\left(V_{2} X_{2}^{*}\right)^{\top} \odot\left(V_{3} X_{3}^{*}\right)^{\top}\right)-A_{1}\right)\right\|_{1} \\
\leq & \left\|V_{1} X_{1}^{*} \cdot\left(\left(V_{2} X_{2}^{*}\right)^{\top} \odot\left(V_{3} X_{3}^{*}\right)^{\top}\right)-A_{1}\right\|_{1} \\
= & \beta_{1}\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{1} . \tag{21.36}
\end{align*}
$$

Then, we have

$$
\begin{aligned}
& \left\|V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-A\right\|_{1} \\
\leq & \beta_{2}\left\|S V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-S A\right\|_{1}+\left(1+\beta_{1} \beta_{2}\right)\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{1} \\
\leq & \beta_{2} \beta\left\|S V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-S A\right\|_{1}+\left(1+\beta_{1} \beta_{2}\right)\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{1} \\
\leq & \beta_{1} \beta_{2} \beta\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{1}+\left(1+\beta_{1} \beta_{2}\right)\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{1} \\
\leq & \beta\left(1+2 \beta_{1} \beta_{2}\right)\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{1} .
\end{aligned}
$$

The first inequality follows by Equation (21.35). The second inequality follows by

$$
\left\|S V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-S A\right\|_{1} \leq \beta \min _{X_{1}, X_{2}, X_{3}}\left\|S V_{1} X_{1} \otimes V_{2} X_{2} \otimes V_{3} X_{3}-S A\right\|_{1}
$$

The third inequality follows by Equation (21.36). The final inequality follows using that $\beta \geq 1$.

Lemma 21.7.5. Let $\min \left(b_{1}, b_{2}, b_{3}\right) \geq k$. Given three matrices $V_{1} \in \mathbb{R}^{n \times b_{1}}, V_{2} \in \mathbb{R}^{n \times b_{2}}$, and $V_{3} \in \mathbb{R}^{n \times b_{3}}$, there exists an algorithm that takes $O(\mathrm{nnz}(A))+n \operatorname{poly}\left(b_{1}, b_{2}, b_{3}\right)$ time and outputs a tensor $C \in \mathbb{R}^{c_{1} \times c_{2} \times c_{3}}$ and three matrices $\widehat{V}_{1} \in \mathbb{R}^{c_{1} \times b_{1}}, \widehat{V}_{2} \in \mathbb{R}^{c_{2} \times b_{2}}$ and $\widehat{V}_{3} \in \mathbb{R}^{c_{3} \times b_{3}}$

```
Algorithm 21.21 Reducing the Size of the Objective Function to poly \((k)\)
    procedure L1PolyKSizeReduction \(\left(A, V_{1}, V_{2}, V_{3}, n, b_{1}, b_{2}, b_{3}, k\right) \quad \triangleright\) Lemma 21.7.5
        for \(i=1 \rightarrow 3\) do
                        \(c_{i} \leftarrow \widetilde{O}\left(b_{i}\right)\).
            Choose sampling and rescaling matrices \(T_{i} \in \mathbb{R}^{c_{i} \times n}\) according to the Lewis weights
    of \(V_{i}\).
            \(\widehat{V}_{i} \leftarrow T_{i} V_{i} \in \mathbb{R}^{c_{i} \times b_{i}}\).
        end for
        \(C \leftarrow A\left(T_{1}, T_{2}, T_{3}\right) \in \mathbb{R}^{c_{1} \times c_{2} \times c_{3}}\).
        return \(\widehat{V}_{1}, \widehat{V}_{2}, \widehat{V}_{3}\) and \(C\).
    end procedure
```

with $c_{1}=c_{2}=c_{3}=\operatorname{poly}\left(b_{1}, b_{2}, b_{3}\right)$, such that with probability 0.99 , for any $\alpha \geq 1$, if $X_{1}^{\prime}, X_{2}^{\prime}, X_{3}^{\prime}$ satisfy that,

$$
\left\|\sum_{i=1}^{k}\left(\widehat{V}_{1} X_{1}^{\prime}\right)_{i} \otimes\left(\widehat{V}_{2} X_{2}^{\prime}\right)_{i} \otimes\left(\widehat{V}_{3} X_{3}^{\prime}\right)_{i}-C\right\|_{1} \leq \alpha_{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k}\left(\widehat{V}_{1} X_{1}\right)_{i} \otimes\left(\widehat{V}_{2} X_{2}\right)_{i} \otimes\left(\widehat{V}_{3} X_{3}\right)_{i}-C\right\|_{1},
$$

then,

$$
\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}^{\prime}\right)_{i} \otimes\left(V_{2} X_{2}^{\prime}\right)_{i} \otimes\left(V_{3} X_{3}^{\prime}\right)_{i}-A\right\|_{1} \lesssim \min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}\right)_{i} \otimes\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}-A\right\|_{1} .
$$

Proof. For simplicity, we define OPT to be

$$
\min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}\right)_{i} \otimes\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}-A\right\|_{1} .
$$

Let $T_{1} \in \mathbb{R}^{c_{1} \times n}$ sample according to the Lewis weights of $V_{1} \in \mathbb{R}^{n \times b_{1}}$, where $c_{1}=\widetilde{O}\left(b_{1}\right)$. Let $T_{2} \in \mathbb{R}^{c_{2} \times n}$ sample according to the Lewis weights of $V_{2} \in \mathbb{R}^{n \times b_{2}}$, where $c_{2}=\widetilde{O}\left(b_{2}\right)$. Let $T_{3} \in \mathbb{R}^{c_{3} \times n}$ sample according to the Lewis weights of $V_{3} \in \mathbb{R}^{n \times b_{3}}$, where $c_{3}=\widetilde{O}\left(b_{3}\right)$.

For any $\alpha \geq 1$, let $X_{1}^{\prime} \in \mathbb{R}^{b_{1} \times k}, X_{2}^{\prime} \in \mathbb{R}^{b_{2} \times k}, X_{3}^{\prime} \in \mathbb{R}^{b_{3} \times k}$ satisfy

$$
\begin{aligned}
& \left\|T_{1} V_{1} X_{1}^{\prime} \otimes T_{2} V_{2} X_{2}^{\prime} \otimes T_{3} V_{3} X_{3}^{\prime}-A\left(T_{1}, T_{2}, T_{3}\right)\right\|_{1} \\
\leq & \alpha_{X_{1} \in \mathbb{R}^{b_{1} \times k}, X_{2} \in \mathbb{R}^{b_{2} \times k}, X_{3} \in \mathbb{R}^{b_{3} \times k}}\left\|T_{1} V_{1} X_{1} \otimes T_{2} V_{2} X_{2} \otimes T_{3} V_{3} X_{3}-A\left(T_{1}, T_{2}, T_{3}\right)\right\|_{1} .
\end{aligned}
$$

First, we regard $T_{1}$ as the sketching matrix for the remainder. Then by Lemma D. 11 in [SWZ17] and Theorem 21.7.4, we have

$$
\begin{aligned}
& \left\|V_{1} X_{1}^{\prime} \otimes T_{2} V_{2} X_{2}^{\prime} \otimes T_{3} V_{3} X_{3}^{\prime}-A\left(I, T_{2}, T_{3}\right)\right\|_{1} \\
\lesssim & \alpha_{X_{1} \in \mathbb{R}^{b_{1} \times k}, X_{2} \in \mathbb{R}^{b_{2} \times k}, X_{3} \in \mathbb{R}^{b_{3} \times k}}\left\|V_{1} X_{1} \otimes T_{2} V_{2} X_{2} \otimes T_{3} V_{3} X_{3}-A\left(I, T_{2}, T_{3}\right)\right\|_{1} .
\end{aligned}
$$

Second, we regard $T_{2}$ as a sketching matrix for $V_{1} X_{1} \otimes V_{2} X_{2} \otimes T_{3} V_{3} X_{3}-A\left(I, I, T_{3}\right)$. Then by Lemma D. 11 in [SWZ17] and Theorem 21.7.4, we have

$$
\begin{aligned}
& \left\|V_{1} X_{1}^{\prime} \otimes V_{2} X_{2}^{\prime} \otimes T_{3} V_{3} X_{3}^{\prime}-A\left(I, I, T_{3}\right)\right\|_{1} \\
\lesssim & \alpha_{X_{1} \in \mathbb{R}^{b_{1} \times k}, X_{2} \in \mathbb{R}^{b_{2} \times k}, X_{3} \in \mathbb{R}^{b_{3} \times k}}\left\|V_{1} X_{1} \otimes V_{2} X_{2} \otimes T_{3} V_{3} X_{3}-A\left(I, I, T_{3}\right)\right\|_{1} .
\end{aligned}
$$

Third, we regard $T_{3}$ as a sketching matrix for $V_{1} X_{1} \otimes V_{2} X_{2} \otimes V_{3} X_{3}-A$. Then by Lemma D. 11 in [SWZ17] and Theorem 21.7.4, we have

$$
\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}^{\prime}\right)_{i} \otimes\left(V_{2} X_{2}^{\prime}\right)_{i} \otimes\left(V_{3} X_{3}^{\prime}\right)_{i}-A\right\|_{1} \lesssim \alpha \min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}\right)_{i} \otimes\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}-A\right\|_{1}
$$

Lemma 21.7.6. Given tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$, and two matrices $U \in \mathbb{R}^{n_{1} \times s}, V \in \mathbb{R}^{n_{2} \times s}$ with $\operatorname{rank}(U)=r$, let $T \in \mathbb{R}^{t \times n_{1}}$ be a sampling/rescaling matrix according to the Lewis weights of $U$ with $t=\widetilde{O}(r)$. Then with probability at least 0.99 , for all $X^{\prime} \in \mathbb{R}^{n_{3} \times s}, \alpha \geq 1$ which satisfy

$$
\left\|T_{1} U \otimes V \otimes X^{\prime}-T_{1} A\right\|_{1} \leq \alpha \cdot \min _{X \in \mathbb{R}^{n} \times s}\left\|T_{1} U \otimes V \otimes X-T_{1} A\right\|_{1}
$$

it holds that

$$
\left\|U \otimes V \otimes X^{\prime}-A\right\|_{1} \lesssim \alpha \cdot \min _{X \in \mathbb{R}^{n_{3} \times s}}\|U \otimes V \otimes X-A\|_{1}
$$

The proof is similar to the proof of Lemma 21.7.5.

Proof. Let $X^{*}=\underset{X \in \mathbb{R}^{n_{3} \times s}}{\arg \min }\|U \otimes V \otimes X-A\|_{1}$. Then according to Lemma D. 11 in [SWZ17], $T$ has at most constant dilation (Definition 21.7.2) on $U \cdot\left(V^{\top} \odot\left(X^{*}\right)^{\top}\right)-A_{1}$, and has at most constant contraction (Definition 21.7.3) on $U$. We first look at

$$
\begin{aligned}
& \left\|T U \otimes V \otimes X^{\prime}-T A\right\|_{1} \\
= & \left\|T U \cdot\left(V^{\top} \odot\left(X^{\prime}\right)^{\top}\right)-T A_{1}\right\|_{1} \\
\geq & \left\|T U \cdot\left(\left(V^{\top} \odot\left(X^{\prime}\right)^{\top}\right)-\left(V^{\top} \odot\left(X^{*}\right)^{\top}\right)\right)\right\|_{1}-\left\|T U \cdot\left(V^{\top} \odot\left(X^{*}\right)^{\top}\right)-T A_{1}\right\|_{1} \\
\geq & \frac{1}{\beta_{2}} \| U \cdot\left(\left(V^{\top} \odot\left(X^{\prime}\right)^{\top}\right)-A_{1}\left\|_{1}-\left(\frac{1}{\beta_{2}}+\beta_{1}\right)\right\| U \cdot\left(V^{\top} \odot\left(X^{*}\right)^{\top}\right)-A_{1} \|_{1},\right.
\end{aligned}
$$

where $\beta_{1} \geq 1, \beta_{2} \geq 1$ are two constants. Then we have:

$$
\begin{aligned}
& \left\|U \otimes V \otimes X^{\prime}-A\right\|_{1} \\
\leq & \beta_{2}\left\|T U \cdot\left(V^{\top} \odot\left(X^{\prime}\right)^{\top}\right)-T A_{1}\right\|_{1}+\left(1+\beta_{1} \beta_{2}\right)\left\|U \cdot\left(V^{\top} \odot\left(X^{*}\right)^{\top}\right)-A_{1}\right\|_{1} \\
\leq & \alpha \beta_{2}\left\|T U \cdot\left(V^{\top} \odot\left(X^{*}\right)^{\top}\right)-T A_{1}\right\|_{1}+\left(1+\beta_{1} \beta_{2}\right)\left\|U \cdot\left(V^{\top} \odot\left(X^{*}\right)^{\top}\right)-A_{1}\right\|_{1} \\
\leq & \alpha \beta_{1} \beta_{2}\left\|U \cdot\left(V^{\top} \odot\left(X^{*}\right)^{\top}\right)-A_{1}\right\|_{1}+\left(1+\beta_{1} \beta_{2}\right)\left\|U \cdot\left(V^{\top} \odot\left(X^{*}\right)^{\top}\right)-A_{1}\right\|_{1} \\
\lesssim & \alpha\left\|U \otimes V \otimes X^{*}-A\right\|_{1 .} .
\end{aligned}
$$

Corollary 21.7.7. Given tensor $A \in \mathbb{R}^{n \times n \times n}$, and two matrices $U \in \mathbb{R}^{n \times s}, V \in \mathbb{R}^{n \times s}$ with $\operatorname{rank}(U)=r_{1}, \operatorname{rank}(V)=r_{2}$, let $T_{1} \in \mathbb{R}^{t_{1} \times n}$ be a sampling/rescaling matrix according to
the Lewis weights of $U$, and let $T_{2} \in \mathbb{R}^{t_{2} \times n}$ be a sampling/rescaling matrix according to the Lewis weights of $V$ with $t_{1}=\widetilde{O}\left(r_{1}\right), t_{2}=\widetilde{O}\left(r_{2}\right)$. Then with probability at least 0.99 , for all $X^{\prime} \in \mathbb{R}^{n \times s}, \alpha \geq 1$ which satisfy

$$
\left\|T_{1} U \otimes T_{2} V \otimes X^{\prime}-A\left(T_{1}, T_{2}, I\right)\right\|_{1} \leq \alpha \cdot \min _{X \in \mathbb{R}^{n \times s}}\left\|T_{1} U \otimes T_{2} V \otimes X-A\left(T_{1}, T_{2}, I\right)\right\|_{1}
$$

it holds that

$$
\left\|U \otimes V \otimes X^{\prime}-A\right\|_{1} \lesssim \alpha \cdot \min _{X \in \mathbb{R}^{n \times s}}\|U \otimes V \otimes X-A\|_{1}
$$

Proof. We apply Lemma 21.7.6 twice: if

$$
\left\|T_{1} U \otimes T_{2} V \otimes X^{\prime}-A\left(T_{1}, T_{2}, I\right)\right\|_{1} \leq \alpha \cdot \min _{X \in \mathbb{R}^{n \times s}}\left\|T_{1} U \otimes T_{2} V \otimes X-A\left(T_{1}, T_{2}, I\right)\right\|_{1}
$$

then

$$
\left\|U \otimes T_{2} V \otimes X^{\prime}-A\left(I, T_{2}, I\right)\right\|_{1} \lesssim \alpha \cdot \min _{X \in \mathbb{R}^{n \times s}}\left\|U \otimes T_{2} V \otimes X-A\left(I, T_{2}, I\right)\right\|_{1}
$$

Then, we have

$$
\left\|U \otimes V \otimes X^{\prime}-A\right\|_{1} \lesssim \alpha \cdot \min _{X \in \mathbb{R}^{n \times s}}\|U \otimes V \otimes X-A\|_{1}
$$

### 21.7.4 Solving small problems

Theorem 21.7.8. Let $\max _{i}\left\{t_{i}, d_{i}\right\} \leq n$. Given $a t_{1} \times t_{2} \times t_{3}$ tensor $A$ and three matrices: a $t_{1} \times d_{1}$ matrix $T_{1}$, a $t_{2} \times d_{2}$ matrix $T_{2}$, and a $t_{3} \times d_{3}$ matrix $T_{3}$, if for $\delta>0$ there exists a solution to

$$
\min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k}\left(T_{1} X_{1}\right)_{i} \otimes\left(T_{2} X_{2}\right)_{i} \otimes\left(T_{3} X_{3}\right)_{i}-A\right\|_{1}:=\mathrm{OPT}
$$

such that each entry of $X_{i}$ can be expressed using $O\left(n^{\delta}\right)$ bits, then there exists an algorithm that takes $n^{O(\delta)} \cdot 2^{O\left(d_{1} k+d_{2} k+d_{3} k\right)}$ time and outputs three matrices: $\widehat{X}_{1}, \widehat{X}_{2}$, and $\widehat{X}_{3}$ such that $\left\|\left(T_{1} \widehat{X}_{1}\right) \otimes\left(T_{2} \widehat{X}_{2}\right) \otimes\left(T_{3} \widehat{X}_{3}\right)-A\right\|_{1}=\mathrm{OPT}$.

Proof. For each $i \in[3]$, we can create $t_{i} \times d_{i}$ variables to represent matrix $X_{i}$. Let $x$ denote the list of these variables. Let $B$ denote tensor $\sum_{i=1}^{k}\left(T_{1} X_{1}\right)_{i} \otimes\left(T_{2} X_{2}\right)_{i} \otimes\left(T_{3} X_{3}\right)_{i}$. Then we can write the following objective function,

$$
\min _{x} \sum_{i=1}^{t_{1}} \sum_{j=1}^{t_{2}} \sum_{l=1}^{t_{3}}\left|B_{i, j, l}(x)-A_{i, j, l}\right| .
$$

To remove the $|\cdot|$, we create $t_{1} t_{2} t_{3}$ extra variables $\sigma_{i, j, l}$. Then we obtain the objective function:

$$
\begin{aligned}
\min _{x, \sigma} & \sum_{i=1}^{t_{1}} \sum_{j=1}^{t_{2}} \sum_{l=1}^{t_{3}} \sigma_{i, j, l}\left(B_{i, j, l}(x)-A_{i, j, l}\right) \\
\text { s.t. } & \sigma_{i, j, l}^{2}=1, \\
& \sigma_{i, j, l}\left(B_{i, j, l}(x)-A_{i, j, l}\right) \geq 0 \\
& \|x\|_{2}^{2}+\|\sigma\|_{2}^{2} \leq 2^{O\left(n^{\delta}\right)}
\end{aligned}
$$

where the last constraint is unharmful, because there exists a solution that can be written using $O\left(n^{\delta}\right)$ bits. Note that the number of inequality constraints in the above system is
$O\left(t_{1} t_{2} t_{3}\right)$, the degree is $O(1)$, and the number of variables is $v=\left(t_{1} t_{2} t_{3}+d_{1} k+d_{2} k+d_{3} k\right)$. Thus by Theorem 21.3.2, we know that the minimum nonzero cost is at least

$$
\left(2^{O\left(n^{\delta}\right)}\right)^{-2^{\tilde{O}(v)}} .
$$

It is immediate that the upper bound on cost is at most $2^{O\left(n^{\delta}\right)}$, and thus the number of binary search steps is at most $\log \left(2^{O\left(n^{\delta}\right)}\right) 2^{\widetilde{O}(v)}$. In each step of the binary search, we need to choose a cost $C$ between the lower bound and the upper bound, and write down the polynomial system,

$$
\begin{aligned}
& \sum_{i=1}^{t_{1}} \sum_{j=1}^{t_{2}} \sum_{l=1}^{t_{3}} \sigma_{i, j, l}\left(B_{i, j, l}(x)-A_{i, j, l}\right) \leq C, \\
& \sigma_{i, j, l}^{2}=1, \\
& \sigma_{i, j, l}\left(B_{i, j, l}(x)-A_{i, j, l}\right) \geq 0, \\
& \|x\|_{2}^{2}+\|\sigma\|_{2}^{2} \leq 2^{O\left(n^{\delta}\right)} .
\end{aligned}
$$

Using Theorem 21.3.1, we can determine if there exists a solution to the above polynomial system. Since the number of variables is $v$, and the degree is $O(1)$, the number of inequality constraints is $t_{1} t_{2} t_{2}$. Thus, the running time is

$$
\operatorname{poly}(\text { bitsize }) \cdot(\# \text { constraints } \cdot \text { degree })^{\# \text { variables }}=n^{O(\delta)} 2^{\widetilde{O}(v)}
$$

### 21.7.5 Bicriteria algorithms

We present several bicriteria algorithms with different tradeoffs. We first present an algorithm that runs in nearly linear time and outputs a solution with rank $\widetilde{O}\left(k^{3}\right)$ in Theorem 21.7.9. Then we show an algorithm that runs in $n n z(A)$ time but outputs a solution with rank poly $(k)$ in Theorem 21.7.10. Then we explain an idea which is able to decrease the cubic rank to quadratic rank, and thus we can obtain Theorem 21.7.11 and Theorem 21.7.12.

### 21.7.5.1 Input sparsity time

Algorithm $21.22 \ell_{1}$-Low Rank Approximation, Bicriteria Algorithm, rank- $\widetilde{O}\left(k^{3}\right)$, Nearly Input Sparsity Time
procedure L1BicriteriaAlgorithm $(A, n, k) \quad \triangleright$ Theorem 21.7.9
$s_{1} \leftarrow s_{2} \leftarrow s_{3} \leftarrow \widetilde{O}(k)$.
For each $i \in[3]$, choose $S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}$ to be a dense Cauchy transform. $\triangleright$ Part (I) of Theorem 21.7.2

Compute $A_{1} \cdot S_{1}, A_{2} \cdot S_{2}, A_{3} \cdot S_{3}$.
$Y_{1}, Y_{2}, Y_{3}, C \leftarrow$ L1PolyKSizeReduction $\left(A, A_{1} S_{1}, A_{2} S_{2}, A_{3} S_{3}, n, s_{1}, s_{2}, s_{3}, k\right) \quad \triangleright$
Algorithm 21.21
Form objective function

$$
\min _{X \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} X_{i, j, l}\left(Y_{1}\right)_{i} \otimes\left(Y_{2}\right)_{j} \otimes\left(Y_{3}\right)_{l}-C\right\|_{1} .
$$

Run $\ell_{1}$-regression solver to find $X$.
return $A_{1} S_{1}, A_{2} S_{2}, A_{3} S_{3}$ and $X$.
end procedure

Theorem 21.7.9. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1, \epsilon \in(0,1)$, let $r=\widetilde{O}\left(k^{3}\right)$. There exists an algorithm which takes $\mathrm{nnz}(A) \cdot \widetilde{O}(k)+O(n) \operatorname{poly}(k)+\operatorname{poly}(k)$
time and outputs three matrices $U, V, W \in \mathbb{R}^{n \times r}$ such that

$$
\left\|\sum_{i=1}^{r} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{1} \leq \widetilde{O}\left(k^{3 / 2}\right) \log ^{3} n_{\operatorname{rank}-k} \min _{A_{k}}\left\|A_{k}-A\right\|_{1}
$$

holds with probability 9/10.

Proof. We first choose three dense Cauchy transforms $S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}$. According to Section 21.3.7, for each $i \in[3], A_{i} S_{i}$ can be computed in $n n z(A) \cdot \widetilde{O}(k)$ time. Then we apply Lemma 21.7.5 (Algorithm 21.21). We obtain three matrices $Y_{1}, Y_{2}, Y_{3}$ and a tensor $C$. Note that for each $i \in[3], Y_{i}$ can be computed in $n \operatorname{poly}(k)$ time. Because $C=A\left(T_{1}, T_{2}, T_{3}\right)$ and $T_{1}, T_{2}, T_{3} \in \mathbb{R}^{n \times \widetilde{O}(k)}$ are three sampling and rescaling matrices, $C$ can be computed in $n n z(A)+\widetilde{O}\left(k^{3}\right)$ time. At the end, we just need to run an $\ell_{1}$-regression solver to find the solution to the problem,

$$
\min _{X \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} X_{i, j, l}\left(Y_{1}\right)_{i} \otimes\left(Y_{2}\right)_{j} \otimes\left(Y_{3}\right)_{j}\right\|_{1},
$$

where $\left(Y_{1}\right)_{i}$ denotes the $i$-th column of matrix $Y_{1}$. Since the size of the above problem is only $\operatorname{poly}(k)$, this can be solved in $\operatorname{poly}(k)$ time.

Theorem 21.7.10. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1, \epsilon \in(0,1)$, let $r=\widetilde{O}\left(k^{15}\right)$. There exists an algorithm that takes $n n z(A)+O(n) \operatorname{poly}(k)+\operatorname{poly}(k)$ time and outputs three matrices $U, V, W \in \mathbb{R}^{n \times r}$ such that

$$
\left\|\sum_{i=1}^{r} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{1} \leq \operatorname{poly}(k, \log n) \min _{\operatorname{rank}-k A_{k}}\left\|A_{k}-A\right\|_{1}
$$

holds with probability 9/10.

Algorithm $21.23 \ell_{1}$-Low Rank Approximation, Bicriteria Algorithm, rank-poly $(k)$, Input Sparsity Time
procedure L1BicriteriaAlgorithm $(A, n, k)$
$\triangleright$ Theorem 21.7.10
$s_{1} \leftarrow s_{2} \leftarrow s_{3} \leftarrow \widetilde{O}\left(k^{5}\right)$.
For each $i \in[3]$, choose $S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}$ to be a sparse Cauchy transform. $\triangleright$ Part (II) of Theorem 21.7.3

Compute $A_{1} \cdot S_{1}, A_{2} \cdot S_{2}, A_{3} \cdot S_{3}$.
$Y_{1}, Y_{2}, Y_{3}, C \leftarrow$ L1PolyKSizeReduction $\left(A, A_{1} S_{1}, A_{2} S_{2}, A_{3}, S_{3}, n, s_{1}, s_{2}, s_{3}, k\right) \quad \triangleright$ Algorithm 21.21

Form objective function

$$
\min _{X \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} X_{i, j, l}\left(Y_{1}\right)_{i} \otimes\left(Y_{2}\right)_{j} \otimes\left(Y_{3}\right)_{l}-C\right\|_{1} .
$$

Run $\ell_{1}$-regression solver to find $X$.
return $A_{1} S_{1}, A_{2} S_{2}, A_{3} S_{3}$ and $X$.
end procedure

Proof. We first choose three dense Cauchy transforms $S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}$. According to Section 21.3.7, for each $i \in[3], A_{i} S_{i}$ can be computed in $O(\mathrm{nnz}(A))$ time. Then we apply Lemma 21.7.5 (Algorithm 21.21), and can obtain three matrices $Y_{1}, Y_{2}, Y_{3}$ and a tensor $C$. Note that for each $i \in[3], Y_{i}$ can be computed in $O(n) \operatorname{poly}(k)$ time. Because $C=A\left(T_{1}, T_{2}, T_{3}\right)$ and $T_{1}, T_{2}, T_{3} \in \mathbb{R}^{n \times \widetilde{O}(k)}$ are three sampling and rescaling matrices, $C$ can be computed in $n n z(A)+\widetilde{O}\left(k^{3}\right)$ time. At the end, we just need to run an $\ell_{1}$-regression solver to find the solution to the problem,

$$
\min _{X \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} X_{i, j, l}\left(Y_{1}\right)_{i} \otimes\left(Y_{2}\right)_{j} \otimes\left(Y_{3}\right)_{l}-C\right\|_{1},
$$

where $\left(Y_{1}\right)_{i}$ denotes the $i$-th column of matrix $Y_{1}$. Since the size of the above problem is only $\operatorname{poly}(k)$, it can be solved in $\operatorname{poly}(k)$ time.

### 21.7.5.2 Improving cubic rank to quadratic rank

Algorithm 21.24 $\ell_{1}$-Low Rank Approximation, Bicriteria Algorithm, rank- $\widetilde{O}\left(k^{2}\right)$, Nearly Input Sparsity Time
procedure L1BicriteriaAlgorithm $(A, n, k) \quad \triangleright$ Theorem 21.7.11
$s_{1} \leftarrow s_{2} \leftarrow s_{3} \leftarrow \widetilde{O}(k)$.
For each $i \in[3]$, choose $S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}$ to be a dense Cauchy transform. $\triangleright$ Part (I) of
Theorem 21.7.2
Compute $A_{1} \cdot S_{1}, A_{2} \cdot S_{2}$.
For each $i \in[2]$, choose $T_{i}$ to be a sampling and rescaling diagonal matrix according
to the Lewis weights of $A_{i} S_{i}$, with $t_{i}=\widetilde{O}(k)$ nonzero entries.
$C \leftarrow A\left(T_{1}, T_{2}, I\right)$.
$B^{i+(j-1) s_{1}} \leftarrow \operatorname{vec}\left(\left(T_{1} A_{1} S_{1}\right)_{i} \otimes\left(T_{2} A_{2} S_{2}\right)_{j}\right), \forall i \in\left[s_{1}\right], j \in\left[s_{2}\right]$.
Form objective function $\min _{W}\left\|W B-C_{3}\right\|_{1}$
Run $\ell_{1}$-regression solver to find $\widehat{W}$.
Construct $\widehat{U}$ by using $A_{1} S_{1}$ according to Equation (21.38).
Construct $\widehat{V}$ by using $A_{2} S_{2}$ according to Equation (21.39).
return $\widehat{U}, \widehat{V}, \widehat{W}$.
end procedure

Theorem 21.7.11. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1, \epsilon \in(0,1)$, let $r=\widetilde{O}\left(k^{2}\right)$. There exists an algorithm which takes $\mathrm{nnz}(A) \cdot \widetilde{O}(k)+O(n) \operatorname{poly}(k)+\operatorname{poly}(k)$ time and outputs three matrices $U, V, W \in \mathbb{R}^{n \times r}$ such that

$$
\left\|\sum_{i=1}^{r} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{1} \leq \widetilde{O}\left(k^{3 / 2}\right) \log ^{3} n \min _{\operatorname{rank}-k A_{k}}\left\|A_{k}-A\right\|_{1}
$$

holds with probability 9/10.

Proof. Let OPT $=\min _{A_{k} \in \mathbb{R}^{n \times n \times n}}\left\|A_{k}-A\right\|_{1}$. We first choose three dense Cauchy transforms $S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}, \forall i \in[3]$. According to Section 21.3.7, for each $i \in[3], A_{i} S_{i}$ can be computed in $\mathrm{nnz}(A) \cdot \widetilde{O}(k)$ time. Then we choose $T_{i}$ to be a sampling and rescaling diagonal matrix according to the Lewis weights of $A_{i} S_{i}, \forall i \in[2]$.

According to Theorem 21.7.3, we have
$\min _{X_{1} \in \mathbb{R}^{s_{1} \times k}, X_{2} \in \mathbb{R}^{s_{2} \times k}, X_{3} \in \mathbb{R}^{s_{3} \times k}}\left\|\sum_{l=1}^{k}\left(A_{1} S_{1} X_{1}\right)_{l} \otimes\left(A_{2} S_{2} X_{2}\right)_{l} \otimes\left(A_{3} S_{3} X_{3}\right)_{l}-A\right\|_{1} \leq \widetilde{O}\left(k^{1.5}\right) \log ^{3} n \mathrm{OPT}$
Now we fix an $l$ and we have:

$$
\begin{aligned}
& \left(A_{1} S_{1} X_{1}\right)_{l} \otimes\left(A_{2} S_{2} X_{2}\right)_{l} \otimes\left(A_{3} S_{3} X_{3}\right)_{l} \\
= & \left(\sum_{i=1}^{s_{1}}\left(A_{1} S_{1}\right)_{i}\left(X_{1}\right)_{i, l}\right) \otimes\left(\sum_{j=1}^{s_{2}}\left(A_{2} S_{2}\right)_{j}\left(X_{2}\right)_{j, l}\right) \otimes\left(A_{3} S_{3} X_{3}\right)_{l} \\
= & \sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}}\left(A_{1} S_{1}\right)_{i} \otimes\left(A_{2} S_{2}\right)_{j} \otimes\left(A_{3} S_{3} X_{3}\right)_{l}\left(X_{1}\right)_{i, l}\left(X_{2}\right)_{j, l}
\end{aligned}
$$

Thus, we have
$\min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}}\left(A_{1} S_{1}\right)_{i} \otimes\left(A_{2} S_{2}\right)_{j} \otimes\left(\sum_{l=1}^{k}\left(A_{3} S_{3} X_{3}\right)_{l}\left(X_{1}\right)_{i, l}\left(X_{2}\right)_{j, l}\right)-A\right\|_{1} \leq \widetilde{O}\left(k^{1.5}\right) \log ^{3} n$ OPT.

We create matrix $\widehat{U} \in \mathbb{R}^{n \times s_{1} s_{2}}$ by copying matrix $A_{1} S_{1} s_{2}$ times, i.e.,

$$
\widehat{U}=\left[\begin{array}{llll}
A_{1} S_{1} & A_{1} S_{1} & \cdots & A_{1} S_{1} \tag{21.38}
\end{array}\right] .
$$

We create matrix $\widehat{V} \in \mathbb{R}^{n \times s_{1} s_{2}}$ by copying the $i$-th column of $A_{2} S_{2}$ a total of $s_{1}$ times into the columns $(i-1) s_{1}, \cdots, i s_{1}$ of $\widehat{V}$, for each $i \in\left[s_{2}\right]$, i.e.,

$$
\widehat{V}=\left[\begin{array}{llllllll}
\left(A_{2} S_{2}\right)_{1} & \cdots & \left(A_{2} S_{2}\right)_{1} & \left(A_{2} S_{2}\right)_{2} & \cdots & \left(A_{2} S_{2}\right)_{2} & \cdots\left(A_{2} S_{2}\right)_{s_{2}} & \cdots \tag{21.39}
\end{array}\left(A_{2} S_{2}\right)_{s_{2}} .\right]
$$

According to Equation (21.37), we have:

$$
\min _{W \in \mathbb{R}^{n \times s_{1} s_{2}}}\|\widehat{U} \otimes \widehat{V} \otimes W-A\|_{1} \leq \widetilde{O}\left(k^{1.5}\right) \log ^{3} n \cdot \mathrm{OPT}
$$

Let

$$
\widehat{W}=\underset{W \in \mathbb{R}^{n \times s_{1} s_{2}}}{\arg \min }\left\|T_{1} \widehat{U} \otimes T_{2} \widehat{V} \otimes W-A\left(T_{1}, T_{2}, I\right)\right\|_{1} .
$$

Due to Corollary 21.7.7, we have

$$
\|\widehat{U} \otimes \widehat{V} \otimes \widehat{W}-A\|_{1} \leq \widetilde{O}\left(k^{1.5}\right) \log ^{3} n \cdot \mathrm{OPT}
$$

Putting it all together, we have that $\widehat{U}, \widehat{V}, \widehat{W}$ gives a rank- $\widetilde{O}\left(k^{2}\right)$ bicriteria algorithm to the original problem.

```
Algorithm \(21.25 \ell_{1}\)-Low Rank Approximation, Bicriteria Algorithm, rank-poly ( \(k\) ), Input
Sparsity Time
procedure L1BicriteriaAlgorithm \((A, n, k) \quad \triangleright\) Theorem 21.7.12
    \(s_{1} \leftarrow s_{2} \leftarrow s_{3} \leftarrow \widetilde{O}\left(k^{5}\right)\).
    For each \(i \in[3]\), choose \(S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}\) to be a sparse Cauchy transform. \(\triangleright\) Part (II) of
Theorem 21.7.2
    Compute \(A_{1} \cdot S_{1}, A_{2} \cdot S_{2}\).
    For each \(i \in[2]\), choose \(T_{i}\) to be a sampling and rescaling diagonal matrix according
to the Lewis weights of \(A_{i} S_{i}\), with \(t_{i}=\widetilde{O}(k)\) nonzero entries.
    \(C \leftarrow A\left(T_{1}, T_{2}, I\right)\).
    \(B^{i+(j-1) s_{1}} \leftarrow \operatorname{vec}\left(\left(T_{1} A_{1} S_{1}\right)_{i} \otimes\left(T_{2} A_{2} S_{2}\right)_{j}\right), \forall i \in\left[s_{1}\right], j \in\left[s_{2}\right]\).
    Form objective function \(\min _{W}\left\|W B-C_{3}\right\|_{1}\).
    Run \(\ell_{1}\)-regression solver to find \(\widehat{W}\).
    Construct \(\widehat{U}\) by using \(A_{1} S_{1}\) according to Equation (21.38).
    Construct \(\widehat{V}\) by using \(A_{2} S_{2}\) according to Equation (21.39).
    return \(\widehat{U}, \widehat{V}, \widehat{W}\).
end procedure
```

Theorem 21.7.12. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1, \epsilon \in(0,1)$, let $r=\widetilde{O}\left(k^{10}\right)$. There exists an algorithm which takes $n n z(A)+O(n) \operatorname{poly}(k)+\operatorname{poly}(k)$ time
and outputs three matrices $U, V, W \in \mathbb{R}^{n \times r}$ such that

$$
\left\|\sum_{i=1}^{r} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{1} \leq \operatorname{poly}(k, \log n) \min _{\mathrm{rank}-k}\left\|A_{k}-A\right\|_{1}
$$

holds with probability 9/10.

Proof. The proof is similar to the proof of Theorem 21.7.11. The only difference is that instead of choosing dense Cauchy matrices $S_{1}, S_{2}$, we choose sparse Cauchy matrices.

Notice that if we firstly apply a sparse Cauchy transform, we can reduce the rank of the matrix to $\operatorname{poly}(k)$. Then we apply a dense Cauchy transform and can further reduce the dimension while only incurring another $\operatorname{poly}(k)$ factor in the approximation ratio. By combining a sparse Cauchy transform and a dense Cauchy transform, we can improve the running time from $n n z(A) \cdot \widetilde{O}(k)$ to $n n z(A)$.

Corollary 21.7.13. Given a 3rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1, \epsilon \in(0,1)$, let $r=\widetilde{O}\left(k^{2}\right)$. There exists an algorithm which takes $\mathrm{nnz}(A)+O(n) \operatorname{poly}(k)+\operatorname{poly}(k)$ time and outputs three matrices $U, V, W \in \mathbb{R}^{n \times r}$ such that

$$
\left\|\sum_{i=1}^{r} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{1} \leq \operatorname{poly}(k, \log n) \min _{\operatorname{rank}-k}\left\|A_{k}-A\right\|_{1}
$$

holds with probability 9/10.

```
Algorithm \(21.26 \ell_{1}\)-Low Rank Approximation, Bicriteria Algorithm, rank- \(\widetilde{O}\left(k^{2}\right)\), Input
Sparsity Time
procedure L1BicriteriaAlgorithm \((A, n, k) \quad\) Corollary 21.7.13
    \(s_{1} \leftarrow s_{2} \leftarrow s_{3} \leftarrow \widetilde{O}(k)\).
    For each \(i \in[3]\), choose \(S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}\) to be the composition of a sparse Cauchy
transform and a dense Cauchy transform. \(\triangleright\) Part (I,II) of Theorem 21.7.2
    Compute \(A_{1} \cdot S_{1}, A_{2} \cdot S_{2}\).
    For each \(i \in[2]\), choose \(T_{i}\) to be a sampling and rescaling diagonal matrix according
to the Lewis weights of \(A_{i} S_{i}\), with \(t_{i}=\widetilde{O}(k)\) nonzero entries.
    \(C \leftarrow A\left(T_{1}, T_{2}, I\right)\).
    \(B^{i+(j-1) s_{1}} \leftarrow \operatorname{vec}\left(\left(T_{1} A_{1} S_{1}\right)_{i} \otimes\left(T_{2} A_{2} S_{2}\right)_{j}\right), \forall i \in\left[s_{1}\right], j \in\left[s_{2}\right]\).
    Form objective function \(\min _{W}\left\|W B-C_{3}\right\|_{1}\).
    Run \(\ell_{1}\)-regression solver to find \(\widehat{W}\).
    Construct \(\widehat{U}\) by using \(A_{1} S_{1}\) according to Equation (21.38).
    Construct \(\widehat{V}\) by using \(A_{2} S_{2}\) according to Equation (21.39).
    return \(\widehat{U}, \widehat{V}, \widehat{W}\).
end procedure
```


### 21.7.6 Algorithms

In this section, we show two different algorithms by using different kind of sketches.
One is shown in Theorem 21.7.14 which gives a fast running time. Another one is shown in Theorem 21.7.16 which gives the best approximation ratio.

### 21.7.6.1 Input sparsity time algorithm

Theorem 21.7.14. Given a $3 r d$ tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, there exists an algorithm that takes $\mathrm{nnz}(A) \cdot \widetilde{O}(k)+O(n) \operatorname{poly}(k)+2^{\widetilde{O}\left(k^{2}\right)}$ time and outputs three matrices $U, V, W \in$ $\mathbb{R}^{n \times k}$ such that,

$$
\|U \otimes V \otimes W-A\|_{1} \leq \operatorname{poly}(k, \log n) \min _{\text {rank }-k}\left\|A^{\prime}-A\right\|_{1} .
$$

```
Algorithm \(21.27 \ell_{1}\)-Low Rank Approximation, Input sparsity Time Algorithm
    procedure L1TensorLowRankApproxInputSparsity \((A, n, k) \triangleright\) Theorem 21.7.14
        \(s_{1} \leftarrow s_{2} \leftarrow s_{3} \leftarrow \widetilde{O}\left(k^{5}\right)\).
        Choose \(S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}\) to be a dense Cauchy transform, \(\forall i \in[3] . \triangleright\) Part (I) of Theorem
    21.7.3
        Compute \(A_{1} \cdot S_{1}, A_{2} \cdot S_{2}\), and \(A_{3} \cdot S_{3}\).
        \(Y_{1}, Y_{2}, Y_{3}, C \leftarrow\) L1PolyKSizeREDUCTion \(\left(A, A_{1} S_{1}, A_{2} S_{2}, A_{3} S_{3}, n, s_{1}, s_{2}, s_{3}, k\right) . \quad \triangleright\)
    Algorithm 21.21
        Create variables \(s_{1} \times k+s_{2} \times k+s_{3} \times k\) variables for each entry of \(X_{1}, X_{2}, X_{3}\).
        Form objective function \(\left\|\left(Y_{1} X_{1}\right) \otimes\left(Y_{2} X_{2}\right) \otimes\left(Y_{3} X_{3}\right)-C\right\|_{F}^{2}\).
        Run polynomial system verifier.
        return \(A_{1} S_{1} X_{1}, A_{2} S_{2} X_{2}, A_{3} S_{3} X_{3}\).
    end procedure
```

holds with probability at least 9/10.

Proof. First, we apply part (II) of Theorem 21.7.3. Then $A_{i} S_{i}$ can be computed in $O(\mathrm{nnz}(A))$ time. Second, we use Lemma 21.7.5 to reduce the size of the objective function from $O\left(n^{3}\right)$ to $\operatorname{poly}(k)$ in $n \operatorname{poly}(k)$ time by only losing a constant factor in approximation ratio. Third, we use Claim 21.3.6 to relax the objective function from entry-wise $\ell_{1}$-norm to Frobenius norm, and this step causes us to lose some other $\operatorname{poly}(k)$ factors in approximation ratio. As a last step, we use Theorem 21.4.44 to solve the Frobenius norm objective function.

Notice again that if we first apply a sparse Cauchy transform, we can reduce the rank of the matrix to poly $(k)$. Then as before we can apply a dense Cauchy transform to further reduce the dimension while only incurring another $\operatorname{poly}(k)$ factor in the approximation ratio. By combining a sparse Cauchy transform and a dense Cauchy transform, we can improve the running time from $n n z(A) \cdot \widetilde{O}(k)$ to $n n z(A)$, while losing some additional poly $(k)$ factors in approximation ratio.

Corollary 21.7.15. Given a $3 r d$ tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, there exists an algorithm that takes $\mathrm{nnz}(A)+O(n) \operatorname{poly}(k)+2^{\widetilde{O}\left(k^{2}\right)}$ time and outputs three matrices $U, V, W \in \mathbb{R}^{n \times k}$ such that,

$$
\|U \otimes V \otimes W-A\|_{1} \leq \operatorname{poly}(k, \log n) \min _{\operatorname{rank}-k}\left\|A^{\prime}-A\right\|_{1} .
$$

holds with probability at least 9/10.

### 21.7.6.2 $\widetilde{O}\left(k^{3 / 2}\right)$-approximation algorithm

```
Algorithm \(21.28 \ell_{1}\)-Low Rank Approximation Algorithm, \(\widetilde{O}\left(k^{3 / 2}\right)\)-approximation
    procedure L1TensorLowRankApproxK \((A, n, k) \quad \triangleright\) Theorem 21.7.16
        \(s_{1} \leftarrow s_{2} \leftarrow s_{3} \leftarrow \widetilde{O}(k)\).
        Guess diagonal matrices \(S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}\) with \(s_{i}\) nonzero entries, \(\forall i \in[3]\). \(\triangleright\) Part (III) of
    Theorem 21.7.3
        Compute \(A_{1} \cdot S_{1}, A_{2} \cdot S_{2}\), and \(A_{3} \cdot S_{3}\).
        \(Y_{1}, Y_{2}, Y_{3}, C \leftarrow\) L1PolyKSizeReduction \(\left(A, A_{1} S_{1}, A_{2} S_{2}, A_{3} S_{3}, n, s_{1}, s_{2}, s_{3}, k\right) . \quad \triangleright\)
    Algorithm 21.21
        Create \(s_{1} \times k+s_{2} \times k+s_{3} \times k\) variables for each entry of \(X_{1}, X_{2}, X_{3}\).
        Form objective function \(\left\|\left(Y_{1} X_{1}\right) \otimes\left(Y_{2} X_{2}\right) \otimes\left(Y_{3} X_{3}\right)-C\right\|_{1}\).
        Run polynomial system verifier.
        return \(U, V, W\).
    end procedure
```

Theorem 21.7.16. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, there exists an algorithm that takes $n^{\widetilde{O}(k)} 2^{\widetilde{O}\left(k^{3}\right)}$ time and output three matrices $U, V, W \in \mathbb{R}^{n \times k}$ such that,

$$
\|U \otimes V \otimes W-A\|_{1} \leq \widetilde{O}\left(k^{3 / 2}\right) \min _{\operatorname{rank}-k A^{\prime}}\left\|A^{\prime}-A\right\|_{1}
$$

holds with probability at least 9/10.

Proof. First, we apply part (III) of Theorem 21.7.3. Then, guessing $S_{i}$ requires $n^{\widetilde{O}(k)}$ time. Second, we use Lemma 21.7.5 to reduce the size of the objective from $O\left(n^{3}\right)$ to poly $(k)$ in polynomial time while only losing a constant factor in approximation ratio. Third, we use Theorem 21.7.8 to solve the entry-wise $\ell_{1}$-norm objective function directly.

### 21.7.7 CURT decomposition

```
Algorithm \(21.29 \ell_{1}\)-CURT Decomposition Algorithm
    procedure L1CURT \(\left(A, U_{B}, V_{B}, W_{B}, n, k\right) \quad \triangleright\) Theorem 21.7.17
        Form \(B_{1}=V_{B}^{\top} \odot W_{B}^{\top} \in \mathbb{R}^{k \times n^{2}}\).
        Let \(D_{1}^{\top} \in \mathbb{R}^{n^{2} \times n^{2}}\) be the sampling and rescaling diagonal matrix corresponding to
    the Lewis weights of \(B_{1}^{\top}\), and let \(D_{1}\) have \(d_{1}=O(k \log k)\) nonzero entries.
        Form \(\widehat{U}=A_{1} D_{1}\left(B_{1} D_{1}\right)^{\dagger} \in \mathbb{R}^{n \times k}\).
        Form \(B_{2}=\widehat{U}^{\top} \odot W_{B}^{\top} \in \mathbb{R}^{k \times n^{2}}\).
        Let \(D_{2}^{\top} \in \mathbb{R}^{n^{2} \times n^{2}}\) be the sampling and rescaling diagonal matrix corresponding to
    the Lewis weights of \(B_{2}^{\top}\), and let \(D_{2}\) have \(d_{2}=O(k \log k)\) nonzero entries.
        Form \(\widehat{V}=A_{2} D_{2}\left(B_{2} D_{2}\right)^{\dagger} \in \mathbb{R}^{n \times k}\).
        Form \(B_{3}=\widehat{U}^{\top} \odot \widehat{V}^{\top} \in \mathbb{R}^{k \times n^{2}}\).
        Let \(D_{3}^{\top} \in \mathbb{R}^{n^{2} \times n^{2}}\) be the sampling and rescaling diagonal matrix corresponding to
    the Lewis weights of \(B_{3}^{\top}\), and let \(D_{3}\) have \(d_{3}=O(k \log k)\) nonzero entries.
        \(C \leftarrow A_{1} D_{1}, R \leftarrow A_{2} D_{2}, T \leftarrow A_{3} D_{3}\).
        \(U \leftarrow \sum_{i=1}^{k}\left(\left(B_{1} D_{1}\right)^{\dagger}\right)_{i} \otimes\left(\left(B_{2} D_{2}\right)^{\dagger}\right)_{i} \otimes\left(\left(B_{3} D_{3}\right)^{\dagger}\right)_{i}\).
        return \(C, R, T\) and \(U\).
    end procedure
```

Theorem 21.7.17. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, let $k \geq 1$, let $U_{B}, V_{B}, W_{B} \in$ $\mathbb{R}^{n \times k}$ denote a rank-k, $\alpha$-approximation to $A$. Then there exists an algorithm which takes $O(\operatorname{nnz}(A))+O\left(n^{2}\right) \operatorname{poly}(k)$ time and outputs three matrices: $C \in \mathbb{R}^{n \times c}$ with columns from $A, R \in \mathbb{R}^{n \times r}$ with rows from $A, T \in \mathbb{R}^{n \times t}$ with tubes from $A$, and a tensor $U \in \mathbb{R}^{c \times r \times t}$ with $\operatorname{rank}(U)=k$ such that $c=r=t=O(k \log k)$, and

$$
\left\|\sum_{i=1}^{c} \sum_{j=1}^{r} \sum_{l=1}^{t} U_{i, j, l} \cdot C_{i} \otimes R_{j} \otimes T_{l}-A\right\|_{1} \leq \widetilde{O}\left(k^{1.5}\right) \alpha \min _{\operatorname{rank}-k A^{\prime}}\left\|A^{\prime}-A\right\|_{1}
$$

holds with probability 9/10.

Proof. We define

$$
\mathrm{OPT}:=\min _{\text {rank }-k}\left\|A^{\prime}-A\right\|_{1} .
$$

We already have three matrices $U_{B} \in \mathbb{R}^{n \times k}, V_{B} \in \mathbb{R}^{n \times k}$ and $W_{B} \in \mathbb{R}^{n \times k}$ and these three matrices provide a rank- $k, \alpha$ approximation to $A$, i.e.,

$$
\begin{equation*}
\left\|\sum_{i=1}^{k}\left(U_{B}\right)_{i} \otimes\left(V_{B}\right)_{i} \otimes\left(W_{B}\right)_{i}-A\right\|_{1} \leq \alpha \mathrm{OPT} \tag{21.40}
\end{equation*}
$$

Let $B_{1}=V_{B}^{\top} \odot W_{B}^{\top} \in \mathbb{R}^{k \times n^{2}}$ denote the matrix where the $i$-th row is the vectorization of $\left(V_{B}\right)_{i} \otimes\left(W_{B}\right)_{i}$. By Section B.3, we can compute $D_{1} \in \mathbb{R}^{n^{2} \times n^{2}}$ which is a sampling and rescaling matrix corresponding to the Lewis weights of $B_{1}^{\top}$ in $O\left(n^{2} \operatorname{poly}(k)\right)$ time, and there are $d_{1}=O(k \log k)$ nonzero entries on the diagonal of $D_{1}$. Let $A_{i} \in \mathbb{R}^{n \times n^{2}}$ denote the matrix obtained by flattening $A$ along the $i$-th direction, for each $i \in[3]$.

Define $U^{*} \in \mathbb{R}^{n \times k}$ to be the optimal solution to $\min _{U \in \mathbb{R}^{n \times k}}\left\|U B_{1}-A_{1}\right\|_{1}, \widehat{U}=A_{1} D_{1}\left(B_{1} D_{1}\right)^{\dagger} \in$ $\mathbb{R}^{n \times k}, V_{0} \in \mathbb{R}^{n \times k}$ to be the optimal solution to $\min _{V \in \mathbb{R}^{n \times k}}\left\|V \cdot\left(\widehat{U}^{\top} \odot W_{B}^{\top}\right)-A_{2}\right\|_{1}$, and $U^{\prime}$ to be the optimal solution to $\min _{U \in \mathbb{R}^{n \times k}}\left\|U B_{1} D_{1}-A_{1} D_{1}\right\|_{1}$.

By Claim 21.3.4, we have

$$
\left\|\widehat{U} B_{1} D_{1}-A_{1} D_{1}\right\|_{1} \leq \sqrt{d_{1}}\left\|U^{\prime} B_{1} D_{1}-A_{1} D_{1}\right\|_{1}
$$

Due to Lemma D. 11 and Lemma D. 8 (in [SWZ17]) with constant probability, we have

$$
\begin{equation*}
\left\|\widehat{U} B_{1}-A_{1}\right\|_{1} \leq \sqrt{d_{1}} \alpha_{D_{1}}\left\|U^{*} B_{1}-A_{1}\right\|_{1}, \tag{21.41}
\end{equation*}
$$

where $\alpha_{D_{1}}=O(1)$.

Recall that $\left(\widehat{U}^{\top} \odot W_{B}^{\top}\right) \in \mathbb{R}^{k \times n^{2}}$ denotes the matrix where the $i$-th row is the vectorization of $\widehat{U}_{i} \otimes\left(W_{B}\right)_{i}, \forall i \in[k]$. Now, we can show,

$$
\begin{array}{rlr}
\left\|V_{0} \cdot\left(\widehat{U}^{\top} \odot W_{B}^{\top}\right)-A_{2}\right\|_{1} & \leq\left\|\widehat{U} B_{1}-A_{1}\right\|_{1} & \text { by } V_{0}=\underset{V \in \mathbb{R}^{n \times k}}{\arg \min }\left\|V \cdot\left(\widehat{U}^{\top} \odot W_{B}^{\top}\right)-A_{2}\right\|_{1} \\
& \lesssim \sqrt{d_{1}}\left\|U^{*} B_{1}-A_{1}\right\|_{1} & \text { by Equation }(21.41) \\
& \leq \sqrt{d_{1}}\left\|U_{B} B_{1}-A_{1}\right\|_{1} & \text { by } U^{*}=\underset{U \in \mathbb{R}^{n \times k}}{\arg \min }\left\|U B_{1}-A_{1}\right\|_{1} \\
& \leq O\left(\sqrt{d_{1}}\right) \alpha \mathrm{OPT} & \text { by Equation (21.40) } \tag{21.42}
\end{array}
$$

We define $B_{2}=\widehat{U}^{\top} \odot W_{B}^{\top}$. We can compute $D_{2} \in \mathbb{R}^{n^{2} \times n^{2}}$ which is a sampling and rescaling matrix corresponding to the Lewis weights of $B_{2}^{\top}$ in $O\left(n^{2}\right.$ poly $\left.(k)\right)$ time, and there are $d_{2}=O(k \log k)$ nonzero entries on the diagonal of $D_{2}$.

Define $V^{*} \in \mathbb{R}^{n \times k}$ to be the optimal solution of $\min _{V \in \mathbb{R}^{n \times k}}\left\|V B_{2}-A_{2}\right\|_{1}, \widehat{V}=$ $A_{2} D_{2}\left(B_{2} D_{2}\right)^{\dagger} \in \mathbb{R}^{n \times k}, W_{0} \in \mathbb{R}^{n \times k}$ to be the optimal solution of $\min _{W \in \mathbb{R}^{n \times k}}\left\|W \cdot\left(\widehat{U}^{\top} \odot \widehat{V}^{\top}\right)-A_{3}\right\|_{1}$, and $V^{\prime}$ to be the optimal solution of $\min _{V \in \mathbb{R}^{n \times k}}\left\|V B_{2} D_{2}-A_{2} D_{2}\right\|_{1}$.

By Claim 21.3.4, we have

$$
\left\|\widehat{V} B_{2} D_{2}-A_{2} D_{2}\right\|_{1} \leq \sqrt{d_{2}}\left\|V^{\prime} B_{2} D_{2}-A_{2} D_{2}\right\|_{1}
$$

Due to Lemma D. 11 and Lemma D.8(in [SWZ17]) with constant probability, we have

$$
\begin{equation*}
\left\|\widehat{V} B_{2}-A_{2}\right\|_{1} \leq \sqrt{d_{2}} \alpha_{D_{2}}\left\|V^{*} B_{2}-A_{2}\right\|_{1} \tag{21.43}
\end{equation*}
$$

where $\alpha_{D_{2}}=O(1)$.
Recall that $\left(\widehat{U}^{\top} \odot \widehat{V}^{\top}\right) \in \mathbb{R}^{k \times n^{2}}$ denotes the matrix for which the $i$-th row is the
vectorization of $\widehat{U}_{i} \otimes \widehat{V}_{i}, \forall i \in[k]$. Now, we can show,

$$
\begin{array}{rlr}
\left\|W_{0} \cdot\left(\widehat{U}^{\top} \odot \widehat{V}^{\top}\right)-A_{3}\right\|_{1} & \leq\left\|\widehat{V} B_{2}-A_{2}\right\|_{1} & \text { by } W_{0}=\underset{W \in \mathbb{R}^{n \times k}}{\arg \min }\left\|W \cdot\left(\widehat{U}^{\top} \odot \widehat{V}^{\top}\right)-A_{3}\right\|_{1} \\
& \lesssim \sqrt{d_{2}}\left\|V^{*} B_{2}-A_{2}\right\|_{1} & \text { by Equation (21.43) } \\
& \leq \sqrt{d_{2}}\left\|V_{0} B_{2}-A_{2}\right\|_{1} & \text { by } V^{*}=\underset{V \in \mathbb{R}^{n \times k}}{\arg \min }\left\|V B_{2}-A_{2}\right\|_{1}
\end{array}
$$

$$
\begin{equation*}
\leq O\left(\sqrt{d_{1} d_{2}}\right) \alpha \mathrm{OPT} \tag{21.42}
\end{equation*}
$$

We define $B_{3}=\widehat{U}^{\top} \odot \widehat{V}^{\top}$. We can compute $D_{3} \in \mathbb{R}^{n^{2} \times n^{2}}$ which is a sampling and rescaling matrix corresponding to the Lewis weights of $B_{3}^{\top}$ in $O\left(n^{2}\right.$ poly $\left.(k)\right)$ time, and there are $d_{3}=O(k \log k)$ nonzero entries on the diagonal of $D_{3}$.

Define $W^{*} \in \mathbb{R}^{n \times k}$ to be the optimal solution to $\min _{W \in \mathbb{R}^{n \times k}}\left\|W B_{3}-A_{3}\right\|_{1}, \widehat{W}=$ $A_{3} D_{3}\left(B_{3} D_{3}\right)^{\dagger} \in \mathbb{R}^{n \times k}$, and $W^{\prime}$ to be the optimal solution to $\min _{W \in \mathbb{R}^{n \times k}}\left\|W B_{3} D_{3}-A_{3} D_{3}\right\|_{1}$.

By Claim 21.3.4, we have

$$
\left\|\widehat{W} B_{3} D_{3}-A_{3} D_{3}\right\|_{1} \leq \sqrt{d_{3}}\left\|W^{\prime} B_{3} D_{3}-A_{3} D_{3}\right\|_{1}
$$

Due to Lemma D. 11 and Lemma D.8(in [SWZ17]) with constant probability, we have

$$
\begin{equation*}
\left\|\widehat{W} B_{3}-A_{3}\right\|_{1} \leq \sqrt{d_{3}} \alpha_{D_{3}}\left\|W^{*} B_{3}-A_{3}\right\|_{1} \tag{21.45}
\end{equation*}
$$

where $\alpha_{D_{3}}=O(1)$. Now we can show,

$$
\begin{array}{rlr}
\left\|\widehat{W} B_{3}-A_{3}\right\|_{1} & \lesssim \sqrt{d_{3}}\left\|W^{*} B_{3}-A_{3}\right\|_{1}, & \text { by Equation }(21.45) \\
& \leq \sqrt{d_{3}}\left\|W_{0} B_{3}-A_{3}\right\|_{1}, & \text { by } W^{*}=\underset{W \in \mathbb{R}^{n \times k}}{\arg \min }\left\|W B_{3}-A_{3}\right\|_{1} \\
& \text { by Equation }(21.44)
\end{array}
$$

Thus, it implies,

$$
\left\|\sum_{i=1}^{k} \widehat{U}_{i} \otimes \widehat{V}_{i} \otimes \widehat{W}_{i}-A\right\|_{1} \leq \operatorname{poly}(k, \log n) \mathrm{OPT}
$$

where $\widehat{U}=A_{1} D_{1}\left(B_{1} D_{1}\right)^{\dagger}, \widehat{V}=A_{2} D_{2}\left(B_{2} D_{2}\right)^{\dagger}, \widehat{W}=A_{3} D_{3}\left(B_{3} D_{3}\right)^{\dagger}$.

```
Algorithm \(21.30 \ell_{1}\)-CURT decomposition algorithm
    procedure \(\mathrm{L1CURT}^{+}(A, n, k)\)
        \(U_{B}, V_{B}, W_{B} \leftarrow\) L1LowRankApproximation \((A, n, k)\).
        \(C, R, T, U \leftarrow \operatorname{L1CURT}\left(A, U_{B}, V_{B}, W_{B}, n, k\right)\).
                                Theorem 21.7.18
                                \(\triangleright\) Corollary 21.7.15
        return \(C, R, T\) and \(U\).
    end procedure
```

Theorem 21.7.18. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, there exists an algorithm which takes $O(\mathrm{nnz}(A))+O\left(n^{2}\right) \operatorname{poly}(k)+2^{\widetilde{O}\left(k^{2}\right)}$ time and outputs three matrices $C \in \mathbb{R}^{n \times c}$ with columns from $A, R \in \mathbb{R}^{n \times r}$ with rows from $A, T \in \mathbb{R}^{n \times t}$ with tubes from $A$, and a tensor $U \in \mathbb{R}^{c \times r \times t}$ with $\operatorname{rank}(U)=k$ such that $c=r=t=O(k \log k)$, and

$$
\left\|\sum_{i=1}^{c} \sum_{j=1}^{r} \sum_{l=1}^{t} U_{i, j, l} \cdot C_{i} \otimes R_{j} \otimes T_{l}-A\right\|_{1} \leq \operatorname{poly}(k, \log n)_{\operatorname{mank}-k} \min _{A^{\prime}}\left\|A^{\prime}-A\right\|_{1},
$$

holds with probability 9/10.

Proof. This follows by combining Corollary 21.7.15 and Theorem 21.7.17.

### 21.8 Weighted Frobenius Norm for Arbitrary Tensors

This section presents several tensor algorithms for the weighted case. For notational purposes, instead of using $U, V, W$ to denote the ground truth factorization, we use $U_{1}, U_{2}, U_{3}$ to denote the ground truth factorization. We use $A$ to denote the input tensor, and $W$ to denote the tensor of weights. Combining our new tensor techniques with existing weighted low rank approximation algorithms [RSW16] allows us to obtain several interesting new results. We provide some necessary definitions and facts in Section 21.8.1. Section 21.8.2 provides an algorithm when $W$ has at most $r$ distinct faces in each dimension. Section 21.8.3 studies relationships between $r$ distinct faces and $r$ distinct columns. Finally, we provides an algorithm with a similar running time but weaker assumption, where $W$ has at most $r$ distinct columns and $r$ distinct rows in Section 21.8.4. The result in Theorem 21.8.2 is fairly similar to Theorem 21.8.5, except for the running time. We only put a very detailed discussion in the statement of Theorem 21.8.5. Note that Theorem 21.8.2 also has other versions which are similar to the Frobnius norm rank- $k$ algorithms described in Section 21.1. For simplicity of presentation, we only present one clean and simple version (which assumes $A_{k}$ exists and has factor norms which are not too large).

### 21.8.1 Definitions and Facts

For a matrix $A \in \mathbb{R}^{n \times m}$ and a weight matrix $W \in \mathbb{R}^{n \times m}$, we define $\|W \circ A\|_{F}$ as follows,

$$
\|W \circ A\|_{F}=\left(\sum_{i=1}^{n} \sum_{j=1}^{m} W_{i, j}^{2} A_{i, j}^{2}\right)^{\frac{1}{2}}
$$

For a tensor $A \in \mathbb{R}^{n \times n \times n}$ and a weight tensor $W \in \mathbb{R}^{n \times n \times n}$, we define $\|W \circ A\|_{F}$ as follows,

$$
\|W \circ A\|_{F}=\left(\sum_{i=1}^{n} \sum_{j=1}^{n} \sum_{l=1}^{n} W_{i, j, l}^{2} A_{i, j, l}^{2}\right)^{\frac{1}{2}}
$$

For three matrices $A \in \mathbb{R}^{n \times m}, U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times m}$ and a weight matrix $W$, from one perspective, we have

$$
\|(U V-A) \circ W\|_{F}^{2}=\sum_{i=1}^{n}\left\|\left(U^{i} V-A^{i}\right) \circ W^{i}\right\|_{2}^{2}=\sum_{i=1}^{n}\left\|\left(U^{i} V-A^{i}\right) D_{W^{i}}\right\|_{2}^{2}
$$

where $W^{i}$ denote the $i$-th row of matrix $W$, and $D_{W^{i}} \in \mathbb{R}^{m \times m}$ denotes a diagonal matrix where the $j$-th entry on diagonal is the $j$-th entry of vector $W^{i}$. From another perspective, we have

$$
\|(U V-A) \circ W\|_{F}^{2}=\sum_{j=1}^{m}\left\|\left(U V_{j}-A_{j}\right) \circ W_{j}\right\|_{2}^{2}=\sum_{j=1}^{m}\left\|\left(U V_{j}-A_{j}\right) D_{W_{j}}\right\|_{2}^{2}
$$

where $W_{j}$ denotes the $j$-th column of matrix $W$, and $D_{W_{j}} \in \mathbb{R}^{n \times n}$ denotes a diagonal matrix where the $i$-th entry on the diagonal is the $i$-th entry of vector $W_{j}$.

One of the key tools we use in this section is,
Lemma 21.8.1 (Cramer's rule). Let $R$ be an $n \times n$ invertible matrix. Then, for each $i \in[n], j \in[n]$,

$$
\left(R^{-1}\right)_{i}^{j}=\operatorname{det}\left(R_{\neg j}^{\neg i}\right) / \operatorname{det}(R)
$$

where $R_{\neg j}^{\neg i}$ is the matrix $R$ with the $i$-th row and the $j$-th column removed.

### 21.8.2 $r$ distinct faces in each dimension

Notice that in the matrix case, it is sufficient to assume that $\left\|A^{\prime}\right\|_{F}$ is upper bounded [RSW16]. Once we have that $\left\|A^{\prime}\right\|_{F}$ is bounded, without loss of generality, we can assume that $U_{1}^{*}$ is an orthonormal basis[CW15a, RSW16]. If $U_{1}^{*}$ is not an orthonormal basis, then let $U_{1}^{\prime} R$ denote a QR factorization of $U_{1}^{*}$, and then write $U_{2}^{\prime}=R U_{2}^{*}$. However, in the case of tensors we have to assume that each factor $\left\|U_{i}^{*}\right\|_{F}$ is upper bounded due to border rank issues (see, e.g., [DSL08]).

```
Algorithm 21.31 Weighted Tensor Low-rank Approximation Algorithm when the Weighted
Tensor has \(r\) Distinct Faces in Each of the Three Dimensions.
    procedure WeightedRDistinctFacesIn3Dimensions \((A, W, n, r, k, \epsilon) \quad \triangleright\)
    Theorem 21.8.2
        for \(j=1 \rightarrow 3\) do
        \(s_{j} \leftarrow O(k / \epsilon)\).
        Choose a sketching matrix \(S_{j} \in \mathbb{R}^{n^{2} \times s_{j}}\).
        for \(i=1 \rightarrow r\) do
            Create \(k \times s_{1}\) variables for matrix \(P_{i, j} \in \mathbb{R}^{k \times s_{j}}\).
        end for
        for \(i=1 \rightarrow n\) do
            Write down \(\left(\widehat{U}_{j}\right)^{i}=A_{i}^{j} D_{W_{1}^{j}} S_{j} P_{j, i}^{\top}\left(P_{j, i} P_{j, i}^{\top}\right)^{-1}\).
        end for
        end for
        Form \(\left\|W \circ\left(\widehat{U}_{1} \otimes \widehat{U}_{2} \otimes \widehat{U}_{3}-A\right)\right\|_{F}^{2}\).
        Run polynomial system verifier.
        return \(U_{1}, U_{2}, U_{3}\)
    end procedure
```

Theorem 21.8.2. Given a 3 rd order $n \times n \times n$ tensor $A$ and an $n \times n \times n$ tensor $W$ of weights with $r$ distinct faces in each of the three dimensions for which each entry can be written using $O\left(n^{\delta}\right)$ bits, for $\delta>0$, define $\mathrm{OPT}=\inf _{\text {rank }-k A_{k}}\left\|W \circ\left(A_{k}-A\right)\right\|_{F}^{2}$. Let $k \geq 1$ be an integer and let $0<\epsilon<1$.

If $\mathrm{OPT}>0$, and there exists a rank- $k A_{k}=U_{1}^{*} \otimes U_{2}^{*} \otimes U_{3}^{*}$ tensor (with size $n \times n \times n$ ) such that $\left\|W \circ\left(A_{k}-A\right)\right\|_{F}^{2}=\mathrm{OPT}$, and $\max _{i \in[3]}\left\|U_{i}^{*}\right\|_{F} \leq 2^{O\left(n^{\delta}\right)}$, then there exists an algorithm that takes $\left(\mathrm{nnz}(A)+\mathrm{nnz}(W)+n 2^{\widetilde{O}\left(r k^{2} / \epsilon\right)}\right) n^{O(\delta)}$ time in the unit cost RAM model with words of size $O(\log n)$ bits ${ }^{10}$ and outputs three $n \times k$ matrices $U_{1}, U_{2}, U_{3}$ such that

$$
\begin{equation*}
\left\|W \circ\left(U_{1} \otimes U_{2} \otimes U_{3}-A\right)\right\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT} \tag{21.46}
\end{equation*}
$$

holds with probability 9/10.

Proof. Note that $W$ has $r$ distinct columns, rows, and tubes. Hence, each of the matrices $W_{1}, W_{2}, W_{3} \in \mathbb{R}^{n \times n^{2}}$ has at most $r$ distinct columns, and at most $r$ distinct rows. Let $U_{1}^{*}, U_{2}^{*}, U_{3}^{*} \in \mathbb{R}^{n \times k}$ denote the matrices satisfying $\left\|W \circ\left(U_{1}^{*} \otimes U_{2}^{*} \otimes U_{3}^{*}-A\right)\right\|_{F}^{2}=$ OPT. We fix $U_{2}^{*}$ and $U_{3}^{*}$, and consider a flattening of the tensor along the first dimension,

$$
\min _{U_{1} \in \mathbb{R}^{n \times k}}\left\|\left(U_{1} Z_{1}-A_{1}\right) \circ W_{1}\right\|_{F}^{2}=\mathrm{OPT},
$$

where matrix $Z_{1}=U_{2}^{* \top} \odot U_{3}^{* \top}$ has size $k \times n^{2}$ and for each $i \in[k]$ the $i$-th row of $Z_{1}$ is $\operatorname{vec}\left(\left(U_{2}^{*}\right)_{i} \otimes\left(U_{3}^{*}\right)_{i}\right)$. For each $i \in[n]$, let $W_{1}^{i}$ denote the $i$-th row of $n \times n^{2}$ matrix $W_{1}$. For each $i \in[n]$, let $D_{W_{1}^{i}}$ denote the diagonal matrix of size $n^{2} \times n^{2}$, where each diagonal entry is from the vector $W_{1}^{i} \in \mathbb{R}^{n^{2}}$. Without loss of generality, we can assume the first $r$ rows of $W_{1}$ are distinct. We can rewrite the objective function along the first dimension as a sum of multiple regression problems. For any $n \times k$ matrix $U_{1}$,

$$
\begin{equation*}
\left\|\left(U_{1} Z_{1}-A_{1}\right) \circ W_{1}\right\|_{F}^{2}=\sum_{i=1}^{n}\left\|U_{1}^{i} Z_{1} D_{W_{1}^{i}}-A_{1}^{i} D_{W_{1}^{i}}\right\|_{2}^{2} \tag{21.47}
\end{equation*}
$$

[^73]Based on the observation that $W_{1}$ has $r$ distinct rows, we can group the $n$ rows of $W^{1}$ into $r$ groups. We use $g_{1,1}, g_{1,2}, \cdots, g_{1, r}$ to denote $r$ sets of indices such that, for each $i \in g_{1, j}$, $W_{1}^{i}=W_{1}^{j}$. Thus we can rewrite Equation (21.47),

$$
\begin{aligned}
\left\|\left(U_{1} Z_{1}-A_{1}\right) \circ W_{1}\right\|_{F}^{2} & =\sum_{i=1}^{n}\left\|U_{1}^{i} Z_{1} D_{W_{1}^{i}}-A_{1}^{i} D_{W_{1}^{i}}\right\|_{2}^{2} \\
& =\sum_{j=1}^{r} \sum_{i \in g_{1, j}}\left\|U_{1}^{i} Z_{1} D_{W_{1}^{i}}-A_{1}^{i} D_{W_{1}^{i}}\right\|_{2}^{2}
\end{aligned}
$$

We can sketch the objective function by choosing Gaussian matrices $S_{1} \in \mathbb{R}^{n^{2} \times s_{1}}$ with $s_{1}=O(k / \epsilon)$.

$$
\sum_{i=1}^{n}\left\|U_{1}^{i} Z_{1} D_{W_{1}^{i}} S_{1}-A_{1}^{i} D_{W_{1}^{i}} S_{1}\right\|_{2}^{2}
$$

Let $\widehat{U}_{1}$ denote the optimal solution of the sketch problem,

$$
\widehat{U}_{1}=\underset{U_{1} \in \mathbb{R}^{n \times k}}{\arg \min } \sum_{i=1}^{n}\left\|U_{1}^{i} Z_{1} D_{W_{1}^{i}} S_{1}-A_{1}^{i} D_{W_{1}^{i}} S_{1}\right\|_{2}^{2}
$$

By properties of $S_{1}\left([\right.$ RSW16] $)$, plugging $\widehat{U} \in \mathbb{R}^{n \times k}$ into the original problem, we obtain,

$$
\sum_{i=1}^{n}\left\|\widehat{U}_{1}^{i} Z_{1} D_{W_{1}^{i}}-A_{1}^{i} D_{W_{1}^{i}}\right\|_{2}^{2} \leq(1+\epsilon) \mathrm{OPT}
$$

Note that $\widehat{U}_{1} \in \mathbb{R}^{n \times k}$ also has the following form. For each $i \in[n]$,

$$
\begin{aligned}
\widehat{U}_{1}^{i} & =A_{1}^{i} D_{W_{1}^{i}} S_{1}\left(Z_{1} D_{W_{1}^{i}} S_{1}\right)^{\dagger} \\
& =A_{1}^{i} D_{W_{1}^{i}} S_{1}\left(Z_{1} D_{W_{1}^{i}} S_{1}\right)^{\top}\left(\left(Z_{1} D_{W_{1}^{i}} S_{1}\right)\left(Z_{1} D_{W_{1}^{i}} S_{1}\right)^{\top}\right)^{-1} .
\end{aligned}
$$

Note that $W_{1}$ has $r$ distinct rows. Thus, we only have $r$ distinct $D_{W_{1}^{i}}$. This implies that there are $r$ distinct matrices $Z_{1} D_{W_{1}^{i}} S_{1} \in \mathbb{R}^{k \times s_{1}}$. Using the definition of $g_{1, j}$, for $j \in[r]$, for
each $i \in g_{1, j} \subset[n]$, we have

$$
\begin{array}{rlr}
\widehat{U}_{1}^{i} & =A_{1}^{i} D_{W_{1}^{i}} S_{1}\left(Z_{1} D_{W_{1}^{i}} S_{1}\right)^{\dagger} & \\
& =A_{1}^{i} D_{W_{1}^{j}} S_{1}\left(Z_{1} D_{W_{1}^{j}} S_{1}\right)^{\dagger} & \text { by } W_{1}^{i}=W_{1}^{j}
\end{array}
$$

which means we only need to write down $r$ different $Z_{1} D_{W_{1}^{j}} S_{1}$. For each $k \times s_{1}$ matrix $Z_{1} D_{W_{1}^{j}} S_{1}$, we create $k \times s_{1}$ variables to represent it. Thus, we need to create $r k s_{1}$ variables to represent $r$ matrices,

$$
\left\{Z_{1} D_{W_{1}^{1}} S_{1}, Z_{1} D_{W_{1}^{2}} S_{1}, \cdots, Z_{1} D_{W_{1}^{r}} S_{1}\right\}
$$

For simplicity, let $P_{1, i} \in \mathbb{R}^{k \times s_{1}}$ denote $Z_{1} D_{W_{1}^{i}} S_{1}$. Then we can rewrite $\widehat{U}^{i} \in \mathbb{R}^{k}$ as follows,

$$
\widehat{U}_{1}^{i}=A_{1}^{i} D_{W_{1}^{i}} S_{1} P_{1, i}^{\top}\left(P_{1, i} P_{1, i}^{\top}\right)^{-1}
$$

If $P_{1, i} P_{1, i}^{\top} \in \mathbb{R}^{k \times k}$ has rank $k$, then we can use Cramer's rule (Lemma 21.8.1) to write down the inverse of $P_{1, i} P_{1, i}^{\top}$. However, vector $W_{1}^{i}$ could have many zero entries. Then the rank of $P_{1, i} P_{1, i}^{\top}$ can be smaller than $k$. There are two different ways to solve this issue.

One way is by using the argument from [RSW16], which allows us to assume that $P_{1, i} P_{1, i}^{\top} \in \mathbb{R}^{k \times k}$ has rank $k$.

The other way is straightforward: we can guess the rank. There are $k$ possibilities. Let $t_{i} \leq k$ denote the rank of $P_{1, i}$. Then we need to figure out a maximal linearly independent subset of rows of $P_{1, i}$. There are $2^{O(k)}$ possibilities. Next, we need to figure out a maximal linearly independent subset of columns of $P_{1, i}$. We can also guess all the possibilities, which is at most $2^{O(k)}$. Because we have $r$ different $P_{1, i}$, the total number of guesses we have is at most $2^{O(r k)}$. Thus, we can write down $\left(P_{1, i} P_{1, i}^{\top}\right)^{-1}$ according to Cramer's rule.

After $\widehat{U}_{1}$ is obtained, we will fix $\widehat{U}_{1}$ and $U_{3}^{*}$ in the next round. We consider the flattening of the tensor along the second direction,

$$
\min _{U_{2} \in \mathbb{R}^{n \times k}}\left\|\left(U_{2} Z_{2}-A_{2}\right) \circ W_{2}\right\|_{F}^{2},
$$

where $n \times n^{2}$ matrix $A_{2}$ is obtained by flattening tensor $A$ along the second dimension, $k \times n^{2}$ matrix $Z_{2}$ denotes $\widehat{U}_{1}^{\top} \odot U_{3}^{* \top}$, and $n \times n^{2}$ matrix $W_{2}$ is obtained by flattening tensor $W$ along the second dimension. For each $i \in[n]$, let $W_{2}^{i}$ denote the $i$-th row of $n \times n^{2}$ matrix $W_{2}$. For each $i \in[n]$, let $D_{W_{1}^{i}}$ denote the diagonal matrix which has size $n^{2} \times n^{2}$ and for which each entry is from vector $W_{2}^{i} \in \mathbb{R}^{n^{2}}$. Without loss of generality, we can assume the first $r$ rows of $W_{2}$ are distinct. We can rewrite the objective function along the second dimension as a sum of multiple regression problems. For any $n \times k$ matrix $U_{2}$,

$$
\begin{equation*}
\left\|\left(U_{2} Z_{2}-A_{2}\right) \circ W_{2}\right\|_{F}^{2}=\sum_{i=1}^{n}\left\|U_{2}^{i} Z_{2} D_{W_{2}^{i}}-A_{2}^{i} D_{W_{2}^{i}}\right\|_{2}^{2} \tag{21.48}
\end{equation*}
$$

Based on the observation that $W_{2}$ has $r$ distinct rows, we can group the $n$ rows of $W^{2}$ into $r$ groups. We use $g_{2,1}, g_{2,2}, \cdots, g_{2, r}$ to denote $r$ sets of indices such that, for each $i \in g_{2, j}$, $W_{2}^{i}=W_{2}^{j}$. Thus we obtain,

$$
\begin{aligned}
\left\|\left(U_{2} Z_{2}-A_{2}\right) \circ W_{2}\right\|_{F}^{2} & =\sum_{i=1}^{n}\left\|U_{2}^{i} Z_{2} D_{W_{2}^{i}}-A_{2}^{i} D_{W_{2}^{i}}\right\|_{2}^{2} \\
& =\sum_{j=1}^{r} \sum_{i \in g_{2, j}}\left\|U_{2}^{i} Z_{2} D_{W_{2}^{i}}-A_{2}^{i} D_{W_{2}^{i}}\right\|_{2}^{2}
\end{aligned}
$$

We can sketch the objective function by choosing a Gaussian sketch $S_{2} \in \mathbb{R}^{n^{2} \times s_{2}}$ with $s_{2}=O(k / \epsilon)$. Let $\widehat{U}_{2}$ denote the optimal solution to the sketch problem. Then $\widehat{U}_{2}$ has the form, for each $i \in[n]$,

$$
\widehat{U}_{2}^{i}=A_{2}^{i} D_{W_{2}^{i}} S_{2}\left(Z_{2} D_{W_{2}^{i}} S_{2}\right)^{\dagger}
$$

Similarly as before, we only need to write down $r$ different matrices $Z_{2} D_{W_{2}^{i}} S_{1}$, and for each of them, create $k \times s_{2}$ variables. Let $P_{2, i} \in \mathbb{R}^{k \times s_{2}}$ denote $Z_{2} D_{W_{2}^{i}} S_{2}$. By our guessing argument, we can obtain $\widehat{U}_{2}$.

In the last round, we fix $\widehat{U}_{1}$ and $\widehat{U}_{2}$. We then write down $\widehat{U}_{3}$. Overall, by creating $l=O\left(r k^{2} / \epsilon\right)$ variables, we have rational polynomials $\widehat{U}_{1}(x), \widehat{U}_{2}(x), \widehat{U}_{3}(x)$. Putting it all together, we can write this objective function,

$$
\begin{aligned}
\min _{x \in \mathbb{R}^{l}} & \left\|\left(\widehat{U}_{1}(x) \otimes \widehat{U}_{2}(x) \otimes \widehat{U}_{3}(x)-A\right) \circ W\right\|_{F}^{2} . \\
\text { s.t. } & h_{1, i}(x) \neq 0, \forall i \in[r] . \\
& h_{2, i}(x) \neq 0, \forall i \in[r] . \\
& h_{3, i}(x) \neq 0, \forall i \in[r] .
\end{aligned}
$$

where $h_{1, i}(x)$ denotes the denominator polynomial related to a full rank sub-block of $P_{1, i}(x)$. By a perturbation argument in Section 4 in [RSW16], we know that the $h_{1, i}(x)$ are nonzero. By a similar argument as in Section 5 in [RSW16], we can show a lower bound on the cost of the denominator polynomial $h_{1, i}(x)$. Thus we can create new bounded variables $x_{l+1}, \cdots, x_{3 r+l}$ to rewrite the objective function,

$$
\begin{array}{rl}
\min _{x \in \mathbb{R}^{l+3 r}} & q(x) / p(x) . \\
\text { s.t. } & h_{1, i}(x) x_{l+i}=0, \forall i \in[r] . \\
& h_{2, i}(x) x_{l+r+i}=0, \forall i \in[r] . \\
& h_{3, i}(x) x_{l+2 r+i}=0, \forall i \in[r] . \\
& p(x)=\prod_{i=1}^{r} h_{1, i}^{2}(x) h_{2, i}^{2}(x) h_{3, i}^{2}(x)
\end{array}
$$

Note that the degree of the above system is poly $(k r)$ and all the equality constraints can be merged into one single constraint. Thus, the number of constraints is $O(1)$. The number of variables is $O\left(r k^{2} / \epsilon\right)$.

Using Theorem 21.3.2 and a similar argument from Section 5 of [RSW16], we have that the minimum nonzero cost is at least $2^{-n^{\delta} 2^{\tilde{O}\left(r k^{2} / \epsilon\right)}}$. Combining the binary search explained in Section 21.4(similar techniques also can be found in Section 6 of [RSW16]) with the lower bound we obtained, we can find the solution for the original problem in time,

$$
\left(\mathrm{nnz}(A)+\operatorname{nnz}(W)+n 2^{\widetilde{O}\left(r k^{2} / \epsilon\right)}\right) n^{O(\delta)} .
$$



$$
\left(W_{1}\right)_{i,(j-1) n+l}=W_{i, j, l}
$$


$\left(W_{2}\right)_{j,(l-1) n+i}=W_{i, j, l}$


$$
\left(W_{3}\right)_{l,(i-1) n+j}=W_{i, j, l}
$$

Figure 21.7: Let $W$ denote a tensor that has columns(red), rows(green) and tubes(blue). For each $i \in[3]$, let $W_{i}$ denote the matrix obtained by flattening tensor $W$ along the $i$-th dimension.

### 21.8.3 $r$ distinct columns, rows and tubes

Lemma 21.8.3. Let $W \in \mathbb{R}^{n \times n \times n}$ denote a tensor that has $r$ distinct columns and $r$ distinct rows, then $W$ has
(I) $r$ distinct column-tube faces .
(II) $r$ distinct row-tube faces.

Proof. Proof of Part (I). Without loss of generality, we consider the first (which is the bottom one) column-row face. Assume it has $r$ distinct rows and $r$ distinct columns. We can re-order all the column-tube faces to make sure that all the $n$ columns in the bottom face have been split into $r$ continuous disjoint groups $C_{i}$, e.g., $\left\{C_{1}, C_{2}, \cdots, C_{r}\right\}=[n]$. Next, we can re-order all the row-tube faces to make sure that all the $n$ rows in the bottom face have been split into $r$ continuous disjoint groups $R_{i}$, e.g., $\left\{R_{1}, R_{2}, \cdots, R_{r}\right\}=[n]$. Thus, the new bottom face can be regarded as $r \times r$ groups, and the number in each position of the same group is the same.

Suppose that the tensor has $r+1$ distinct column-tube faces. By the pigeonhole principle there exist two different column-tube faces belonging to the same group $C_{i}$, for some $i \in[r]$. Note that these two column-tube faces are the same by looking at the bottom (column-row) face. Since they are distinct faces, there must exist one row vector $v$ which is not in the bottom (column-row) face, and it has a different value in coordinates belong to group $C_{i}$. Note that, considering the bottom face, for each row vector, it has the same value over coordinates belonging to group $C_{i}$. But $v$ has different values in coordinates belong to group $C_{i}$. Also, note that the bottom (column-row) face also has $r$ distinct rows, and $v$ is not one of them. This means there are at least $r+1$ distinct rows, which contradicts that there are $r$ distinct rows in total. Thus, there are at most $r$ distinct column-tube faces.

Proof of Part (II). It is similar to Part (I).

Corollary 21.8.4. Let $W \in \mathbb{R}^{n \times n \times n}$ denote $a$ tensor that has $r$ distinct columns, $r$ distinct rows, and $r$ distinct rubes. Then $W$ has $r$ distinct column-tube faces, $r$ distinct row-tube faces, and $r$ distinct column-row faces.


Figure 21.8: Each face $W_{*, *, i}$ is a column-row face. $W_{*, *, 1}$ is the bottom column-row face. $r=3$. The blue blocks represent column-tube faces, the red blocks represent column-tube faces.

Proof. This follows by applying Lemma 21.8.3 twice.

Thus, we obtain the same result as in Theorem 21.8.2 by changing the assumption from $r$ distinct faces in each dimension to $r$ distinct columns, $r$ distinct rows and $r$ distinct tubes.

### 21.8.4 $r$ distinct columns and rows

The main difference between Theorem 21.8.2 and Theorem 21.8.5 is the running time. The first one takes $2^{\widetilde{O}\left(r k^{2} / \epsilon\right)}$ time and the second one is slightly longer, $2^{\widetilde{O}\left(r^{2} k^{2} / \epsilon\right)}$. By Lemma 21.8.3, $r$ distinct columns in two dimensions implies $r$ distinct faces in two of the three kinds of faces. Thus, the following theorem also holds for $r$ distinct columns in two dimensions.

```
Algorithm 21.32 Weighted Tensor Low-rank Approximation Algorithm when the Weighted
Tensor has \(r\) Distinct Faces in Each of the Two Dimensions.
    procedure WeightedRDistinctFacesIn2Dimensions \((A, W, n, r, k, \epsilon) \quad \triangleright\)
    Theorem 21.8.5
        for \(j=1 \rightarrow 3\) do
        \(s_{j} \leftarrow O(k / \epsilon)\).
        Choose a sketching matrix \(S_{j} \in \mathbb{R}^{n^{2} \times s_{j}}\).
        if \(j \neq 3\) then
            for \(i=1 \rightarrow r\) do
                Create \(k \times s_{1}\) variables for matrix \(P_{i, j} \in \mathbb{R}^{k \times s_{j}}\).
            end for
        end if
        for \(i=1 \rightarrow n\) do
            Write down \(\left(\widehat{U}_{j}\right)^{i}=A_{i}^{j} D_{W_{1}^{j}} S_{j} P_{j, i}^{\top}\left(P_{j, i} P_{j, i}^{\top}\right)^{-1}\).
        end for
    end for
    Form \(\left\|W \circ\left(\widehat{U}_{1} \otimes \widehat{U}_{2} \otimes \widehat{U}_{3}-A\right)\right\|_{F}^{2}\).
    Run polynomial system verifier.
    return \(U_{1}, U_{2}, U_{3}\)
    end procedure
```

Theorem 21.8.5. Given a 3 rd order $n \times n \times n$ tensor $A$ and an $n \times n \times n$ tensor $W$ of weights with $r$ distinct faces in two dimensions (out of three dimensions) such that each entry can be written using $O\left(n^{\delta}\right)$ bits for some $\delta>0$, define $\mathrm{OPT}=\inf _{\text {rank }-k A_{k}}\left\|W \circ\left(A_{k}-A\right)\right\|_{F}^{2}$.

For any $k \geq 1$ and any $0<\epsilon<1$.
(I) If OPT $>0$, and there exists a rank-k $A_{k}=U_{1}^{*} \otimes U_{2}^{*} \otimes U_{3}^{*}$ tensor (with size $n \times n \times n)$ such that $\left\|W \circ\left(A_{k}-A\right)\right\|_{F}^{2}=\mathrm{OPT}$, and $\max _{i \in[3]}\left\|U_{i}^{*}\right\|_{F} \leq 2^{O\left(n^{\delta}\right)}$, then there exists an algorithm that takes $\left(\mathrm{nnz}(A)+\mathrm{nnz}(W)+n 2^{\widetilde{O}\left(r^{2} k^{2} / \epsilon\right)}\right) n^{O(\delta)}$ time in the unit cost RAM model with words of size $O(\log n)$ bits ${ }^{11}$ and outputs three $n \times k$ matrices $U_{1}, U_{2}, U_{3}$ such that

$$
\begin{equation*}
\left\|W \circ\left(U_{1} \otimes U_{2} \otimes U_{3}-A\right)\right\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT} \tag{21.49}
\end{equation*}
$$

holds with probability 9/10.
(II) If $\mathrm{OPT}>0, A_{k}$ does not exist, and there exist three $n \times k$ matrices $U_{1}^{\prime}, U_{2}^{\prime}, U_{3}^{\prime}$ where each entry can be written using $O\left(n^{\delta}\right)$ bits and $\left\|W \circ\left(U_{1}^{\prime} \otimes U_{2}^{\prime} \otimes U_{3}^{\prime}-A\right)\right\|_{F}^{2} \leq(1+\epsilon / 2)$ OPT, then we can find $U, V, W$ such that (21.49) holds.
(III) If $\mathrm{OPT}=0, A_{k}$ exists, and there exists a solution $U_{1}^{*}, U_{2}^{*}, U_{3}^{*}$ such that each entry of the matrix can be written using $O\left(n^{\delta}\right)$ bits, then we can obtain (21.49).
(IV) If $\mathrm{OPT}=0$, and there exist three $n \times k$ matrices $U_{1}, U_{2}, U_{3}$ such that $\max _{i \in[3]}\left\|U_{i}^{*}\right\|_{F} \leq$ $2^{O\left(n^{\delta}\right)}$ and

$$
\begin{equation*}
\left\|W \circ\left(U_{1} \otimes U_{2} \otimes U_{3}-A\right)\right\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}+2^{-\Omega\left(n^{\delta}\right)} \tag{21.50}
\end{equation*}
$$

then we can output $U_{1}, U_{2}, U_{3}$ such that (21.50) holds.
(V) Further if $A_{k}$ exists, we can output a number $Z$ for which $\mathrm{OPT} \leq Z \leq(1+\epsilon) \mathrm{OPT}$.

For all the cases, the algorithm succeeds with probability at least 9/10.

[^74]Proof. By Lemma 21.8.3, we have $W$ has $r$ distinct column-tube faces and $r$ distinct row-tube faces. By Claim 21.8.6, we know that $W$ has $R=2^{O(r \log r)}$ distinct column-row faces.

We use the same approach as in proof of Theorem 21.8.2 (which is also similar to Section 8 of [RSW16]) to create variables, write down the polynomial systems and add not equal constraints. Instead of having $3 r$ distinct denominators as in the proof of Theorem 21.8.2, we have $2 r+R$.

We create $l=O\left(r k^{2} / \epsilon\right)$ variables for $\left\{Z_{1} D_{W_{1}^{1}} S_{1}, Z_{1} D_{W_{1}^{2}} S_{1}, \cdots, Z_{1} D_{W_{1}^{r}} S_{1}\right\}$. Then we can write down $\widehat{U}_{1}$ with $r$ distinct denominators $g_{i}(x)$. Each $g_{i}(x)$ is non-zero in an optimal solution using the perturbation argument in Section 4 in [RSW16]. We create new variables $x_{2 l+i}$ to remove the denominators $g_{i}(x), \forall i \in[r]$. Then the entries of $\widehat{U}_{1}$ are polynomials as opposed to rational functions.

We create $l=O\left(r k^{2} / \epsilon\right)$ variables for $\left\{Z_{2} D_{W_{2}^{1}} S_{2}, Z_{2} D_{W_{2}^{2}} S_{2}, \cdots, Z_{2} D_{W_{2}^{r}} S_{2}\right\}$. Then we can write down $\widehat{U}_{2}$ with $r$ distinct denominators $g_{r+i}(x)$. Each $g_{r+i}(x)$ is non-zero in an optimal solution using the perturbation argument in Section 4 in [RSW16]. We create new variables $x_{2 l+r+i}$ to remove the denominators $g_{r+i}(x), \forall i \in[r]$. Then the entries of $\widehat{U}_{2}$ are polynomials as opposed to rational functions.

Using $\widehat{U}_{1}$ and $\widehat{U}_{2}$ we can express $\widehat{U}_{3}$ with $R$ distinct denominators $f_{i}(x)$, which are also non-zero by using the perturbation argument in Section 4 in [RSW16], and using that $W_{3}$ has at most this number of distinct rows. Finally we can write the following optimization
problem,

$$
\begin{array}{rl}
\min _{x \in \mathbb{R}^{2 l+2 r}} & p(x) / q(x) \\
\text { s.t. } & g_{i}(x) x_{2 l+i}-1=0, \forall i \in[r] \\
& g_{r+i}(x) x_{2 l+r+i}-1=0, \forall i \in[r] \\
& f_{j}^{2}(x) \neq 0, \forall j \in[R] \\
& q(x)=\prod_{j=1}^{R} f_{j}^{2}(x)
\end{array}
$$

We then determine if there exists a solution to the above semi-algebraic set in time

$$
(\operatorname{poly}(k, r) R)^{O\left(r k^{2} / \epsilon\right)}=2^{\widetilde{O}\left(r^{2} k^{2} / \epsilon\right)}
$$

Using similar techniques from Section 5 of [RSW16], we can show a lower bound on the cost similar to Section 8.3 of [RSW16], namely, the minimum nonzero cost is at least

$$
2^{-n^{\delta} 2^{\tilde{O}\left(r^{2} k^{2} / \epsilon\right)}}
$$

Combining the binary search explained in Section 21.4 (a similar techniques also can be found in Section 6 of [RSW16]) with the lower bound we obtained, we can find a solution for the original problem in time

$$
\left(\operatorname{nnz}(A)+\operatorname{nnz}(W)+n 2^{\widetilde{O}\left(r^{2} k^{2} / \epsilon\right)}\right) n^{O(\delta)} .
$$

Remark 21.8.1. Note that the running time for the Frobenius norm and for the $\ell_{1}$ norm are of the form $\operatorname{poly}(n)+\exp (\operatorname{poly}(k / \epsilon))$ rather than $\operatorname{poly}(n) \cdot \exp (k / \epsilon)$. The reason is, we can use an input sparsity reduction to reduce the size of the objective function from poly $(n)$ to $\operatorname{poly}(k)$.


Figure 21.9: Each face $W_{*, *, i}$ is a column-row face. $W_{*, *, 1}$ is the bottom column-row face. $r=3$. The blue blocks represent $\left|C_{3}\right|$ column-tube faces. The green blocks represet $\left|R_{3}\right|$ row-tube faces. In each column-row face, the intersection between blue faces and green faces is a size $\left|R_{3}\right| \times\left|C_{3}\right|$ block, and all the entries in this block are the same.

Claim 21.8.6. Let $W \in \mathbb{R}$ denote a third order tensor that has $r$ distinct columns and $r$ distinct rows. Then it has $2^{O(r \log r)}$ distinct column-row faces.

Proof. By similar arguments as in the proof of Lemma 21.8.3, the bottom (column-row) face can be split into $r$ groups $C_{1}, C_{2}, \cdots, C_{r}$ based on $r$ columns, and split into $r$ groups $R_{1}, R_{2}, \cdots, R_{r}$ based on rows. Thus, the bottom (column-row) face can be regarded as having $r \times r$ groups, and the number in each position of the same group is the same.

We can assume that all the $r^{2}$ blocks in the bottom column-row face have the same size. Otherwise, we can expand the tensor to the situation that all the $r^{2}$ blocks have the same size. Because this small tensor is a sub-tensor of the big tensor, if the big tensor has at most $t$ distinct column-row faces, then the small tensor has at most $t$ distinct column-row faces.

By Lemma 21.8.3, we know that the tensor $W$ has at most $r$ distinct column-tube faces and row-tube faces. Because it has $r$ distinct column-tube faces, then all the faces belonging to coordinates in $C_{r}$ are the same. Thus, all the columns belonging to $C_{r}$ and in the second column-row face are the same. Similarly, we have that all the rows belonging to $R_{r}$ and in the second column-row face are the same. Thus we have that all the entries in block $C_{R} \cup R_{r}$ and in the second column-row faces are the same. Further, we can conclude, for every column-row face, for every $C_{i} \cup R_{j}$ block, all the entries in the same block are the same.

The next observation is, if there exist $r^{2}+1$ different values in the tensor, then there exist either $r$ distinct columns or $r$ distinct rows. Indeed, otherwise since we have $r$ distinct columns, each column has at most $r$ distinct entries given our bound on the nunber of distinct rows. Thus, the $r$ distinct columns could have at most $r^{2}$ distinct entries in total, a contradiction.

For each column-row face, there are at most $r^{2}$ blocks, and the value in each block can have at most $r^{2}$ possibilities. Thus, overall we have at most $\left(r^{2}\right)^{r^{2}}=2^{O\left(r^{2} \log r\right)}$ column-row faces.

By using different argument, we can improve the above bound. Note that we already show in each column-row face of a tensor, it has $r^{2}$ blocks, and all the values in each block have to be the same. Since we have $r$ distinct rows, we can fix the those $r$ distinct rows. If we copy row $v$ into one row of $R_{i}$, then we have to copy row $v$ into every row of $R_{i}$. This is because if $R_{i}$ contains two distinct rows, then there must exist a block $C_{j}$ for which the entries in block $R_{i} \cup C_{j}$ are not all the same. Thus, for each row group, all the rows in that group are the same.

Now, for each column-row face, consider the leftmost $r$ blocks, $R_{1} \cup C_{1}, R_{2} \cup C_{1}$, $\cdots, R_{r} \cup C_{1}$. There are at most $r$ possible values in each block, because we have $r$ distinct rows in total. Overall the total number of possibilities for the leftmost $r$ blocks is at most $(r)^{r}=2^{O(r \log r)}$. Once the leftmost $r$ blocks are determined, the remaining $r(r-1)$ are also determined. This completes the proof.

Also, notice that there is an example that has $2^{\Omega(r \log r)}$ distinct column-row faces. For the bottom column-row faces, there are $r \times r$ blocks for which all the blocks have the same size, the blocks on the diagonal have all 1s, and all the other blocks contain 0s everywhere. For the later column-row faces, we can arbitrarily permute this block diagonal matrix, and the total number of possibilities is $\Omega(r!) \geq 2^{\Omega(r \log r)}$.

### 21.9 Hardness

We first provide definitions and results for some fundamental problems in Section 21.9.1. Section 21.9.2 presents our hardness result for the symmetric tensor eigenvalue problem. Section 21.9.3 presents our hardness results for symmetric tensor singular value problems, computing tensor spectral norm, and rank-1 approximation. We improve Håstad's NPhardness[Hås90] result for tensor rank in Section 21.9.4.

### 21.9.1 Definitions

We first provide the definitions for 3SAT, ETH, MAX-3SAT, MAX-E3SAT and then state some fundamental results related to those definitions.

Definition 21.9.1 (3SAT problem). Given $n$ variables and $m$ clauses in a conjunctive normal form CNF formula with the size of each clause at most 3, the goal is to decide whether there exists an assignment to the $n$ Boolean variables to make the CNF formula be satisfied.

Hypothesis 21.9.1 (Exponential Time Hypothesis (ETH) [IPZ98]). There is a $\delta>0$ such that the 3SAT problem defined in Definition 21.9.1 cannot be solved in $O\left(2^{\delta n}\right)$ time.

Definition 21.9.2 (MAX-3SAT). Given $n$ variables and $m$ clauses, a conjunctive normal form CNF formula with the size of each clause at most 3, the goal is to find an assignment that satisfies the largest number of clauses.

We use MAX-E3SAT to denote the version of MAX-3SAT where each clause contains exactly 3 literals.

Theorem 21.9.2 ([Hås01]). For every $\delta>0$, it is NP-hard to distinguish a satisfiable instance of MAX-E3SAT from an instance where at most a $7 / 8+\delta$ fraction of the clauses can be simultaneously satisfied.

Theorem 21.9.3 ([Hås01, MR10]). Assume ETH holds. For every $\delta>0$, there is no $2^{o\left(n^{1-o(1)}\right)}$ time algorithm to distinguish a satisfiable instance of MAX-E3SAT from an instance where at most a fraction $7 / 8+\delta$ of the clauses can be simultaneously satisfied.

We use MAX-E3SAT(B) to denote the restricted special case of MAX-3SAT where every variable occurs in at most $B$ clauses. Håstad [Hås00] proved that the problem is
approximable to within a factor $7 / 8+1 /(64 B)$ in polynomial time, and that it is hard to approximate within a factor $7 / 8+1 /(\log B)^{\Omega(1)}$. In 2001, Trevisan improved the hardness result,

Theorem 21.9.4 ([Tre01]). Unless $\mathbf{R P}=\mathbf{N P}$, there is no polynomial time $(7 / 8+5 / \sqrt{B})$ approximate algorithm for MAX-E3SAT(B) .

Theorem 21.9.5 ([Hås01, Tre01, MR10]). Unless ETH fails, there is no $2^{o\left(n^{1-o(1)}\right)}$ time $(7 / 8+5 / \sqrt{B})$-approximate algorithm for MAX-E3SAT(B).

Theorem 21.9.6 ([LMS11]). Unless ETH fails, there is no $2^{o(n)}$ time algorithm for the Independent Set problem.

Definition 21.9.3 (MAX-CUT decision problem). Given a positive integer $c^{*}$ and an unweighted graph $G=(V, E)$ where $V$ is the set of vertices of $G$ and $E$ is the set of edges of $G$, the goal is to determine whether there is a cut of $G$ that has at least $c^{*}$ edges.

Note that Feige's original assumption[Fei02] states that there is no polynomial time algorithm for the problem in Assumption 21.9.7. We do not know of any better algorithm for the problem in Assumption 21.9.7 and have consulted several experts ${ }^{12}$ about the assumption who do not know a counterexample to it.

Assumption 21.9.7 (Random Exponential Time Hypothesis). Let $c>\ln 2$ be a constant. Consider a random 3SAT formula on $n$ variables in which each clause has 3 literals, and in which each of the $8 n^{3}$ clauses is picked independently with probability $c / n^{2}$. Then any

[^75]algorithm which always outputs 1 when the random formula is satisfiable, and outputs 0 with probability at least $1 / 2$ when the random formula is unsatisfiable, must run in $2^{c^{\prime} n}$ time on some input, where $c^{\prime}>0$ is an absolute constant.

The 4SAT-version of the above random-ETH assumption has been used in [GL04] and [RSW16] (Assumption 1.3).

### 21.9.2 Symmetric tensor eigenvalue

Definition 21.9.4 (Tensor Eigenvalue [HL13]). An eigenvector of a tensor $A \in \mathbb{R}^{n \times n \times n}$ is a nonzero vector $x \in \mathbb{R}^{n}$ such that

$$
\sum_{i=1}^{n} \sum_{j=1}^{n} A_{i, j, k} x_{i} x_{j}=\lambda x_{k}, \forall k \in[n]
$$

for some $\lambda \in \mathbb{R}$, which is called an eigenvalue of $A$.

Theorem 21.9.8 $\left(\left[\mathrm{N}^{+} 03\right]\right)$. Let $G=(V, E)$ on $v$ vertices have stability number (the size of a maximum independent set) $\alpha(G)$. Let $n=v+\frac{v(v-1)}{2}$ and $\mathbb{S}^{n-1}=\left\{(x, y) \in \mathbb{R}^{v} \times \mathbb{R}^{v(v-1) / 2}\right.$ : $\left.\|x\|_{2}^{2}+\|y\|_{2}^{2}=1\right\}$. Then,

$$
\sqrt{1-\frac{1}{\alpha(G)}}=3 \sqrt{3 / 2} \max _{(x, y) \in \mathbb{S}^{n-1}} \sum_{i<j,(i, j) \notin E} x_{i} x_{j} y_{i, j}
$$

For any graph $G(V, E)$, we can construct a symmetric tensor $A \in \mathbb{R}^{n \times n \times n}$. For any $1 \leq i<j<k \leq v$, let

$$
A_{i, j, k}= \begin{cases}1 & 1 \leq i<j \leq v, k=v+\phi(i, j),(i, j) \notin E \\ 0 & \text { otherwise }\end{cases}
$$

where $\phi(i, j)=(i-1) v-i(i-1) / 2+j-i$ is a lexicographical enumeration of the $v(v-1) / 2$ pairs $i<j$. For the other cases $i<k<j, \cdots, k<j<i$, we set

$$
A_{i, j, k}=A_{i, k, j}=A_{j, i, k}=A_{j, k, i}=A_{k, i, j}=A_{k, j, i}
$$

If two or more indices are equal, we set $A_{i, j, k}=0$. Thus tensor $T$ has the following property,

$$
A(z, z, z)=6 \sum_{i<j,(i, j) \notin E} x_{i} x_{j} y_{i, j},
$$

where $z=(x, y) \in \mathbb{R}^{n}$.
Thus, we have

$$
\lambda=\max _{z \in \mathbb{S}^{n-1}} A(z, z, z)=\max _{(x, y) \in \mathbb{S}^{n-1}} 6 \sum_{i<j,(i, j) \notin E} x_{i} x_{j} y_{i, j} .
$$

Furthermore, $\lambda$ is the maximum eigenvalue of $A$.
Theorem 21.9.9. Unless ETH fails, there is no $2^{o(\sqrt{n})}$ time to approximate the largest eigenvalue of an n-dimensional symmetric tensor within $(1 \pm \Theta(1 / n))$ relative error.

Proof. The additive error is at least

$$
\sqrt{1-1 / v}-\sqrt{1-1 /(v-1)}=\frac{1 /(v-1)-1 / v}{\sqrt{1-1 / v}+\sqrt{1-1 /(v-1)}} \gtrsim 1 /(v-1)-1 / v \geq 1 / v^{2}
$$

Thus, the relative error is $\left(1 \pm \Theta\left(1 / v^{2}\right)\right)$. By the definition of $n$, we know $n=\Theta\left(v^{2}\right)$. Assuming ETH, there is no $2^{o(v)}$ time algorithm to compute the clique number of $\bar{G}$. Because the clique number of $\bar{G}$ is $\alpha(G)$, there is no $2^{o(v)}$ time algorithm to compute $\alpha(G)$. Furthermore, there is no $2^{o(v)}$ time algorithm to approximate the maximum eigenvalue within $\left(1 \pm \Theta\left(1 / v^{2}\right)\right)$ relative error. Thus, we complete the proof.

Corollary 21.9.10. Unless ETH fails, there is no polynomial running time algorithm to approximate the largest eigenvalue of an $n$-dimensional tensor within $\left(1 \pm \Theta\left(1 / \log ^{2+\gamma}(n)\right)\right)$ relative-error, where $\gamma>0$ is an arbitrarily small constant.

Proof. We can apply a padding argument here. According to Theorem 21.9.9, there is a $d$-dimensional tensor such that there is no $2^{o(\sqrt{d})}$ time algorithm that can give a $(1+\Theta(1 / d))$ relative error approximation. If we pad 0 s everywhere to extend the size of the tensor
to $n=2^{d^{\left(1-\gamma^{\prime}\right) / 2}}$, where $\gamma^{\prime}>0$ is a sufficiently small constant, then $\operatorname{poly}(n)=2^{o(\sqrt{d})}$, so $d=\log ^{2+O\left(\gamma^{\prime}\right)}(n)$. Thus, it means that there is no polynomial running time algorithm which can output a $\left(1+1 /\left(\log ^{2+\gamma}\right)\right)$-relative approximation to the tensor which has size $n$.

### 21.9.3 Symmetric tensor singular value, spectral norm and rank-1 approximation

[HL13] defines two kinds of singular values of a tensor. In this paper, we only consider the following kind:

Definition 21.9.5 ( $\ell_{2}$ singular value in [HL13]). Given a 3 rd order tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$, the number $\sigma \in \mathbb{R}$ is called a singular value and the nonzero $u \in \mathbb{R}^{n_{1}}, v \in \mathbb{R}^{n_{2}}, w \in \mathbb{R}^{n_{3}}$ are called singular vectors of $A$ if

$$
\begin{aligned}
& \sum_{j=1}^{n_{2}} \sum_{k=1}^{n_{3}} A_{i, j, k} v_{j} w_{k}=\sigma u_{i}, \forall i \in\left[n_{1}\right] \\
& \sum_{i=1}^{n_{1}} \sum_{k=1}^{n_{3}} A_{i, j, k} u_{i} w_{k}=\sigma v_{j}, \forall j \in\left[n_{2}\right] \\
& \sum_{i=1}^{n_{1}} \sum_{j=1}^{n_{2}} A_{i, j, k} u_{i} v_{j}=\sigma w_{k}, \forall k \in\left[n_{3}\right] .
\end{aligned}
$$

Definition 21.9.6 (Spectral norm [HL13]). The spectral norm of a tensor $A$ is:

$$
\|A\|_{2}=\sup _{x, y, z \neq 0} \frac{|A(x, y, z)|}{\|x\|_{2}\|y\|_{2}\|z\|_{2}}
$$

Notice that the spectral norm is the absolute value of either the maximum value of $\frac{A(x, y, z)}{\|x\|_{2}\|y\|_{2}\left\|_{z}\right\|_{2}}$ or the minimum value of it. Thus, it is an $\ell_{2}$-singular value of $A$. Furthermore, it is the maximum $\ell_{2}$-singular value of $A$.

Theorem 21.9.11 ([Ban38]). Let $A \in \mathbb{R}^{n \times n \times n}$ be a symmetric 3 rd order tensor. Then,

$$
\|A\|_{2}=\sup _{x, y, z \neq 0} \frac{A(x, y, z)}{\|x\|_{2}\|y\|_{2}\|z\|_{2}}=\sup _{x \neq 0} \frac{|A(x, x, x)|}{\|x\|_{2}^{3}} .
$$

It means that if a tensor is symmetric, then its largest eigenvalue is the same as its largest singular value and its spectral norm. Then, by combining with Theorem 21.9.9, we have the following corollary:

Corollary 21.9.12. Unless ETH fails,

1. There is no $2^{o(\sqrt{n})}$ time algorithm to approximate the largest singular value of an $n$ dimensional symmetric tensor within $(1+\Theta(1 / n))$ relative-error.
2. There is no $2^{o(\sqrt{n})}$ time algorithm to approximate the spectral norm of an n-dimensional symmetric tensor within $(1+\Theta(1 / n))$ relative-error.

By Corollary 21.9.10, we have:

Corollary 21.9.13. Unless ETH fails,

1. There is no polynomial time algorithm to approximate the largest singular value of an $n$-dimensional tensor within $\left(1+\Theta\left(1 / \log ^{2+\gamma}(n)\right)\right)$ relative-error, where $\gamma>0$ is an arbitrarily small constant.
2. There is no polynomial time algorithm to approximate the spectral norm of an $n$ dimensional tensor within $\left(1+\Theta\left(1 / \log ^{2+\gamma}(n)\right)\right)$ relative-error, where $\gamma>0$ is an arbitrarily small constant.

Now, let us consider Frobenius norm rank-1 approximation.

Theorem 21.9.14 ([Ban38]). Let $A \in \mathbb{R}^{n \times n \times n}$ be a symmetric 3 rd order tensor. Then,

$$
\min _{\sigma \geq 0,\|u\|_{2}=\|v\|_{2}=\|w\|_{2}=1}\|A-\sigma u \otimes v \otimes w\|_{F}=\min _{\lambda \geq 0,\|v\|_{2}=1}\|A-\lambda v \otimes v \otimes v\|_{F} .
$$

Furthermore, the optimal $\sigma$ and $\lambda$ may be chosen to be equal.

Notice that

$$
\|A-\sigma u \otimes v \otimes w\|_{F}^{2}=\|A\|_{F}^{2}-2 \sigma A(u, v, w)+\sigma^{2}\|u \otimes v \otimes w\|_{F}^{2}
$$

Then, if $\|u\|_{2}=\|v\|_{2}=\|w\|_{2}=1$, we have:

$$
\|A-\sigma u \otimes v \otimes w\|_{F}^{2}=\|A\|_{F}^{2}-2 \sigma A(u, v, w)+\sigma^{2}
$$

When $A(u, v, w)=\sigma$, then the above is minimized.
Thus, we have:

$$
\min _{\sigma \geq 0,\|u\|_{2}=\|v\|_{2}=\|w\|_{2}=1}\|A-\sigma u \otimes v \otimes w\|_{F}^{2}+\|A\|_{2}^{2}=\|A\|_{F}^{2} .
$$

It is sufficient to prove the following theorem:

Theorem 21.9.15. Given $A \in \mathbb{R}^{n \times n \times n}$, unless ETH fails, there is no $2^{o(\sqrt{n})}$ time algorithm to compute $u^{\prime}, v^{\prime}, w^{\prime} \in \mathbb{R}^{n}$ such that

$$
\left\|A-u^{\prime} \otimes v^{\prime} \otimes w^{\prime}\right\|_{F}^{2} \leq(1+\epsilon) \min _{u, v, w \in \mathbb{R}^{n}}\|A-u \otimes v \otimes w\|_{F}^{2}
$$

where $\epsilon=O\left(1 / n^{2}\right)$.

Proof. Let $A \in \mathbb{R}^{n \times n \times n}$ be the same hard instance mentioned in Theorem 21.9.8. Notice that each entry of $A$ is either 0 or 1 . Thus, $\min _{u, v, w \in \mathbb{R}^{n}}\|A-u \otimes v \otimes w\|_{F}^{2} \leq\|A\|_{F}^{2}$. Notice that Theorem 21.9.8 also implies that it is hard to distinguish the two cases $\|A\|_{2} \leq 2 \sqrt{2 / 3}$. $\sqrt{1-1 / c}$ or $\|A\|_{2} \geq 2 \sqrt{2 / 3} \cdot \sqrt{1-1 /(c+1)}$ where $c$ is an integer which is no greater than $\sqrt{n}$. So the difference between $(2 \sqrt{2 / 3} \cdot \sqrt{1-1 / c})^{2}$ and $(2 \sqrt{2 / 3} \cdot \sqrt{1-1 /(c+1)})^{2}$ is at
least $\Theta(1 / n)$. Since $\|A\|_{F}^{2}$ is at most $n$ (see construction of $A$ in the proof of Lemma 21.9.8), $\Theta(1 / n)$ is an $\epsilon=O\left(1 / n^{2}\right)$ fraction of $\min _{u, v, w \in \mathbb{R}^{n}}\|A-u \otimes v \otimes w\|_{F}^{2}$. Because

$$
\min _{u, v, w \in \mathbb{R}^{n}}\|A-u \otimes v \otimes w\|_{F}^{2}+\|A\|_{2}^{2}=\|A\|_{F}^{2}
$$

if we have a $2^{o(\sqrt{n})}$ time algorithm to compute $u^{\prime}, v^{\prime}, w^{\prime} \in \mathbb{R}^{n}$ such that

$$
\left\|A-u^{\prime} \otimes v^{\prime} \otimes w^{\prime}\right\|_{F}^{2} \leq(1+\epsilon) \min _{u, v, w \in \mathbb{R}^{n}}\|A-u \otimes v \otimes w\|_{F}^{2}
$$

for $\epsilon=O\left(1 / n^{2}\right)$, it will contradict the fact that we cannot distinguish whether $\|A\|_{2} \leq$ $2 \sqrt{2 / 3} \cdot \sqrt{1-1 / c}$ or $\|A\|_{2} \geq 2 \sqrt{2 / 3} \cdot \sqrt{1-1 /(c+1)}$.

Corollary 21.9.16. Given $A \in \mathbb{R}^{n \times n \times n}$, unless ETH fails, for any $\epsilon$ for which $\frac{1}{2} \geq \epsilon \geq c / n^{2}$ where $c$ is any constant, there is no $2^{o\left(\epsilon^{-1 / 4}\right)}$ time algorithm to compute $u^{\prime}, v^{\prime}, w^{\prime} \in \mathbb{R}^{n}$ such that

$$
\left\|A-u^{\prime} \otimes v^{\prime} \otimes w^{\prime}\right\|_{F}^{2} \leq(1+\epsilon) \min _{u, v, w \in \mathbb{R}^{n}}\|A-u \otimes v \otimes w\|_{F}^{2} .
$$

Proof. If $\epsilon=\Omega\left(1 / n^{2}\right)$, it means that $n=\Omega(1 / \sqrt{\epsilon})$. Then, we can construct a hard instance $B$ with size $m \times m \times m$ where $m=\Theta(1 / \sqrt{\epsilon})$, and we can put $B$ into $A$, and let $A$ have zero entries elsewhere. Since $B$ is hard, i.e., there is no $2^{o\left(m^{-1 / 2}\right)}=2^{o\left(\epsilon^{-1 / 4}\right)}$ running time to compute a rank-1 approximation to $B$, this means there is no $2^{o\left(\epsilon^{-1 / 4}\right)}$ running time algorithm to find an approximate rank-1 approximation to $A$.

Corollary 21.9.17. Unless ETH fails, there is no polynomial time algorithm to approximate the best rank-1 approximation of an $n$-dimensional tensor within $\left(1+\Theta\left(1 / \log ^{2+\gamma}(n)\right)\right)$ relative-error, where $\gamma>0$ is an arbitrarily small constant.

Proof. We can apply a padding argument here. According to Theorem 21.9.15, there is a $d$ dimensional tensor such that there is no $2^{o(\sqrt{d})}$ time algorithm which can give a $\left(1+\Theta\left(1 / d^{4}\right)\right)$ relative approximation. Then, if we pad with 0s everywhere to extend the size of the tensor to $n=2^{d^{\left(1-\gamma^{\prime}\right) / 2}}$ where $\gamma^{\prime}>0$ is a sufficiently small constant, then poly $(n)=2^{o(\sqrt{d})}$, and $d^{4}=\log ^{2+O\left(\gamma^{\prime}\right)}(n)$. Thus, it means that there is no polynomial time algorithm which can output a $\left(1+1 /\left(\log ^{2+\gamma}\right)\right)$-relative error approximation to the tensor which has size $n$.

### 21.9.4 Tensor rank is hard to approximate

This section presents the hardness result for approximating tensor rank under ETH. According to our new result, we notice that not only deciding the tensor rank is a hard problem, but also approximating the tensor rank is a hard problem. This therefore strengthens Håstad's NP-Hadness [Hås90] for computing tensor rank.

### 21.9.4.1 Cover number

Before getting into the details of the reduction, we provide a definition of an important concept called the "cover number" and discuss the cover number for the MAX-E3SAT(B) problem.

Definition 21.9.7 (Cover number). For any 3SAT instance $S$ with $n$ variables and $m$ clauses, we are allowed to assign one of three values $\{0,1, *\}$ to each variable. For each clause, if one of the literals outputs true, then the clause outputs true. For each clause, if the corresponding variable of one of the literals is assigned to $*$, then the clause outputs true. We say $y \in\{0,1\}^{n}$ is a string, and $z \in\{0,1, *\}^{n}$ is a star string. For an instance $S$, if there exists a string $y \in\{0,1\}^{n}$ that causes all the clauses to be true, then we say that $S$ is satisfiable, otherwise it is unsatisfiable. For an instance $S$, let $Z_{S}$ denote the set of star strings which cause all of the clauses of $S$ to be true. For each star string $z \in\{0,1, *\}^{n}$, let $\operatorname{star}(z)$ denote the number of $*$ s in the star-string $z$. We define the "cover number" of instance $S$ to be

$$
\text { cover-number }(S)=\min _{z \in Z_{S}} \operatorname{star}(z)
$$

Notice that for a satisfiable 3SAT instance $S$, the cover number $p$ is 0 . Also, for any


Figure 21.10: Cover number. For a 3SAT instance with $n$ variables and $m$ clauses, we can draw a bipartite graph which has $n$ nodes on the left and $m$ nodes on the right. Each node (blue) on the left corresponds to a variable $x_{i}$, each node (green) on the right corresponds to a clause $C_{j}$. If either $x_{i}$ or $\bar{x}_{i}$ belongs to clause $C_{j}$, then we draw a line between these two nodes. Consider an input string $y \in\{0,1\}^{7}$. There exists some unsatisfied clauses with respect to this input string $y$. For for example, let $C_{1}, C_{2}$ and $C_{3}$ denote those unsatisfied clauses. We want to pick a smallest set of nodes on the left partition of the graph to guarantee that for each unsatisfied clause in the right partition, there exists a node on the left to cover it. The cover number is defined to be the smallest such number over all possible input strings.
unsatisfiable 3SAT instance $S$, the cover number $p$ is at least 1 . This is because for any input string, there exists at least one clause which cannot be satisfied. To fix that clause, we have to assign $*$ to a variable belonging to that clause. (Assigning $*$ to a variable can be regarded as assigning both 0 and 1 to a variable)

Lemma 21.9.18. Let $S$ denote a MAX-E3SAT(B) instance with $n$ variables and $m$ clauses and $S$ suppose $S$ is at most $7 / 8+A$ satisfiable, where $A \in(0,1 / 8)$. Then the cover number
of $S$ is at least $(1 / 8-A) m / B$.

Proof. For any input string $y \in\{0,1\}^{n}$, there exists at least $(1 / 8-A) m$ clauses which are not satisfied. Since each variable appears in at most $B$ clauses, we need to assign $*$ to at least $(1 / 8-A) m / B$ variables. Thus, the cover number of $S$ is at least $(1 / 8-A) m / B$.

We say $x_{1}, x_{2}, \cdots, x_{n}$ are variables and $x_{1}, \bar{x}_{1}, x_{2}, \bar{x}_{2}, \cdots, x_{n}, \bar{x}_{n}$ are literals.

Definition 21.9.8. For a list of clauses $C$ and a set of variables $P$, if for each clause, there exists at least one literal such that the corresponding variable of that literal belongs to $P$, then we say $P$ covers $L$.

### 21.9.4.2 Properties of 3SAT instances

Fact 21.9.19. For any 3SAT instance $S$ with $n$ variables and $m=\Theta(n)$ clauses, let $c>0$ denote a constant. If $S$ is $(1-c) m$ satisfiable, then let $y \in\{0,1\}^{n}$ denote a string for which $S$ has the smallest number of unsatisfiable clauses. Let $T$ denote the set of unsatisfiable clauses and let $b$ denote the number of variables in $T$. Then $\Omega\left((c m)^{1 / 3}\right) \leq b \leq O(c m)$.

Proof. Note that in $S$, there is no duplicate clause. Let $T$ denote the set of unsatisfiable clauses by assigning string $y$ to $S$. First, we can show that any two literals $x_{i}, \bar{x}_{i}$ cannot belong to $T$ at the same time. If $x_{i}$ and $\bar{x}_{i}$ belong to the same clause, then that clause must be an "always" satisfiable clause. If $x_{i}$ and $\bar{x}_{i}$ belong to different clauses, then one of the clauses must be satisfiable. This contradicts the fact that that clause belongs to $T$. Thus, we can assume that literals $x_{1}, x_{2}, \cdots, x_{b}$ belong to $T$.

There are two extreme cases: one is that each clause only contains three literals and each literal appears in exactly one clause in $T$. Then $b=3 \mathrm{~cm}$. The other case is that each clause contains 3 literals, and each literal appears in as many clauses as possible. Then $\binom{b}{3}=c m$, which gives $b=\Theta\left((c m)^{1 / 3}\right)$.

Lemma 21.9.20. For a random 3SAT instance, with probability $1-2^{-\Omega(\log n \log \log n)}$ there is no literal appearing in at least $\log n$ clauses.

Proof. By the property of random 3SAT, for any literal $x$ and any clause $C$, the probability that $x$ appears in $C$ is $\frac{3}{2 n}$, i.e., $\operatorname{Pr}[x \in C]=\frac{3}{2 n}=\Theta(1 / n)$. Let $p$ denote this probability. For any literal $x$, the probability of $x$ appearing in at least $\log n$ clauses (out of $m$ clauses) is

$$
\begin{aligned}
& \operatorname{Pr}[x \text { appearing in } \geq \log n \text { clauses }] \\
= & \sum_{i=\log n}^{m}\binom{m}{i} p^{i}(1-p)^{m-i} \\
= & \sum_{i=\log n}^{m / 2}\binom{m}{i} p^{i}(1-p)^{m-i}+\sum_{i=m / 2}^{m}\binom{m}{i} p^{i}(1-p)^{m-i} \\
\leq & \sum_{i=\log n}^{m / 2}(e m / i)^{i} p^{i}+\sum_{i=m / 2}^{m}\binom{m}{i} p^{i} \quad \text { by }(1-p) \leq 1,\binom{m}{i} \leq(e m / i)^{i} \\
\leq & (\Theta(1 / \log n))^{\log n}+2 \cdot(2 e)^{m / 2} \cdot \Theta(1 / n)^{m / 2} \quad \\
\leq & 2^{-\Omega(\log n \cdot \log \log n)} .
\end{aligned}
$$

Taking a union bound over all the literals, we complete the proof,

$$
\operatorname{Pr}[\nexists x \text { appearing in } \geq \log n \text { clauses }] \geq 1-2^{-\Omega(\log n \log \log n)}
$$

Lemma 21.9.21. For a sufficiently large constant $c^{\prime}>0$ and a constant $c>0$, for any random 3SAT instance which has $n$ variables and $m=c^{\prime} n$ clauses, suppose it is $(1-c) m$ satisfiable. Then with probability $1-2^{-\Omega(\log n \log \log n)}$, for all input strings $y$, among the unsatisfied clauses, each literal appears in $O(\log n)$ places.

Proof. This follows by Lemma 21.9.20.

Next, we show how to reduce the $O(\log n)$ to $O(1)$.

Lemma 21.9.22. For a sufficiently large constant $c$, for any random 3SAT instance that has $n$ variables and $m=c n$ clauses, for any constant $B \geq 1, b \in(0,1)$, with probability at least $1-\frac{9 m}{B b n}$, there exist at least $(1-b) m$ clauses such that each variable (in these $(1-b) m$ clauses) only appears in at most $B$ clauses (out of these $(1-b) m$ clauses).

Proof. For each $i \in[m]$, we use $z_{i}$ to denote the indicator variable such that it is 1 , if for each variable in the $i$ th clause, it appears in at most $a$ clauses. Let $B \in[1, \infty)$ denote a sufficiently large constant, which we will decide upon later.

For each variable $x$, the probability of it appearing in the $i$-th clause is $\frac{3}{n}$. Then we have

$$
\mathbb{E}[\# \text { clauses that contain } x]=\sum_{i=1}^{m} \mathbb{E}[i \text {-th clause contains } x]=\frac{3 m}{n}
$$

By Markov's inequality,

$$
\operatorname{Pr}[\# \text { clauses that contain } x \geq a] \leq \mathbb{E}[\# \text { clauses that contain } x] / B=\frac{3 m}{B n}
$$

By a union bound, we can compute $\mathbb{E}\left[z_{i}\right]$,

$$
\begin{aligned}
\mathbb{E}\left[z_{i}\right] & =\operatorname{Pr}\left[z_{i}=1\right] \\
& \geq 1-3 \operatorname{Pr}[\text { one variable in } i \text {-th clause appearing } \geq B \text { clauses }] \\
& \geq 1-\frac{9 m}{B n}
\end{aligned}
$$

Furthermore, we have

$$
\mathbb{E}[z]=\mathbb{E}\left[\sum_{i=1}^{m} z_{i}\right]=\sum_{i=1}^{m} \mathbb{E}\left[z_{i}\right] \geq\left(1-\frac{9 m}{B n}\right) m
$$

Note that $z \leq m$. Thus $\mathbb{E}[z] \leq m$. Let $b \in(0,1)$ denote a sufficiently small constant. We can show

$$
\begin{aligned}
\operatorname{Pr}[m-z \geq b m] & \leq \frac{\mathbb{E}[m-z]}{b m} \\
& =\frac{m-\mathbb{E}[z]}{b m} \\
& \leq \frac{m-\left(1-\frac{9 m}{B n}\right) m}{b m} \\
& =\frac{9 m}{B b n}
\end{aligned}
$$

This implies that with probability at least $1-\frac{9 m}{B b n}$, we have $m-z \leq b m$. Notice that in random-ETH, $m=c n$ for a constant $c$. Thus, by choosing a sufficiently large constant $B$ (which is a function of $c, b$ ), we can obtain arbitrarily large constant success probability.

### 21.9.4.3 Reduction

We reduce 3SAT to tensor rank by following the same construction in [Hås90]. To obtain a stronger hardness result, we use the property that each variable only appears in at most $B$ (some constant) clauses and that the cover number of an unsatisfiable 3SAT instance
is large. Note that both MAX-E3SAT(B) instances and random-ETH instances have that property. Also each MAX-E3SAT(B) is also a 3SAT instance. Thus if the reduction holds for 3SAT, it also holds for MAX-E3SAT(B), and similarly for random-ETH .

Recall the definition of 3SAT : 3SAT is the problem of given a Boolean formula of $n$ variables in CNF form with at most 3 variables in each of the $m$ clauses, is it possible to find a satisfying assignment to the formula? We say $x_{1}, x_{2}, \cdots, x_{n}$ are variables and $x_{1}, \bar{x}_{1}, x_{2}, \bar{x}_{2}, \cdots, x_{n}, \bar{x}_{n}$ are literals. We transform this to the problem of computing the rank of a tensor of size $n_{1} \times n_{2} \times n_{3}$ where $n_{1}=2+n+2 m, n_{2}=3 n$ and $n_{3}=3 n+m$. T has the following $n_{3}$ column-row faces, where each of the faces is an $m_{1} \times n_{2}$ matrix,

- $n$ variable matrices $V_{i} \in \mathbb{R}^{n_{1} \times n_{2}}$. It has a 1 in positions $(1,2 i-1)$ and $(2,2 i)$ while all other elements are 0.
- $n$ help matrices $S_{i} \in \mathbb{R}^{n_{1} \times n_{2}}$. It has a 1 position in $(1,2 n+i)$ and is 0 otherwise.
- $n$ help matrices $M_{i} \in \mathbb{R}^{n_{1} \times n_{2}}$. It has a 1 in positions $(1,2 i-1),(2+i, 2 i)$ and $(2+i, 2 n+i)$ and is 0 otherwise.
- $m$ clause matrices $C_{l} \in \mathbb{R}^{n_{1} \times n_{2}}$. Suppose the clause $c_{l}$ contains the literals $u_{l, 1}, u_{l, 2}$ and $u_{l, 3}$. For each $j \in[3], u_{l, j} \in\left\{x_{1}, x_{2}, \cdots, x_{n}, \bar{x}_{1}, \bar{x}_{2}, \cdots, \bar{x}_{n}\right\}$. Note that $x_{i}, \bar{x}_{i}$ are the literals of the 3SAT formula. We can also think of $x_{i}, \bar{x}_{i}$ as length $3 n$ vectors. Let $x_{i}$ denote the vector that has a 1 in position $2 i-1$, i.e., $x_{i}=e_{2 i-1}$. Let $\bar{x}_{i}$ denote the vector that has a 1 in positions $2 i-1$ and $2 i, \bar{x}_{i}=e_{2 i-1}+e_{2 i}$.
- Row 1 is the vector $u_{l, 1} \in \mathbb{R}^{3 n}$,
- Row $2+n+2 l-1$ is the vector $u_{l, 1}-u_{l, 2} \in \mathbb{R}^{3 n}$,


Figure 21.11: There are $3 n+m$ column-row faces, $V_{i}, \forall i \in[n], S_{i}, \forall i \in[n], M_{i}, \forall i \in[n]$, $C_{l}, \forall l \in[m]$. In face $C_{l}$, each $u_{l, j}$ is either $x_{i}$ or $\bar{x}_{i}$ where $x_{i}=e_{2 i-1}$ and $\bar{x}_{i}=e_{2 i-1}+e_{2 i}$.

- Row $2+n+2 l$ is the vector $u_{l, 1}-u_{l, 3} \in \mathbb{R}^{3 n}$.

First, we can obtain Lemma 21.9.23 which follows by Lemma 2 in [Hås90]. For completeness, we provide a proof.

Lemma 21.9.23. If the formula is satisfiable, then the constructed tensor has rank at most $4 n+2 m$.

Proof. We will construct $4 n+2 m$ rank-1 matrices $V_{i}^{(1)}, V_{i}^{(2)}, S_{i}^{(1)}, M_{i}^{(1)}, C_{l}^{(1)}$ and $C_{l}^{(2)}$. Then the goal is to show that for each matrix in the set

$$
\left\{V_{1}, V_{2}, \cdots, V_{n}, S_{1}, S_{2}, \cdots, S_{n}, M_{1}, M_{2}, \cdots, M_{n}, C_{1}, C_{2}, \cdots, C_{m}\right\}
$$

it can be written as a linear combination of these constructed matrices.

- Matrices $V_{i}^{(1)}$ and $V_{i}^{(2)} . V_{i}^{(1)}$ has the first row equal to $x_{i}$ iff $\alpha_{i}=1$ and otherwise $\bar{x}_{i}$.

All the other rows are 0 . We set $V_{i}^{(2)}=V_{i}-V_{i}^{(1)}$.

- Matrices $S_{i}^{(1)} . S_{i}^{(1)}=S_{i}$.
- Matrices $M_{i}^{(1)}$.

$$
M_{i}^{(1)}= \begin{cases}M_{i}-V_{i}^{(1)} & \text { if } \alpha_{i}=1 \\ M_{i}-V_{i}^{(1)}-S_{i} & \text { if } \alpha_{i}=0\end{cases}
$$

- Matrices $C_{l}^{(1)}$ and $C_{l}^{(2)}$. Let $x_{i}=\alpha_{i}$ be the assignment that makes the clause $c_{l}$ true. Then $C_{l}-V_{i}^{(1)}$ has rank 2, since either it has just two nonzero rows (in the case where $x_{i}$ is the first variable in the clause) or it has three nonzero rows of which two are equal. In both cases we just need two additional rank 1 matrices.


Figure 21.12: Two possibilities for $V_{i}^{(1)}, \forall i \in[n], V^{(2)}, \forall i \in[n], M_{i}^{(1)}, \forall i \in[n]$.

Once the 3SAT instance $S$ is unsatisfiable, then its cover number is at least 1. For each unsatisfiable 3SAT instance $S$ with cover number $p$, we can show that the constructed tensor has rank at most $4 n+2 m+O(p)$ and also has rank at least $4 n+2 m+\Omega(p)$. We first prove an upper bound,

Lemma 21.9.24. For a 3SAT instance $S$, let $y \in\{0,1\}$ denote a string such that $S(y)$ has
a set $L$ that contains unsatisfiable clauses. Let $p$ denote the smallest number of variables that cover all clauses in $L$. Then the constructed tensor $T$ has rank at most $4 n+2 m+p$.

Proof. Let $y$ denote a length- $n$ Boolean string $\left(\alpha_{1}, \alpha_{2}, \cdots, \alpha_{n}\right)$. Based on the assignment $y$, all the clauses of $S$ can be split into two sets: $L$ contains all the unsatisfied clauses and $\bar{L}$ contains all the satisfied clauses. We use set $P$ to denote a set of variables that covers all the clauses in set $L$. Let $p=|P|$. We will construct $4 n+2 m+p$ rank- 1 matrices $V_{i}^{(1)}, V_{i}^{(2)}$, $S_{i}^{(1)}, M_{i}^{(1)}, \forall i \in[n], C_{l}^{(1)}, C_{l}^{(2)}, \forall l \in[m]$, and $V_{j}^{(3)}, \forall j \in P$. Then the goal is to show that the $V_{i}, S_{i}, M_{i}$ and $C_{l}$ can be written as linear combinations of these constructed matrices.

- Matrices $V_{i}^{(1)}$ and $V_{i}^{(2)}$. $V_{i}^{(1)}$ has first row equal to $x_{i}$ iff $\alpha_{i}=1$ and otherwise $\bar{x}_{i}$. All the other rows are 0 . We set $V_{i}^{(2)}=V_{i}-V_{i}^{(1)}$.
- Matrices $V_{j}^{(3)}$. For each $j \in P, V_{j}^{(3)}$ has the first row equal to $x_{i}$ iff $\alpha_{i}=0$ and otherwise $\bar{x}_{i}$.
- Matrices $S_{i}^{(1)} . S_{i}^{(1)}=S_{i}$.
- Matrices $M_{i}^{(1)}$.

$$
M_{i}^{(1)}= \begin{cases}M_{i}-V_{i}^{(1)} & \text { if } \alpha_{i}=1 \\ M_{i}-V_{i}^{(1)}-S_{i} & \text { if } \alpha_{i}=0\end{cases}
$$

- Matrices $C_{l}^{(1)}$ and $C_{l}^{(2)}$.
- For each $l \notin L$, clause $c_{l}$ is satisfied according to assignment $y$. Let $x_{i}=\alpha_{i}$ be the assignment that makes the clause $c_{l}$ true. Then $C_{l}-V_{i}^{(1)}$ has rank 2, since either it has just two nonzero rows (in the case where $x_{i}$ is the first variables in
the clause) or it has three nonzero rows of which two are equal. In both cases we just need two additional rank 1 matrices.
- For each $l \in L$. It means clause $c_{l}$ is unsatisfied according to assignment $y$. Let $x_{j_{1}}=\alpha_{j_{1}}, x_{j_{2}}=\alpha_{j_{2}}, x_{j_{3}}=\alpha_{j_{3}}$ be an assignment that makes the clause $c_{l}$ false. In other words, one of $j_{1}, j_{2}, j_{3}$ must be $P$ according to the definition that $P$ covers L. Then matrix $C_{l}-V_{j_{1}}^{(3)}$ has rank 2, since either it has just two nonzero rows (in the case where $x_{j_{1}}$ is the first variables in the clause) or it has three nonzero rows of which two are equal. In both cases we just need two additional rank 1 matrices.

We finish the proof by taking the $P$ that has the smallest size.

Further, we have:

Corollary 21.9.25. For a 3SAT instance $S$, let $p$ denote the cover number of $S$, then the constructed tensor $T$ has rank at most $4 n+2 m+p$.

Proof. This follows by applying Lemma 21.9.24 to all the input strings and the definition of cover number (Definition 21.9.7).

We can split the tensor $T \in \mathbb{R}^{(2+n+3 m) \times 3 n \times(3 n+m)}$ into two sub-tensors, one is $T_{1} \in$ $\mathbb{R}^{2 \times 3 n \times(3 n+m)}$ (that contains the first two row-tube faces of $T$ and linear combination of the remaining $2 m$ row-tube faces of $T$ ), and the other is $T_{2} \in \mathbb{R}^{(n+2 m) \times 3 n \times(3 n+m)}$ (that contains the next $n+2 m$ row-tube faces of $T$ ). We first analyze the rank of $T_{1}$ and then analyze the rank of $T_{2}$.


Figure 21.13: $\widetilde{V}_{i}, \widetilde{S}_{i}, \widetilde{M}_{i}, \widetilde{C}_{l}$.

Claim 21.9.26. The rank of $T_{2}$ is $n+2 m$.

Proof. According to Figure 21.11, the nonzero rows are distributed in $n+m$ fully separated sub-tensors. It is obvious that the rank of each one of those $n$ sub-tensors is 1 , and the rank of each of those $m$ sub-tensors is 2 . Thus, overall, the rank $T_{2}$ is $n+2 m$.

To make sure $\operatorname{rank}(T)=\operatorname{rank}\left(T_{1}\right)+\operatorname{rank}\left(T_{2}\right)$, the $T_{1} \in \mathbb{R}^{2 \times 3 n \times(3 n+m)}$ can be described as the following $3 n+m$ column-row faces, and each of the faces is a $2 \times 3 n$ matrix.

- Matrices $\widetilde{V}_{i}, \forall i \in[n]$. The two rows are from the first two rows of $V_{i}$ in Figure 21.11, i.e., the first row is $e_{2 i-1}$ and the second row is $e_{2 i}$.
- Matrices $\widetilde{S}_{i}, \forall i \in[n]$. The two rows are from the first two rows of $S_{i}$ in Figure 21.11, i.e., the first row is $e_{2 n+i}$ and the second row is zero everywhere else.
- Matrices $\widetilde{M}_{i}, \forall i \in[n]$. The first row is $e_{2 i-1}+\beta_{i, 1}\left(e_{2 i}+e_{2 n+i}\right)$, while the second row is $\beta_{i, 2}\left(e_{2 i}+e_{2 n+i}\right)$.


Figure 21.14: There are $n+p$ matrices $A_{i} \in \mathbb{R}^{2 \times(2 n+p)}, \forall i \in[n+p]$ and $2 n+p$ matrices $B_{i} \in \mathbb{R}^{2 \times(n+p)}, \forall i \in[2 n+p]$. Tensor $A$ and tensor $B$ represet the same tensor, and for each $i \in[n+p], j \in[2], l \in[2 n+p],\left(A_{i}\right)_{j, l}=\left(B_{l}\right)_{j, i}$.

- Matrices $\widetilde{C}_{l}, \forall i \in[m]$. The first row is $\left(1+\gamma_{l, 1}+\gamma_{l, 2}\right) u_{l, 1}-\gamma_{l, 1} u_{l, 2}-\gamma_{l, 2} u_{l, 3}$ and the second is $\left(\gamma_{l, 3}+\gamma_{l, 4}\right) u_{l, 1}-\gamma_{l, 3} u_{l, 2}-\gamma_{l, 4} u_{l, 3}$,
where for each $i \in[3 n]$, we use vector $e_{i}$ to denote a length $3 n$ vector such that it only has a 1 in position $i$ and 0 otherwise. $\beta, \gamma$ are variables. The goal is to show a lower bound for,

$$
\operatorname{rank}_{\beta, \gamma}\left(T_{1}\right) .
$$

Lemma 21.9.27. Let $P$ denote the set $\left\{i \mid\right.$ the second row of matrix $\widetilde{M}_{i}$ is nonzero, $\forall i \in$ $[n]\}$. Then the rank of $T_{1}$ is at least $3 n+|P|$.

Proof. We define $p=|P|$. Without loss of generality, we assume that for each $i \in[p]$, the second row of matrix $\widetilde{M}_{i}$ is nonzero.

Notice that matrices $\widetilde{V}_{i}, \widetilde{S}_{i}, \widetilde{M}_{i}$ have size $2 \times 3 n$, but we only focus on the first $2 n+p$ columns. Thus, we have $n+p$ column-row faces (from the 3rd dimension) $A_{j} \in \mathbb{R}^{2 \times(2 n+p)}$,

- $A_{j}, 1 \leq j \leq n, A_{j}$ is the first $2 n+p$ columns of $\widetilde{V}_{j}-\sum_{i=1}^{n} \alpha_{i, j} \widetilde{S}_{i} \in \mathbb{R}^{2 \times 3 n}$, where $\alpha_{i, j}$ are some coefficients.
- $A_{n+j}, 1 \leq j \leq p, A_{j}$ is the first $2 n+p$ columns of $\widetilde{M}_{j}-\sum_{i=1}^{n} \alpha_{i, n+j} \widetilde{S}_{i} \in \mathbb{R}^{2 \times 3 n}$, where $\alpha_{i, j}$ are some coefficients.

Consider the first $2 n+p$ column-tube faces (from 2nd dimension), $B_{j}, \forall j \in[2 n+p]$, of $T_{1}$. Notice that these matrices have size $2 \times(n+p)$.

- $B_{2 i-1}, 1 \leq i \leq p$, it has a 1 in positions $(1, i)$ and $(1, n+i)$.
- $B_{2 i}, 1 \leq i \leq p$, it has $\beta_{i, 1}$ in position $(1, n+i), 1$ in position $(2, i)$ and $\beta_{i, 2}$ in position $(2, n+i)$.
- $B_{2 i-1}, p+1 \leq i \leq n$, it has 1 in position $(1, i)$.
- $B_{2 i}, p+1 \leq i \leq n$, it has 1 in position $(2, i)$.
- $B_{2 n+i}, 1 \leq i \leq p$, the first row is unknown, the second row has $\beta_{i, 2}$ in position in $(2, n+i)$.

It is obvious that the first $2 n$ matrices are linearly independent, thus the rank is at least $2 n$. We choose the first $2 n$ matrices as our basis. For $B_{2 n+1}$, we try to write it as a linear combination of the first $2 n$ matrices $\left\{B_{i}\right\}_{i \in[2 n]}$. Consider the second row of $B_{2 n+1}$. The first $n$ positions are all 0 . The matrices $B_{2 i}$ all have disjoint support for the second row of the
first $n$ columns. Thus, the matrices $B_{2 i}$ should not be used. Consider the second row of $B_{2 i-1}, \forall i \in[n]$. None of them has a nonzero value in position $n+1$. Thus $B_{2 n+1}$ cannot be written as a linear combination of of the first $2 n$ matrices. Thus, we can show for any $i \in[p], B_{2 n+i}$ cannot be written as a linear combination of matrices $\left\{B_{i}\right\}_{i \in[2 n]}$. Consider the $p$ matrices $\left\{B_{2 n+i}\right\}_{i \in[p]}$. Each of them has a different nonzero position in the second row. Thus these matrices are all linearly independent. Putting it all together, we know that the rank of matrices $\left\{B_{i}\right\}_{i \in[2 n+p]}$ is at least $2 n+p$.

Next, we consider another special case when $\beta_{i, 2}=0$, for all $i \in[n]$. If we subtract $\beta_{i, 1}$ times $\widetilde{S}_{i}$ from $\widetilde{M}_{i}$ and leave the other column-row faces (from the 3rd dimension) as they are, and we make all column-tube faces(from the 2 nd dimension) for $j>2 n$ identically 0 , then all other choices do not change the first $2 n$ column-tube faces (from the 2 nd dimension) and make some other column-tube faces (from the 2nd dimension) nonzero. Such a choice could clearly only increase the rank of $T$. Thus, we obtain,

$$
\operatorname{rank}(T)=2 n+2 m+\min \operatorname{rank}\left(T_{3}\right)
$$

where $T_{3}$ is a tensor of size $2 \times 2 n \times(2 n+m)$ given by the following column-row faces (from 3rd dimension) $A_{i}, \forall i \in[2 n+m]$ and each matrix has size $2 \times 2 n$ (shown in Figure 21.15).

- $A_{i}, i \in[n]$, the first $2 n$ columns of $\tilde{V}_{i}$.
- $A_{n+i}, i \in[n]$, the first $2 n$ columns of $\widetilde{M}_{i}$. The first row is $e_{2 i-1}+\beta_{i, 1} e_{2 i}$, and the second row is 0 .
- $A_{2 n+l}, l \in[m]$, the first $2 n$ columns of $\widetilde{C}_{l}$. The first row is $\left(1+\gamma_{l, 1}+\gamma_{l, 2}\right) u_{l, 1}-\gamma_{l, 1} u_{l, 2}-$ $\gamma_{l, 2} u_{l, 3}$, and the second row is $\left(\gamma_{l, 3}+\gamma_{l, 4}\right) u_{l, 1}-\gamma_{l, 3} u_{l, 2}-\gamma_{l, 4} u_{l, 3}$.

We can show

Lemma 21.9.28. Let $p$ denote the cover number of the 3SAT instance. $T_{3}$ has rank at least $2 n+\Omega(p)$.

Proof. First, we can show that all matrices $A_{n+i}-A_{i}$ and $A_{n+i}$ (for all $i \in[n]$ ) are in the expansion of tensor $T_{3}$. Thus, the rank of $T_{3}$ is at least $2 n$.

We need the following claim:
Claim 21.9.29. For any $l \in[m]$, if $A_{2 n+l}$ can be written as a linear combination of $\left\{A_{n+i}-\right.$ $\left.A_{i}\right\}_{i \in[n]}$ and $\left\{A_{n+i}\right\}_{i \in[n]}$, then the second row of $A_{2 n+l}$ is 0 , and the first row of one of the $A_{n+i}$ is $u_{i}$ where $u_{i}$ is one of the literals appearing in clause $c_{l}$.

Proof. We prove this for the second row first. For each $l \in[m]$, we consider the possibility of using all matrices $A_{n+i}-A_{i}$ and $A_{n+i}$ to express matrix $A_{2 n+l}$. If the second row of $A_{2 n+l}$ is nonzero, then it must have a nonzero entry in an odd position. But there is no nonzero in an odd position of the second row of any of matrices $A_{n+i}-A_{i}$ and $A_{n+i}$.

For the first row. It is obvious that the first row of $A_{2 n+l}$ must have at least one nonzero position, for any $\gamma_{l, 1}, \gamma_{l, 2}$. Let $u_{j}$ be a literal belonging to the variable $x_{i}$ which appears in the first row of $A_{2 n+l}$ with a nonzero coefficient. Since only $A_{n+i}$ of all the other $A_{n+s}, \forall s \in[n]$ matrices has nonzero elements in either of the positions $(1,2 i-1)$ or $(1,2 i)$, then $A_{n+i}$ must be used to cancel these elements. Thus, the first row of $A_{n+i}$ must be a multiple of $u_{j}$ and since the element in position $(1,2 i-1)$ of $A_{n+i}$ is 1 , this multiple must be 1 .


Figure 21.15: For any $i \in[n], \beta_{i, 1} \in \mathbb{R}$, for any $l \in[m], \gamma_{l, 1}, \gamma_{l, 2} \in \mathbb{R}$, for any $l \in[m]$, if the first literal of clause $l$ is $x_{j}$, then row vector $u_{l, 1}=e_{2 i-1} \in \mathbb{R}^{2 n}$; if the first literal of clause $l$ is $\bar{x}_{j}$, then row vector $u_{l, 1}=e_{2 i-1}+e_{2 i} \in \mathbb{R}^{2 n}$.

Note that matrices $A_{i}, \forall i \in[n]$ have the property that, for any matrix in $\left\{A_{n+1}, \cdots, A_{2 n+m}\right\}$, it cannot be written as the linear combination of matrices $A_{i}, \forall i \in[n]$. Let $\widetilde{A} \in \mathbb{R}^{(n+m) \times 2 n}$ denote a matrix that consists of the first rows of $\left\{A_{n+1}, \cdots, A_{2 n+m}\right\}$. According to the property of matrices $A_{i}, \forall i \in[n]$, and that the rank of a tensor is always greater than or equal to the rank of any sub-tensor, we know that

$$
\operatorname{rank}\left(T_{3}\right) \geq n+\min \operatorname{rank}(\widetilde{A})
$$

Claim 21.9.30. For a 3SAT instance $S$, for any input string $y \in\{0,1\}^{n}$, set $\beta_{*, 1}$ to be the entry-wise flipping of $y$, (I) if the clause $l$ is satisfied, then the $(n+l)$-th row of $\widetilde{A} \in \mathbb{R}^{(n+m) \times 2 n}$ can be written as a linear combination of the first $n$ rows of $\widetilde{A}$. (II) if the clause $l$ is unsatisfied, then the $(n+l)$-th row of $\widetilde{A}$ cannot be written as a linear combination of the first $n$ rows of $\widetilde{A}$.

Proof. Part (I), consider a clause $l$ which is satisfied with input string $y$. Then there must exist a variable $x_{i}$ belonging to clause $l$ (either literal $x_{i}$ or literal $\bar{x}_{i}$ ) and one of the following
holds: if $x_{i}$ belongs to clause $l$, then $\alpha_{i}=1$; if $\bar{x}_{i}$ belongs to clause $l$, then $\alpha_{i}=0$. Suppose clause $l$ contains literal $x_{i}$. The other case can be proved in a similar way. We consider the $(n+l)$-th row. One of the following assignments $(0,0),(-1,0),(0,-1)$ to $\gamma_{l, 1}, \gamma_{l, 1}$ is going to set the $(n+l)$-th row of $\widetilde{A}$ to be vector $e_{2 i-1}$. We consider the $i$-th row of $\widetilde{A}$. Since we set $\alpha_{i}=1$, then we set $\beta_{i, 1}=0$, it follows that the $i$-th row of $A$ becomes $e_{2 i-1}$. Therefore, the $(n+l)$-th row of $\widetilde{A}$ can be written as a linear combination of $\widetilde{A}$.

Part (II), consider a clause $l$ which is unsatisfied with input string $y$. Suppose that clause contains three literals $x_{i_{1}}, x_{i_{2}}, x_{i_{3}}$ (the other seven possibilities can be proved in a similar way). Then for input string $y$, we have $\alpha_{i_{1}}=0, \alpha_{i_{2}}=0$ and $\alpha_{i_{3}}=0$, otherwise this clause $l$ is satisfied. Consider $i_{1}$-th row of $\widetilde{A}$. It becomes $e_{2 i_{1}-1}+e_{2 i_{1}}$. Similarly for the $i_{2}$-th row and $i_{3}$-th row. Consider the $(n+l)$-th row. We can observe that all of positions $2 i_{1}, 2 i_{2}, 2 i_{3}$ must be 0 . Any linear combination formed by the $i_{1}, i_{2}, i_{3}$-th row of $\widetilde{A}$ must have one nonzero in one of positions $2 i_{1}, 2 i_{2}, 2 i_{3}$. However, if we consider the $(n+l)$-th row of $\widetilde{A}$, one of the positions $2 i_{1}, 2 i_{2}, 2 i_{3}$ must be 0 . Also, the remaining $n-3$ of the first $n$ rows of $\widetilde{A}$ also have 0 in positions $2 i_{1}, 2 i_{2}, 2 i_{3}$. Thus, we can show that the $(n+l)$-th row of $\widetilde{A}$ cannot be written as a linear combination of the first $n$ rows. Similarly, for the other seven cases.

Note that in order to make sure as many as possible rows in $n+1, \cdots, n+m$ can be written as linear combinations of the first $n$ rows of $\widetilde{A}$, the $\beta_{i, 1}$ should be set to either 0 or 1 . Also each possibility of input string $y$ is corresponding to a choice of $\beta_{i, 1}$. According to the above Claim 21.9.30, let $l_{0}$ denote the smallest number of unsatisfied clauses over the choices of all the $2^{n}$ input strings. Then over all choices of $\beta, \gamma$, there must exist at least $l_{0}$ rows of
$\widetilde{A}_{n+1}, \cdots \widetilde{A}_{n+m}$, such that each of those rows cannot be written as the linear combination of the first $n$ rows.

Claim 21.9.31. Let $\widetilde{A} \in \mathbb{R}^{(n+m) \times 2 n}$ denote a matrix that consists of the first rows of $A_{n+i}, \forall i \in[n]$ and $A_{n+l}, \forall l \in[m]$. Let $p$ denote the cover number of 3SAT instance. Then $\min \operatorname{rank}(\widetilde{A}) \geq n+\Omega(p)$.

Proof. For any choices of $\left\{\beta_{i, 1}\right\}_{i \in[n]}$, there must exist a set of rows out of the next $m$ rows such that, each of those rows cannot be written as a linear combination of the first $n$ rows. Let $L$ denote the set of those rows. Let $t$ denote the maximum size set of disjoint rows from $L$. Since those $t$ rows in $L$ all have disjoint support, they are always linearly independent. Thus the rank is at least $n+t$.

Note that each row corresponds to a unique clause and each clause corresponds to a unique row. We can just pick an arbitrary clause $l$ in $L$, then remove the clauses that are using the same literal as clause $l$ from $L$. Because each variable occurs in at most $B$ clauses, we only need to remove at most $3 B$ clauses from $L$. We repeat the procedure until there is no clause $L$. The corresponding rows of all the clauses we picked have disjoint supports, thus we can show a lower bound for $t$,

$$
t \geq|L| /(3 B) \geq l_{0} /(3 B) \geq p /(9 B) \gtrsim p
$$

where the second step follows by $|L| \geq l_{0}$, the third step follows $3 l_{0} \geq p$, and the last step follows by $B$ is some constant.

Thus, putting it all together, we complete the proof.


Figure 21.16: For any $i \in[n], \beta_{i, 1} \in \mathbb{R}$. For any $i \in[q], \beta_{i, 2} \in \mathbb{R}$. For any $l \in[m], \gamma_{l, 1}, \gamma_{l, 2} \in \mathbb{R}$. For any $l \in[m]$, if the first literal of clause $l$ is $x_{j}$, then row vector $u_{l, 1}=e_{2 i-1} \in \mathbb{R}^{2 n}$; if the first literal of clause $l$ is $\bar{x}_{j}$, then row vector $u_{l, 1}=e_{2 i-1}+e_{2 i} \in \mathbb{R}^{2 n}$.

Now, we consider a general case when there are $q$ different $i \in[n]$ satisfying that $\beta_{i, 2} \neq 0$. Similar to tensor $T_{3}$, we can obtain $T_{4}$ such that,

$$
\operatorname{rank}(T)=2 n+2 m+\min \operatorname{rank}\left(T_{4}\right)
$$

where $T_{4}$ is a tensor of size $2 \times 2 n \times(2 n+m)$ given by the following column-row faces (from 3rd dimension) $A_{i}, \forall i \in[2 n+m]$ and each matrix has size $2 \times 2 n$ (shown in Figure 21.16).

- $A_{i}, i \in[n]$, the first $2 n$ columns of $\widetilde{V}_{i}$.
- $A_{n+i}, i \in[q]$, the first $2 n$ columns of $\widetilde{M}_{i}$. The first row is $e_{2 i-1}+\beta_{i, 1} e_{2 i}$, and the second row is $\beta_{i, 2} e_{2 i}$.
- $A_{n+i}, i \in\{q+1, \cdots, n\}$, the first $2 n$ columns of $\widetilde{M}_{i}$. The first row is $e_{2 i-1}+\beta_{i, 1} e_{2 i}$, and the second row is 0 .
- $A_{2 n+l}, l \in[m]$, the first $2 n$ columns of $\widetilde{C}_{l}$. The first row is $\left(1+\gamma_{l, 1}+\gamma_{l, 2}\right) u_{l, 1}-\gamma_{l, 1} u_{l, 2}-$ $\gamma_{l, 2} u_{l, 3}$, and the second row is $\left(\gamma_{l, 3}+\gamma_{l, 4}\right) u_{l, 1}-\gamma_{l, 3} u_{l, 2}-\gamma_{l, 4} u_{l, 3}$.

Note that modifying $q$ entries(from Figure 21.15 to Figure 21.16) of a tensor can only decrease the rank by $q$, thus we obtain

Lemma 21.9.32. Let $q$ denote the number of $i$ such that $\beta_{i, 2} \neq 0$, and let $p$ denote the cover number of the 3SAT instance. Then $T_{4}$ has rank at least $2 n+\Omega(p)-q$.

Combining the two perspectives we have

Lemma 21.9.33. Let $p$ denote the cover number of an unsatisfiable 3SAT instance. Then the tensor has rank at least $4 n+2 m+\Omega(p)$.

Proof. Let $q$ denote the $q$ in Figure 21.16. From one perspective, we know that the tensor has rank at least $4 n+2 m+\Omega(p)-q$. From another perspective, we know that the tensor has rank at least $4 n+2 m+q$. Combining them together, we obtain the rank is at least $4 n+2 m+\Omega(p) / 2$, which is still $4 n+2 m+\Omega(p)$.

Theorem 21.9.34. Unless ETH fails, there is a $\delta>0$ and an absolute constant $c_{0}>1$ such that the following holds. For the problem of deciding if the rank of a $q$-th order tensor, $q \geq 3$, with each dimension $n$, is at most $k$ or at least $c_{0} k$, there is no $2^{\delta k^{1-o(1)}}$ time algorithm.

Proof. The reduction can be split into three parts. ${ }^{13}$ The first part reduces the MAX3SAT problem to the MAX-E3SAT problem by [MR10]. For each MAX-3SAT instance with

[^76]size $n$, the corresponding MAX-E3SAT instance has size $n^{1+o(1)}$. The second part is by reducing the MAX-E3SAT problem to MAX-E3SAT(B) by [Tre01]. For each MAX-E3SAT instance with size $n$, the corresponding MAX-E3SAT(B) instance has size $\Theta(n)$ when $B$ is a constant. The third part is by reducing the MAX-E3SAT(B) problem to the tensor problem. Combining Theorem 21.9.5, Lemma 21.9.18 with this reduction, we complete the proof.

Theorem 21.9.35. Unless random-ETH fails, there is an absolute constant $c_{0}>1$ for which any deterministic algorithm for deciding if the rank of a q-th order tensor is at most $k$ or at least $c_{0} k$, requires $2^{\Omega(k)}$ time.

Proof. This follows by combining the reduction with random-ETH and Lemma 21.9.22.

Note that, if $\mathbf{B P P}=\mathbf{P}$ then it also holds for randomized algorithms which succeed with probability $2 / 3$.

Indeed, we know that any deterministic algorithm requires $2^{\Omega(n)}$ running time on tensors that have size $n \times n \times n$. Let $g(n)$ denote a fixed function of $n$, and $g(n)=o(n)$. We change the original tensor from size $n \times n \times n$ to $2^{g(n)} \times 2^{g(n)} \times 2^{g(n)}$ by adding zero entries. Then the number of entries in the new tensor is $2^{3 g(n)}$ and the deterministic algorithm still requires $2^{\Omega(n)}$ running time on this new tensor. Assume there is a randomized algorithm that runs in $2^{c g(n)}$ time, for some constant $c>3$. Then considering the size of this new tensor, the deterministic algorithm is a super-polynomial time algorithm, but the randomized algorithm is a polynomial time algorithm. Thus, by assuming $\mathbf{B P P}=\mathbf{P}$, we can rule out randomized algorithms, which means Theorem 21.9.35 also holds for randomized algorithms which succeed with probability $2 / 3$.

We provide some some motivation for the $\mathbf{B P P}=\mathbf{P}$ assumption: this is a standard conjecture in complexity theory, as it is implied by the existence of strong pseudorandom generators or if any problem in deterministic exponential time has exponential size circuits [IW97].

### 21.10 Extension to Other Tensor Ranks

The tensor rank studied in the previous sections is also called the CP rank or canonical rank. The tensor rank can be thought of as a direct extension of the matrix rank. We would like to point out that there are other definitions of tensor rank, e.g., the tucker rank and train rank. In this section we explain how to extend our proofs to other notions of tensor rank. Section 21.10.1 provides the extension to tucker rank, and Section 21.10.2 provides the extension to train rank.

### 21.10.1 Tensor Tucker rank

Tensor Tucker rank has been studied in a number of works [KC07, PC08, MH09, ZW13, YC14]. We provide the formal definition here:

### 21.10.1.1 Definitions

Definition 21.10.1 (Tucker rank). Given a third order tensor $A \in \mathbb{R}^{n \times n \times n}$, we say $A$ has tucker rank $k$ if $k$ is the smallest integer such that there exist three matrices $U, V, W \in \mathbb{R}^{n \times k}$ and a (small) tensor $C \in \mathbb{R}^{k \times k \times k}$ satisfying

$$
A_{i, j, l}=\sum_{i^{\prime}=1}^{k} \sum_{j^{\prime}=1}^{k} \sum_{l^{\prime}=1}^{k} C_{i^{\prime}, j^{\prime}, l^{\prime}} U_{i, i^{\prime}} V_{j, j^{\prime}} W_{l, l^{\prime}}, \forall i, j, l \in[n] \times[n] \times[n],
$$

or equivalently,

$$
A=C(U, V, W)
$$

### 21.10.1.2 Algorithm

Theorem 21.10.1. Given a third order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$ and $\epsilon \in(0,1)$, there exists an algorithm which takes $O(\operatorname{nnz}(A))+n \operatorname{poly}(k, 1 / \epsilon)+2^{O\left(k^{2} / \epsilon+k^{3}\right)}$ time and outputs three matrices $U, V, W \in \mathbb{R}^{n \times k}$, and a tensor $C \in \mathbb{R}^{k \times k \times k}$ for which

$$
\|C(U, V, W)-A\|_{F}^{2} \leq(1+\epsilon) \min _{\text {tucker rank }-k A_{k}}\left\|A_{k}-A\right\|_{F}^{2}
$$

holds with probability 9/10.

Proof. We define OPT to be

$$
\mathrm{OPT}=\min _{\text {tucker rank }-k}\left\|A^{\prime}-A\right\|_{F}^{2} .
$$

```
Algorithm 21.33 Frobenius Norm Low (Tucker) Rank Approximation
    procedure FLowTuckerRankApprox \((A, n, k, \epsilon) \quad \triangleright\) Theorem 21.10.1
        \(s_{1} \leftarrow s_{2} \leftarrow s_{3} \leftarrow O(k / \epsilon)\).
        \(t_{1} \leftarrow t_{2} \leftarrow t_{3} \leftarrow \operatorname{poly}(k, 1 / \epsilon)\).
        Choose sketching matrices \(S_{1} \in \mathbb{R}^{n^{2} \times s_{1}}, S_{2} \in \mathbb{R}^{n^{2} \times s_{2}}, S_{3} \in \mathbb{R}^{n^{2} \times s_{3}}\). \(\quad\) Definition
    21.3.12
        Choose sketching matrices \(T_{1} \in \mathbb{R}^{t_{1} \times n}, T_{2} \in \mathbb{R}^{t_{2} \times n}, T_{3} \in \mathbb{R}^{t_{3} \times n}\).
        Compute \(A_{i} S_{i}, \forall i \in[3]\).
        Compute \(T_{i} A_{i} S_{i}, \forall i \in[3]\).
        Compute \(B \leftarrow A\left(T_{1}, T_{2}, T_{3}\right)\).
        Create variables for \(X_{i} \in \mathbb{R}^{s_{i} \times k}, \forall i \in[3]\).
        Create variables for \(C \in \mathbb{R}^{k \times k \times k}\).
        Run a polynomial system verifier for \(\left\|C\left(\left(Y_{1} X_{1}\right),\left(Y_{2} X_{2}\right),\left(Y_{3} X_{3}\right)\right)-B\right\|_{F}^{2}\).
        return \(C, A_{1} S_{1} X_{1}, A_{2} S_{2} X_{2}\), and \(A_{3} S_{3} X_{3}\).
    end procedure
```

Suppose the optimal $A_{k}=C^{*}\left(U^{*}, V^{*}, W^{*}\right)$. We fix $C^{*} \in \mathbb{R}^{k \times k \times k}, V^{*} \in \mathbb{R}^{n \times k}$ and $W^{*} \in \mathbb{R}^{n \times k}$. We use $V_{1}^{*}, V_{2}^{*}, \cdots, V_{k}^{*}$ to denote the columns of $V^{*}$ and $W_{1}^{*}, W_{2}^{*}, \cdots, W_{k}^{*}$ to denote the columns of $W^{*}$.

We consider the following optimization problem,

$$
\min _{U_{1}, \cdots, U_{k} \in \mathbb{R}^{n}}\left\|C^{*}\left(U, V^{*}, W^{*}\right)-A\right\|_{F}^{2},
$$

which is equivalent to

$$
\min _{U_{1}, \cdots, U_{k} \in \mathbb{R}^{n}}\left\|U \cdot C^{*}\left(I, V^{*}, W^{*}\right)-A\right\|_{F}^{2}
$$

because $C^{*}\left(U, V^{*}, W^{*}\right)=U \cdot C^{*}\left(I, V^{*}, W^{*}\right)$ according to Definition 21.2.6.
Recall that $C^{*}\left(I, V^{*}, W^{*}\right)$ denotes a $k \times n \times n$ tensor. Let $\left(C^{*}\left(I, V^{*}, W^{*}\right)\right)_{1}$ denote the matrix obtained by flattening $C^{*}\left(I, V^{*}, W^{*}\right)$ along the first dimension. We use matrix $Z_{1}$
to denote $\left(C^{*}\left(I, V^{*}, W^{*}\right)\right)_{1} \in \mathbb{R}^{k \times n^{2}}$. Then we can obtain the following equivalent objective function,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1}-A_{1}\right\|_{F}^{2}
$$

Notice that $\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1}-A_{1}\right\|_{F}^{2}=$ OPT, since $A_{k}=U^{*} Z_{1}$.
Let $S_{1}^{\top} \in \mathbb{R}^{s_{1} \times n^{2}}$ be the sketching matrix defined in Definition 21.3.12, where $s_{1}=$ $O(k / \epsilon)$. We obtain the following optimization problem,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1} S_{1}-A_{1} S_{1}\right\|_{F}^{2}
$$

Let $\widehat{U} \in \mathbb{R}^{n \times k}$ denote the optimal solution to the above optimization problem. Then $\widehat{U}=$ $A_{1} S_{1}\left(Z_{1} S_{1}\right)^{\dagger}$. By Lemma 21.3.10 and Theorem 21.3.11, we have

$$
\left\|\widehat{U} Z_{1}-A_{1}\right\|_{F}^{2} \leq(1+\epsilon) \min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1}-A_{1}\right\|_{F}^{2}=(1+\epsilon) \mathrm{OPT},
$$

which implies

$$
\left\|C^{*}\left(\widehat{U}, V^{*}, W^{*}\right)-A\right\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}
$$

To write down $\widehat{U}_{1}, \cdots, \widehat{U}_{k}$, we use the given matrix $A_{1}$, and we create $s_{1} \times k$ variables for matrix $\left(Z_{1} S_{1}\right)^{\dagger}$.

As our second step, we fix $\widehat{U} \in \mathbb{R}^{n \times k}$ and $W^{*} \in \mathbb{R}^{n \times k}$, and we convert tensor $A$ into matrix $A_{2}$. Let matrix $Z_{2}$ denote $\left(C^{*}\left(\widehat{U}, I, W^{*}\right)\right)_{2} \in \mathbb{R}^{k \times n^{2}}$. We consider the following objective function,

$$
\min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2}-A_{2}\right\|_{F}^{2}
$$

for which the optimal cost is at most $(1+\epsilon)$ OPT.
Let $S_{2}^{\top} \in \mathbb{R}^{s_{2} \times n^{2}}$ be a sketching matrix defined in Definition 21.3.12, where $s_{2}=$ $O(k / \epsilon)$. We sketch $S_{2}$ on the right of the objective function to obtain a new objective function,

$$
\min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2} S_{2}-A_{2} S_{2}\right\|_{F}^{2}
$$

Let $\widehat{V} \in \mathbb{R}^{n \times k}$ denote the optimal solution to the above problem. Then $\widehat{V}=A_{2} S_{2}\left(Z_{2} S_{2}\right)^{\dagger}$. By Lemma 21.3.10 and Theorem 21.3.11, we have,

$$
\left\|\widehat{V} Z_{2}-A_{2}\right\|_{F}^{2} \leq(1+\epsilon) \min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2}-A_{2}\right\|_{F}^{2} \leq(1+\epsilon)^{2} \mathrm{OPT}
$$

which implies

$$
\left\|C^{*}\left(\widehat{U}, \widehat{V}, W^{*}\right)-A\right\|_{F}^{2} \leq(1+\epsilon)^{2} \mathrm{OPT}
$$

To write down $\widehat{V}_{1}, \cdots, \widehat{V}_{k}$, we need to use the given matrix $A_{2} \in \mathbb{R}^{n^{2} \times n}$, and we need to create $s_{2} \times k$ variables for matrix $\left(Z_{2} S_{2}\right)^{\dagger}$.

As our third step, we fix the matrices $\widehat{U} \in \mathbb{R}^{n \times k}$ and $\widehat{V} \in \mathbb{R}^{n \times k}$. We convert tensor $A \in \mathbb{R}^{n \times n \times n}$ into matrix $A_{3} \in \mathbb{R}^{n^{2} \times n}$. Let matrix $Z_{3}$ denote $\left(C^{*}(\widehat{U}, \widehat{V}, I)\right)_{3} \in \mathbb{R}^{k \times n^{2}}$. We consider the following objective function,

$$
\min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3}-A_{3}\right\|_{F}^{2},
$$

which has optimal cost at most $(1+\epsilon)^{2}$ OPT.
Let $S_{3}^{\top} \in \mathbb{R}^{s_{3} \times n^{2}}$ be a sketching matrix defined in Definition 21.3.12, where $s_{3}=$ $O(k / \epsilon)$. We sketch $S_{3}$ on the right of the objective function to obtain a new objective
function,

$$
\min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3} S_{3}-A_{3} S_{3}\right\|_{F}^{2}
$$

Let $\widehat{W} \in \mathbb{R}^{n \times k}$ denote the optimal solution of the above problem. Then $\widehat{W}=A_{3} S_{3}\left(Z_{3} S_{3}\right)^{\dagger}$. By Lemma 21.3.10 and Theorem 21.3.11, we have,

$$
\left\|\widehat{W} Z_{3}-A_{3}\right\|_{F}^{2} \leq(1+\epsilon) \min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3}-A_{3}\right\|_{F}^{2} \leq(1+\epsilon)^{3} \mathrm{OPT}
$$

Thus, we have

$$
\min _{X_{1}, X_{2}, X_{3}}\left\|C^{*}\left(\left(A_{1} S_{1} X_{1}\right),\left(A_{2} S_{2} X_{2}\right),\left(A_{3} S_{3} X_{3}\right)\right)-A\right\|_{F}^{2} \leq(1+\epsilon)^{3} \mathrm{OPT} .
$$

Let $V_{1}=A_{1} S_{1}, V_{2}=A_{2} S_{2}$, and $V_{3}=A_{3} S_{3}$. We then apply Lemma 21.4.3, and we obtain $\widehat{V}_{1}, \widehat{V}_{2}, \widehat{V}_{3}, B$. We then apply Theorem 21.4.44. Correctness follows by rescaling $\epsilon$ by a constant factor.

Running time. Due to Definition 21.3.12, the running time of line 7 (Algorithm 21.33) is $O(\operatorname{nnz}(A))+n \operatorname{poly}(k, 1 / \epsilon)$. Due to Lemma 21.4.3, line 7 and 8 can be executed in $n n z(A)+$ $n$ poly $(k, 1 / \epsilon)$ time. The running time of line 11 is given by Theorem 21.4.44. (For simplicity, we ignore the bit complexity in the running time.)

### 21.10.2 Tensor Train rank

### 21.10.2.1 Definitions

The tensor train rank has been studied in several works [Ose11, OTZ11, ZWZ16, PTBD16]. We provide the formal definition here.

Definition 21.10.2 (Tensor Train rank). Given a third order tensor $A \in \mathbb{R}^{n \times n \times n}$, we say $A$ has train rank $k$ if $k$ is the smallest integer such that there exist three tensors $U \in \mathbb{R}^{1 \times n \times k}$, $V \in \mathbb{R}^{k \times n \times k}, W \in \mathbb{R}^{k \times n \times 1}$ satisfying:

$$
A_{i, j, l}=\sum_{i_{1}=1}^{1} \sum_{i_{2}=1}^{k} \sum_{i_{3}=1}^{k} \sum_{i_{4}=1}^{1} U_{i_{1}, i_{2}} V_{i_{2}, j, i_{3}} W_{i_{3}, l, i_{4}}, \forall i, j, l \in[n] \times[n] \times[n],
$$

or equivalently,

$$
A_{i, j, l}=\sum_{i_{2}=1}^{k} \sum_{i_{3}=1}^{k}\left(U_{2}\right)_{i, i_{2}}\left(V_{2}\right)_{j, i_{2}+k\left(i_{3}-1\right)}\left(W_{2}\right)_{l, i_{3}},
$$

where $V_{2} \in \mathbb{R}^{n \times k^{2}}$ denotes the matrix obtained by flattening the tensor $U$ along the second dimension, and $\left(V_{2}\right)_{i, i_{1}+k\left(i_{2}-1\right)}$ denotes the entry in the $i$-th row and $i_{1}+k\left(i_{2}-1\right)$-th column of $V_{2}$. We similarly define $U_{2}, W_{2} \in \mathbb{R}^{n \times k}$.

### 21.10.2.2 Algorithm

Theorem 21.10.2. Given a third order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1, \epsilon \in(0,1)$, there exists an algorithm which takes $O(\mathrm{nnz}(A))+n \operatorname{poly}(k, 1 / \epsilon)+2^{O\left(k^{4} / \epsilon\right)}$ time and outputs three tensors $U \in \mathbb{R}^{1 \times n \times k}, V \in \mathbb{R}^{k \times n \times k}, W \in \mathbb{R}^{k \times n \times 1}$ such that

$$
\left\|\sum_{i=1}^{k} \sum_{j=1}^{k}\left(U_{2}\right)_{i} \otimes\left(V_{2}\right)_{i+k(j-1)} \otimes\left(W_{2}\right)_{j}-A\right\|_{F}^{2} \leq(1+\epsilon)_{\operatorname{train} \operatorname{rank}-k A_{k}}\left\|A_{k}-A\right\|_{F}^{2}
$$

holds with probability 9/10.

```
Algorithm 21.34 Frobenius Norm Low (Train) rank Approximation
    procedure FLowTrainRankApprox \((A, n, k, \epsilon) \quad \triangleright\) Theorem 21.10.2
        \(s_{1} \leftarrow s_{3} \leftarrow O(k / \epsilon)\).
        \(s_{2} \leftarrow O\left(k^{2} / \epsilon\right)\).
        \(t_{1} \leftarrow t_{2} \leftarrow t_{3} \leftarrow \operatorname{poly}(k, 1 / \epsilon)\).
        Choose sketching matrices \(S_{1} \in \mathbb{R}^{n^{2} \times s_{1}}, S_{2} \in \mathbb{R}^{n^{2} \times s_{2}}, S_{3} \in \mathbb{R}^{n^{2} \times s_{3}}\). \(\quad\) Definition
    21.3.12
        Choose sketching matrices \(T_{1} \in \mathbb{R}^{t_{1} \times n}, T_{2} \in \mathbb{R}^{t_{2} \times n}, T_{3} \in \mathbb{R}^{t_{3} \times n}\).
        Compute \(A_{i} S_{i}, \forall i \in[3]\).
        Compute \(T_{i} A_{i} S_{i}, \forall i \in[3]\).
        Compute \(B \leftarrow A\left(T_{1}, T_{2}, T_{3}\right)\).
        Create variables for \(X_{1} \in \mathbb{R}^{s_{1} \times k}\).
        Create variables for \(X_{3} \in \mathbb{R}^{s_{3} \times k}\).
        Create variables for \(X_{2} \in \mathbb{R}^{s_{2} \times k^{2}}\).
        Create variables for \(C \in \mathbb{R}^{k \times k \times k}\).
        Run polynomial system verifier for \(\| \sum_{i_{2}=1}^{k} \sum_{i_{3}=1}^{k}\left(Y_{1} X_{1}\right)_{i_{2}}\left(Y_{2} X_{2}\right)_{i_{2}+k\left(i_{3}-1\right)}\left(Y_{3} X_{3}\right)_{i_{3}}-\)
    \(B \|_{F}^{2}\).
        return \(A_{1} S_{1} X_{1}, A_{2} S_{2} X_{2}\), and \(A_{3} S_{3} X_{3}\).
    end procedure
```

Proof. We define OPT as

$$
\mathrm{OPT}=\min _{\text {train }}\left\|A^{\prime}-A\right\|_{F}^{2}
$$

Suppose the optimal

$$
A_{k}=\sum_{i=1}^{k} \sum_{j=1}^{k} U_{i}^{*} \otimes V_{i+k(j-1)}^{*} \otimes W_{j}^{*}
$$

We fix $V^{*} \in \mathbb{R}^{n \times k^{2}}$ and $W^{*} \in \mathbb{R}^{n \times k}$. We use $V_{1}^{*}, V_{2}^{*}, \cdots, V_{k^{2}}^{*}$ to denote the columns of $V^{*}$, and $W_{1}^{*}, W_{2}^{*}, \cdots, W_{k}^{*}$ to denote the columns of $W^{*}$.

We consider the following optimization problem,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|\sum_{i=1}^{k} \sum_{j=1}^{k} U_{i} \otimes V_{i+k(j-1)}^{*} \otimes W_{j}^{*}-A\right\|_{F}^{2}
$$

which is equivalent to

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U \cdot\left[\begin{array}{c}
\sum_{j=1}^{k} V_{1+k(j-1)}^{*} \otimes W_{j}^{*} \\
\sum_{j=1}^{k} V_{2+k(j-1)}^{*} \otimes W_{j}^{*} \\
\cdots \\
\sum_{j=1}^{k} V_{k+k(j-1)}^{*} \otimes W_{j}^{*}
\end{array}\right]-A\right\|_{F}^{2}
$$

Let $A_{1} \in \mathbb{R}^{n \times n^{2}}$ denote the matrix obtained by flattening the tensor $A$ along the first dimension. We use matrix $Z_{1} \in \mathbb{R}^{k \times n^{2}}$ to denote

$$
\left[\begin{array}{c}
\sum_{j=1}^{k} \operatorname{vec}\left(V_{1+k(j-1)}^{*} \otimes W_{j}^{*}\right) \\
\sum_{j=1}^{k} \operatorname{vec}\left(V_{2+k(j-1)}^{*} \otimes W_{j}^{*}\right) \\
\cdots \\
\sum_{j=1}^{k} \operatorname{vec}\left(V_{k+k(j-1)}^{*} \otimes W_{j}^{*}\right)
\end{array}\right] .
$$

Then we can obtain the following equivalent objective function,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1}-A_{1}\right\|_{F}^{2} .
$$

Notice that $\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1}-A_{1}\right\|_{F}^{2}=$ OPT, since $A_{k}=U^{*} Z_{1}$.
Let $S_{1}^{\top} \in \mathbb{R}^{s_{1} \times n^{2}}$ be a sketching matrix defined in Definition 21.3.12, where $s_{1}=$ $O(k / \epsilon)$. We obtain the following optimization problem,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1} S_{1}-A_{1} S_{1}\right\|_{F}^{2}
$$

Let $\widehat{U} \in \mathbb{R}^{n \times k}$ denote the optimal solution to the above optimization problem. Then $\widehat{U}=$ $A_{1} S_{1}\left(Z_{1} S_{1}\right)^{\dagger}$. By Lemma 21.3.10 and Theorem 21.3.11, we have

$$
\left\|\widehat{U} Z_{1}-A_{1}\right\|_{F}^{2} \leq(1+\epsilon) \min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1}-A_{1}\right\|_{F}^{2}=(1+\epsilon) \mathrm{OPT},
$$

which implies

$$
\left\|\sum_{i=1}^{k} \sum_{j=1}^{k} \widehat{U}_{i} \otimes V_{i+k(j-1)}^{*} \otimes W_{j}^{*}-A\right\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}
$$

To write down $\widehat{U}_{1}, \cdots, \widehat{U}_{k}$, we use the given matrix $A_{1}$, and we create $s_{1} \times k$ variables for matrix $\left(Z_{1} S_{1}\right)^{\dagger}$.

As our second step, we fix $\widehat{U} \in \mathbb{R}^{n \times k}$ and $W^{*} \in \mathbb{R}^{n \times k}$, and we convert the tensor $A$ into matrix $A_{2}$. Let matrix $Z_{2} \in \mathbb{R}^{k^{2} \times n^{2}}$ denote the matrix where the $(i, j)$-th row is the vectorization of $\widehat{U}_{i} \otimes W_{j}^{*}$. We consider the following objective function,

$$
\min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2}-A_{2}\right\|_{F}^{2}
$$

for which the optimal cost is at most $(1+\epsilon)$ OPT.
Let $S_{2}^{\top} \in \mathbb{R}^{s_{2} \times n^{2}}$ be a sketching matrix defined in Definition 21.3.12, where $s_{2}=$ $O\left(k^{2} / \epsilon\right)$. We sketch $S_{2}$ on the right of the objective function to obtain the new objective function,

$$
\min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2} S_{2}-A_{2} S_{2}\right\|_{F}^{2}
$$

Let $\widehat{V} \in \mathbb{R}^{n \times k}$ denote the optimal solution of the above problem. Then $\widehat{V}=A_{2} S_{2}\left(Z_{2} S_{2}\right)^{\dagger}$. By Lemma 21.3.10 and Theorem 21.3.11, we have,

$$
\left\|\widehat{V} Z_{2}-A_{2}\right\|_{F}^{2} \leq(1+\epsilon) \min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2}-A_{2}\right\|_{F}^{2} \leq(1+\epsilon)^{2} \mathrm{OPT}
$$

which implies

$$
\left\|\sum_{i=1}^{k} \sum_{j=1}^{k} \widehat{U}_{i} \otimes \widehat{V}_{i+k(j-1)} \otimes W^{*}-A\right\|_{F}^{2} \leq(1+\epsilon)^{2} \mathrm{OPT}
$$

To write down $\widehat{V}_{1}, \cdots, \widehat{V}_{k}$, we need to use the given matrix $A_{2} \in \mathbb{R}^{n^{2} \times n}$, and we need to create $s_{2} \times k$ variables for matrix $\left(Z_{2} S_{2}\right)^{\dagger}$.

As our third step, we fix the matrices $\widehat{U} \in \mathbb{R}^{n \times k}$ and $\widehat{V} \in \mathbb{R}^{n \times k}$. We convert tensor $A \in \mathbb{R}^{n \times n \times n}$ into matrix $A_{3} \in \mathbb{R}^{n^{2} \times n}$. Let matrix $Z_{3} \in \mathbb{R}^{k \times n^{2}}$ denote

$$
\left[\begin{array}{c}
\sum_{i=1}^{k} \operatorname{vec}\left(\widehat{U}_{i} \otimes \widehat{V}_{i+k \cdot 0}\right) \\
\sum_{i=1}^{k} \operatorname{vec}\left(\widehat{U}_{i} \otimes \widehat{V}_{i+k \cdot 1}\right) \\
\cdots \\
\sum_{i=1}^{k} \operatorname{vec}\left(\widehat{U}_{i} \otimes \widehat{V}_{i+k \cdot(k-1)}\right)
\end{array}\right]
$$

We consider the following objective function,

$$
\min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3}-A_{3}\right\|_{F}^{2},
$$

which has optimal cost at most $(1+\epsilon)^{2}$ OPT.
Let $S_{3}^{\top} \in \mathbb{R}^{s_{3} \times n^{2}}$ be a sketching matrix defined in Definition 21.3.12, where $s_{3}=$ $O(k / \epsilon)$. We sketch $S_{3}$ on the right of the objective function to obtain a new objective function,

$$
\min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3} S_{3}-A_{3} S_{3}\right\|_{F}^{2} .
$$

Let $\widehat{W} \in \mathbb{R}^{n \times k}$ denote the optimal solution of the above problem. Then $\widehat{W}=A_{3} S_{3}\left(Z_{3} S_{3}\right)^{\dagger}$. By Lemma 21.3.10 and Theorem 21.3.11, we have,

$$
\left\|\widehat{W} Z_{3}-A_{3}\right\|_{F}^{2} \leq(1+\epsilon) \min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3}-A_{3}\right\|_{F}^{2} \leq(1+\epsilon)^{3} \mathrm{OPT}
$$

Thus, we have

$$
\min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k} \sum_{j=1}^{k}\left(A_{1} S_{1} X_{1}\right)_{i} \otimes\left(A_{2} S_{2} X_{2}\right)_{i+k(j-1)} \otimes\left(A_{3} S_{3} X_{3}\right)_{j}-A\right\|_{F}^{2} \leq(1+\epsilon)^{3} \mathrm{OPT} .
$$

Let $V_{1}=A_{1} S_{1}, V_{2}=A_{2} S_{2}$, and $V_{3}=A_{3} S_{3}$. We then apply Lemma 21.4.3, and we obtain $\widehat{V}_{1}, \widehat{V}_{2}, \widehat{V}_{3}, B$. We then apply Theorem 21.4.44. Correctness follows by rescaling $\epsilon$ by a constant factor.

Running time. Due to Definition 21.3.12, the running time of line 7 (Algorithm 21.34) is $O(\mathrm{nnz}(A))+n \operatorname{poly}(k, 1 / \epsilon)$. Due to Lemma 21.4.3, lines 8 and 9 can be executed in $\mathrm{nnz}(A)+$ $n$ poly $(k, 1 / \epsilon)$ time. The running time of $2^{O\left(k^{4} / \epsilon\right)}$ comes from running Theorem 21.4.44 (For simplicity, we ignore the bit complexity in the running time.)

### 21.11 Entry-wise $\ell_{p}$ Norm for Arbitrary Tensors, $1<p<2$

There is a long line of research dealing with $\ell_{p}$ norm-related problems $\left[\mathrm{DDH}^{+} 09\right.$, MM13, $\left.\mathrm{CDMI}^{+} 13, \mathrm{CP} 15, \mathrm{BCKY16}, \mathrm{YCRM} 16, \mathrm{BBC}^{+} 17\right]$.

In this section, we provide several different algorithms for tensor $\ell_{p}$-low rank approximation. Section 21.11 .1 formally states the $\ell_{p}$ version of Theorem C. 1 in [SWZ17]. Section 21.11.2 presents several existence results. Section 21.11 .3 describes a tool that is able to reduce the size of the objective function from $\operatorname{poly}(n)$ to $\operatorname{poly}(k)$. Section 21.11.4 discusses the case when the problem size is small. Section 21.11 .5 provides several bicriteria algorithms. Section 21.11 .6 summarizes a batch of algorithms. Section 21.11 .7 provides an algorithm for $\ell_{p}$ norm CURT decomposition.

Notice that if the rank- $k$ solution does not exist, then every bicriteria algorithm in Section 21.11.5 can be stated in the form as Theorem 21.1.1, and every algorithm which can output a rank- $k$ solution in Section 21.11.6 can be stated in the form as Theorem 21.1.2. See Section 21.1 for more details.

### 21.11.1 Existence results for matrix case

Theorem 21.11.1 ([SWZ17]). Let $1 \leq p<2$. Given $V \in \mathbb{R}^{k \times n}, A \in \mathbb{R}^{d \times n}$. Let $S \in \mathbb{R}^{n \times s}$ be a proper random sketching matrix. Let

$$
\widehat{U}=\arg \min _{U \in \mathbb{R}^{d \times k}}\|U V S-A S\|_{F}^{2}
$$

i.e.,

$$
\widehat{U}=A S(V S)^{\dagger}
$$

Then with probability at least 0.999,

$$
\|\widehat{U} V-A\|_{p}^{p} \leq \alpha \cdot \min _{U \in \mathbb{R}^{d \times k}}\|U V-A\|_{p}^{p}
$$

(I). $S$ denotes a dense p-stable transform, $s=\widetilde{O}(k), \alpha=\widetilde{O}\left(k^{1-p / 2}\right) \log d$.
(II). $S$ denotes a sparse p-stable transform,
$s=\widetilde{O}\left(k^{5}\right), \alpha=\widetilde{O}\left(k^{5-5 p / 2+2 / p}\right) \log d$.
(III). $S^{\top}$ denotes a sampling/rescaling matrix according to the $\ell_{p}$ Lewis weights of $V^{\top}$,
$s=\widetilde{O}(k), \alpha=\widetilde{O}\left(k^{1-p / 2}\right)$.

We give the proof for completeness.

Proof. Let $S \in \mathbb{R}^{n \times s}$ be a sketching matrix which satisfies the property ( $*$ ): $\forall c \geq 1, \widetilde{U} \in \mathbb{R}^{d \times k}$ which satisfy

$$
\|\widetilde{U} V S-A S\|_{p}^{p} \leq c \cdot \min _{U \in \mathbb{R}^{d \times k}}\|U V S-A S\|_{p}^{p}
$$

we have

$$
\|\widetilde{U} V-A\|_{p}^{p} \leq c \beta_{S} \cdot \min _{U \in \mathbb{R}^{d \times k}}\|U V-A\|_{p}^{p}
$$

where $\beta_{S} \geq 1$ only depends on the sketching matrix $S$. Let

$$
\forall i \in[d],\left(\widehat{U}^{i}\right)^{\top}=\arg \min _{x \in \mathbb{R}^{k}}\left\|x^{\top} V S-A^{i} S\right\|_{2}^{2},
$$

i.e.,

$$
\widehat{U}=A S(V S)^{\dagger}
$$

Let

$$
\widetilde{U}=\arg \min _{U \in \mathbb{R}^{d \times k}}\|U V S-A S\|_{p}^{p}
$$

Then, we have:

$$
\begin{aligned}
& \|\widehat{U} V S-A S\|_{p}^{p} \\
= & \sum_{i=1}^{d}\left\|\widehat{U}^{i} V S-A^{i} S\right\|_{p}^{p} \\
\leq & \sum_{i=1}^{d}\left(s^{1 / p-1 / 2}\left\|\widehat{U}^{i} V S-A^{i} S\right\|_{2}\right)^{p} \\
\leq & \sum_{i=1}^{d}\left(s^{1 / p-1 / 2}\left\|\widetilde{U}^{i} V S-A^{i} S\right\|_{2}\right)^{p} \\
\leq & \sum_{i=1}^{d}\left(s^{1 / p-1 / 2}\left\|\widetilde{U}^{i} V S-A^{i} S\right\|_{p}\right)^{p} \\
\leq & s^{1-p / 2}\|\widetilde{U} V S-A S\|_{p}^{p} .
\end{aligned}
$$

The first inequality follows using $\forall x \in \mathbb{R}^{s},\|x\|_{p} \leq s^{1 / p-1 / 2}\|x\|_{2}$ since $p<2$. The third inequality follows using $\forall x \in \mathbb{R}^{s},\|x\|_{2} \leq\|x\|_{p}$ since $p<2$. Thus, according to the property $(*)$ of $S$,

$$
\|\widehat{U} V-A\|_{p}^{p} \leq s^{1-p / 2} \beta_{S} \min _{U \in \mathbb{R}^{d \times k}}\|U V-A\|_{p}^{p} .
$$

Due to Lemma E. 8 and Lemma E. 11 of [SWZ17], we have:

$$
\begin{aligned}
& \text { for }(\mathrm{I}), s=\widetilde{O}(k), \beta_{S}=O(\log d), \alpha=s^{1-p / 2} \beta_{S}=\widetilde{O}\left(k^{1-p / 2}\right) \log d, \\
& \text { for }(\mathrm{II}), s=\widetilde{O}\left(k^{5}\right), \beta_{S}=\widetilde{O}\left(k^{2 / p} \log d\right), \alpha=s^{1-p / 2} \beta_{S}=\widetilde{O}\left(k^{5-5 p / 2+2 / p}\right) \log d, \\
& \text { for }(\mathrm{III}), s=\widetilde{O}(k), \beta_{S}=O(1), \alpha=s^{1-p / 2} \beta_{S}=\widetilde{O}\left(k^{1-p / 2}\right) .
\end{aligned}
$$

### 21.11.2 Existence results

Theorem 21.11.2. Given a 3rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, there exist three matrices $S_{1} \in \mathbb{R}^{n^{2} \times s_{1}}, S_{2} \in \mathbb{R}^{n^{2} \times s_{2}}, S_{3} \in \mathbb{R}^{n^{2} \times s_{3}}$ such that

$$
\min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k}\left(A_{1} S_{1} X_{1}\right)_{i} \otimes\left(A_{2} S_{2} X_{2}\right)_{i} \otimes\left(A_{3} S_{3} X_{3}\right)_{i}-A\right\|_{p}^{p} \leq \alpha_{\mathrm{rank}-k} \min _{A_{k} \in \mathbb{R}^{n \times n \times n}}\left\|A_{k}-A\right\|_{p}^{p},
$$

holds with probability 99/100.
(I). Using a dense p-stable transform,
$s_{1}=s_{2}=s_{3}=\widetilde{O}(k), \alpha=\widetilde{O}\left(k^{3-1.5 p}\right) \log ^{3} n$.
(II). Using a sparse p-stable transform, $s_{1}=s_{2}=s_{3}=\widetilde{O}\left(k^{5}\right), \alpha=\widetilde{O}\left(k^{15-7.5 p+6 / p}\right) \log ^{3} n$.
(III). Guessing Lewis weights, $s_{1}=s_{2}=s_{3}=\widetilde{O}(k), \alpha=\widetilde{O}\left(k^{3-1.5 p}\right)$.

Proof. We use OPT to denote

$$
\mathrm{OPT}:=\min _{\text {rank }-k} A_{k} \in \mathbb{R}^{n \times n \times n} \mid ~\left\|A_{k}-A\right\|_{p}^{p}
$$

Given a tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$, we define three matrices $A_{1} \in \mathbb{R}^{n_{1} \times n_{2} n_{3}}, A_{2} \in \mathbb{R}^{n_{2} \times n_{3} n_{1}}, A_{3} \in$ $\mathbb{R}^{n_{3} \times n_{1} n_{2}}$ such that, for any $i \in\left[n_{1}\right], j \in\left[n_{2}\right], l \in\left[n_{3}\right]$

$$
A_{i, j, l}=\left(A_{1}\right)_{i,(j-1) \cdot n_{3}+l}=\left(A_{2}\right)_{j,(l-1) \cdot n_{1}+i}=\left(A_{3}\right)_{l,(i-1) \cdot n_{2}+j} .
$$

We fix $V^{*} \in \mathbb{R}^{n \times k}$ and $W^{*} \in \mathbb{R}^{n \times k}$, and use $V_{1}^{*}, V_{2}^{*}, \cdots, V_{k}^{*}$ to denote the columns of $V^{*}$ and $W_{1}^{*}, W_{2}^{*}, \cdots, W_{k}^{*}$ to denote the columns of $W^{*}$.

We consider the following optimization problem,

$$
\min _{U_{1}, \cdots, U_{k} \in \mathbb{R}^{n}}\left\|\sum_{i=1}^{k} U_{i} \otimes V_{i}^{*} \otimes W_{i}^{*}-A\right\|_{p}^{p}
$$

which is equivalent to

$$
\min _{U_{1}, \cdots, U_{k} \in \mathbb{R}^{n}}\left\|\left[\begin{array}{llll}
U_{1} & U_{2} & \cdots & U_{k}
\end{array}\right]\left[\begin{array}{c}
V_{1}^{*} \otimes W_{1}^{*} \\
V_{2}^{*} \otimes W_{2}^{*} \\
\cdots \\
V_{k}^{*} \otimes W_{k}^{*}
\end{array}\right]-A\right\|_{p}^{p}
$$

We use matrix $Z_{1}$ to denote $V^{* T} \odot W^{* \top} \in \mathbb{R}^{k \times n^{2}}$ and matrix $U$ to denote $\left[\begin{array}{llll}U_{1} & U_{2} & \cdots & U_{k}\end{array}\right]$. Then we can obtain the following equivalent objective function,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1}-A_{1}\right\|_{p}^{p}
$$

Choose a sketching matrix (a dense $p$-stable, a sparse $p$-stable or an $\ell_{p}$ Lewis weight sampling/rescaling matrix to $Z_{1}$ ) $S_{1} \in \mathbb{R}^{n^{2} \times s_{1}}$. We can obtain the optimization problem,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1} S_{1}-A_{1} S_{1}\right\|_{p}^{p}=\min _{U \in \mathbb{R}^{n \times k}} \sum_{i=1}^{n}\left\|U^{i} Z_{1} S_{1}-\left(A_{1} S_{1}\right)^{i}\right\|_{p}^{p}
$$

where $U^{i}$ denotes the $i$-th row of matrix $U \in \mathbb{R}^{n \times k}$ and $\left(A_{1} S_{1}\right)^{i}$ denotes the $i$-th row of matrix $A_{1} S_{1}$. Instead of solving it under the $\ell_{p}$-norm, we consider the $\ell_{2}$-norm relaxation,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1} S_{1}-A_{1} S_{1}\right\|_{F}^{2}=\min _{U \in \mathbb{R}^{n \times k}} \sum_{i=1}^{n}\left\|U^{i} Z_{1} S_{1}-\left(A_{1} S_{1}\right)^{i}\right\|_{2}^{2}
$$

Let $\widehat{U} \in \mathbb{R}^{n \times k}$ denote the optimal solution of the above optimization problem. Then, $\widehat{U}=$ $A_{1} S_{1}\left(Z_{1} S_{1}\right)^{\dagger}$. We plug $\widehat{U}$ into the objective function under the $\ell_{p}$-norm. By choosing $s_{1}$ and by the properties of sketching matrices (a dense $p$-stable, a sparse $p$-stable or an $\ell_{p}$ Lewis weight sampling/rescaling matrix to $\left.Z_{1}\right) S_{1} \in \mathbb{R}^{n^{2} \times s_{1}}$, we have

$$
\left\|\widehat{U} Z_{1}-A_{1}\right\|_{p}^{p} \leq \alpha \min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1}-A_{1}\right\|_{p}^{p}=\alpha \mathrm{OPT}
$$

This implies

$$
\left\|\widehat{U} \otimes V^{*} \otimes W^{*}-A\right\|_{p}^{p} \leq \alpha \mathrm{OPT}
$$

As a second step, we fix $\widehat{U} \in \mathbb{R}^{n \times k}$ and $W^{*} \in \mathbb{R}^{n \times k}$, and convert tensor $A$ into matrix $A_{2}$. Let matrix $Z_{2}$ denote $\widehat{U}^{\top} \odot W^{* \top}$. We consider the following objective function,

$$
\min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2}-A_{2}\right\|_{p}^{p}
$$

and the optimal cost of it is at most $\alpha$ OPT.
We choose a sketching matrix (a dense $p$-stable, a sparse $p$-stable or an $\ell_{p}$ Lewis weight sampling/rescaling matrix to $\left.Z_{2}\right) S_{2} \in \mathbb{R}^{n^{2} \times s_{2}}$ and sketch on the right of the objective function to obtain the new objective function,

$$
\min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2} S_{2}-A_{2} S_{2}\right\|_{p}^{p}=\min _{V \in \mathbb{R}^{n \times k}} \sum_{i=1}^{n}\left\|V^{i} Z_{2} S_{2}-\left(A_{2} S_{2}\right)^{i}\right\|_{p}^{p}
$$

where $V^{i}$ denotes the $i$-th row of matrix $V$ and $\left(A_{2} S_{2}\right)^{i}$ denotes the $i$-th row of matrix $A_{2} S_{2}$. Instead of solving this under the $\ell_{p}$-norm, we consider the $\ell_{2}$-norm relaxation,

$$
\min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2} S_{2}-A_{2} S_{2}\right\|_{F}^{2}=\min _{V \in \mathbb{R}^{n \times k}} \sum_{i=1}^{n}\left\|V^{i}\left(Z_{2} S_{2}\right)-\left(A_{2} S_{2}\right)^{i}\right\|_{2}^{2}
$$

Let $\widehat{V} \in \mathbb{R}^{n \times k}$ denote the optimal solution of the above problem. Then $\widehat{V}=$ $A_{2} S_{2}\left(Z_{2} S_{2}\right)^{\dagger}$. By properties of sketching matrix $S_{2} \in \mathbb{R}^{n^{2} \times s_{2}}$, we have,

$$
\left\|\widehat{V} Z_{2}-A_{2}\right\|_{p}^{p} \leq \alpha \min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2}-A_{2}\right\|_{p}^{p} \leq \alpha^{2} \mathrm{OPT}
$$

which implies

$$
\left\|\widehat{U} \otimes \widehat{V} \otimes W^{*}-A\right\|_{p}^{p} \leq \alpha^{2} \mathrm{OPT}
$$

As a third step, we fix the matrices $\widehat{U} \in \mathbb{R}^{n \times k}$ and $\widehat{V} \in \mathbb{R}^{n \times k}$. We can convert tensor $A \in \mathbb{R}^{n \times n \times n}$ into matrix $A_{3} \in \mathbb{R}^{n^{2} \times n}$. Let matrix $Z_{3}$ denote $\widehat{U}^{\top} \odot \widehat{V}^{\top} \in \mathbb{R}^{k \times n^{2}}$. We consider the following objective function,

$$
\min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3}-A_{3}\right\|_{p}^{p}
$$

and the optimal cost of it is at most $\alpha^{2}$ OPT.
We choose sketching matrix (a dense $p$-stable, a sparse $p$-stable or an $\ell_{p}$ Lewis weight sampling/rescaling matrix to $\left.Z_{3}\right) S_{3} \in \mathbb{R}^{n^{2} \times s_{3}}$ and sketch on the right of the objective function to obtain the new objective function,

$$
\min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3} S_{3}-A_{3} S_{3}\right\|_{p}^{p}
$$

Instead of solving this under the $\ell_{p}$-norm, we consider the $\ell_{2}$-norm relaxation,

$$
\min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3} S_{3}-A_{3} S_{3}\right\|_{F}^{2}=\min _{W \in \mathbb{R}^{n \times k}} \sum_{i=1}^{n}\left\|W^{i}\left(Z_{3} S_{3}\right)-\left(A_{3} S_{3}\right)^{i}\right\|_{2}^{2}
$$

Let $\widehat{W} \in \mathbb{R}^{n \times k}$ denote the optimal solution of the above problem. Then $\widehat{W}=A_{3} S_{3}\left(Z_{3} S_{3}\right)^{\dagger}$. By properties of sketching matrix $S_{3} \in \mathbb{R}^{n^{2} \times s_{3}}$, we have,

$$
\left\|\widehat{W} Z_{3}-A_{3}\right\|_{p}^{p} \leq \alpha \min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3}-A_{3}\right\|_{p}^{p} \leq \alpha^{3} \mathrm{OPT}
$$

Thus, we obtain,

$$
\min _{X_{1} \in \mathbb{R}^{s_{1} \times k}, X_{2} \in \mathbb{R}^{s_{2} \times k}, X_{3} \in \mathbb{R}^{s_{3} \times k}}\left\|\sum_{i=1}^{k}\left(A_{1} S_{1} X_{1}\right)_{i} \otimes\left(A_{2} S_{2} X_{2}\right)_{i} \otimes\left(A_{3} S_{3} X_{3}\right)_{i}-A\right\|_{p}^{p} \leq \alpha^{3} \mathrm{OPT}
$$

According to Theorem 21.11.1, we let $s=s_{1}=s_{2}=s_{3}$ and take the corresponding $\alpha$. We can directly get the results for (I), (II) and (III).

### 21.11.3 Polynomial in $k$ size reduction

Definition 21.11.1 (Definition E. 1 in [SWZ17]). Given a matrix $M \in \mathbb{R}^{n \times d}$, if matrix $S \in \mathbb{R}^{m \times n}$ satisfies

$$
\|S M\|_{p}^{p} \leq \beta\|M\|_{p}^{p}
$$

then $S$ has at most $\beta$ dilation on $M$ in the $\ell_{p}$ case.
Definition 21.11.2 (Definition E. 2 in [SWZ17]). Given a matrix $U \in \mathbb{R}^{n \times k}$, if matrix $S \in \mathbb{R}^{m \times n}$ satisfies

$$
\forall x \in \mathbb{R}^{k},\|S U x\|_{p}^{p} \geq \frac{1}{\beta}\|U x\|_{p}^{p}
$$

then $S$ has at most $\beta$ contraction on $U$ in the $\ell_{p}$ case.

Theorem 21.11.3. Given a tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$ and three matrices $V_{1} \in \mathbb{R}^{n_{1} \times b_{1}}, V_{2} \in$ $\mathbb{R}^{n_{2} \times b_{2}}, V_{3} \in \mathbb{R}^{n_{3} \times b_{3}}$, let $X_{1}^{*} \in \mathbb{R}^{b_{1} \times k}, X_{2}^{*} \in \mathbb{R}^{b_{2} \times k}, X_{3}^{*} \in \mathbb{R}^{b_{3} \times k}$ satisfy

$$
X_{1}^{*}, X_{2}^{*}, X_{3}^{*}=\underset{X_{1} \in \mathbb{R}^{b_{1} \times k}, X_{2} \in \mathbb{R}^{b_{2} \times k}, X_{3} \in \mathbb{R}^{b_{3} \times k}}{\arg \min }\left\|V_{1} X_{1} \otimes V_{2} X_{2} \otimes V_{3} X_{3}-A\right\|_{p}^{p}
$$

Let $S \in \mathbb{R}^{m \times n}$ have at most $\beta_{1} \geq 1$ dilation on $V_{1} X_{1}^{*} \cdot\left(\left(V_{2} X_{2}^{*}\right)^{\top} \odot\left(V_{3} X_{3}^{*}\right)^{\top}\right)-A_{1}$ and $S$ have at most $\beta_{2} \geq 1$ contraction on $V_{1}$ in the $\ell_{p}$ case. If $\widehat{X}_{1} \in \mathbb{R}^{b_{1} \times k}, \widehat{X}_{2} \in \mathbb{R}^{b_{2} \times k}, \widehat{X}_{3} \in \mathbb{R}^{b_{3} \times k}$ satisfy
$\left\|S V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-S A\right\|_{p}^{p} \leq \beta_{X_{1} \in \mathbb{R}^{b_{1} \times k}, X_{2} \in \mathbb{R}^{b_{2} \times k}, X_{3} \in \mathbb{R}^{b_{3} \times k}}\left\|S V_{1} X_{1} \otimes V_{2} X_{2} \otimes V_{3} X_{3}-S A\right\|_{p}^{p}$, where $\beta \geq 1$, then

$$
\left\|V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-A\right\|_{p}^{p} \lesssim \beta_{1} \beta_{2} \beta \min _{X_{1}, X_{2}, X_{3}}\left\|V_{1} X_{1} \otimes V_{2} X_{2} \otimes V_{3} X_{3}-A\right\|_{p}^{p}
$$

The proof is essentially the same as the proof of Theorem 21.7.4:

Proof. Let $A, V_{1}, V_{2}, V_{3}, S, X_{1}^{*}, X_{2}^{*}, X_{3}^{*}, \beta_{1}, \beta_{2}$ be as stated in the theorem. Let $\widehat{X}_{1} \in \mathbb{R}^{b_{1} \times k}, \widehat{X}_{2} \in$ $\mathbb{R}^{b_{2} \times k}, \widehat{X}_{3} \in \mathbb{R}^{b_{3} \times k}$ satisfy
$\left\|S V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-S A\right\|_{p}^{p} \leq \beta_{X_{1} \in \mathbb{R}^{b_{1} \times k}, X_{2} \in \mathbb{R}^{b_{2} \times k}, X_{3} \in \mathbb{R}^{b_{3} \times k}}\left\|S V_{1} X_{1} \otimes V_{2} X_{2} \otimes V_{3} X_{3}-S A\right\|_{p}^{p}$.
Similar to the proof of Theorem 21.7.4, we have,

$$
\begin{aligned}
& \left\|S V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-S A\right\|_{p}^{p} \\
= & 2^{2-2 p} \frac{1}{\beta_{2}}\left\|V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-A\right\|_{p}^{p}-\left(2^{1-p} \frac{1}{\beta_{2}}+\beta_{1}\right)\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{p}^{p}
\end{aligned}
$$

The only difference from the proof of Theorem 21.7.4 is that instead of using triangle inequality, we actually use $\|x+y\|_{p}^{p} \leq 2^{p-1}\|x\|_{p}^{p}+\|y\|_{p}^{p}$. Then, we have

$$
\begin{aligned}
& \left\|V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-A\right\|_{p}^{p} \\
\leq & 2^{2 p-2} \beta_{2}\left\|S V_{1} \widehat{X}_{1} \otimes V_{2} \widehat{X}_{2} \otimes V_{3} \widehat{X}_{3}-S A\right\|_{p}^{p}+\left(2^{p-1}+2^{2 p-2} \beta_{1} \beta_{2}\right)\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{p}^{p} \\
\leq & 2^{2 p-2} \beta_{2} \beta\left\|S V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-S A\right\|_{p}^{p}+\left(2^{p-1}+2^{2 p-2} \beta_{1} \beta_{2}\right)\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{p}^{p} \\
\leq & 2^{2 p-2} \beta_{1} \beta_{2} \beta\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{p}^{p}+\left(2^{p-1}+2^{2 p-2} \beta_{1} \beta_{2}\right)\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{p}^{p} \\
\leq & 2^{p-1} \beta\left(1+2 \beta_{1} \beta_{2}\right)\left\|V_{1} X_{1}^{*} \otimes V_{2} X_{2}^{*} \otimes V_{3} X_{3}^{*}-A\right\|_{p}^{p} .
\end{aligned}
$$

Lemma 21.11.4. Let $\min \left(b_{1}, b_{2}, b_{3}\right) \geq k$. Given three matrices $V_{1} \in \mathbb{R}^{n \times b_{1}}, V_{2} \in \mathbb{R}^{n \times b_{2}}$, and $V_{3} \in \mathbb{R}^{n \times b_{3}}$, there exists an algorithm which takes $O(\mathrm{nnz}(A))+n \operatorname{poly}\left(b_{1}, b_{2}, b_{3}\right)$ time and outputs a tensor $C \in \mathbb{R}^{c_{1} \times c_{2} \times c_{3}}$ and three matrices $\widehat{V}_{1} \in \mathbb{R}^{c_{1} \times b_{1}}, \widehat{V}_{2} \in \mathbb{R}^{c_{2} \times b_{2}}$ and $\widehat{V}_{3} \in \mathbb{R}^{c_{3} \times b_{3}}$
with $c_{1}=c_{2}=c_{3}=\operatorname{poly}\left(b_{1}, b_{2}, b_{3}\right)$, such that with probability 0.99 , for any $\alpha \geq 1$, if $X_{1}^{\prime}, X_{2}^{\prime}, X_{3}^{\prime}$ satisfy that,

$$
\left\|\sum_{i=1}^{k}\left(\widehat{V}_{1} X_{1}^{\prime}\right)_{i} \otimes\left(\widehat{V}_{2} X_{2}^{\prime}\right)_{i} \otimes\left(\widehat{V}_{3} X_{3}^{\prime}\right)_{i}-C\right\|_{p}^{p} \leq \min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k}\left(\widehat{V}_{1} X_{1}\right)_{i} \otimes\left(\widehat{V}_{2} X_{2}\right)_{i} \otimes\left(\widehat{V}_{3} X_{3}\right)_{i}-C\right\|_{p}^{p}
$$

then,

$$
\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}^{\prime}\right)_{i} \otimes\left(V_{2} X_{2}^{\prime}\right)_{i} \otimes\left(V_{3} X_{3}^{\prime}\right)_{i}-A\right\|_{p}^{p} \lesssim \min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}\right)_{i} \otimes\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}-A\right\|_{p}^{p}
$$

Proof. For simplicity, we define OPT to be

$$
\min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}\right)_{i} \otimes\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}-A\right\|_{p}^{p} .
$$

Let $T_{1} \in \mathbb{R}^{c_{1} \times n}$ correspond to sampling according to the $\ell_{p}$ Lewis weights of $V_{1} \in \mathbb{R}^{n \times b_{1}}$, where $c_{1}=\widetilde{b_{1}}$. Let $T_{2} \in \mathbb{R}^{c_{2} \times n}$ be sampling according to the $\ell_{p}$ Lewis weights of $V_{2} \in \mathbb{R}^{n \times b_{2}}$, where $c_{2}=\widetilde{b_{2}}$. Let $T_{3} \in \mathbb{R}^{c_{3} \times n}$ be sampling according to the $\ell_{p}$ Lewis weights of $V_{3} \in \mathbb{R}^{n \times b_{3}}$, where $c_{3}=\widetilde{b_{3}}$.

For any $\alpha \geq 1$, let $X_{1}^{\prime} \in \mathbb{R}^{b_{1} \times k}, X_{2}^{\prime} \in \mathbb{R}^{b_{2} \times k}, X_{3}^{\prime} \in \mathbb{R}^{b_{3} \times k}$ satisfy

$$
\begin{aligned}
& \left\|T_{1} V_{1} X_{1}^{\prime} \otimes T_{2} V_{2} X_{2}^{\prime} \otimes T_{3} V_{3} X_{3}^{\prime}-A\left(T_{1}, T_{2}, T_{3}\right)\right\|_{p}^{p} \\
\leq & \alpha{ }_{X_{1} \in \mathbb{R}^{b_{1} \times k}, X_{2} \in \mathbb{R}^{b_{2} \times k}, X_{3} \in \mathbb{R}^{b_{3} \times k}}\left\|T_{1} V_{1} X_{1} \otimes T_{2} V_{2} X_{2} \otimes T_{3} V_{3} X_{3}-A\left(T_{1}, T_{2}, T_{3}\right)\right\|_{p}^{p} .
\end{aligned}
$$

First, we regard $T_{1}$ as the sketching matrix for the remainder. Then by Lemma D. 11 in [SWZ17] and Theorem 21.7.4, we have

$$
\begin{aligned}
& \left\|V_{1} X_{1}^{\prime} \otimes T_{2} V_{2} X_{2}^{\prime} \otimes T_{3} V_{3} X_{3}^{\prime}-A\left(I, T_{2}, T_{3}\right)\right\|_{p}^{p} \\
\lesssim & \alpha \min _{X_{1} \in \mathbb{R}^{b_{1} \times k}, X_{2} \in \mathbb{R}^{b_{2} \times k}, X_{3} \in \mathbb{R}^{b_{3} \times k}}\left\|V_{1} X_{1} \otimes T_{2} V_{2} X_{2} \otimes T_{3} V_{3} X_{3}-A\left(I, T_{2}, T_{3}\right)\right\|_{p}^{p}
\end{aligned}
$$

Second, we regard $T_{2}$ as the sketching matrix for $V_{1} X_{1} \otimes V_{2} X_{2} \otimes T_{3} V_{3} X_{3}-A\left(I, I, T_{3}\right)$. Then by Lemma D. 11 in [SWZ17] and Theorem 21.7.4, we have

$$
\begin{aligned}
& \left\|V_{1} X_{1}^{\prime} \otimes V_{2} X_{2}^{\prime} \otimes T_{3} V_{3} X_{3}^{\prime}-A\left(I, I, T_{3}\right)\right\|_{p}^{p} \\
\lesssim & \alpha{ }_{X_{1} \in \mathbb{R}^{b_{1} \times k}, X_{2} \in \mathbb{R}^{b_{2} \times k}, X_{3} \in \mathbb{R}^{b_{3} \times k}}\left\|V_{1} X_{1} \otimes V_{2} X_{2} \otimes T_{3} V_{3} X_{3}-A\left(I, I, T_{3}\right)\right\|_{p}^{p} .
\end{aligned}
$$

Third, we regard $T_{3}$ as the sketching matrix for $V_{1} X_{1} \otimes V_{2} X_{2} \otimes V_{3} X_{3}-A$. Then by Lemma D. 11 in [SWZ17] and Theorem 21.7.4, we have

$$
\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}^{\prime}\right)_{i} \otimes\left(V_{2} X_{2}^{\prime}\right)_{i} \otimes\left(V_{3} X_{3}^{\prime}\right)_{i}-A\right\|_{p}^{p} \lesssim \alpha \min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k}\left(V_{1} X_{1}\right)_{i} \otimes\left(V_{2} X_{2}\right)_{i} \otimes\left(V_{3} X_{3}\right)_{i}-A\right\|_{p}^{p}
$$

### 21.11.4 Solving small problems

Combining Section B. 5 in [SWZ17] and the proof of Theorem 21.7.4, for any $p=a / b$ with $a, b$ are integers, we can obtain the $\ell_{p}$ version of Theorem 21.7.4.

### 21.11.5 Bicriteria algorithm

We present several bicriteria algorithms with different tradeoffs. We first present an algorithm that runs in nearly linear time and outputs a solution with rank $\widetilde{O}\left(k^{3}\right)$ in Theorem 21.11.5. Then we show an algorithm that runs in $n n z(A)$ time but outputs a solution with rank poly $(k)$ in Theorem 21.11.6. Then we explain an idea which is able to decrease the cubic rank to quadratic, and thus we can obtain Theorem 21.11.7.

Theorem 21.11.5. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, let $r=\widetilde{O}\left(k^{3}\right)$. There exists an algorithm which takes $\mathrm{nnz}(A) \cdot \widetilde{O}(k)+n \operatorname{poly}(k)+\operatorname{poly}(k)$ time and outputs three matrices $U, V, W \in \mathbb{R}^{n \times r}$ such that

$$
\left\|\sum_{i=1}^{r} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{p}^{p} \leq \widetilde{O}\left(k^{3-p / 2}\right) \log ^{3} n \min _{\operatorname{rank}-k A_{k}}\left\|A_{k}-A\right\|_{p}^{p}
$$

holds with probability 9/10.

Proof. We first choose three dense Cauchy transforms $S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}$. According to Section 21.3.7, for each $i \in[3], A_{i} S_{i}$ can be computed in $n n z(A) \cdot \widetilde{O}(k)$ time. Then we apply Lemma 21.11.4. We obtain three matrices $Y_{1}=T_{1} A_{1} S_{1}, Y_{2}=T_{2} A_{2} S_{2}, Y_{3}=T_{3} A_{3} S_{3}$ and a tensor $C=A\left(T_{1}, T_{2}, T_{3}\right)$. Note that for each $i \in[3], Y_{i}$ can be computed in $n \operatorname{poly}(k)$ time. Because $C=A\left(T_{1}, T_{2}, T_{3}\right)$ and $T_{1}, T_{2}, T_{3} \in \mathbb{R}^{n \times \widetilde{O}(k)}$ are three sampling and rescaling matrices, $C$ can be computed in $n n z(A)+\widetilde{O}\left(k^{3}\right)$ time. At the end, we just need to run an $\ell_{p}$-regression solver to find the solution for the problem:

$$
\min _{X \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} X_{i, j, l}\left(Y_{1}\right)_{i} \otimes\left(Y_{2}\right)_{j} \otimes\left(Y_{3}\right)_{j}\right\|_{p}^{p},
$$

where $\left(Y_{1}\right)_{i}$ denotes the $i$-th column of matrix $Y_{1}$. Since the size of the above problem is only $\operatorname{poly}(k)$, this can be solved in $\operatorname{poly}(k)$ time.

Theorem 21.11.6. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, let $r=\widetilde{O}\left(k^{15}\right)$. There exists an algorithm that takes $n n z(A)+n \operatorname{poly}(k)+\operatorname{poly}(k)$ time and outputs three matrices $U, V, W \in \mathbb{R}^{n \times r}$ such that

$$
\left\|\sum_{i=1}^{r} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{p}^{p} \leq \operatorname{poly}(k, \log n) \min _{\operatorname{rank}-k A_{k}}\left\|A_{k}-A\right\|_{p}^{p}
$$

holds with probability 9/10.

Proof. We first choose three sparse $p$-stable transforms $S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}$. According to Section 21.3.7, for each $i \in[3], A_{i} S_{i}$ can be computed in $O(\mathrm{nnz}(A))$ time. Then we apply Lemma 21.11.4, and can obtain three matrices $Y_{1}=T_{1} A_{1} S_{1}, Y_{2}=T_{2} A_{2} S_{2}, Y_{3}=T_{3} A_{3} S_{3}$ and a tensor $C=A\left(T_{1}, T_{2}, T_{3}\right)$. Note that for each $i \in[3], Y_{i}$ can be computed in $n$ poly $(k)$ time. Because $C=A\left(T_{1}, T_{2}, T_{3}\right)$ and $T_{1}, T_{2}, T_{3} \in \mathbb{R}^{n \times \widetilde{O}(k)}$ are three sampling and rescaling matrices, $C$ can be computed in $n n z(A)+\widetilde{O}\left(k^{3}\right)$ time. At the end, we just need to run an $\ell_{p}$-regression solver to find the solution to the problem,

$$
\min _{X \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} X_{i, j, l}\left(Y_{1}\right)_{i} \otimes\left(Y_{2}\right)_{j} \otimes\left(Y_{3}\right)_{l}-C\right\|_{p}^{p}
$$

where $\left(Y_{1}\right)_{i}$ denotes the $i$-th column of matrix $Y_{1}$. Since the size of the above problem is only $\operatorname{poly}(k)$, it can be solved in poly $(k)$ time.

Theorem 21.11.7. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, $\epsilon \in(0,1)$, let $r=\widetilde{O}\left(k^{2}\right)$. There exists an algorithm which takes $\mathrm{nnz}(A) \cdot \widetilde{O}(k)+n \operatorname{poly}(k)+\operatorname{poly}(k)$ time and outputs three matrices $U, V, W \in \mathbb{R}^{n \times r}$ such that

$$
\left\|\sum_{i=1}^{r} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{p}^{p} \leq \widetilde{O}\left(k^{3-1.5 p}\right) \log ^{3} n_{\text {rank }-k} \min _{A_{k}}\left\|A_{k}-A\right\|_{p}^{p}
$$

holds with probability 9/10.

Proof. The proof is similar to Theorem 21.7.11.

```
Algorithm \(21.35 \ell_{p}\)-Low Rank Approximation, Bicriteria Algorithm, rank- \(\widetilde{O}\left(k^{2}\right)\), Input
Sparsity Time
    procedure LpBicriteriaAlgorithm \((A, n, k) \quad\) Corollary 21.11.8
    \(s_{1} \leftarrow s_{2} \leftarrow s_{3} \leftarrow \widetilde{O}(k)\).
    For each \(i \in[3]\), choose \(S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}\) to be the composition of a sparse \(p\)-stable
    transform and a dense \(p\)-stable transform.
                                    \(\triangleright\) Part (I,II) of Theorem 21.11.2
    Compute \(A_{1} \cdot S_{1}, A_{2} \cdot S_{2}\).
    For each \(i \in[2]\), choose \(T_{i}\) to be a sampling and rescaling diagonal matrix according
    to the Lewis weights of \(A_{i} S_{i}\), with \(t_{i}=\widetilde{O}(k)\) nonzero entries.
    \(C \leftarrow A\left(T_{1}, T_{2}, I\right)\).
    \(B^{i+(j-1) s_{1}} \leftarrow \operatorname{vec}\left(\left(T_{1} A_{1} S_{1}\right)_{i} \otimes\left(T_{2} A_{2} S_{2}\right)_{j}\right), \forall i \in\left[s_{1}\right], j \in\left[s_{2}\right]\).
    Form objective function \(\min _{W}\left\|W B-C_{3}\right\|_{1}\).
    Run \(\ell_{p}\)-regression solver to find \(\widehat{W}\).
    Construct \(\widehat{U}\) by copying \(\left(A_{1} S_{1}\right)_{i}\) to the \((i, j)\)-th column of \(\widehat{U}\).
    Construct \(\widehat{V}\) by copying \(\left(A_{2} S_{2}\right)_{j}\) to the \((i, j)\)-th column of \(\widehat{V}\).
    return \(\widehat{U}, \widehat{V}, \widehat{W}\).
end procedure
```

As for $\ell_{1}$, notice that if we first apply a sparse Cauchy transform, we can reduce the rank of the matrix to poly $(k)$. Theyn we can apply a dense Cauchy transform and further reduce the dimension, while only incurring another $\operatorname{poly}(k)$ factor in the approximation ratio. By combining sparse $p$-stable and dense $p$-stable transforms, we can improve the running time from $n n z(A) \cdot \widetilde{O}(k)$ to be $n n z(A)$ by losing some additional poly $(k)$ factors in the approximation ratio.

Corollary 21.11.8. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1, \epsilon \in(0,1)$, let $r=\widetilde{O}\left(k^{2}\right)$. There exists an algorithm which takes $\mathrm{nnz}(A)+n \operatorname{poly}(k)+\operatorname{poly}(k)$ time and
outputs three matrices $U, V, W \in \mathbb{R}^{n \times r}$ such that

$$
\left\|\sum_{i=1}^{r} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{p}^{p} \leq \operatorname{poly}(k, \log n) \min _{\operatorname{rank}-k A_{k}}\left\|A_{k}-A\right\|_{p}^{p}
$$

holds with probability 9/10.

### 21.11.6 Algorithms

In this section, we show two different algorithms by using different kind of sketches. One is shown in Theorem 21.11 .9 which gives a fast running time. Another one is shown in Theorem 21.11.10 which gives the best approximation ratio.

Theorem 21.11.9. Given a 3 rd tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, there exists an algorithm which takes $O(\mathrm{nnz}(A))+n$ poly $(k)+2^{\widetilde{O}\left(k^{2}\right)}$ time and outputs three matrices $U, V, W \in \mathbb{R}^{n \times k}$ such that,

$$
\|U \otimes V \otimes W-A\|_{p}^{p} \leq \operatorname{poly}(k, \log n) \min _{\text {rank }-k}\left\|A^{\prime}-A\right\|_{p}^{p} .
$$

holds with probability at least 9/10.

Proof. First, we apply part (II) of Theorem 21.11.2. Then $A_{i} S_{i}$ can be computed in $O(\mathrm{nnz}(A))$ time. Second, we use Lemma 21.11.4 to reduce the size of the objective function from $O\left(n^{3}\right)$ to poly $(k)$ in $n$ poly $(k)$ time by only losing a constant factor in approximation ratio. Third, we use Claim 21.3.6 to relax the objective function from entry-wise $\ell_{p}$-norm to Frobenius norm, and this step causes us to lose some other $\operatorname{poly}(k)$ factors in approximation ratio. As a last step, we use Theorem 21.4.44 to solve the Frobenius norm objective function.

Theorem 21.11.10. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, there exists an algorithm that takes $n^{\widetilde{O}(k)} 2^{\widetilde{O}\left(k^{3}\right)}$ time and output three matrices $U, V, W \in \mathbb{R}^{n \times k}$ such that,

$$
\|U \otimes V \otimes W-A\|_{p}^{p} \leq \widetilde{O}\left(k^{3-1.5 p}\right) \min _{\operatorname{rank}-k A^{\prime}}\left\|A^{\prime}-A\right\|_{p}^{p}
$$

holds with probability at least 9/10.

Proof. First, we apply part (III) of Theorem 21.11.2. Then, guessing $S_{i}$ requires $n^{\widetilde{O}(k)}$ time. Second, we use Lemma 21.11.4 to reduce the size of the objective from $O\left(n^{3}\right)$ to $\operatorname{poly}(k)$ in polynomial time while only losing a constant factor in approximation ratio. Third, we solve the small optimization problem.

### 21.11.7 CURT decomposition

Theorem 21.11.11. Given a $3 r d$ order tensor $A \in \mathbb{R}^{n \times n \times n}$, let $k \geq 1$, and let $U_{B}, V_{B}, W_{B} \in$ $\mathbb{R}^{n \times k}$ denote a rank-k, $\alpha$-approximation to $A$. Then there exists an algorithm which takes $O(\mathrm{nnz}(A))+O\left(n^{2}\right)$ poly $(k)$ time and outputs three matrices $C \in \mathbb{R}^{n \times c}$ with columns from $A$, $R \in \mathbb{R}^{n \times r}$ with rows from $A, T \in \mathbb{R}^{n \times t}$ with tubes from $A$, and a tensor $U \in \mathbb{R}^{c \times r \times t}$ with $\operatorname{rank}(U)=k$ such that $c=r=t=O(k \log k \log \log k)$, and

$$
\left\|\sum_{i=1}^{c} \sum_{j=1}^{r} \sum_{l=1}^{t} U_{i, j, l} \cdot C_{i} \otimes R_{j} \otimes T_{l}-A\right\|_{p}^{p} \leq \widetilde{O}\left(k^{3-1.5 p}\right) \alpha \min _{\operatorname{rank}-k A^{\prime}}\left\|A^{\prime}-A\right\|_{p}^{p}
$$

holds with probability 9/10.

Proof. We define

$$
\mathrm{OPT}:=\min _{\text {rank }-k}\left\|A^{\prime}-A\right\|_{p}^{p}
$$

We already have three matrices $U_{B} \in \mathbb{R}^{n \times k}, V_{B} \in \mathbb{R}^{n \times k}$ and $W_{B} \in \mathbb{R}^{n \times k}$ and these three matrices provide a rank- $k, \alpha$ approximation to $A$, i.e.,

$$
\begin{equation*}
\left\|\sum_{i=1}^{k}\left(U_{B}\right)_{i} \otimes\left(V_{B}\right)_{i} \otimes\left(W_{B}\right)_{i}-A\right\|_{p}^{p} \leq \alpha \mathrm{OPT} \tag{21.51}
\end{equation*}
$$

Let $B_{1}=V_{B}^{\top} \odot W_{B}^{\top} \in \mathbb{R}^{k \times n^{2}}$ denote the matrix where the $i$-th row is the vectorization of $\left(V_{B}\right)_{i} \otimes\left(W_{B}\right)_{i}$. By Section B. 3 in [SWZ17], we can compute $D_{1} \in \mathbb{R}^{n^{2} \times n^{2}}$ which is a sampling and rescaling matrix corresponding to the Lewis weights of $B_{1}^{\top}$ in $O\left(n^{2} \operatorname{poly}(k)\right)$ time, and there are $d_{1}=O(k \log k \log \log k)$ nonzero entries on the diagonal of $D_{1}$. Let $A_{i} \in \mathbb{R}^{n \times n^{2}}$ denote the matrix obtained by flattening $A$ along the $i$-th direction, for each $i \in[3]$.

Define $U^{*} \in \mathbb{R}^{n \times k}$ to be the optimal solution to $\min _{U \in \mathbb{R}^{n \times k}}\left\|U B_{1}-A_{1}\right\|_{p}^{p}, \widehat{U}=A_{1} D_{1}\left(B_{1} D_{1}\right)^{\dagger} \in$ $\mathbb{R}^{n \times k}, V_{0} \in \mathbb{R}^{n \times k}$ to be the optimal solution to $\min _{V \in \mathbb{R}^{n \times k}}\left\|V \cdot\left(\widehat{U}^{\top} \odot W_{B}^{\top}\right)-A_{2}\right\|_{p}^{p}$, and $U^{\prime}$ to be the optimal solution to $\min _{U \in \mathbb{R}^{n \times k}}\left\|U B_{1} D_{1}-A_{1} D_{1}\right\|_{p}^{p}$.

By Claim 21.3.4, we have

$$
\left\|\widehat{U} B_{1} D_{1}-A_{1} D_{1}\right\|_{p}^{p} \leq d_{1}^{1-p / 2}\left\|U^{\prime} B_{1} D_{1}-A_{1} D_{1}\right\|_{p}^{p}
$$

Due to Lemma E. 11 and Lemma E. 8 in [SWZ17], with constant probability, we have

$$
\begin{equation*}
\left\|\widehat{U} B_{1}-A_{1}\right\|_{p}^{p} \leq d_{1}^{1-p / 2} \alpha_{D_{1}}\left\|U^{*} B_{1}-A_{1}\right\|_{p}^{p} \tag{21.52}
\end{equation*}
$$

where $\alpha_{D_{1}}=O(1)$.
Recall that $\left(\widehat{U}^{\top} \odot W_{B}^{\top}\right) \in \mathbb{R}^{k \times n^{2}}$ denotes the matrix where the $i$-th row is the vectorization of $\widehat{U}_{i} \otimes\left(W_{B}\right)_{i}, \forall i \in[k]$. Now, we can show,

$$
\begin{array}{rlr}
\left\|V_{0} \cdot\left(\widehat{U}^{\top} \odot W_{B}^{\top}\right)-A_{2}\right\|_{p}^{p} \leq\left\|\widehat{U} B_{1}-A_{1}\right\|_{p}^{p} & \text { by } V_{0}=\underset{V \in \mathbb{R}^{n \times k}}{\arg \min }\left\|V \cdot\left(\widehat{U}^{\top} \odot W_{B}^{\top}\right)-A_{2}\right\|_{p}^{p} \\
& \lesssim d_{1}^{1-p / 2}\left\|U^{*} B_{1}-A_{1}\right\|_{p}^{p} & \text { by Equation (21.52) } \\
& \leq d_{1}^{1-p / 2}\left\|U_{B} B_{1}-A_{1}\right\|_{p}^{p} & \text { by } U^{*}=\underset{U \in \mathbb{R}^{n \times k}}{\arg \min }\left\|U B_{1}-A_{1}\right\|_{p}^{p} \\
& \leq O\left(d_{1}^{1-p / 2}\right) \alpha \mathrm{OPT} . & \text { by Equation (21.51) } \tag{21.51}
\end{array}
$$

We define $B_{2}=\widehat{U}^{\top} \odot W_{B}^{\top}$. We can compute $D_{2} \in \mathbb{R}^{n^{2} \times n^{2}}$ which is a sampling and rescaling matrix corresponding to the $\ell_{p}$ Lewis weights of $B_{2}^{\top}$ in $O\left(n^{2}\right.$ poly $\left.(k)\right)$ time, and there are $d_{2}=O(k \log k \log \log k)$ nonzero entries on the diagonal of $D_{2}$.

Define $V^{*} \in \mathbb{R}^{n \times k}$ to be the optimal solution of $\min _{V \in \mathbb{R}^{n \times k}}\left\|V B_{2}-A_{2}\right\|_{p}^{p}$, $\widehat{V}=$ $A_{2} D_{2}\left(B_{2} D_{2}\right)^{\dagger} \in \mathbb{R}^{n \times k}, W_{0} \in \mathbb{R}^{n \times k}$ to be the optimal solution of $\min _{W \in \mathbb{R}^{n \times k}}\left\|W \cdot\left(\widehat{U}^{\top} \odot \hat{V}^{\top}\right)-A_{3}\right\|_{p}^{p}$, and $V^{\prime}$ to be the optimal solution of $\min _{V \in \mathbb{R}^{n \times k}}\left\|V B_{2} D_{2}-A_{2} D_{2}\right\|_{p}^{p}$.

By Claim 21.3.4, we have

$$
\left\|\widehat{V} B_{2} D_{2}-A_{2} D_{2}\right\|_{p}^{p} \leq d_{2}^{1-p / 2}\left\|V^{\prime} B_{2} D_{2}-A_{2} D_{2}\right\|_{p}^{p}
$$

Due to Lemma E. 11 and Lemma E. 8 in [SWZ17], with constant probability, we have

$$
\begin{equation*}
\left\|\widehat{V} B_{2}-A_{2}\right\|_{p}^{p} \leq d_{2}^{1-p / 2} \alpha_{D_{2}}\left\|V^{*} B_{2}-A_{2}\right\|_{p}^{p} \tag{21.54}
\end{equation*}
$$

where $\alpha_{D_{2}}=O(1)$.
Recall that $\left(\widehat{U}^{\top} \odot \widehat{V}^{\top}\right) \in \mathbb{R}^{k \times n^{2}}$ denotes the matrix for which the $i$-th row is the vectorization of $\widehat{U}_{i} \otimes \widehat{V}_{i}, \forall i \in[k]$. Now, we can show,

$$
\begin{array}{rr} 
& \left\|W_{0} \cdot\left(\widehat{U}^{\top} \odot \widehat{V}^{\top}\right)-A_{3}\right\|_{p}^{p} \\
\leq & \left\|\widehat{V} B_{2}-A_{2}\right\|_{p}^{p} \\
\vdots & \text { by } W_{0}=\underset{W \in \mathbb{R}^{n \times k}}{\arg \min }\left\|W \cdot\left(\widehat{U}^{\top} \odot \widehat{V}^{\top}\right)-A_{3}\right\|_{p}^{p} \\
\leq d_{2}^{1-p / 2}\left\|V^{*} B_{2}-A_{2}\right\|_{p}^{p} & \text { by Equation }(21.54) \\
\leq d_{2}^{1-p / 2}\left\|V_{0} B_{2}-A_{2}\right\|_{p}^{p} & \text { by } V^{*}=\underset{V \in \mathbb{R}^{n \times k}}{\arg \min }\left\|V B_{2}-A_{2}\right\|_{p}^{p}  \tag{21.55}\\
\leq O\left(\left(d_{1} d_{2}\right)^{1-p / 2}\right) \alpha \mathrm{OPT} . & \text { by Equation }(21.53)
\end{array}
$$

We define $B_{3}=\widehat{U}^{\top} \odot \widehat{V}^{\top}$. We can compute $D_{3} \in \mathbb{R}^{n^{2} \times n^{2}}$ which is a sampling and rescaling matrix corresponding to the $\ell_{p}$ Lewis weights of $B_{3}^{\top}$ in $O\left(n^{2}\right.$ poly $\left.(k)\right)$ time, and there are $d_{3}=O(k \log k \log \log k)$ nonzero entries on the diagonal of $D_{3}$.

Define $W^{*} \in \mathbb{R}^{n \times k}$ to be the optimal solution to $\min _{W \in \mathbb{R}^{n \times k}}\left\|W B_{3}-A_{3}\right\|_{p}^{p}, \widehat{W}=$ $A_{3} D_{3}\left(B_{3} D_{3}\right)^{\dagger} \in \mathbb{R}^{n \times k}$, and $W^{\prime}$ to be the optimal solution to $\min _{W \in \mathbb{R}^{n \times k}}\left\|W B_{3} D_{3}-A_{3} D_{3}\right\|_{p}^{p}$.

By Claim 21.3.4, we have

$$
\left\|\widehat{W} B_{3} D_{3}-A_{3} D_{3}\right\|_{p}^{p} \leq d_{3}^{1-p / 2}\left\|W^{\prime} B_{3} D_{3}-A_{3} D_{3}\right\|_{p}^{p}
$$

Due to Lemma E. 11 and Lemma E. 8 in [SWZ17], with constant probability, we have

$$
\begin{equation*}
\left\|\widehat{W} B_{3}-A_{3}\right\|_{p}^{p} \leq d_{3}^{1-p / 2} \alpha_{D_{3}}\left\|W^{*} B_{3}-A_{3}\right\|_{p}^{p} \tag{21.56}
\end{equation*}
$$

where $\alpha_{D_{3}}=O(1)$. Now we can show,

$$
\begin{array}{rlr}
\left\|\widehat{W} B_{3}-A_{3}\right\|_{p}^{p} \lesssim d_{3}^{1-p / 2}\left\|W^{*} B_{3}-A_{3}\right\|_{p}^{p}, & \text { by Equation }(21.56) \\
\leq d_{3}^{1-p / 2}\left\|W_{0} B_{3}-A_{3}\right\|_{p}^{p}, & \text { by } W^{*}=\underset{W \in \mathbb{R}^{n \times k}}{\arg \min }\left\|W B_{3}-A_{3}\right\|_{p}^{p} \\
\leq O\left(\left(d_{1} d_{2} d_{3}\right)^{1-p / 2}\right) \alpha \text { OPT } & \text { by Equation }(21.55)
\end{array}
$$

Thus, it implies,

$$
\left\|\sum_{i=1}^{k} \widehat{U}_{i} \otimes \widehat{V}_{i} \otimes \widehat{W}_{i}-A\right\|_{p}^{p} \leq \operatorname{poly}(k, \log n) \mathrm{OPT}
$$

where $\widehat{U}=A_{1} D_{1}\left(B_{1} D_{1}\right)^{\dagger}, \widehat{V}=A_{2} D_{2}\left(B_{2} D_{2}\right)^{\dagger}, \widehat{W}=A_{3} D_{3}\left(B_{3} D_{3}\right)^{\dagger}$.

### 21.12 Robust Subspace Approximation (Asymmetric Norms for Arbitrary Tensors)

Recently, [CW15b] and [CW15a] study the linear regression problem and low-rank approximation problem under M-Estimator loss functions. In this section, we extend the matrix version of the low rank approximation problem to tensors, i.e., in particular focusing on tensor low-rank approximation under M-Estimator norms. Note that M-Estimators are very different from Frobenius norm and Entry-wise $\ell_{1}$ norm, which are symmetric norms. Namely, flattening the tensor objective function along any of the dimensions does not change the cost if the norm is Frobenius or Entry-wise $\ell_{1}$-norm. However, for M-Estimator norms, we cannot flatten the tensor along all three dimensions. This property makes the tensor low-rank approximation problem under M-Estimator norms more difficult. This section can be split into two independent parts. Section 21.12 .2 studies the $\ell_{1}-\ell_{2}-\ell_{2}$ norm setting, and Section 21.12.3 studies the $\ell_{1}-\ell_{1}-\ell_{2}$ norm setting.

### 21.12.1 Preliminaries

Definition 21.12.1 (Nice functions for $M$-Estimators, $\mathcal{M}_{2}, \mathcal{L}_{p}$, [CW15a]). We say an $M$ Estimator is nice if $M(x)=M(-x), M(0)=0, M$ is non-decreasing in $|x|$, there is a constant $C_{M}>0$ and a constant $p \geq 1$ so that for all $a, b \in \mathbb{R}_{>0}$ with $a \geq b$, we have

$$
C_{m} \frac{|a|}{|b|} \leq \frac{M(a)}{M(b)} \leq\left(\frac{a}{b}\right)^{p},
$$

and also that $M(x)^{\frac{1}{p}}$ is subadditive, that is, $M(x+y)^{\frac{1}{p}} \leq M(x)^{\frac{1}{p}}+M(y)^{\frac{1}{p}}$.
Let $\mathcal{M}_{2}$ denote the set of such nice $M$-estimators, for $p=2$. Let $\mathcal{L}_{p}$ denote $M$ Estimators with $M(x)=|x|^{p}$ and $p \in[1,2)$.

### 21.12.2 $\quad \ell_{1}$-Frobenius (a.k.a $\ell_{1}-\ell_{2}-\ell_{2}$ ) norm

Section 21.12.2.1 presents basic definitions and facts for the $\ell_{1}-\ell_{2}-\ell_{2}$ norm setting. Section 21.12.2.2 introduces some useful tools. Section 21.12.2.3 presents the "no dilation" and "no contraction" bounds, which are the key ideas for reducing the problem to a "generalized" Frobenius norm low rank approximation problem. Finally, we provide our algorithms in Section 21.12.2.6.

### 21.12.2.1 Definitions

We first give the definition for the $v$-norm of a tensor, and then give the definition of the $v$-norm for a matrix and a weighted version of the $v$-norm for a matrix.

Definition 21.12.2 (Tensor $v$-norm). For an $n \times n \times n$ tensor $A$, we define the $v$-norm of $A$, denoted $\|A\|_{v}$, to be

$$
\left(\sum_{i=1}^{n} M\left(\left\|A_{i, *, *}\right\|_{F}\right)\right)^{1 / p}
$$

where $A_{i, *, *}$ is the $i$-th face of $A$ (along the 1 st direction), and $p$ is a parameter associated with the function $M()$, which defines a nice $M$-Estimator.

Definition 21.12.3 (Matrix $v$-norm). For an $n \times d$ matrix $A$, we define the $v$-norm of $A$, denoted $\|A\|_{v}$, to be

$$
\sum_{i=1}^{n} M\left(\left\|A_{i, *}\right\|_{2}\right)^{1 / p}
$$

where $A_{i, *}$ is the $i$-th row of $A$, and $p$ is a parameter associated with the function $M()$, which defines a nice $M$-Estimator.

Definition 21.12.4. Given matrix $A \in \mathbb{R}^{n \times d}$, let $A_{i, *}$ denote the $i$-th row of $A$. Let $T_{S} \subset[n]$ denote the indices $i$ such that $e_{i}$ is chosen for $S$. Using a probability vector $q$ and a sampling and rescaling matrix $S \in \mathbb{R}^{n \times n}$ from $q$, we will estimate $\|A\|_{v}$ using $S$ and a re-weighted version, $\|S \cdot\|_{v, w^{\prime}}$ of $\|\cdot\|_{v}$, with

$$
\|S A\|_{v, w^{\prime}}=\left(\sum_{i \in T_{S}} w_{i}^{\prime} M\left(\left\|A_{i, *}\right\|_{2}\right)\right)^{1 / p}
$$

where $w_{i}^{\prime}=w_{i} / q_{i}$. Since $w^{\prime}$ is generally understood, we will usually just write $\|S A\|_{v}$. We will also need an "entrywise row-weighted" version :

$$
\||S A|\|=\left(\sum_{i \in T_{S}} \frac{w_{i}}{q_{i}}\left\|A_{i, *}\right\|_{M}^{p}\right)^{1 / p}=\left(\sum_{i \in T_{S}, j \in[d]} \frac{w_{i}}{q_{i}} M\left(A_{i, j}\right)\right)^{1 / p}
$$

where $A_{i, j}$ denotes the entry in the $i$-th row and $j$-th column of $A$.
Fact 21.12.1. For $p=1$, for any two matrices $A$ and $B$, we have $\|A+B\|_{v} \leq\|A\|_{v}+\|B\|_{v}$. For any two tensors $A$ and $B$, we have $\|A+B\|_{v} \leq\|A\|_{v}+\|B\|_{v}$.

### 21.12.2.2 Sampling and rescaling sketches

Note that Lemmas 42 and 44 in [CW15a] are stronger than stated. In particular, we do not need to assume $X$ is a square matrix. For any $m \geq z$, if $X \in \mathbb{R}^{d \times m}$, then we have the same result.

Lemma 21.12.2 (Lemma 42 in [CW15a]). Let $\rho>0$ and integer $z>0$. For sampling matrix $S$, suppose for a given $y \in \mathbb{R}^{d}$ with failure probability $\delta$ it holds that $\|S A y\|_{M}=$ $(1 \pm 1 / 10)\|A y\|_{M}$. There is $K_{1}=O\left(z^{2} / C_{M}\right)$ so that with failure probability $\delta\left(K_{\mathcal{N}} / C_{M}\right)^{(1+p) d}$, for a constant $K_{\mathcal{N}}$, any rank-z matrix $X \in \mathbb{R}^{d \times m}$ has the property that if $\|A X\|_{v} \geq K_{1} \rho$, then $\|S A X\|_{v} \geq \rho$, and that if $\|A X\|_{v} \leq \rho / K_{1}$, then $\|S A X\|_{v} \leq \rho$.

Lemma 21.12.3 (Lemma 44 in [CW15a]). Let $\delta, \rho>0$ and integer $z>0$. Given matrix $A \in$ $\mathbb{R}^{n \times d}$, there exists a sampling and rescaling matrix $S \in \mathbb{R}^{n \times n}$ with $r=O\left(\gamma(A, M, w) \epsilon^{-2} d z^{2} \log (z / \epsilon) \log (1 / \delta)\right.$, nonzero entries such that, with probability at least $1-\delta$, for any rank-z matrix $X \in \mathbb{R}^{d \times m}$, we have either

$$
\|S A X\|_{v} \geq \rho
$$

or

$$
(1-\epsilon)\|A X\|_{v}-\epsilon \rho \leq\|S A X\|_{v} \leq(1+\epsilon)\|A X\|_{v}+\epsilon \rho .
$$

Lemma 21.12.4 (Lemma 43 in [CW15a]). For $r>0$, let $\widehat{r}=r / \gamma(A, M, w)$, and let $q \in \mathbb{R}^{n}$ have

$$
q_{i}=\min \left\{1, \widehat{r} \gamma_{i}(A, M, w)\right\} .
$$

Let $S$ be a sampling and rescaling matrix generated using $q$, with weights as usual $w_{i}^{\prime}=w_{i} / q_{i}$. Let $W \in \mathbb{R}^{d \times z}$, and $\delta>0$. There is an absolute constant $C$ so that for $\widehat{r} \geq C z \log (1 / \delta) / \epsilon^{2}$, with probability at least $1-\delta$, we have

$$
(1-\epsilon)\|A W\|_{v, w} \leq\|S A W\|_{v, w^{\prime}} \leq(1+\epsilon)\|A W\|_{v, w}
$$

### 21.12.2.3 No dilation and no contraction

Lemma 21.12.5. Given matrices $A \in \mathbb{R}^{n \times m}, U \in \mathbb{R}^{n \times d}$, let $V^{*}=\underset{\operatorname{rank}-k V \in \mathbb{R}^{d \times m}}{\arg \min }\|U V-A\|_{v}$. If $S \in \mathbb{R}^{s \times n}$ has at most $c_{1}$-dilation on $U V^{*}-A$, i.e.,

$$
\left\|S\left(U V^{*}-A\right)\right\|_{v} \leq c_{1}\left\|U V^{*}-A\right\|_{v}
$$

and it has at most $c_{2}$-contraction on $U$, i.e.,

$$
\forall x \in \mathbb{R}^{d},\|S U x\|_{v} \geq \frac{1}{c_{2}}\|U x\|_{v}
$$

then $S$ has at most $\left(c_{2}, c_{1}+\frac{1}{c_{2}}\right)$-contraction on $(U, A)$, i.e.,

$$
\forall \operatorname{rank}-k V \in \mathbb{R}^{d \times m},\|S U V-S A\|_{v} \geq \frac{1}{c_{2}}\|U V-A\|_{v}-\left(c_{1}+\frac{1}{c_{2}}\right)\left\|U V^{*}-A\right\|_{v} .
$$

Proof. Let $A \in \mathbb{R}^{n \times m}, U \in \mathbb{R}^{n \times d}$ and $S \in \mathbb{R}^{s \times n}$ be the same as that described in the lemma. Let $\left(V-V^{*}\right)_{j}$ denote the $j$-th column of $V-V^{*}$. Then $\forall \operatorname{rank}-k V \in \mathbb{R}^{d \times m}$,

$$
\begin{aligned}
\|S U V-S A\|_{v} & \geq\left\|S U V-S U V^{*}\right\|_{v}-\left\|S U V^{*}-S A\right\|_{v} \\
& \geq\left\|S U V-S U V^{*}\right\|_{v}-c_{1}\left\|U V^{*}-A\right\|_{v} \\
& =\left\|S U\left(V-V^{*}\right)\right\|_{v}-c_{1}\left\|U V^{*}-A\right\|_{v} \\
& =\sum_{j=1}^{m}\left\|S U\left(V-V^{*}\right)_{j}\right\|_{v}-c_{1}\left\|U V^{*}-A\right\|_{v} \\
& \geq \sum_{j=1}^{m} \frac{1}{c_{2}}\left\|U\left(V-V^{*}\right)_{j}\right\|_{v}-c_{1}\left\|U V^{*}-A\right\|_{v} \\
& =\frac{1}{c_{2}}\left\|U V-U V^{*}\right\|_{v}-c_{1}\left\|U V^{*}-A\right\|_{v} \\
& \geq \frac{1}{c_{2}}\|U V-A\|_{v}-\frac{1}{c_{2}}\left\|U V^{*}-A\right\|_{v}-c_{1}\left\|U V^{*}-A\right\|_{v} \\
& =\frac{1}{c_{2}}\|U V-A\|_{v}-\left(\left(\frac{1}{c_{2}}+c_{2}\right)\left\|U V^{*}-A\right\|_{v}\right),
\end{aligned}
$$

where the first inequality follows by the triangle inequality, the second inequality follows since $S$ has at most $c_{1}$ dilation on $U V^{*}-A$, the third inequality follows since $S$ has at most $c_{2}$ contraction on $U$, and the fourth inequality follows by the triangle inequality.

Claim 21.12.6. Given matrix $A \in \mathbb{R}^{n \times m}$, for any distribution $p=\left(p_{1}, p_{2}, \cdots, p_{n}\right)$ define random variable $X$ such that $X=\left\|A_{i}\right\|_{2} / p_{i}$ with probability $p_{i}$ where $A_{i}$ is the $i$-th row of matrix $A$. Then take $m$ independent samples $X^{1}, X^{2}, \cdots, X^{m}$, and let $Y=\frac{1}{m} \sum_{j=1}^{m} X^{j}$. We have

$$
\operatorname{Pr}\left[Y \leq 1000\|A\|_{v}\right] \geq .999
$$

Proof. We can compute the expectation of $X^{j}$, for any $j \in[m]$,

$$
\mathbb{E}\left[X^{j}\right]=\sum_{i=1}^{n} \frac{\left\|A_{i}\right\|_{2}}{p_{i}} \cdot p_{i}=\|A\|_{v}
$$

Then $\mathbb{E}[Y]=\frac{1}{m} \sum_{j=1}^{m} \mathbb{E}\left[X^{j}\right]=\|A\|_{v}$. Using Markov's inequality, we have

$$
\operatorname{Pr}\left[Y \geq\|A\|_{v}\right] \leq .001
$$

Lemma 21.12.7. For any fixed $U^{*} \in \mathbb{R}^{n \times d}$ and rank- $k V^{*} \in \mathbb{R}^{d \times m}$ with $d=\operatorname{poly}(k)$, there exists an algorithm that takes poly $(n, d)$ time to compute a sampling and rescaling diagonal matrix $S \in \mathbb{R}^{n \times n}$ with $s=\operatorname{poly}(k)$ nonzero entries such that, with probability at least .999, we have: for all rank-k $V \in \mathbb{R}^{d \times m}$,

$$
\left\|U^{*} V^{*}-U^{*} V\right\|_{v} \lesssim\left\|S U^{*} V^{*}-S U^{*} V\right\|_{v} \lesssim\left\|U^{*} V^{*}-U^{*} V\right\|_{v}
$$

Lemma 21.12.8 (No dilation). Given matrices $A \in \mathbb{R}^{n \times m}, U^{*} \in \mathbb{R}^{n \times d}$ with $d=\operatorname{poly}(k)$, define $V^{*} \in \mathbb{R}^{d \times m}$ to be the optimal solution $\min _{\operatorname{rank}-k \in \mathbb{R}^{d \times m}}\left\|U^{*} V-A\right\|_{v}$. Choose a sampling and rescaling diagonal matrix $S \in \mathbb{R}^{n \times n}$ with $s=\operatorname{poly}(k)$ according to Lemma 21.12.4. Then with probability at least .99, we have: for all rank-k $V \in \mathbb{R}^{d \times m}$,

$$
\left\|S U^{*} V-S A\right\|_{v} \lesssim\left\|U^{*} V^{*}-U^{*} V\right\|_{v}+O(1)\left\|U^{*} V^{*}-A\right\|_{v} \lesssim\left\|U^{*} V-A\right\|_{v}
$$

Proof. Using Claim 21.12.6 and Lemma 21.12.7, we have with probability at least .99, for all rank- $k V \in \mathbb{R}^{d \times m}$,

$$
\begin{aligned}
& \left\|S U^{*} V-S A\right\|_{v} \\
\leq & \text { by triangle inequality } \\
\lesssim & \left\|S U^{*} V-S U^{*} V-S U^{*} V^{*}\right\|_{v}+O(1)\left\|U^{*} V^{*}-A\right\|_{v} \\
\lesssim\left\|U^{*} V-U^{*} V^{*}\right\|_{v}+O(1)\left\|U^{*} V^{*}-A\right\|_{v} & \text { by Claim } 21.12 .6 \\
\lesssim\left\|U^{*} V-A\right\|_{v}+\left\|U^{*} V^{*}-A\right\|_{v}+O(1)\left\|U^{*} V^{*}-A\right\|_{v} & \text { by triangle inequality } \\
\lesssim & \left\|U^{*} V-A\right\|_{v}
\end{aligned}
$$

Lemma 21.12.9 (No contraction). Given matrices $A \in \mathbb{R}^{n \times m}, U^{*} \in \mathbb{R}^{n \times d}$ with $d=\operatorname{poly}(k)$, define $V^{*} \in \mathbb{R}^{d \times m}$ to be the optimal solution $\min _{\text {rank }-k \in \mathbb{R}^{d \times m}}\left\|U^{*} V-A\right\|_{v}$. Choose a sampling and rescaling diagonal matrix $S \in \mathbb{R}^{n \times n}$ with $s=\operatorname{poly}(k)$ according to Lemma 21.12.4. Then with probability at least .99, we have: for all rank- $k V \in \mathbb{R}^{d \times m}$,

$$
\left\|U^{*} V-A\right\|_{v} \lesssim\left\|S U^{*} V-S A\right\|_{v}+O(1)\left\|U^{*} V^{*}-A\right\|_{v}
$$

Proof. This follows by Lemma 21.12.5, Claim 21.12.6 and Lemma 21.12.8.

### 21.12.2.4 Oblivious sketches, MSketch

In this section, we recall a concept called $M$-sketches for $M$-estimators which is defined in [CW15b]. $M$-sketch is an oblivious sketch for matrices.

Theorem 21.12.10 (Theorem 3.1 in [CW15b]). Let OPT denote $\min _{x \in \mathbb{R}^{d}}\|A x-b\|_{G}$. There is an algorithm that in $O(\mathrm{nnz}(A))+\operatorname{poly}(d \log n)$ time, with constant probability finds $x^{\prime}$ such that $\left\|A x^{\prime}-b\right\|_{G} \leq O(1) \mathrm{OPT}$.

Definition 21.12.5 (M-Estimator sketches or MSketch [CW15b]). Given parameters $N, n, m, b>1$, define $h_{\max }=\left\lfloor\log _{b}(n / m)\right\rfloor, \beta=\left(b-b^{-h_{\max }}\right) /(b-1)$ and $s=N h_{\max }$. For each $p \in[n], \sigma_{p}, g_{p}, h_{p}$ are generated (independently) in the following way,

$$
\begin{array}{lr}
\sigma_{p} \leftarrow \pm 1, & \text { chosen with equal probability, } \\
g_{p} \in[N], & \text { chosen with equal probability }, \\
h_{p} \leftarrow t, & \text { chosen with probability } 1 /\left(\beta b^{t}\right) \text { for } t \in\left\{0,1, \cdots h_{\max }\right\} .
\end{array}
$$

For each $p \in[n]$, we define $j_{p}=g_{p}+N h_{p}$. Let $w \in \mathbb{R}^{s}$ denote the scaling vector such that, for each $j \in[s]$,

$$
w_{j}= \begin{cases}\beta b^{h_{p}}, & \text { if there exists } p \in[n] \text { s.t. } j=j_{p} \\ 0 & \text { otherwise }\end{cases}
$$

Let $\bar{S} \in \mathbb{R}^{N h_{\max } \times n}$ be such that, for each $j \in[s]$,for each $p \in[n]$,

$$
\bar{S}_{j, p}= \begin{cases}\sigma_{p}, & \text { if } j=g_{p}+N \cdot h_{p} \\ 0, & \text { otherwise }\end{cases}
$$

Let $D_{w}$ denote the diagonal matrix where the $i$-th entry on the diagonal is the $i$-th entry of $w$. Let $S=D_{w} \bar{S}$. We say $(\bar{S}, w)$ or $S$ is an MSketch.

Definition 21.12.6 (Tensor $\left\|\|_{v, w^{-}}\right.$-norm). For a tensor $A \in \mathbb{R}^{d \times n_{1} \times n_{2}}$ and a vector $w \in$, we define

$$
\|A\|_{v, w}=\sum_{i=1}^{d} w_{i}\left\|A_{i, *, *}\right\|_{F}
$$

Let $(\bar{S}, w)$ denote an MSketch, and let $S=D_{w} \bar{S}$. If $v$ corresponds to a scaleinvariant M-Estimator, then for any three matrices $U, V, W$, we have the following,

$$
\|(\bar{S} U) \otimes V \otimes W\|_{v, w}=\left\|\left(D_{w} \bar{S} U\right) \otimes V \otimes W\right\|_{v}=\|(S U) \otimes V \otimes W\|_{v}
$$

Fact 21.12.11. For a tensor $A \in \mathbb{R}^{n \times n \times n}$, let $S \in \mathbb{R}^{s \times n}$ denote an MSketch (defined in 21.12.5) with $s=\operatorname{poly}(k, \log n)$. Then $S A$ can be computed in $O(\operatorname{nnz}(A))$ time.

Lemma 21.12.12. For any fixed $U^{*} \in \mathbb{R}^{n \times d}$ and rank- $k V^{*} \in \mathbb{R}^{d \times m}$ with $d=\operatorname{poly}(k)$, let $S \in \mathbb{R}^{s \times n}$ denote an MSketch (defined in Definition 21.12.5) with $s=\operatorname{poly}(k, \log n)$ rows. Then with probability at least .999, we have: for all rank- $k V \in \mathbb{R}^{d \times m}$,

$$
\left\|U^{*} V^{*}-U^{*} V\right\|_{v} \lesssim\left\|S U^{*} V^{*}-S U^{*} V\right\|_{v} \lesssim\left\|U^{*} V^{*}-U^{*} V\right\|_{v}
$$

Lemma 21.12.13 (No dilation, Theorem 3.4 in [CW15b]). Given matrices $A \in \mathbb{R}^{n \times m}, U^{*} \in$ $\mathbb{R}^{n \times d}$ with $d=\operatorname{poly}(k)$, define $V^{*} \in \mathbb{R}^{d \times m}$ to be the optimal solution to $\min _{\text {rank }-k V \in \mathbb{R}^{d \times m}} \| U^{*} V-$ $A \|_{v}$. Choose an MSketch $S \in \mathbb{R}^{s \times n}$ with $s=\operatorname{poly}(k, \log n)$ according to Definition 21.12.5. Then with probability at least .99, we have: for all rank-k $V \in \mathbb{R}^{d \times m}$,

$$
\left\|S U^{*} V-S A\right\|_{v} \lesssim\left\|U^{*} V^{*}-U^{*} V\right\|_{v}+O(1)\left\|U^{*} V^{*}-A\right\|_{v} \lesssim\left\|U^{*} V-A\right\|_{v}
$$

Lemma 21.12.14 (No contraction). Given matrices $A \in \mathbb{R}^{n \times m}, U^{*} \in \mathbb{R}^{n \times d}$ with $d=$ $\operatorname{poly}(k)$, define $V^{*} \in \mathbb{R}^{d \times m}$ to be the optimal solution to $\min _{\text {rank }-k V \in \mathbb{R}^{d \times m}}\left\|U^{*} V-A\right\|_{v}$. Choose an MSketch $S \in \mathbb{R}^{s \times n}$ with $s=\operatorname{poly}(k, \log n)$ according to Definition 21.12.5. Then with probability at least .99, we have: for all rank- $k V \in \mathbb{R}^{d \times m}$,

$$
\left\|U^{*} V-A\right\|_{v} \lesssim\left\|S U^{*} V-S A\right\|_{v}+O(1)\left\|U^{*} V^{*}-A\right\|_{v}
$$

### 21.12.2.5 Running time analysis

Lemma 21.12.15. Given a tensor $A \in \mathbb{R}^{n \times d \times d}$, let $S \in \mathbb{R}^{s \times n}$ denote an MSkETCH with $s$ rows. Let $S A$ denote a tensor that has size $s \times d \times d$. For each $i \in\{2,3\}$, let $(S A)_{i} \in \mathbb{R}^{d \times d s}$ denote a matrix obtained by flattening tensor SA along the $i$-th dimension. For each $i \in$ $\{2,3\}$, let $S_{i} \in \mathbb{R}^{d s \times s_{i}}$ denote a CountSketch transform with $s_{i}$ columns. For each $i \in\{2,3\}$, let $T_{i} \in \mathbb{R}^{t_{i} \times d}$ denote a CountSketch transform with $t_{i}$ rows. Then
(I) For each $i \in\{2,3\},(S A)_{i} S_{i}$ can be computed in $O(\mathrm{nnz}(A))$ time.
(II) For each $i \in\{2,3\}, T_{i}(S A)_{i} S_{i}$ can be computed in $O(\mathrm{nnz}(A))$ time.

Proof. Proof of Part (I). First note that $(S A)_{2} S_{2}$ has size $n \times S_{2}$. Thus for each $i \in[d], j \in\left[s_{2}\right]$, we have,

$$
\begin{array}{rlr}
\left((S A)_{2} S_{2}\right)_{i, j} & =\sum_{x^{\prime}=1}^{d s}\left((S A)_{2}\right)_{i, x^{\prime}}\left(S_{2}\right)_{x^{\prime}, j} & \text { by }(S A)_{2} \in \mathbb{R}^{d \times d s}, S_{2} \in \mathbb{R}^{d s \times s_{2}} \\
& =\sum_{y=1}^{d} \sum_{z=1}^{s}\left((S A)_{2}\right)_{i,(y-1) s+z}\left(S_{2}\right)_{(y-1) s+z, j} & \quad \text { by unflattening } \\
& =\sum_{y=1}^{d} \sum_{z=1}^{s}(S A)_{z, i, y}\left(S_{2}\right)_{(y-1) s+z, j} & \\
& =\sum_{y=1}^{d} \sum_{z=1}^{s}\left(\sum_{x=1}^{n} S_{z, x} A_{x, i, y}\right)\left(S_{2}\right)_{(y-1) s+z, j} \\
& =\sum_{y=1}^{d} \sum_{z=1}^{s} \sum_{x=1}^{n} S_{z, x} \cdot A_{x, i, y} \cdot\left(S_{2}\right)_{(y-1) s+z, j}
\end{array}
$$

For each nonzero entry $A_{x, i, y}$, there is only one $z$ such that $S_{z, x}$ is nonzero. Thus there is only one $j$ such that $\left(S_{2}\right)_{(y-1) s+z, j}$ is nonzero. It means that $A_{x, i, y}$ can only affect one entry of $\left((S A)_{2} S_{2}\right)_{i, j}$. Thus, $(S A)_{2} S_{2}$ can be computed in $O(\operatorname{nnz}(A))$ time. Similarly, we can compute $(S A)_{3} S_{3}$ in $O(\mathrm{nnz}(A))$ time.

Proof of Part (II). Note that $T_{2}(S A)_{2} S_{2}$ has size $t_{2} \times s_{2}$. Thus for each $i \in\left[t_{2}\right], j \in\left[s_{2}\right]$, we have,

$$
\begin{array}{rlr}
\left(T_{2}(S A)_{2} S_{2}\right)_{i, j} & =\sum_{x=1}^{d} \sum_{y^{\prime}=1}^{d s}\left(T_{2}\right)_{i, x}\left((S A)_{2}\right)_{x, y^{\prime}}\left(S_{2}\right)_{y^{\prime}, j} & \text { by }(S A)_{2} \in \mathbb{R}^{d \times d s} \\
& =\sum_{x=1}^{d} \sum_{y=1}^{d} \sum_{z=1}^{s}\left(T_{2}\right)_{i, x}\left((S A)_{2}\right)_{x,(y-1) s+z}\left(S_{2}\right)_{(y-1) s+z, j} & \\
& =\sum_{x=1}^{d} \sum_{y=1}^{d} \sum_{z=1}^{s}\left(T_{2}\right)_{i, x}(S A)_{z, x, y}\left(S_{2}\right)_{(y-1) s+z, j} & \text { by unflattening } \\
& =\sum_{x=1}^{d} \sum_{y=1}^{d} \sum_{z=1}^{s}\left(T_{2}\right)_{i, x}\left(\sum_{w=1}^{n} S_{z, w} A_{w, x, y}\right)\left(S_{2}\right)_{(y-1) s+z, j} \\
& =\sum_{x=1}^{d} \sum_{y=1}^{d} \sum_{z=1}^{s} \sum_{w=1}^{n}\left(T_{2}\right)_{i, x} \cdot S_{z, w} \cdot A_{w, x, y} \cdot\left(S_{2}\right)_{(y-1) s+z, j}
\end{array}
$$

For each nonzero entry $A_{w, x, y}$, there is only one $z$ such that $S_{z, w}$ is nonzero. There is only one $i$ such that $\left(T_{2}\right)_{i, x}$ is nonzero. Since there is only one $z$ to make $S_{z, w}$ nonzero, there is only one $j$, such that $\left(S_{2}\right)_{(y-1) s+z, j}$ is nonzero. Thus, $T_{2}(S A)_{2} S_{2}$ can be computed in $O(\mathrm{nnz}(A))$ time. Similarly, we can compute $T_{3}(S A)_{3} S_{3}$ in $O(\mathrm{nnz}(A))$ time.

### 21.12.2.6 Algorithms

We first give a "warm-up" algorithm in Theorem 21.12.16 by using a sampling and rescaling matrix. Then we improve the running time to be polynomial in all the parameters by using an oblivious sketch, and thus we obtain Theorem 21.12.17.

Theorem 21.12.16. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, let $r=O\left(k^{2}\right)$. There exists an algorithm which takes $n^{\mathrm{poly}(k)}$ time and outputs three matrices $U, V, W \in \mathbb{R}^{n \times r}$

```
Algorithm \(21.36 \ell_{1}\)-Frobenius \(\left(\ell_{1}-\ell_{2}-\ell_{2}\right)\) Low-rank Approximation Algorithm, poly \((k)\) ap-
proximation
procedure L122TensorLowRankApprox \((A, n, k) \quad \triangleright\) Theorem 21.12.16
        \(\epsilon \leftarrow 0.1\).
        \(s \leftarrow \operatorname{poly}(k, 1 / \epsilon)\).
        Guess a sampling and rescaling matrix \(S \in \mathbb{R}^{s \times n}\).
        \(s_{2} \leftarrow s_{3} \leftarrow O(k / \epsilon)\).
        \(r \leftarrow s_{2} s_{3}\).
        Choose sketching matrices \(S_{2} \in \mathbb{R}^{n s \times s_{2}}, S_{3} \in \mathbb{R}^{n s \times s_{3}}\).
        Compute \((S A)_{2} S_{2},(S A)_{3} S_{3}\).
        Form \(\widetilde{V} \in \mathbb{R}^{n \times r}\) by repeating \((S A)_{2} S_{2} s_{3}\) times according to Equation (21.64).
        Form \(\widetilde{W} \in \mathbb{R}^{n \times r}\) by repeating \((S A)_{3} S_{3} s_{2}\) times according to Equation (21.65).
        Form objective function \(\min _{U \in \mathbb{R}^{n \times r}}\left\|U \cdot\left(\widetilde{V}^{\top} \odot \widetilde{W}^{\top}\right)-A_{1}\right\|_{F}\).
        Use a linear regression solver to find a solution \(\widetilde{U}\).
        Take the best solution found over all guesses.
        return \(\widetilde{U}, \widetilde{V}, \widetilde{W}\).
    end procedure
```

such that

$$
\|U \otimes V \otimes W-A\|_{v} \leq \operatorname{poly}(k) \min _{\operatorname{rank}-k A^{\prime}}\left\|A^{\prime}-A\right\|_{v},
$$

holds with probability at least 9/10.

Proof. We define OPT as follows,

$$
\mathrm{OPT}=\min _{U, V, W \in \mathbb{R}^{n \times k}}\|U \otimes V \otimes W-A\|_{v}=\min _{U, V, W \in \mathbb{R}^{n \times k}}\left\|\sum_{i=1}^{k} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{v} .
$$

Let $A_{1} \in \mathbb{R}^{n \times n^{2}}$ denote the matrix obtained by flattening tensor $A$ along the 1 st dimension. Let $U^{*} \in \mathbb{R}^{n \times k}$ denote the optimal solution. We fix $U^{*} \in \mathbb{R}^{n \times k}$, and consider this objective function,

$$
\begin{equation*}
\min _{V, W \in \mathbb{R}^{n \times k}}\left\|U^{*} \otimes V \otimes W-A\right\|_{v} \equiv \min _{V, W \in \mathbb{R}^{n \times k}}\left\|U^{*} \cdot\left(V^{\top} \odot W^{\top}\right)-A_{1}\right\|_{v} \tag{21.57}
\end{equation*}
$$

which has cost at most OPT, and where $V^{\top} \odot W^{\top} \in \mathbb{R}^{k \times n^{2}}$ denotes the matrix for which the $i$-th row is a vectorization of $V_{i} \otimes W_{i}, \forall i \in[k]$. (Note that $V_{i} \in \mathbb{R}^{n}$ is the $i$-th column of matrix $V \in \mathbb{R}^{n \times k}$ ). Choose a sampling and rescaling diagonal matrix $S \in \mathbb{R}^{n \times n}$ according to $U^{*}$, which has $s=\operatorname{poly}(k)$ non-zero entries. Using $S$ to sketch on the left of the objective function when $U^{*}$ is fixed (Equation (21.57)), we obtain a smaller problem,

$$
\begin{equation*}
\min _{V, W \in \mathbb{R}^{n \times k}}\left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{v} \equiv \min _{V, W \in \mathbb{R}^{n \times k}}\left\|S U^{*} \cdot\left(V^{\top} \odot W^{\top}\right)-S A_{1}\right\|_{v} \tag{21.58}
\end{equation*}
$$

Let $V^{\prime}, W^{\prime}$ denote the optimal solution to the above problem, i.e.,

$$
V^{\prime}, W^{\prime}=\underset{V, W \in \mathbb{R}^{n \times k}}{\arg \min }\left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{v}
$$

Then using properties (no dilation Lemma 21.12 .8 and no contraction Lemma 21.12.9) of $S$, we have

$$
\left\|U^{*} \otimes V^{\prime} \otimes W^{\prime}-A\right\|_{v} \leq \alpha \mathrm{OPT}
$$

where $\alpha$ is an approximation ratio determined by $S$.
By definition of $\|\cdot\|_{v}$ and $\|\cdot\|_{2} \leq\|\cdot\|_{1} \leq \sqrt{\operatorname{dim}}\|\cdot\|_{2}$, we can rewrite Equation (21.58) in the following way,

$$
\begin{align*}
& \left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{v} \\
= & \sum_{i=1}^{s}\left(\sum_{j=1}^{n} \sum_{l=1}^{n}\left(\left(\left(S U^{*}\right) \otimes V \otimes W\right)_{i, j, l}-(S A)_{i, j, l}\right)^{2}\right)^{\frac{1}{2}} \\
\leq & \sqrt{s}\left(\sum_{i=1}^{s} \sum_{j=1}^{n} \sum_{l=1}^{n}\left(\left(\left(S U^{*}\right) \otimes V \otimes W\right)_{i, j, l}-(S A)_{i, j, l}\right)^{2}\right)^{\frac{1}{2}} \\
= & \sqrt{s}\left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{F} . \tag{21.59}
\end{align*}
$$

Given the above properties of $S$ and Equation (21.59), for any $\beta \geq 1$, let $V^{\prime \prime}, W^{\prime \prime}$ denote a $\beta$-approximate solution of $\min _{V, W \in \mathbb{R}^{n \times k}}\left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{F}$, i.e.,

$$
\begin{equation*}
\left\|\left(S U^{*}\right) \otimes V^{\prime \prime} \otimes W^{\prime \prime}-S A\right\|_{F} \leq \beta \cdot \min _{V, W \in \mathbb{R}^{n \times k}}\left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{F} \tag{21.60}
\end{equation*}
$$

Then,

$$
\begin{equation*}
\left\|U^{*} \otimes V^{\prime \prime} \otimes W^{\prime \prime}-A\right\|_{v} \leq \sqrt{s} \alpha \beta \cdot \mathrm{OPT} \tag{21.61}
\end{equation*}
$$

In the next few paragraphs we will focus on solving Equation (21.60). We start by fixing $W^{*} \in \mathbb{R}^{n \times k}$ to be the optimal solution of

$$
\min _{V, W \in \mathbb{R}^{n \times k}}\left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{F}
$$

We use $(S A)_{2} \in \mathbb{R}^{n \times n s}$ to denote the matrix obtained by flattening the tensor $S A \in \mathbb{R}^{s \times n \times n}$ along the second direction. We use $Z_{2}=\left(S U^{*}\right)^{\top} \odot\left(W^{*}\right)^{\top} \in \mathbb{R}^{k \times n s}$ to denote the matrix where the $i$-th row is the vectorization of $\left(S U^{*}\right)_{i} \otimes W_{i}^{*}$. We can consider the following objective function,

$$
\min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2}-(S A)_{2}\right\|_{F} .
$$

Choosing a sketching matrix $S_{2} \in \mathbb{R}^{n s \times s_{2}}$ with $s_{2}=O(k / \epsilon)$ gives a smaller problem,

$$
\min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2} S_{2}-(S A)_{2} S_{2}\right\|_{F} .
$$

Letting $\widehat{V}=(S A)_{2} S_{2}\left(Z_{2} S_{2}\right)^{\dagger} \in \mathbb{R}^{n \times k}$, then

$$
\begin{array}{rlr}
\left\|\widehat{V} Z_{2}-(S A)_{2}\right\|_{F} & \leq(1+\epsilon) \min _{V \in \mathbb{R}^{n \times k}}\left\|V Z_{2}-(S A)_{2}\right\|_{F} & \\
& =(1+\epsilon) \min _{V \in \mathbb{R}^{n \times k}}\left\|V\left(\left(S U^{*}\right)^{\top} \odot\left(W^{*}\right)^{\top}\right)-(S A)_{2}\right\|_{F} & \\
& =(1+\epsilon) \min _{V \in \mathbb{R}^{n \times k}}\left\|\left(S U^{*}\right) \otimes V \otimes W^{*}-S A\right\|_{F} & \text { by unflattening } \\
& =(1+\epsilon) \min _{V, W \in \mathbb{R}^{n \times k}}\left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{F} . & \text { by definition of } W^{*} \tag{21.62}
\end{array}
$$

We define $D_{2} \in \mathbb{R}^{n^{2} \times n^{2}}$ to be a diagonal matrix obrained by copying the $n \times n$ identity matrix $s$ times on $n$ diagonal blocks of $D_{2}$. Then it has $n s$ nonzero entries. Thus, $D_{2}$ also can be thought of as a matrix that has size $n^{2} \times n s$.

We can think of $(S A)_{2} S_{2} \in \mathbb{R}^{n \times s_{2}}$ as follows,

$$
\begin{aligned}
(S A)_{2} S_{2} & =(A(S, I, I))_{2} S_{2} \\
& =\underbrace{A_{2}}_{n \times n^{2}} \cdot \underbrace{D_{2}}_{n^{2} \times n^{2}} \cdot \underbrace{S_{2}}_{n s \times s_{2}} \text { by } D_{2} \text { can be thought of as having size } n^{2} \times n s \\
& =A_{2} \cdot\left[\begin{array}{llll}
c_{2,1} I_{n} & & & \\
& c_{2,2} I_{n} & & \\
& & \ddots & \\
& & & c_{2, n} I_{n}
\end{array}\right] \cdot S_{2}
\end{aligned}
$$

where $I_{n}$ is an $n \times n$ identity matrix, $c_{2, i} \geq 0$ for each $i \in[n]$, and the number of nonzero $c_{2, i}$ is $s$.

For the last step, we fix $S U^{*}$ and $\widehat{V}$. We use $(S A)_{3} \in \mathbb{R}^{n \times n s}$ to denote the matrix obtained by flattening the tensor $S A \in \mathbb{R}^{s \times n \times n}$ along the third direction. We use $Z_{3}=$ $\left(S U^{*}\right)^{\top} \odot \widehat{V}^{\top} \in \mathbb{R}^{k \times n s}$ to denote the matrix where the $i$-th row is the vectorization of
$\left(S U^{*}\right)_{i} \otimes \widehat{V}_{i}$. We can consider the following objective function,

$$
\min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3}-(S A)_{3}\right\|_{F}
$$

Choosing a sketching matrix $S_{3} \in \mathbb{R}^{n s \times s_{3}}$ with $s_{3}=O(k / \epsilon)$ gives a smaller problem,

$$
\min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3} S_{3}-(S A)_{3} S_{3}\right\|_{F}
$$

Let $\widehat{W}=(S A)_{3} S_{3}\left(Z_{3} S_{3}\right)^{\dagger} \in \mathbb{R}^{n \times k}$. Then

$$
\begin{aligned}
\left\|\widehat{W} Z_{3}-(S A)_{3}\right\|_{F} & \leq(1+\epsilon) \min _{W \in \mathbb{R}^{n \times k}}\left\|W Z_{3}-(S A)_{3}\right\|_{F} & \text { by property of } S_{3} \\
& =(1+\epsilon) \min _{W \in \mathbb{R}^{n \times k}}\left\|W\left(\left(S U^{*}\right)^{\top} \odot \widehat{V}^{\top}\right)-(S A)_{3}\right\|_{F} & \text { by definition } Z_{3} \\
& =(1+\epsilon) \min _{W \in \mathbb{R}^{n \times k}}\left\|\left(S U^{*}\right) \otimes \widehat{V} \otimes W-S A\right\|_{F} & \text { by unflattening } \\
& \leq(1+\epsilon)^{2}\left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{F} . & \text { by Equation }(21.62)
\end{aligned}
$$

We define $D_{3} \in \mathbb{R}^{n^{2} \times n^{2}}$ to be a diagonal matrix formed by copying the $n \times n$ identity matrix $s$ times on $n$ diagonal blocks of $D_{3}$. Then it has $n s$ nonzero entries. Thus, $D_{3}$ also can be thought of as a matrix that has size $n^{2} \times n s$ and $D_{3}$ is uniquely determined by $S$.

Similarly as to the 2 nd dimension, for the 3 rd dimension, we can think of $(S A)_{3} S_{3}$ as follows,

$$
\begin{aligned}
(S A)_{3} S_{3} & =(A(S, I, I))_{3} S_{3} \\
& =\underbrace{A_{3}}_{n \times n^{2}} \cdot \underbrace{D_{3}}_{n^{2} \times n^{2}} \cdot \underbrace{S_{3}}_{n s \times s_{3}} \\
& =A_{3} \cdot\left[\begin{array}{llll}
c_{3,1} I_{n} & & \\
& c_{3,2} I_{n} & & \\
& & \ddots & \\
& & & c_{3, n} I_{n}
\end{array}\right] \cdot S_{3}
\end{aligned}
$$

$$
=\underbrace{A_{3}}_{n \times n^{2}} \cdot \underbrace{D_{3}}_{n^{2} \times n^{2}} \cdot \underbrace{S_{3}}_{n s \times s_{3}} \quad \text { by } D_{3} \text { can be thought of as having size } n^{2} \times n s
$$

where $I_{n}$ is an $n \times n$ identity matrix, $c_{3, i} \geq 0$ for each $i \in[n]$ and the number of nonzero $c_{3, i}$ is $s$.

Overall, we have proved that,

$$
\begin{equation*}
\min _{X_{2}, X_{3}}\left\|\left(S U^{*}\right) \otimes\left(A_{2} D_{2} S_{2} X_{2}\right) \otimes\left(A_{3} D_{3} S_{3} X_{3}\right)-S A\right\|_{F} \leq(1+\epsilon)^{2}\left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{F}, \tag{21.63}
\end{equation*}
$$

where diagonal matrix $D_{2} \in \mathbb{R}^{n^{2} \times n^{2}}$ (with $n s$ nonzero entries) and $D_{3} \in \mathbb{R}^{n^{2} \times n^{2}}$ (with $n s$ nonzero entries) are uniquely determined by diagonal matrix $S \in \mathbb{R}^{n \times n}$ ( $s$ nonzero entries). Let $X_{2}^{\prime}$ and $X_{3}^{\prime}$ denote the optimal solution to the above problem (Equation (21.63)). Let $V^{\prime \prime}=\left(A_{2} D_{2} S_{2} X_{2}^{\prime}\right) \in \mathbb{R}^{n \times k}$ and $W^{\prime \prime}=\left(A_{3} D_{3} S_{3} X_{3}^{\prime}\right) \in \mathbb{R}^{n \times k}$. Then we have

$$
\left\|U^{*} \otimes V^{\prime \prime} \otimes W^{\prime \prime}-A\right\|_{v} \leq \sqrt{s} \alpha \beta \mathrm{OPT}
$$

We construct matrix $\widetilde{V} \in \mathbb{R}^{n \times s_{2} s_{3}}$ by copying matrix $(S A)_{2} S_{2} \in \mathbb{R}^{n \times s_{2}} s_{3}$ times,

$$
\tilde{V}=\left[\begin{array}{llll}
(S A)_{2} S_{2} & (S A)_{2} S_{2} & \cdots & (S A)_{2} S_{2} . \tag{21.64}
\end{array}\right]
$$

We construct matrix $\widetilde{W} \in \mathbb{R}^{n \times s_{2} s_{3}}$ by copying the $i$-th column of matrix $(S A)_{3} S_{3} \in \mathbb{R}^{n \times s_{3}}$ into $(i-1) s_{2}+1, \cdots, i s_{2}$ columns of $\widetilde{W}$,

$$
\widetilde{W}=\left[\begin{array}{llll}
\left((S A)_{3} S_{3}\right)_{1} \cdots\left((S A)_{3} S_{3}\right)_{1} & \left((S A)_{3} S_{3}\right)_{2} \cdots\left((S A)_{3} S_{3}\right)_{2} & \cdots & \left((S A)_{3} S_{3}\right)_{s_{3}} \cdots\left((S A)_{3} S_{3}\right)_{s_{3}} \cdot \tag{21.65}
\end{array}\right]
$$

Although we don't know $S$, we can guess all of the possibilities. For each possibility,
we can find a solution $\widetilde{U} \in \mathbb{R}^{n \times s_{2} s_{3}}$ to the following problem,

$$
\begin{aligned}
& \min _{U \in \mathbb{R}^{n \times s_{2} s_{3}}}\left\|\sum_{i=1}^{s_{2}} \sum_{j=1}^{s_{3}} U_{(i-1) s_{3}+j} \otimes\left((S A)_{2} S_{2}\right)_{i} \otimes\left((S A)_{3} S_{3}\right)_{j}-A\right\|_{v} \\
= & \min _{U \in \mathbb{R}^{n \times s_{2} s_{3}}}\left\|\sum_{i=1}^{s_{2}} \sum_{j=1}^{s_{3}} U_{(i-1) s_{3}+j} \cdot \operatorname{vec}\left(\left((S A)_{2} S_{2}\right)_{i} \otimes\left((S A)_{3} S_{3}\right)_{j}\right)-A_{1}\right\|_{v} \\
= & \min _{U \in \mathbb{R}^{n \times s_{2} s_{3}}}\left\|\sum_{i=1}^{s_{2}} \sum_{j=1}^{s_{3}} U_{(i-1) s_{3}+j} \cdot\left(\widetilde{V}^{\top} \odot \widetilde{W}^{\top}\right)^{(i-1) s_{3}+j}-A_{1}\right\|_{v} \\
= & \min _{U \in \mathbb{R}^{n \times s_{2} s_{3}}}\left\|U \cdot\left(\widetilde{V}^{\top} \odot \widetilde{W}^{\top}\right)-A_{1}\right\|_{v} \\
= & \min _{U \in \mathbb{R}^{n \times s_{2} s_{3}}}\left\|U Z-A_{1}\right\|_{v} \\
= & \min _{U \in \mathbb{R}^{n \times s_{2} s_{3}}} \sum_{i=1}^{s_{2} s_{3}}\left\|U^{i} Z-A_{1}^{i}\right\|_{2},
\end{aligned}
$$

where the first step follows by flattening the tensor along the 1 st dimension, $U_{(i-1) s_{3}+j}$ denotes the $(i-1) s_{3}+j$-th column of $U \in \mathbb{R}^{n \times s_{2} s_{3}}, A_{1} \in \mathbb{R}^{n \times n^{2}}$ denotes the matrix obtained by flattening tensor $A$ along the 1st dimension, the second step follows since $\widetilde{V}^{\top} \odot \widetilde{W}^{\top} \in \mathbb{R}^{s_{2} s_{3} \in n^{2}}$ is defined to be the matrix where the $(i-1) s_{3}+j$-th row is vectorization of $\left((S A)_{2} S_{2}\right)_{i} \otimes$ $\left((S A)_{3} S_{3}\right)_{j}$, the fourth step follows by defining $Z$ to be $\widetilde{V}^{\top} \odot \widetilde{W}^{\top}$, and the last step follows by definition of $\|\cdot\|_{v}$ norm. Thus, we obtain a multiple regression problem and it can be solved directly by using [CW13, NN13a].

Finally, we take the best $\widetilde{U}, \widetilde{V}, \widetilde{W}$ over all the guesses. The entire running time is dominated by the number of guesses, which is $n^{\operatorname{poly}(k)}$. This completes the proof.

Theorem 21.12.17. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, let $r=O\left(k^{2}\right)$. There exists an algorithm which takes $O(\mathrm{nnz}(A))+n \operatorname{poly}(k, \log n)$ time and outputs three

```
Algorithm \(21.37 \ell_{1}\)-Frobenius \(\left(\ell_{1}-\ell_{2}-\ell_{2}\right)\) Low-rank Approximation Algorithm, poly \((k, \log n)\)
approximation
    procedure L122TensorLowRankApprox \((A, n, k) \quad \triangleright\) Theorem 21.12.17
        \(\epsilon \leftarrow 0.1\).
        \(s \leftarrow \operatorname{poly}(k, \log n)\).
        Choose \(S \in \mathbb{R}^{s \times n}\) to be an MSketch. \(\triangleright\) Definition 21.12.5
        \(s_{2} \leftarrow s_{3} \leftarrow O(k / \epsilon)\).
        \(t_{2} \leftarrow t_{3} \leftarrow \operatorname{poly}(k / \epsilon)\).
        \(r \leftarrow s_{2} s_{3}\).
        Choose sketching matrices \(S_{2} \in \mathbb{R}^{n s \times s_{2}}, S_{3} \in \mathbb{R}^{n s \times s_{3}}\).
        Choose sketching matrices \(T_{2} \in \mathbb{R}^{t_{2} \times n}, T_{3} \in \mathbb{R}^{t_{3} \times n}\).
        Compute \((S A)_{2} S_{2},(S A)_{3} S_{3}\).
        Compute \(T_{2}(S A)_{2} S_{2}, T_{3}(S A)_{3} S_{3}\).
        Form \(\widetilde{V} \in \mathbb{R}^{n \times r}\) by repeating \((S A)_{2} S_{2} s_{3}\) times according to Equation (21.74).
        Form \(\widetilde{W} \in \mathbb{R}^{n \times r}\) by repeating \((S A)_{3} S_{3} s_{2}\) times according to Equation (21.75).
        Form \(\bar{V} \in \mathbb{R}^{t_{2} \times r}\) by repeating \(T_{2}(S A)_{2} S_{2} s_{3}\) times according to Equation (21.72).
    Form \(\bar{W} \in \mathbb{R}^{t_{3} \times r}\) by repeating \(T_{3}(S A)_{3} S_{3} s_{2}\) times according to Equation (21.73).
    \(C \leftarrow A\left(I, T_{2}, T_{3}\right)\).
    Form objective function \(\min _{U \in \mathbb{R}^{n \times r}}\left\|U \cdot\left(\bar{V}^{\top} \odot \bar{W}^{\top}\right)-C_{1}\right\|_{F}\).
    Use linear regression solver to find a solution \(\widetilde{U}\).
    return \(\widetilde{U}, \widetilde{V}, \widetilde{W}\).
    end procedure
```

matrices $U, V, W \in \mathbb{R}^{n \times r}$ such that

$$
\|U \otimes V \otimes W-A\|_{v} \leq \operatorname{poly}(k, \log n) \min _{\text {rank }-k}\left\|A^{\prime} A^{\prime}-A\right\|_{v}
$$

holds with probability at least 9/10.

Proof. We define OPT as follows,

$$
\mathrm{OPT}=\min _{U, V, W \in \mathbb{R}^{n \times k}}\|U \otimes V \otimes W-A\|_{v}=\min _{U, V, W \in \mathbb{R}^{n \times k}}\left\|\sum_{i=1}^{k} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{v} .
$$

Let $A_{1} \in \mathbb{R}^{n \times n^{2}}$ denote the matrix obtained by flattening tensor $A$ along the 1 st dimension. Let $U^{*} \in \mathbb{R}^{n \times k}$ denote the optimal solution. We fix $U^{*} \in \mathbb{R}^{n \times k}$, and consider the objective function,

$$
\begin{equation*}
\min _{V, W \in \mathbb{R}^{n \times k}}\left\|U^{*} \otimes V \otimes W-A\right\|_{v} \equiv \min _{V, W \in \mathbb{R}^{n \times k}}\left\|U^{*} \cdot\left(V^{\top} \odot W^{\top}\right)-A_{1}\right\|_{v}, \tag{21.66}
\end{equation*}
$$

which has cost at most OPT, and where $V^{\top} \odot W^{\top} \in \mathbb{R}^{k \times n^{2}}$ denotes the matrix for which the $i$-th row is a vectorization of $V_{i} \otimes W_{i}, \forall i \in[k]$. (Note that $V_{i} \in \mathbb{R}^{n}$ is the $i$-th column of matrix $V \in \mathbb{R}^{n \times k}$ ). Choose an (oblivious) MSketch $S \in \mathbb{R}^{s \times n}$ with $s=\operatorname{poly}(k, \log n)$ according to Definition 21.12.5. Using MSketch $S, w$ to sketch on the left of the objective function when $U^{*}$ is fixed (Equation (21.66)), we obtain a smaller problem,

$$
\begin{equation*}
\min _{V, W \in \mathbb{R}^{n \times k}}\left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{v} \equiv \min _{V, W \in \mathbb{R}^{n \times k}}\left\|S U^{*} \cdot\left(V^{\top} \odot W^{\top}\right)-S A_{1}\right\|_{v} \tag{21.67}
\end{equation*}
$$

Let $V^{\prime}, W^{\prime}$ denote the optimal solution to the above problem, i.e.,

$$
V^{\prime}, W^{\prime}=\underset{V, W \in \mathbb{R}^{n \times k}}{\arg \min }\left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{v}
$$

Then using properties (no dilation Lemma 21.12.13 and no contraction Lemma 21.12.14) of $S$, we have

$$
\left\|U^{*} \otimes V^{\prime} \otimes W^{\prime}-A\right\|_{v} \leq \alpha \mathrm{OPT}
$$

where $\alpha$ is an approximation ratio determined by $S$.
By definition of $\|\cdot\|_{v}$ and $\|\cdot\|_{2} \leq\|\cdot\|_{1} \leq \sqrt{\operatorname{dim}}\|\cdot\|_{2}$, we can rewrite Equation (21.67)
in the following way,

$$
\begin{align*}
& \left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{v} \\
= & \sum_{i=1}^{s}\left(\sum_{j=1}^{n} \sum_{l=1}^{n}\left(\left(\left(S U^{*}\right) \otimes V \otimes W\right)_{i, j, l}-(S A)_{i, j, l}\right)^{2}\right)^{\frac{1}{2}} \\
\leq & \sqrt{s}\left(\sum_{i=1}^{s} \sum_{j=1}^{n} \sum_{l=1}^{n}\left(\left(\left(S U^{*}\right) \otimes V \otimes W\right)_{i, j, l}-(S A)_{i, j, l}\right)^{2}\right)^{\frac{1}{2}} \\
= & \sqrt{s}\left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{F} \tag{21.68}
\end{align*}
$$

Using the properties of $S$ and Equation (21.68), for any $\beta \geq 1$, let $V^{\prime \prime}, W^{\prime \prime}$ denote a $\beta$ approximation solution of $\min _{V, W \in \mathbb{R}^{n \times k}}\left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{F}$, i.e.,

$$
\begin{equation*}
\left\|\left(S U^{*}\right) \otimes V^{\prime \prime} \otimes W^{\prime \prime}-S A\right\|_{F} \leq \beta \cdot \min _{V, W \in \mathbb{R}^{n \times k}}\left\|\left(S U^{*}\right) \otimes V \otimes W-S A\right\|_{F} \tag{21.69}
\end{equation*}
$$

Then,

$$
\begin{equation*}
\left\|U^{*} \otimes V^{\prime \prime} \otimes W^{\prime \prime}-A\right\|_{v} \leq \sqrt{s} \alpha \beta \cdot \mathrm{OPT} \tag{21.70}
\end{equation*}
$$

Let $\widehat{A}$ denote $S A$. Choose $S_{i} \in \mathbb{R}^{n s \times s_{i}}$ to be Gaussian matrix with $s_{i}=O(k / \epsilon)$, $\forall i\{2,3\}$. By a similar proof as in Theorem 21.12.16, we have if $X_{2}^{\prime}, X_{3}^{\prime}$ is a $\beta$-approximate solution to

$$
\min _{X_{2}, X_{3}}\left\|\left(S U^{*}\right) \otimes\left(\widehat{A}_{2} S_{2} X_{2}\right) \otimes\left(\widehat{A}_{3} S_{3} X_{3}\right)-S A\right\|_{F},
$$

then,

$$
\left\|U^{*} \otimes\left(\widehat{A}_{2} S_{2} X_{2}\right) \otimes\left(\widehat{A}_{3} S_{3} X_{3}\right)-A\right\|_{v} \leq \sqrt{s} \alpha \beta
$$

To reduce the size of the objective function from $\operatorname{poly}(n)$ to $\operatorname{poly}(k / \epsilon)$, we use perform an "input sparsity reduction" (in Lemma 21.4.3). Note that, we do not need to use this idea to optimize the running time in Theorem 21.12.16. The running time of Theorem 21.12.16 is dominated by guessing sampling and rescaling matrices. (That running time is $\gg \mathrm{nnz}(A)$.) Choose $T_{i} \in \mathbb{R}^{t_{i} \times n}$ to be a sparse subspace embedding matrix (CountSketch transform) with $t_{i}=\operatorname{poly}(k, 1 / \epsilon), \forall i \in\{2,3\}$. Applying the proof of Lemma 21.4.3 here, we obtain, if $X_{2}^{\prime}, X_{3}^{\prime}$ is a $\beta$-approximate solution to

$$
\min _{X_{2}, X_{3}}\left\|\left(S U^{*}\right) \otimes\left(T_{2}(S A)_{2} S_{2} X_{2}\right) \otimes\left(T_{3}(S A)_{3} S_{3} X_{3}\right)-S A\right\|_{F},
$$

then,

$$
\begin{equation*}
\left\|U^{*} \otimes\left((S A)_{2} S_{2} X_{2}\right) \otimes\left((S A)_{3} S_{3} X_{3}\right)-A\right\|_{v} \leq \sqrt{s} \alpha \beta . \tag{21.71}
\end{equation*}
$$

Similar to the bicriteria results in Section 21.4.4, Equation (21.71) indicates that we can construct a bicriteria solution by using two matrices $(S A)_{2} S_{2}$ and $(S A)_{3} S_{3}$. The next question is how to obtain the final results $\widehat{U}, \widehat{V}, \widehat{W}$. We first show how to obtain $\widehat{U}$. Then we show to construct $\widehat{V}$ and $\widehat{W}$.

To obtain $\widehat{U}$, we need to solve a regression problem related to two matrices $\bar{V}, \widehat{W}$ and a tensor $A\left(I, T_{2}, T_{3}\right)$. We construct matrix $\bar{V} \in \mathbb{R}^{t_{2} \times s_{2} s_{3}}$ by copying matrix $T_{2}(S A)_{2} S_{2} \in \mathbb{R}^{t_{2} \times s_{2}}$ $s_{3}$ times,

$$
\bar{V}=\left[\begin{array}{llll}
T_{2}(S A)_{2} S_{2} & T_{2}(S A)_{2} S_{2} & \cdots & T_{2}(S A)_{2} S_{2} \tag{21.72}
\end{array}\right] .
$$

We construct matrix $\bar{W} \in \mathbb{R}^{t_{3} \times s_{2} s_{3}}$ by copying the $i$-th column of matrix $T_{3}(S A)_{3} S_{3} \in \mathbb{R}^{t_{3} \times s_{3}}$ into $(i-1) s_{2}+1, \cdots, i s_{2}$ columns of $\bar{W}$,

$$
\bar{W}=\left[\begin{array}{llll}
F_{1} \cdots F_{1} & F_{2} \cdots F_{2} & \cdots & F_{s_{3}} \cdots F_{s_{3}} \tag{21.73}
\end{array}\right],
$$

where $F=T_{3}(S A)_{3} S_{3}$.
Thus, to obtain $\widetilde{U} \in \mathbb{R}^{s_{2} s_{3}}$, we just need to use a linear regression solver to solve a smaller problem,

$$
\min _{U \in \mathbb{R}^{s^{2} s_{3}}}\left\|U \cdot\left(\bar{V}^{\top} \odot \bar{W}^{\top}\right)-A\left(I, T_{2}, T_{3}\right)\right\|_{F},
$$

which can be solved in $O(\operatorname{nnz}(A))+n \operatorname{poly}(k, \log n)$ time. We will show how to obtain $\widetilde{V}$ and $\widetilde{W}$.

We construct matrix $\tilde{V} \in \mathbb{R}^{n \times s_{2} s_{3}}$ by copying matrix $(S A)_{2} S_{2} \in \mathbb{R}^{n \times s_{2}} s_{3}$ times,

$$
\tilde{V}=\left[\begin{array}{llll}
(S A)_{2} S_{2} & (S A)_{2} S_{2} & \cdots & (S A)_{2} S_{2} . \tag{21.74}
\end{array}\right]
$$

We construct matrix $\widetilde{W} \in \mathbb{R}^{n \times s_{2} s_{3}}$ by copying the $i$-th column of matrix $(S A)_{3} S_{3} \in \mathbb{R}^{n \times s_{3}}$ into $(i-1) s_{2}+1, \cdots, i s_{2}$ columns of $\widetilde{W}$,

$$
\widetilde{W}=\left[\begin{array}{llll}
F_{1} \cdots F_{1} & F_{2} \cdots F_{2} & \cdots & F_{s_{3}} \cdots F_{s_{3}} \tag{21.75}
\end{array}\right]
$$

where $F=(S A)_{3} S_{3}$.

### 21.12.3 $\ell_{1}-\ell_{1}-\ell_{2}$ norm

Section 21.12.3.1 presents some definitions and useful facts for the tensor $\ell_{1}-\ell_{1}-\ell_{2}$ norm. We provide some tools in Section 21.12.3.2. Section 21.12.3.3 presents a key idea which shows we are able to reduce the original problem to a new problem under entrywise $\ell_{1}$ norm. Section 21.12.3.4 presents several existence results. Finally, Section 21.12.3.6 introduces several algorithms with different tradeoffs.

### 21.12.3.1 Definitions

Definition 21.12.7. (Tensor $u$-norm) For an $n \times n \times n$ tensor $A$, we define the $u$-norm of $A$, denoted $\|A\|_{u}$, to be

$$
\left(\sum_{i=1}^{n} \sum_{j=1}^{n} M\left(\left\|A_{i, j, *}\right\|_{2}\right)\right)^{1 / p}
$$

where $A_{i, j, *}$ is the $(i, j)$-th tube of $A$, and $p$ is a parameter associated with the function $M()$, which defines a nice $M$-Estimator.

Definition 21.12.8. (Matrix $u$-norm) For an $n \times n$ matrix $A$, we define $u$-norm of $A$, denoted $\|A\|_{u}$, to be

$$
\left(\sum_{i=1}^{n} M\left(\left\|A_{i, *}\right\|_{2}\right)\right)^{1 / p}
$$

where $A_{i, *}$ is the $i$-th row of $A$, and $p$ is a parameter associated with the function $M()$, which defines a nice $M$-Estimator.

Fact 21.12.18. For $p=1$, for any two matrices $A$ and $B$, we have $\|A+B\|_{u} \leq\|A\|_{u}+\|B\|_{u}$. For any two tensors $A$ and $B$, we have $\|A+B\|_{u} \leq\|A\|_{u}+\|B\|_{u}$.

### 21.12.3.2 Projection via Gaussians

Definition 21.12.9. Let $p \geq 1$. Let $\ell_{p}^{g^{n-1}}$ be an infinite dimensional $\ell_{p}$ metric which consists of a coordinate for each vector $r$ in the unit sphere $\mathfrak{S}^{n-1}$. Define function $f: S^{n-1} \rightarrow \mathbb{R}$. The $\ell_{1}$-norm of any such $f$ is defined as follows:

$$
\|f\|_{1}=\left(\int_{r \in \delta^{n-1}}|f(r)|^{p} \mathrm{~d} r\right)^{1 / p}
$$

Claim 21.12.19. Let $f_{v}(r)=\langle v, r\rangle$. There exists a universal constant $\alpha_{p}$ such that

$$
\left\|f_{v}\right\|_{p}=\alpha_{p}\|v\|_{2}
$$

Proof. We have,

$$
\begin{aligned}
\left\|f_{v}\right\|_{p} & =\left(\int_{r \in \delta^{n-1}}|\langle v, r\rangle|^{p} \mathrm{~d} r\right)^{1 / p} \\
& =\left(\int_{\theta \in \delta^{n-1}}\|v\|_{2}^{p} \cdot|\cos \theta|^{p} \mathrm{~d} \theta\right)^{1 / p} \\
& =\|v\|_{2}\left(\int_{\theta \in \delta^{n-1}}|\cos \theta|^{p} \mathrm{~d} \theta\right)^{1 / p} \\
& =\alpha_{p}\|v\|_{2} .
\end{aligned}
$$

This completes the proof.
Lemma 21.12.20. Let $G \in \mathbb{R}^{k \times n}$ denote i.i.d. random Gaussian matrices with rescaling.
Then for any $v \in \mathbb{R}^{n}$, we have

$$
\operatorname{Pr}\left[(1-\epsilon)\|v\|_{2} \leq\|G v\|_{1} \leq(1+\epsilon)\|v\|_{2}\right] \geq 1-2^{-\Omega\left(k \epsilon^{2}\right)} .
$$

Proof. For each $i \in[k]$, we define $X_{i}=\left\langle v, g_{i}\right\rangle$, where $g_{i} \in \mathbb{R}^{n}$ is the $i$-th row of $G$. Then $X_{i}=\sum_{j=1}^{n} v_{j} g_{i, j}$ and $\mathbb{E}\left[\left|X_{i}\right|\right]=\alpha_{p}\|v\|_{2}$. Define $Y=\sum_{i=1}^{k}\left|X_{i}\right|$. We have $\mathbb{E}[Y]=k \alpha_{1}\|v\|_{2}=$ $k \alpha_{1}$.

We can show

$$
\begin{array}{rlr}
\operatorname{Pr}\left[Y \geq(1+\epsilon) \alpha_{1} k\right] & =\operatorname{Pr}\left[e^{s Y} \geq e^{s(1+\epsilon) \alpha_{1} k}\right] & \text { for all } s>0 \\
& \leq \mathbb{E}\left[e^{s Y}\right] / e^{s(1+\epsilon) \alpha_{1} k} & \text { by Markov's inequality } \\
& =e^{-s(1+\epsilon) \alpha_{1} k} \cdot \mathbb{E}\left[\prod_{i=1}^{k} e^{s\left|X_{i}\right|}\right] & \text { by } Y=\sum_{i=1}^{k}\left|X_{i}\right| \\
& =e^{-s(1+\epsilon) \alpha_{1} k} \cdot\left(\mathbb{E}\left[e^{s\left|X_{1}\right|}\right]\right)^{k} &
\end{array}
$$

It remains to bound $\mathbb{E}\left[e^{s\left|X_{1}\right|}\right]$. Since $X_{1} \sim \mathcal{N}(0,1)$, we have that $X_{1}$ has density function $e^{-t^{2} / 2}$. Thus, we have,

$$
\begin{array}{rlr}
\mathbb{E}\left[e^{s\left|X_{1}\right|}\right] & =\frac{1}{\sqrt{2 \pi}} \int_{-\infty}^{+\infty} e^{s|t|} \cdot e^{-t^{2} / 2} \mathrm{~d} t \\
& =\frac{1}{\sqrt{2 \pi}} \int_{-\infty}^{+\infty} e^{s^{2} / 2} \cdot e^{-(|t|-s)^{2} / 2} \mathrm{~d} t & \\
& =e^{s^{2} / 2}(\operatorname{erf}(s / \sqrt{2})+1) & \\
& \leq e^{s^{2} / 2}\left(\left(1-\exp \left(-2 s^{2} / \pi\right)\right)^{1 / 2}+1\right) & \text { by } 1-\exp \left(-4 x^{2} / \pi\right) \geq \operatorname{erf}(x)^{2} \\
& \leq e^{s^{2} / 2}(\sqrt{2 / \pi} s+1) . & \text { by } 1-e^{-x} \leq x
\end{array}
$$

Thus, we have

$$
\begin{aligned}
\operatorname{Pr}\left[Y \geq(1+\epsilon) \alpha_{1} k\right] & \leq e^{-s(1+\epsilon) k} e^{k s^{2} / 2}(1+s \sqrt{2 / \pi})^{k} \\
& =e^{-s(1+\epsilon) \alpha_{1} k} e^{k s^{2} / 2} e^{k \cdot \log (1+s \sqrt{2 / \pi})} \\
& \leq e^{-s(1+\epsilon) \alpha_{1} k+k s^{2} / 2+k \cdot s \sqrt{2 / \pi}}
\end{aligned}
$$

$$
\leq e^{-\Omega\left(k \epsilon^{2}\right)} . \quad \text { by } \alpha_{1} \geq \sqrt{2 / \pi} \text { and setting } s=\epsilon
$$

Lemma 21.12.21. For any $\epsilon \in(0,1)$, let $k=O\left(n / \epsilon^{2}\right)$. Let $G \in \mathbb{R}^{k \times n}$ denote i.i.d. random Gaussian matrices with rescaling. Then for any $v \in \mathbb{R}^{n}$, with probability at least $1-2^{-\Omega\left(n / \epsilon^{2}\right)}$, we have : for all $v \in \mathbb{R}^{n}$,

$$
(1-\epsilon)\|v\|_{2} \leq\|G v\|_{1} \leq(1+\epsilon)\|v\|_{2} .
$$

Proof. Let $\mathcal{S}$ denote $\left\{y \in \mathbb{R}^{n} \mid\|y\|_{2}=1\right\}$. We construct a $\gamma$-net so that for all $y \in \mathcal{S}$, there exists a vector $w \in \mathcal{N}$ for which $\|y-w\|_{2} \leq \gamma$. We set $\gamma=1 / 2$.

For any unit vector $y$, we can write

$$
y=y^{0}+y^{1}+y^{2}+\cdots
$$

where $\left\|y^{i}\right\|_{2} \leq 1 / 2^{i}$ and $y^{i}$ is a scalar multiple of a vector in $\mathcal{N}$. Thus, we have

$$
\begin{aligned}
\|G y\|_{1} & =\left\|G\left(y^{0}+y^{1}+y^{2}+\cdots\right)\right\|_{1} \\
& \leq \sum_{i=0}^{\infty}\left\|G y^{i}\right\|_{1} \\
& \leq \sum_{i=0}^{\infty}(1+\epsilon)\left\|y^{i}\right\|_{2} \\
& \leq \sum_{i=0}^{\infty}(1+\epsilon) \frac{1}{2^{i}} \\
& \leq 1+\Theta(\epsilon)
\end{aligned}
$$

Similarly, we can lower bound $\|G y\|_{1}$ by $1-\Theta(\epsilon)$. By Lemma 2.2 in [Woo14b], we know that for any $\gamma \in(0,1)$, there exists a $\gamma$-net $\mathcal{N}$ of $\mathcal{S}$ for which $|\mathcal{N}| \leq(1+4 / \gamma)^{n}$.

### 21.12.3.3 Reduction, projection to high dimension

Lemma 21.12.22. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, let $S \in \mathbb{R}^{n \times s}$ denote a Gaussian matrix with $s=O\left(n / \epsilon^{2}\right)$ columns. With probability at least $1-2^{-\Omega\left(n / \epsilon^{2}\right)}$, for any $U, V, W \in$
$\mathbb{R}^{n \times k}$, we have

$$
(1-\epsilon)\|U \otimes V \otimes W-A\|_{u} \leq\|(U \otimes V \otimes W) S-A S\|_{1} \leq(1+\epsilon)\|U \otimes V \otimes W-A\|_{u}
$$

Proof. By definition of the $\otimes$ product between matrices and • product between a tensor and a matrix, we have $(U \otimes V \otimes W) S=U \otimes V \otimes(S W) \in \mathbb{R}^{n \times n \times s}$. We use $A_{i, j, *} \in \mathbb{R}^{n}$ to denote the $(i, j)$-th tube (the column in the 3rd dimension) of tensor $A$. We first prove the upper bound,

$$
\begin{aligned}
\|(U \otimes V \otimes W) S-A S\|_{1} & =\sum_{i=1}^{n} \sum_{j=1}^{n}\left\|\left((U \otimes V \otimes W)_{i, j, *}-A_{i, j, *}\right) S\right\|_{1} \\
& \leq \sum_{i=1}^{n} \sum_{j=1}^{n}(1+\epsilon)\left\|(U \otimes V \otimes W)_{i, j, *}-A_{i, j, *}\right\|_{2} \\
& =(1+\epsilon)\|U \otimes V \otimes W-A\|_{u}
\end{aligned}
$$

where the first step follows by definition of tensor $\|\cdot\|_{u}$ norm, the second step follows by Lemma 21.12.21, and the last step follows by tensor entry-wise $\ell_{1}$ norm. Similarly, we can prove the lower bound,

$$
\begin{aligned}
\|(U \otimes V \otimes W) S-A S\|_{1} & \geq \sum_{i=1}^{n} \sum_{j=1}^{n}(1-\epsilon)\left\|(U \otimes V \otimes W)_{i, j, *}-A_{i, j, *}\right\|_{2} \\
& =(1-\epsilon)\|U \otimes V \otimes W-A\|_{u}
\end{aligned}
$$

This completes the proof.
Corollary 21.12 .23 . For any $\alpha \geq 1$, if $U^{\prime}, V^{\prime}, W^{\prime}$ satisfy

$$
\left\|\left(U^{\prime} \otimes V^{\prime} \otimes W^{\prime}-A\right) S\right\|_{1} \leq \gamma \min _{\text {rank }-k A_{k}}\left\|\left(A_{k}-A\right) S\right\|_{1}
$$

then

$$
\left\|U^{\prime} \otimes V^{\prime} \otimes W^{\prime}-A\right\|_{u} \leq \gamma \frac{1+\epsilon}{1-\epsilon} \min _{\operatorname{rank}-k}\left\|A_{k}-A\right\|_{u}
$$

Proof. Let $\widehat{U}, \widehat{V}, \widehat{W}$ denote the optimal solution to $\min _{\text {rank }-k A_{k}}\left\|\left(A_{k}-A\right) S\right\|_{1}$. Let $U^{*}, V^{*}, W^{*}$ denote the optimal solution to $\min _{\text {rank }-k A_{k}}\left\|A_{k}-A\right\|_{u}$. Then,

$$
\begin{aligned}
\left\|U^{\prime} \otimes V^{\prime} \otimes W^{\prime}-A\right\|_{u} & \leq \frac{1}{1-\epsilon}\left\|\left(U^{\prime} \otimes V^{\prime} \otimes W^{\prime}-A\right) S\right\|_{1} \\
& \leq \gamma \frac{1}{1-\epsilon}\|(\widehat{U} \otimes \widehat{V} \otimes \widehat{W}-A) S\|_{1} \\
& \leq \gamma \frac{1}{1-\epsilon}\left\|\left(U^{*} \otimes V^{*} \otimes W^{*}-A\right) S\right\|_{1} \\
& \leq \gamma \frac{1+\epsilon}{1-\epsilon}\left\|U^{*} \otimes V^{*} \otimes W^{*}-A\right\|_{u}
\end{aligned}
$$

which completes the proof.

### 21.12.3.4 Existence results

Theorem 21.12.24 (Existence results). Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$ and a matrix $S \in \mathbb{R}^{n \times \bar{n}}$, let OPT denote $\min _{\mathrm{rank}-k} A_{k} \in \mathbb{R}^{n \times n \times n}\left\|\left(A_{k}-A\right) S\right\|_{1}$, let $\widehat{A}=A S \in \mathbb{R}^{n \times n \times \bar{n}}$. For any $k \geq 1$, there exist three matrices $S_{1} \in \mathbb{R}^{n \bar{n} \times s_{1}}, S_{2} \in \mathbb{R}^{n \bar{n} \times s_{2}}, S_{3} \in \mathbb{R}^{n^{2} \times s_{3}}$ such that

$$
\min _{X_{1} \in \mathbb{R}^{s_{1} \times k}, X_{2} \in \mathbb{R}^{s_{2} \times k}, X_{3} \in \mathbb{R}^{s_{3} \times k}}\left\|\left(\widehat{A}_{1} S_{1} X_{1}\right) \otimes\left(\widehat{A}_{2} S_{2} X_{2}\right) \otimes\left(\widehat{A}_{3} S_{3} X_{3}\right)-\widehat{A}\right\|_{1} \leq \alpha \mathrm{OPT},
$$

or equivalently,

$$
\min _{X_{1} \in \mathbb{R}^{s_{1} \times k}, X_{2} \in \mathbb{R}^{s_{2} \times k}, X_{3} \in \mathbb{R}^{s_{3} \times k}}\left\|\left(\left(\widehat{A}_{1} S_{1} X_{1}\right) \otimes\left(\widehat{A}_{2} S_{2} X_{2}\right) \otimes\left(A_{3} S_{3} X_{3}\right)-A\right) S\right\|_{1} \leq \alpha \mathrm{OPT},
$$

holds with probability 99/100.
(I). Using a dense Cauchy transform, $s_{1}=s_{2}=s_{3}=\widetilde{O}(k), \alpha=\widetilde{O}\left(k^{1.5}\right) \log ^{3} n$.
(II). Using a sparse Cauchy transform, $s_{1}=s_{2}=s_{3}=\widetilde{O}\left(k^{5}\right), \alpha=\widetilde{O}\left(k^{13.5}\right) \log ^{3} n$.
(III). Guessing Lewis weights, $s_{1}=s_{2}=s_{3}=\widetilde{O}(k), \alpha=\widetilde{O}\left(k^{1.5}\right)$.

Proof. We use OPT to denote the optimal cost,

$$
\mathrm{OPT}:=\min _{\text {rank }-k} A_{A_{k} \in \mathbb{R}^{n \times n \times n}}\left\|\left(A_{k}-A\right) S\right\|_{1} .
$$

We fix $V^{*} \in \mathbb{R}^{n \times k}$ and $W^{*} \in \mathbb{R}^{n \times k}$ to be the optimal solution to

$$
\min _{U, V, W}\|(U \otimes V \otimes W-A) S\|_{1} .
$$

We define $Z_{1} \in \mathbb{R}^{k \times n \bar{n}}$ to be the matrix where the $i$-th row is the vectorization of $V_{i}^{*} \otimes\left(S W_{i}^{*}\right)$. We define tensor

$$
\widehat{A}=A S \in \mathbb{R}^{n \times n \times \bar{n}}
$$

Then we also have $\widehat{A}=A(I, I, S)$ according to the definition of the $\cdot$ product between a tensor and a matrix.

Let $\widehat{A}_{1} \in \mathbb{R}^{n \times n \bar{n}}$ denote the matrix obtained by flattening tensor $\widehat{A}$ along the first direction. We can consider the following optimization problem,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1}-\widehat{A}_{1}\right\|_{1} .
$$

Choosing $S_{1}$ to be one of the following sketching matrices:
(I) a dense Cauchy transform,
(II) a sparse Cauchy transform,
(III) a sampling and rescaling diagonal matrix according to Lewis weights.

Let $\alpha_{S_{1}}$ denote the approximation ratio produced by the sketching matrix $S_{1}$. We use $S_{1} \in \mathbb{R}^{n \bar{n} \times s_{1}}$ to sketch on right of the above problem, and obtain the problem:

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1} S_{1}-\widehat{A}_{1} S_{1}\right\|_{1}=\min _{U \in \mathbb{R}^{n \times k}} \sum_{i=1}^{n}\left\|U^{i} Z_{1} S_{1}-\left(\widehat{A}_{1} S_{1}\right)^{i}\right\|_{1},
$$

where $U^{i}$ denotes the $i$-th row of matrix $U \in \mathbb{R}^{n \times k}$ and $\left(\widehat{A}_{1} S_{1}\right)^{i}$ denotes the $i$-th row of matrix $\widehat{A}_{1} S_{1}$. Instead of solving it under $\ell_{1}$-norm, we consider the $\ell_{2}$-norm relaxation,

$$
\min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1} S_{1}-\widehat{A}_{1} S_{1}\right\|_{F}^{2}=\min _{U \in \mathbb{R}^{n \times k}} \sum_{i=1}^{n}\left\|U^{i} Z_{1} S_{1}-\left(\widehat{A}_{1} S_{1}\right)^{i}\right\|_{2}^{2}
$$

Let $\widehat{U} \in \mathbb{R}^{n \times k}$ denote the optimal solution of the above optimization problem, so that $\widehat{U}=\widehat{A}_{1} S_{1}\left(Z_{1} S_{1}\right)^{\dagger}$. We plug $\widehat{U}$ into the objective function under the $\ell_{1}$-norm. By the property of sketching matrix $S_{1} \in \mathbb{R}^{n \bar{n} \times s_{1}}$, we have,

$$
\left\|\widehat{U} Z_{1}-\widehat{A}_{1}\right\|_{1} \leq \alpha_{S_{1}} \min _{U \in \mathbb{R}^{n \times k}}\left\|U Z_{1}-\widehat{A}_{1}\right\|_{1}=\alpha_{S_{1}} \mathrm{OPT}
$$

which implies that,

$$
\left\|\widehat{U} \otimes V^{*} \otimes\left(S W^{*}\right)-\widehat{A}\right\|_{1}=\left\|\left(\widehat{U} \otimes V^{*} \otimes W^{*}\right) S-\widehat{A}\right\|_{1} \leq \alpha_{S_{1}} \mathrm{OPT}
$$

In the second step, we fix $\widehat{U} \in \mathbb{R}^{n \times k}$ and $W^{*} \in \mathbb{R}^{n \times k}$. Let $\widehat{A}_{2} \in \mathbb{R}^{n \times n \bar{n}}$ denote the matrix obtained by flattening tensor $\widehat{A} \in \mathbb{R}^{n \times n \times n}$ along the second direction. We choose a sketching matrix $S_{2} \in \mathbb{R}^{n \bar{n} \times s_{2}}$. Let $Z_{2}=\widehat{U}^{\top} \odot\left(S W^{*}\right)^{\top} \in \mathbb{R}^{k \times n \bar{n}}$ denote the matrix where the $i$-th row is the vectorization of $\widehat{U}_{i} \otimes\left(S W_{i}^{*}\right)$. Define $\widehat{V}=\widehat{A}_{2} S_{2}\left(Z_{2} S_{2}\right)^{\dagger}$. By the properties of sketching matrix $S_{2}$, we have

$$
\left\|\widehat{V} Z_{2}-\widehat{A}_{2}\right\|_{1} \leq \alpha_{S_{2}} \alpha_{S_{1}} \mathrm{OPT}
$$

In the third step, we fix $\widehat{U} \in \mathbb{R}^{n \times k}$ and $\widehat{V} \in \mathbb{R}^{n \times k}$. Let $\widehat{A}_{3} \in \mathbb{R}^{\bar{n} \times n^{2}}$ denote the matrix obtained by flattening tensor $\widehat{A} \in \mathbb{R}^{n \times n \times \bar{n}}$ along the third direction. We choose a sketching matrix $S_{3} \in \mathbb{R}^{n^{2} \times s_{3}}$. Let $Z_{3} \in \mathbb{R}^{k \times n^{2}}$ denote the matrix where the $i$-th row is the vectorization of $\widehat{U}_{i} \otimes \widehat{V}_{i}$. Define $W^{\prime}=\widehat{A}_{3} S_{3}\left(Z_{3} S_{3}\right)^{\dagger} \in \mathbb{R}^{\bar{n} \times k}$ and $\widehat{W}=A_{3} S_{3}\left(Z_{3} S_{3}\right)^{\dagger} \in \mathbb{R}^{n \times k}$. Then we have,

$$
\begin{aligned}
W^{\prime} & =\widehat{A}_{3} S_{3}\left(Z_{3} S_{3}\right)^{\dagger} \\
& =(A(I, I, S))_{3} S_{3}\left(Z_{3} S_{3}\right)^{\dagger} \\
& =\left(S^{\top} A_{3}\right) S_{3}\left(Z_{3} S_{3}\right)^{\dagger} \\
& =S^{\top} \widehat{W}
\end{aligned}
$$

By properties of sketching matrix $S_{3}$, we have

$$
\left\|W^{\prime} Z_{3}-\widehat{A}_{3}\right\|_{1} \leq \alpha_{S_{3}} \alpha_{S_{2}} \alpha_{S_{1}} \mathrm{OPT} .
$$

Replacing $W^{\prime}$ by $S^{\top} \widehat{W}$, we obtain,

$$
\left\|W^{\prime} Z_{3}-\widehat{A}_{3}\right\|_{1}=\left\|S^{\top} \widehat{W} Z_{3}-\widehat{A}_{3}\right\|_{1}=\left\|S^{\top} \widehat{W} Z_{3}-S^{\top} A_{3}\right\|_{1}=\|(\widehat{U} \otimes \widehat{V} \otimes \widehat{W}-A) S\|_{1} .
$$

Thus, we have

$$
\min _{X_{1} \in \mathbb{R}^{s_{1} \times k}, X_{2} \in \mathbb{R}^{s_{2} \times k}, X_{3} \in \mathbb{R}^{s_{3} \times k}}\left\|\left(\widehat{A}_{1} S_{1} X_{1}\right) \otimes\left(\widehat{A}_{2} S_{2} X_{2}\right) \otimes\left(\widehat{A}_{3} S_{3} X_{3}\right)-\widehat{A}\right\|_{1} \leq \alpha_{S_{1}} \alpha_{S_{2}} \alpha_{S_{3}} \mathrm{OPT} .
$$

### 21.12.3.5 Running time analysis

Fact 21.12.25. Given tensor $A \in \mathbb{R}^{n \times n \times n}$ and a matrix $B \in \mathbb{R}^{n \times d}$ with $d=O(n)$, let $A B$ denote an $n \times n \times d$ size tensor, For each $i \in[3]$, let $(A B)_{i}$ denote a matrix obtained by
flattening tensor $A B$ along the $i$-th dimension, then

$$
(A B)_{1} \in \mathbb{R}^{n \times n d},(A B)_{2} \in \mathbb{R}^{n \times n d},(A B)_{3} \in \mathbb{R}^{d \times n^{2}}
$$

For each $i \in[3]$, let $S_{i} \in \mathbb{R}^{n d \times s_{i}}$ denote a sparse Cauchy transform, $T_{i} \in \mathbb{R}^{t_{i} \times n}$. Then we have,
(I) If $T_{1}$ denotes a sparse Cauchy transform or a sampling and rescaling matrix according to the Lewis weights, $T_{1}(A B)_{1} S_{1}$ can be computed in $O(\mathrm{nnz}(A) d)$ time. Otherwise, it can be computed in $O\left(\mathrm{nnz}(A) d+n s_{1} t_{1}\right)$.
(II) If $T_{2}$ denotes a sparse Cauchy transform or a sampling and rescaling matrix according to the Lewis weights, $T_{2}(A B)_{2} S_{2}$ can be computed in $O(\mathrm{nnz}(A) d)$ time. Otherwise, it can be computed in $O\left(\mathrm{nnz}(A) d+n s_{2} t_{2}\right)$.
(III) If $T_{3}$ denotes a sparse Cauchy transform or a sampling and rescaling matrix according to the Lewis weights, $T_{3}(A B)_{3} S_{3}$ can be computed in $O(\mathrm{nnz}(A) d)$ time. Otherwise, it can be computed in $O\left(\mathrm{nnz}(A) d+d s_{3} t_{3}\right)$.

Proof. Part (I). Note that $T_{1}(A B)_{1} S_{1} \in \mathbb{R}^{t_{1} \times s_{1}}$ and $(A B)_{1} \in \mathbb{R}^{n \times n d}$, for each $i \in\left[t_{1}\right], j \in\left[s_{1}\right]$,

$$
\begin{aligned}
\left(T_{1}(A B)_{1} S_{1}\right)_{i, j} & =\sum_{x=1}^{n} \sum_{y^{\prime}=1}^{n d}\left(T_{1}\right)_{i, x}\left((A B)_{1}\right)_{x, y^{\prime}}\left(S_{1}\right)_{y^{\prime}, j} \\
& =\sum_{x=1}^{n} \sum_{y=1}^{n} \sum_{z=1}^{d}\left(T_{1}\right)_{i, x}\left((A B)_{1}\right)_{x,(y-1) d+z}\left(S_{1}\right)_{(y-1) d+z, j} \\
& =\sum_{x=1}^{n} \sum_{y=1}^{n} \sum_{z=1}^{d}\left(T_{1}\right)_{i, x}(A B)_{x, y, z}\left(S_{1}\right)_{(y-1) d+z, j} \\
& =\sum_{x=1}^{n} \sum_{y=1}^{n} \sum_{z=1}^{d}\left(T_{1}\right)_{i, x} \sum_{w=1}^{n}\left(A_{x, y, w} B_{w, z}\right)\left(S_{1}\right)_{(y-1) d+z, j} \\
& =\sum_{x=1}^{n} \sum_{y=1}^{n}\left(T_{1}\right)_{i, x} \sum_{w=1}^{n} A_{x, y, w} \sum_{z=1}^{d} B_{w, z}\left(S_{1}\right)_{(y-1) d+z, j}
\end{aligned}
$$

We look at a non-zero entry $A_{x, y, w}$ and the entry $B_{w, z}$. If $T_{1}$ denotes a sparse Cauchy transform or a sampling and rescaling matrix according to the Lewis weights, then there is at most one pair $(i, j)$ such that $\left(T_{1}\right)_{i, x} A_{x, y, w} B_{w, z}\left(S_{1}\right)_{(y-1) d+z, j}$ is non-zero. Therefore, computing $T_{1}(A B)_{1} S_{1}$ only needs $\mathrm{nnz}(A) d$ time. If $T_{1}$ is not in the above case, since $S_{1}$ is sparse, we can compute $(A B)_{1} S_{1}$ in $\mathrm{nnz}(A) d$ time by a similar argument. Then, we can compute $T_{1}(A B)_{1} S_{1}$ in $n t_{1} s_{1}$ time.

Part (II). It is as the same as Part (I).
Part (III). Note that $T_{3}(A B)_{3} S_{3} \in \mathbb{R}^{t_{3} \times s_{3}}$ and $(A B)_{3} \in \mathbb{R}^{d \times n^{2}}$. For each $i \in\left[t_{3}\right], j \in$
$\left[s_{3}\right]$,

$$
\begin{aligned}
\left(T_{3}(A B)_{3} S_{3}\right)_{i, j} & =\sum_{x=1}^{d} \sum_{y^{\prime}=1}^{n^{2}}\left(T_{3}\right)_{i, x}\left((A B)_{3}\right)_{x, y^{\prime}}\left(S_{3}\right)_{y^{\prime}, j} \\
& =\sum_{x=1}^{d} \sum_{y=1}^{n} \sum_{z=1}^{n}\left(T_{3}\right)_{i, x}\left((A B)_{3}\right)_{x,(y-1) n+z}\left(S_{3}\right)_{(y-1) n+z, j} \\
& =\sum_{x=1}^{d} \sum_{y=1}^{n} \sum_{z=1}^{n}\left(T_{3}\right)_{i, x}(A B)_{y, z, x}\left(S_{3}\right)_{(y-1) n+z, j} \\
& =\sum_{x=1}^{d} \sum_{y=1}^{n} \sum_{z=1}^{n}\left(T_{3}\right)_{i, x} \sum_{w=1}^{n} A_{y, z, w} B_{w, x}\left(S_{3}\right)_{(y-1) n+z, j}
\end{aligned}
$$

Similar to Part (I), if $T_{1}$ denotes a sparse Cauchy transform or a sampling and rescaling matrix according to the Lewis weights, computing $T_{3}(A B)_{3} S_{3}$ only needs nnz $(A) d$ time. Otherwise, it needs $d t_{3} s_{3}+\mathrm{nnz}(A) d$ running time.

### 21.12.3.6 Algorithms

Theorem 21.12.26. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, there exists an algorithm which takes $O(\mathrm{nnz}(A) n)+\widetilde{O}(n) \operatorname{poly}(k)+n 2^{\widetilde{O}\left(k^{2}\right)}$ time and outputs three matrices $U, V, W \in \mathbb{R}^{n \times k}$ such that,

$$
\|U \otimes V \otimes W-A\|_{u} \leq \operatorname{poly}(k, \log n) \min _{\text {rank }-k A^{\prime}}\left\|A^{\prime}-A\right\|_{u}
$$

holds with probability at least 9/10.

Proof. We first choose a Gaussian matrix $S \in \mathbb{R}^{n \times \bar{n}}$ with $\bar{n}=O(n)$. By applying Corollary 21.12.23, we can reduce the original problem to a "generalized" $\ell_{1}$ low rank approximation

```
Algorithm \(21.38 \ell_{1}-\ell_{1}-\ell_{2}\)-Low Rank Approximation algorithm, input sparsity time
    procedure L112TensorLowRankApproxInputSparsity \((A, n, k)\)
    Theorem 21.12.26
        \(\bar{n} \leftarrow O(n)\).
        \(s_{1} \leftarrow s_{2} \leftarrow s_{3} \leftarrow \widetilde{O}\left(k^{5}\right)\).
        Choose \(S \in \mathbb{R}^{n \times \bar{n}}\) to be a Gaussian matrix.
        Choose \(S_{1} \in \mathbb{R}^{n \bar{n} \times s_{1}}\) to be a sparse Cauchy transform. \(\quad\) Part (II) of
    Theorem 21.12.24
            Choose \(S_{2} \in \mathbb{R}^{n \bar{n} \times s_{2}}\) to be a sparse Cauchy transform.
            Choose \(S_{3} \in \mathbb{R}^{n^{2} \times s_{3}}\) to be a sparse Cauchy transform.
            Form \(\widehat{A}=A S\).
            Compute \(\widehat{A}_{1} S_{1}, \widehat{A}_{2} S_{2}\), and \(\widehat{A}_{3} S_{3}\)
            \(Y_{1}, Y_{2}, Y_{3}, C \leftarrow\) L1PolyKSizeREduction \(\left(\widehat{A}, \widehat{A}_{1} S_{1}, \widehat{A}_{2} S_{2}, \widehat{A}_{3} S_{3}, n, n, \bar{n}, s_{1}, s_{2}, s_{3}, k\right) \triangleright\)
    Algorithm 21.21
            Create \(s_{1} k+s_{2} k+s_{3} k\) variables for each entry of \(X_{1}, X_{2}, X_{3}\).
            Form objective function \(\left\|\left(Y_{1} X_{1}\right) \otimes\left(Y_{2} X_{2}\right) \otimes\left(Y_{3} X_{3}\right)-C\right\|_{F}^{2}\).
            Run polynomial system verifier.
            return \(A_{1} S_{1} X_{1}, A_{2} S_{2} X_{2}, A_{3} S_{3} X_{3}\)
    end procedure
```

problem. Next, we use the existence results (Theorem 21.12.24) and polynomial in $k$ size reduction (Lemma 21.7.5). At the end, we relax the $\ell_{1}$-norm objective function to a Frobenius norm objective function (Fact 21.7.1).

Theorem 21.12.27. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, there exists an algorithm which takes $n^{\widetilde{O}(k)} 2^{\widetilde{O}\left(k^{3}\right)}$ time and outputs three matrices $U, V, W \in \mathbb{R}^{n \times k}$ such that,

$$
\|U \otimes V \otimes W-A\|_{u} \leq O\left(k^{3 / 2}\right) \min _{\operatorname{rank}-k A^{\prime}}\left\|A^{\prime}-A\right\|_{u}
$$

holds with probability at least 9/10.

Proof. We first choose a Gaussian matrix $S \in \mathbb{R}^{n \times \bar{n}}$ with $\bar{n}=O(n)$. By applying Corollary 21.12.23, we can reduce the original problem to a "generalized" $\ell_{1}$ low rank approximation

```
Algorithm \(21.39 \ell_{1}-\ell_{1}-\ell_{2}\)-Low Rank Approximation Algorithm, \(\widetilde{O}\left(k^{2 / 3}\right)\)
    procedure L112TensorLowRankApproxK \((A, n, k) \quad \triangleright\) Theorem 21.12.27
        \(\bar{n} \leftarrow O(n)\).
        \(s_{1} \leftarrow s_{2} \leftarrow s_{3} \leftarrow \widetilde{O}(k)\).
        Choose \(S \in \mathbb{R}^{n \times \bar{n}}\) to be a Gaussian matrix.
        Guess a diagonal matrix \(S_{1} \in \mathbb{R}^{n \bar{n} \times s_{1}}\) with \(s_{1}\) nonzero entries. \(\triangleright\) Part (III) of
    Theorem 21.12.24
        Guess a diagonal matrix \(S_{2} \in \mathbb{R}^{n \bar{n} \times s_{2}}\) with \(s_{2}\) nonzero entries.
        Guess a diagonal matrix \(S_{3} \in \mathbb{R}^{n^{2} \times s_{3}}\) with \(s_{3}\) nonzero entries.
        Form \(\widehat{A}=A S\).
        Compute \(\widehat{A}_{1} S_{1}, \widehat{A}_{2} S_{2}\), and \(\widehat{A}_{3} S_{3}\)
        \(Y_{1}, Y_{2}, Y_{3}, C \leftarrow\) L1PoLYKSizeREDUCTION \(\left(\widehat{A}, \widehat{A}_{1} S_{1}, \widehat{A}_{2} S_{2}, \widehat{A}_{3} S_{3}, n, n, \bar{n}, s_{1}, s_{2}, s_{3}, k\right) \triangleright\)
    Algorithm 21.21
        Create \(s_{1} k+s_{2} k+s_{3} k\) variables for each entry of \(X_{1}, X_{2}, X_{3}\).
        Form objective function \(\left\|\left(Y_{1} X_{1}\right) \otimes\left(Y_{2} X_{2}\right) \otimes\left(Y_{3} X_{3}\right)-C\right\|_{1}\).
        Run polynomial system verifier.
        return \(A_{1} S_{1} X_{1}, A_{2} S_{2} X_{2}, A_{3} S_{3} X_{3}\)
    end procedure
```

problem. Next, we use the existence results (Theorem 21.12.24) and polynomial in $k$ size reduction (Lemma 21.7.5). At the end, we solve an entry-wise $\ell_{1}$ norm objective function directly.

Theorem 21.12.28. Given a 3 rd order tensor $A \in \mathbb{R}^{n \times n \times n}$, for any $k \geq 1$, let $r=\widetilde{O}\left(k^{2}\right)$. There is an algorithm which takes $O(\mathrm{nnz}(A) n)+\widetilde{O}(n) \operatorname{poly}(k)$ time and outputs three matrices $U, V, W \in \mathbb{R}^{n \times r}$ such that

$$
\|U \otimes V \otimes W-A\|_{u} \leq \operatorname{poly}(\log n, k) \min _{\text {rank }-k A_{k}}\left\|A_{k}-A\right\|_{u},
$$

holds with probability at least 9/10.

```
Algorithm \(21.40 \ell_{1}-\ell_{1}-\ell_{2}\)-Low Rank Approximation Algorithm, Bicriteria Algorithm
    procedure L112TensorLowRankApproxBicteriteria \((A, n, k)\)
    Theorem 21.12.28
        \(\bar{n} \leftarrow O(n)\).
        \(s_{2} \leftarrow s_{3} \leftarrow \widetilde{O}\left(k^{5}\right)\).
        \(t_{2} \leftarrow t_{3} \leftarrow \widetilde{O}(k)\).
        \(r \leftarrow s_{2} s_{3}\).
        Choose \(S \in \mathbb{R}^{n \times \bar{n}}\) to be a Gaussian matrix.
        Form \(\widehat{A}=A S \in \mathbb{R}^{n \times n \times \bar{n}}\).
        Choose a sketching matrix \(S_{2} \in \mathbb{R}^{n \bar{n} \times s_{2}}\) with \(s_{2}\) nonzero entries (Sparse Cauchy
    transform), for each \(i \in\{2,3\}\). \(\triangleright\) Part (II) of Theorem 21.12.24
        Choose a sampling and rescaling diagonal matrix \(D_{i}\) according to the Lewis weights
    of \(\widehat{A}_{i} S_{i}\) with \(t_{i}\) nonzero entries, for each \(i \in\{2,3\}\).
        Form \(\widehat{V} \in \mathbb{R}^{n \times r}\) by setting the \((i, j)\)-th column to be \(\left(\widehat{A}_{2} S_{2}\right)_{i}\).
        Form \(\widehat{W} \in \mathbb{R}^{n \times r}\) by setting the \((i, j)\)-th column to be \(\left(A_{3} S_{3}\right)_{j}\).
        Form matrix \(B \in \mathbb{R}^{r \times t_{2} t_{3}}\) by setting the \((i, j)\)-th column to be the vectorization of
    \(\left(T_{2} \widehat{A}_{2} S_{2}\right)_{i} \otimes\left(T_{3} \widehat{A}_{3} S_{3}\right)_{j}\).
        Solve \(\min _{U}\left\|U \cdot B-\left(\widehat{A}\left(I, T_{2}, T_{3}\right)\right)_{1}\right\|_{1}\).
        return \(\widehat{U}, \widehat{V}, \widehat{W}\)
    end procedure
```

Proof. We first choose a Gaussian matrix $S \in \mathbb{R}^{n \times \bar{n}}$ with $\bar{n}=O(n)$. By applying Corollary 21.12 .23 , we can reduce the original problem to a "generalized" $\ell_{1}$ low rank approximation problem. Next, we use the existence results (Theorem 21.12.24) and polynomial in $k$ size reduction (Lemma 21.7.5). At the end, we solve an entry-wise $\ell_{1}$ norm objective function directly.

### 21.13 Streaming Setting

One of the computation models which is closely related to the distributed model of computation is the streaming model. There is a growing line of work in the streaming model. Some problems are very fundamental in the streaming model such like Heavy Hitters [LNNT16, BCI ${ }^{+} 16$, BCIW16], and streaming numerical linear algebra problems [CW09]. Streaming low rank matrix approximation has been extensively studied by previous work like [CW09, KL11, GP14, Lib13, KLM ${ }^{+}$14a, BWZ16, SWZ17]. In this section, we show that there is a streaming algorithm which can compute a low rank tensor approximation.

In the following, we introduce the turnstile streaming model and the turnstile streaming tensor Frobenius norm low rank approximation problem. The following gives a formal definition of the computation model we study.

Definition 21.13.1 (Turnstile model). Initially, tensor $A \in \mathbb{R}^{n \times n \times n}$ is an all zero tensor. In the turnstile streaming model, there is a stream of update operations, and the $i^{\text {th }}$ update operation is in the form $\left(x_{i}, y_{i}, z_{i}, \delta_{i}\right)$ where $x_{i}, y_{i}, z_{i} \in[n]$, and $\delta_{i} \in \mathbb{R}$ has $O(\log n)$ bits. Each $\left(x_{i}, y_{i}, z_{i}, \delta_{i}\right)$ means that $A_{x_{i}, y_{i}, z_{i}}$ should be incremented by $\delta_{i}$. And each entry of $A$ has at most $O(\log n)$ bits at the end of the stream. An algorithm in this computation model is only allowed one pass over the stream. At the end of the stream, the algorithm stores a summary of $A$. The space complexity of the algorithm is the total number of words required to compute and store this summary while scanning the stream. Here, each word has at most $O(\log (n))$ bits.

The following is the formal definition of the problem.

Definition 21.13.2 (Turnstile model Frobenius norm rank- $k$ tensor approximation). Given tensor $A \in \mathbb{R}^{n \times n \times n}, k \in \mathbb{N}_{+}$and an error parameter $1>\epsilon>0$, the goal is to design an algorithm in the streaming model of Definition 21.13 .1 such that

1. Upon termination, the algorithm outputs three matrices $U^{*}, V^{*}, W^{*} \in \mathbb{R}^{n \times k}$.
2. $U^{*}, V^{*}, W^{*}$ satisft that

$$
\left\|\sum_{i=1}^{k} U_{i}^{*} \otimes V_{i}^{*} \otimes W_{i}^{*}-A\right\|_{F}^{2} \leq(1+\epsilon) \min _{\operatorname{rank}-k}\left\|A^{\prime}-A\right\|_{F}^{2} .
$$

3. The space complexity of the algorithm is as small as possible.

Theorem 21.13.1. Suppose tensor $A \in \mathbb{R}^{n \times n \times n}$ is given in the turnstile streaming model (see Definition 21.13.1), there is an streaming algorithm (in Algorithm 21.41) which solves the problem in Definition 21.13.2 with constant success probability. In addition, the space complexity of the algorithm is poly $(k / \epsilon)+O(n k / \epsilon)$ words.

Proof. Correctness. Similar to the distributed protocol, the correctness of this streaming algorithm is also implied by Algorithm 21.2 and Algorithm 21.3 (Theorem 21.4.1.) Notice that at the end of the stream $V_{1}=A_{1} S_{1} \in \mathbb{R}^{n \times s_{1}}, V_{2}=A_{2} S_{2} \in \mathbb{R}^{n \times s_{2}}, V_{3}=A_{3} S_{3} \in$ $\mathbb{R}^{n \times s_{3}}, C=A\left(T_{1}, T_{2}, T_{3}\right) \in \mathbb{R}^{t_{1} \times t_{2} \times t_{3}}$. It also means that

$$
Y_{1}=T_{1} A_{1} S_{1}, Y_{2}=T_{2} A_{2} S_{2}, Y_{3}=T_{3} A_{3} S_{3} .
$$

According to line 26 of procedure TurnstileStreaming,

$$
X_{1}^{*}, X_{2}^{*}, X_{3}^{*}=\underset{X_{1} \in \mathbb{R}^{s_{1} \times k}, X_{2} \in \mathbb{R}^{s_{2} \times k}, X_{3} \in \mathbb{R}^{s_{3} \times k}}{\arg \min }\left\|\sum_{j=1}^{k}\left(Y_{1} X_{1}\right)_{j} \otimes\left(Y_{2} X_{2}\right)_{j} \otimes\left(Y_{3} X_{3}\right)_{j}-C\right\|_{F}
$$

According to Lemma 21.4.3, we have

$$
\begin{aligned}
& \left\|\sum_{j=1}^{k}\left(Y_{1} X_{1}\right)_{j} \otimes\left(Y_{2} X_{2}\right)_{j} \otimes\left(Y_{3} X_{3}\right)_{j}-C\right\|_{F}^{2} \\
= & \left\|\sum_{j=1}^{k}\left(T_{1} A_{1} S_{1} X_{1}^{*}\right)_{j} \otimes\left(T_{2} A_{2} S_{2} X_{2}^{*}\right)_{j} \otimes\left(T_{3} A_{3} S_{3} X_{3}^{*}\right)_{j}-A\left(T_{1}, T_{2}, T_{3}\right)\right\|_{F}^{2} \\
\leq & (1+O(\epsilon)) \min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{j=1}^{k}\left(A_{1} S_{1} X_{1}\right)_{j} \otimes\left(A_{2} S_{2} X_{2}\right)_{j} \otimes\left(A_{3} Y_{3} X_{3}\right)_{j}-A\right\|_{F}^{2} \\
\leq & (1+O(\epsilon)) \min _{U, V, W}\left\|\sum_{i=1}^{k} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{F}^{2}
\end{aligned}
$$

where the last inequality follows by the proof of Theorem 21.4.1. By scaling a constant of $\epsilon$, we complete the proof of correctness.

Space complexity. Since $S_{1}, S_{2}, S_{3}$ are $w_{1}$-wise independent, and $T_{1}, T_{2}, T_{3}$ are $w_{2^{-}}$ wise independent, the space needed to construct these sketching matrices in line 3 and line 5 of procedure TurnstileStreaming is $O\left(w_{1}+w_{2}\right)$ words, where $w_{1}=O(k), w_{2}=O(1)$ (see [KVW14, CW13, Woo14b, KN14]). The cost to maintain $V_{1}, V_{2}, V_{3}$ is $O(n k / \epsilon)$ words, and the cost to maintain $C$ is $\operatorname{poly}(k / \epsilon)$ words.

Notice that, since each entry of $A$ has at most $O(\log (s n))$ bits, each entry of $Y_{1}, Y_{2}, Y_{3}, C$ has at most $O(\log (s n))$ bits. Due to Theorem 21.15.2, each entry of $X_{1}^{*}, X_{2}^{*}, X_{3}^{*}$ has at most $O(\log (s n))$ bits, and the sizes of $X_{1}^{*}, X_{2}^{*}, X_{3}^{*}$ are poly $(k / \epsilon)$ words. Thus the space cost in line 26 is $\operatorname{poly}(k / \epsilon)$ words.

The total space cost is $\operatorname{poly}(k / \epsilon)+O(n k / \epsilon)$ words.

Remark 21.13.1. In the Algorithm 21.41, for each update operation, we need $O(k / \epsilon)$ time to maintain matrices $V_{1}, V_{2}, V_{3}$, and we need $\operatorname{poly}(k / \epsilon)$ time to maintain tensor $C$. Thus the

```
Algorithm 21.41 Turnstile Frobenius Norm Low Rank Approximation Algorithm
    procedure TurnstileStreaming \((k, \mathcal{S})\)
        \(s_{1} \leftarrow s_{2} \leftarrow s_{3} \leftarrow O(k / \epsilon)\).
        Construct sketching matrices \(S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}, \forall i \in\) [3] where entries of \(S_{1}, S_{2}, S_{3}\) are
    \(w_{1}\)-wise independent random \(N\left(0,1 / s_{i}\right)\) Gaussian variables.
            \(t_{1} \leftarrow t_{2} \leftarrow t_{3} \leftarrow \operatorname{poly}(k / \epsilon)\).
            Construct sparse embedding matrices \(T_{i} \in \mathbb{R}^{t_{i} \times n}, \forall i \in[3]\) where entries are \(w_{2}\)-wise
    independent.
        Initialize matrices:
        \(V_{i} \leftarrow\{0\}^{n \times s_{i}}, \forall i \in[3]\).
        \(C \leftarrow\{0\}^{t_{1} \times t_{2} \times t_{3}}\)
        for \(i \in[l]\) do
            Receive update operation \(\left(x_{i}, y_{i}, z_{i}, \delta_{i}\right)\) from the data stream \(\mathcal{S}\).
            for \(r=1 \rightarrow s_{1}\) do
                \(\left(V_{1}\right)_{x_{i}, r} \leftarrow\left(V_{1}\right)_{x_{i}, r}+\delta_{i} \cdot\left(S_{1}\right)_{\left(y_{i}-1\right) n+z_{i}, r}\).
            end for
            for \(r=1 \rightarrow s_{2}\) do
                \(\left(V_{2}\right)_{y_{i}, r} \leftarrow\left(V_{2}\right)_{y_{i}, r}+\delta_{i} \cdot\left(S_{2}\right)_{\left(z_{i}-1\right) n+x_{i}, r}\).
            end for
            for \(r=1 \rightarrow s_{3}\) do
                \(\left(V_{3}\right)_{z_{i}, r} \leftarrow\left(V_{3}\right)_{z_{i}, r}+\delta_{i} \cdot\left(S_{3}\right)_{\left(x_{i}-1\right) n+y_{i}, r}\).
            end for
            for \(r=1 \rightarrow t_{1}, p=1 \rightarrow t_{2}, q=1 \rightarrow t_{3}\) do
                \(C_{r, p, q} \leftarrow C_{r, p, q}+\delta_{i} \cdot\left(T_{1}\right)_{r, x_{i}}\left(T_{2}\right)_{p, y_{i}}\left(T_{3}\right)_{q, z_{i}}\).
            end for
        end for
        Compute \(Y_{1} \leftarrow T_{1} V_{1}, Y_{2} \leftarrow T_{2} V_{2}, Y_{3} \leftarrow T_{3} V_{3}\).
        Compute \(X_{i}^{*} \in \mathbb{R}^{s_{i} \times k}, \forall i \in[3]\) by solving
            \(\min _{X_{1}, X_{2}, X_{3}}\left\|\left(Y_{1} X_{1}\right) \otimes\left(Y_{2} X_{2}\right) \otimes\left(Y_{3} X_{3}\right)-C\right\|_{F}\)
        Compute \(U^{*} \leftarrow V_{1} X_{1}^{*}, V^{*} \leftarrow V_{2} X_{2}^{*}, W^{*} \leftarrow V_{3} X_{3}^{*}\).
        return \(U^{*}, V^{*}, W^{*}\)
    end procedure
```

update time is $\operatorname{poly}(k / \epsilon)$. At the end of the stream, the time to compute

$$
X_{1}^{*}, X_{2}^{*}, X_{3}^{*}=\underset{X_{1}, X_{2}, X_{3} \in \mathbb{R}^{O(k / \epsilon) \times k}}{\arg \min }\left\|\sum_{j=1}^{k}\left(Y_{1} X_{1}\right)_{j} \otimes\left(Y_{2} X_{2}\right)_{j} \otimes\left(Y_{3} X_{3}\right)_{j}-C\right\|_{F}
$$

is exponential in poly $(k / \epsilon)$ running time since it should use a polynomial system solver. Instead of computing the rank- $k$ solution, we can solve the following:

$$
\alpha^{*}=\underset{\alpha \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}}{\arg \min }\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} \alpha_{i, j, l} \cdot\left(Y_{1}\right)_{i} \otimes\left(Y_{2}\right)_{j} \otimes\left(Y_{3}\right)_{l}-C\right\|_{F}
$$

which will then give

$$
\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} \alpha_{i, j, l}^{*} \cdot\left(Y_{1}\right)_{i} \otimes\left(Y_{2}\right)_{j} \otimes\left(Y_{3}\right)_{l}
$$

to be a rank- $O\left(k^{3} / \epsilon^{3}\right)$ bicriteria solution.
Further, similar to Theorem 21.4.8, we can solve

$$
\min _{U \in \mathbb{R}^{n \times s_{2} s_{3}}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} U_{i+s_{1}(j-1)} \otimes\left(Y_{2}\right)_{i} \otimes\left(Y_{3}\right)_{j}-C\right\|_{F}
$$

where $C=\sum_{i} A_{i}\left(I, T_{2}, T_{3}\right)$. Thus, we can obtain a rank- $O\left(k^{2} / \epsilon^{2}\right)$ in polynomial time.
Remark 21.13.2. If we choose $S_{1}, S_{2}, S_{3}, T_{1}, T_{2}, T_{3}$ to be random Cauchy matrices, then we are able to apply the entry-wise $\ell_{1}$ norm low rank tensor approximation algorithm (see Theorem 21.7.14) in turnstile model.

### 21.14 Distributed Setting

Input data to large-scale machine learning and data mining tasks may be distributed across different machines. The communication cost becomes the major bottleneck of distributed protocols, and so there is a growing body of work on low rank matrix approximations in the distributed model [TD99, QOSG02, BCL05, BRB08, MBZ10, FEGK13, PMvdG ${ }^{+}$13, KVW14, BKLW14, BLS ${ }^{+}$16a, BWZ16, WZ16, SWZ17] and also many other machine learning problems such as clustering, boosting, and column subset selection [BBLM14, BLG ${ }^{+} 15$, ABW17]. Thus, it is natural to ask whether our algorithm can be applied in the distributed setting. This section will discuss the distributed Frobenius norm low rank tensor approximation protocol in the so-called arbitrary-partition model (see, e.g. [KVW14, BWZ16]).

In the following, we extend the definition of the arbitrary-partition model [KVW14] to fit our tensor setting.

Definition 21.14.1 (Arbitrary-partition model [KVW14]). There are $s$ machines, and the $i^{\text {th }}$ machine holds a tensor $A_{i} \in \mathbb{R}^{n \times n \times n}$ as its local data tensor. The global data tensor is implicit and is denoted as $A=\sum_{i=1}^{s} A_{i}$. Then, we say that $A$ is arbitrarily partitioned into $s$ matrices distributed in the $s$ machines. In addition, there is also a coordinator. In this model, the communication is only allowed between the machines and the coordinator. The total communication cost is the total number of words delivered between machines and the coordinator. Each word has $O(\log (s n))$ bits.

Now, let us introduce the distributed Frobenius norm low rank tensor approximation problem in the arbitrary partition model:

Definition 21.14.2 (Arbitrary-partition model Frobenius norm rank- $k$ tensor approximation). Tensor $A \in \mathbb{R}^{n \times n \times n}$ is arbitrarily partitioned into $s$ matrices $A_{1}, A_{2}, \cdots, A_{s}$ distributed in $s$ machines respectively, and $\forall i \in[s]$, each entry of $A_{i}$ is at most $O(\log (s n))$ bits. Given tensor $A, k \in \mathbb{N}_{+}$and an error parameter $0<\epsilon<1$, the goal is to find a distributed protocol in the model of Definition 21.15.1 such that

1. Upon termination, the protocol leaves three matrices $U^{*}, V^{*}, W^{*} \in \mathbb{R}^{n \times k}$ on the coordinator.
2. $U^{*}, V^{*}, W^{*}$ satisfies that

$$
\left\|\sum_{i=1}^{k} U_{i}^{*} \otimes V_{i}^{*} \otimes W_{i}^{*}-A\right\|_{F}^{2} \leq(1+\epsilon)_{\operatorname{rank}-k} \min _{A^{\prime}}\left\|A^{\prime}-A\right\|_{F}^{2}
$$

3. The communication cost is as small as possible.

Theorem 21.14.1. Suppose tensor $A \in \mathbb{R}^{n \times n \times n}$ is distributed in the arbitrary partition model (See Definition 21.15.1). There is a protocol( in Algorithm 21.43) which solves the problem in Definition 21.15.2 with constant success probability. In addition, the communication complexity of the protocol is $s(\operatorname{poly}(k / \epsilon)+O(k n))$ words.

Proof. Correctness. The correctness is implied by Algorithm 21.2 and Algorithm 21.3 (Theorem 21.4.1.) Notice that $A_{1}=\sum_{i=1}^{s} A_{i, 1}, A_{2}=\sum_{i=1}^{s} A_{i, 2}, A_{3}=\sum_{i=1}^{s} A_{i, 3}$, which means that

$$
Y_{1}=T_{1} A_{1} S_{1}, Y_{2}=T_{2} A_{2} S_{2}, Y_{3}=T_{3} A_{3} S_{3}
$$

and

$$
C=A\left(T_{1}, T_{2}, T_{3}\right)
$$

According to line 23,

$$
X_{1}^{*}, X_{2}^{*}, X_{3}^{*}=\underset{X_{1}, X_{2}, X_{3}}{\arg \min }\left\|\sum_{j=1}^{k}\left(Y_{1} X_{1}\right)_{j} \otimes\left(Y_{2} X_{2}\right)_{j} \otimes\left(Y_{3} X_{3}\right)_{j}-C\right\|_{F}
$$

According to Lemma 21.4.3, we have

$$
\begin{aligned}
& \left\|\sum_{j=1}^{k}\left(T_{1} A_{1} S_{1} X_{1}^{*}\right)_{j} \otimes\left(T_{2} A_{2} S_{2} X_{2}^{*}\right)_{j} \otimes\left(T_{3} A_{3} S_{3} X_{3}^{*}\right)_{j}-A\left(T_{1}, T_{2}, T_{3}\right)\right\|_{F}^{2} \\
\leq & (1+O(\epsilon))_{X_{1}, X_{2}, X_{3}} \min _{F}\left\|\sum_{j=1}^{k}\left(A_{1} S_{1} X_{1}\right)_{j} \otimes\left(A_{2} S_{2} X_{2}\right)_{j} \otimes\left(A_{3} Y_{3} X_{3}\right)_{j}-A\right\|_{F}^{2} \\
\leq & (1+O(\epsilon)) \min _{U, V, W}\left\|\sum_{i=1}^{k} U_{i} \otimes V_{i} \otimes W_{i}-A\right\|_{F}^{2},
\end{aligned}
$$

where the last inequality follows by the proof of Theorem 21.4.1. By scaling a constant of $\epsilon$, we complete the proof of correctness.

Communication complexity. Since $S_{1}, S_{2}, S_{3}$ are $w_{1}$-wise independent, and $T_{1}, T_{2}, T_{3}$ are $w_{2}$-wise independent, the communication cost of sending random seeds in line 5 is $O\left(s\left(w_{1}+w_{2}\right)\right)$ words, where $w_{1}=O(k), w_{2}=O(1)$ (see [KVW14, CW13, Woo14b, KN14]). The communication cost in line 18 is $s \cdot \operatorname{poly}(k / \epsilon)$ words due to $T_{1} A_{i, 1} S_{1}, T_{2} A_{i, 2} S_{2}, T_{3} A_{i, 3} S_{3} \in$ $\mathbb{R}^{\text {poly }(k / \epsilon) \times O(k / \epsilon)}$ and $C_{i}=A_{i}\left(T_{1}, T_{2}, T_{3}\right) \in \mathbb{R}^{\operatorname{poly}(k / \epsilon) \times \operatorname{poly}(k / \epsilon) \times \operatorname{poly}(k / \epsilon)}$.

Notice that, since $\forall i \in[s]$ each entry of $A_{i}$ has at most $O(\log (s n))$ bits, each entry of $Y_{1}, Y_{2}, Y_{3}, C$ has at most $O(\log (s n))$ bits. Due to Theorem 21.15.2, each entry of $X_{1}^{*}, X_{2}^{*}, X_{3}^{*}$ has at most $O(\log (s n))$ bits, and the sizes of $X_{1}^{*}, X_{2}^{*}, X_{3}^{*}$ are poly $(k / \epsilon)$ words. Thus the communication cost in line 24 is $s \cdot \operatorname{poly}(k / \epsilon)$ words.

Finally, since $\forall i \in[s], U_{i}^{*}, V_{i}^{*}, W_{i}^{*} \in \mathbb{R}^{n \times k}$, the communication here is at most $O(s k n)$ words. The total communication cost is $s(\operatorname{poly}(k / \epsilon)+O(k n))$ words.

Remark 21.14.1. If we slightly change the goal in Definition 21.15 .2 to the following: the coordinator does not need to output $U^{*}, V^{*}, W^{*}$, but each machine $i$ holds $U_{i}^{*}, V_{i}^{*}, W_{i}^{*}$ such that $U^{*}=\sum_{i=1}^{s} U_{i}^{*}, V^{*}=\sum_{i=1}^{s} V_{i}^{*}, W^{*}=\sum_{i=1}^{s} W_{i}^{*}$, then the protocol shown in Algorithm 21.43 does not have to do the line 28 . Thus the total communication cost is at most $s \cdot \operatorname{poly}(k / \epsilon)$ words in this setting.

Remark 21.14.2. Algorithm 21.43 needs exponential in poly $(k / \epsilon)$ running time since it solves a polynomial solver in line 23. Instead of solving line 23, we can solve the following optimization problem:

$$
\alpha^{*}=\underset{\alpha \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}}{\arg \min _{\|}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} \alpha_{i, j, l} \cdot\left(Y_{1}\right)_{i} \otimes\left(Y_{2}\right)_{j} \otimes\left(Y_{3}\right)_{l}-C\right\|_{F}
$$

Since it is actually a regression problem, it only takes polynomial running time to get $\alpha^{*}$. And according to Lemma 21.4.5,

$$
\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} \alpha_{i, j, l}^{*} \cdot\left(Y_{1}\right)_{i} \otimes\left(Y_{2}\right)_{j} \otimes\left(Y_{3}\right)_{l}
$$

gives a rank- $O\left(k^{3} / \epsilon^{3}\right)$ bicriteria solution.
Further, similar to Theorem 21.4.8, we can solve

$$
\min _{U \in \mathbb{R}^{n \times s_{2} s_{3}}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} U_{i+s_{1}(j-1)} \otimes\left(Y_{2}\right)_{i} \otimes\left(Y_{3}\right)_{j}-C\right\|_{F},
$$

where $C=\sum_{i} A_{i}\left(I, T_{2}, T_{3}\right)$. Thus, we can obtain a rank- $O\left(k^{2} / \epsilon^{2}\right)$ in polynomial time.
Remark 21.14.3. If we select sketching matrices $S_{1}, S_{2}, S_{3}, T_{1}, T_{2}, T_{3}$ to be random Cauchy matrices, then we are able to compute distributed entry-wise $\ell_{1}$ norm rank- $k$ tensor approximation (see Theorem 21.7.14). The communication cost is still $s(\operatorname{poly}(k / \epsilon)+O(k n))$ words. If we only require a bicriteria solution, then it only needs polynomial running time.

Using similar techniques as in the proof of Theorem 21.4.44, we can obtain:

Theorem 21.14.2. Let $\max _{i}\left\{t_{i}, d_{i}\right\} \leq n$. Given a $t_{1} \times t_{2} \times t_{3}$ tensor $A$ and three matrices: a $t_{1} \times d_{1}$ matrix $T_{1}$, a $t_{2} \times d_{2}$ matrix $T_{2}$, and a $t_{3} \times d_{3}$ matrix $T_{3}$. For any $\delta>0$, if there exists a solution to

$$
\min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k}\left(T_{1} X_{1}\right)_{i} \otimes\left(T_{2} X_{2}\right)_{i} \otimes\left(T_{3} X_{3}\right)_{i}-A\right\|_{F}^{2}:=\mathrm{OPT}
$$

and each entry of $X_{i}$ can be expressed using $O(\log n)$ bits, then there exists an algorithm that takes $\operatorname{poly}(\log n) \cdot 2^{O\left(d_{1} k+d_{2} k+d_{3} k\right)}$ time and outputs three matrices: $\widehat{X}_{1}, \widehat{X}_{2}$, and $\widehat{X}_{3}$ such that $\left\|\left(T_{1} \widehat{X}_{1}\right) \otimes\left(T_{2} \widehat{X}_{2}\right) \otimes\left(T_{3} \widehat{X}_{3}\right)-A\right\|_{F}^{2}=\mathrm{OPT}$.

### 21.15 Distributed Setting

Input data to large-scale machine learning and data mining tasks may be distributed across different machines. The communication cost becomes the major bottleneck of distributed protocols, and so there is a growing body of work on low rank matrix approximations in the distributed model [TD99, QOSG02, BCL05, BRB08, MBZ10, FEGK13, PMvdG ${ }^{+}$13, KVW14, BKLW14, BLS ${ }^{+}$16a, BWZ16, WZ16, SWZ17] and also many other machine learning problems such as clustering, boosting, and column subset selection [BBLM14, BLG ${ }^{+} 15$, ABW17]. Thus, it is natural to ask whether our algorithm can be applied in the distributed setting. This section will discuss the distributed Frobenius norm low rank tensor approximation protocol in the so-called arbitrary-partition model (see, e.g. [KVW14, BWZ16]).

In the following, we extend the definition of the arbitrary-partition model [KVW14] to fit our tensor setting.

Definition 21.15.1 (Arbitrary-partition model [KVW14]). There are $s$ machines, and the $i^{\text {th }}$ machine holds a tensor $A_{i} \in \mathbb{R}^{n \times n \times n}$ as its local data tensor. The global data tensor is implicit and is denoted as $A=\sum_{i=1}^{s} A_{i}$. Then, we say that $A$ is arbitrarily partitioned into $s$ matrices distributed in the $s$ machines. In addition, there is also a coordinator. In this model, the communication is only allowed between the machines and the coordinator. The total communication cost is the total number of words delivered between machines and the coordinator. Each word has $O(\log (s n))$ bits.

Now, let us introduce the distributed Frobenius norm low rank tensor approximation problem in the arbitrary partition model:

Definition 21.15.2 (Arbitrary-partition model Frobenius norm rank- $k$ tensor approximation). Tensor $A \in \mathbb{R}^{n \times n \times n}$ is arbitrarily partitioned into $s$ matrices $A_{1}, A_{2}, \cdots, A_{s}$ distributed in $s$ machines respectively, and $\forall i \in[s]$, each entry of $A_{i}$ is at most $O(\log (s n))$ bits. Given tensor $A, k \in \mathbb{N}_{+}$and an error parameter $0<\epsilon<1$, the goal is to find a distributed protocol in the model of Definition 21.15.1 such that

1. Upon termination, the protocol leaves three matrices $U^{*}, V^{*}, W^{*} \in \mathbb{R}^{n \times k}$ on the coordinator.
2. $U^{*}, V^{*}, W^{*}$ satisfies that

$$
\left\|\sum_{i=1}^{k} U_{i}^{*} \otimes V_{i}^{*} \otimes W_{i}^{*}-A\right\|_{F}^{2} \leq(1+\epsilon) \min _{\operatorname{rank}-k A^{\prime}}\left\|A^{\prime}-A\right\|_{F}^{2} .
$$

3. The communication cost is as small as possible.

Theorem 21.15.1. Suppose tensor $A \in \mathbb{R}^{n \times n \times n}$ is distributed in the arbitrary partition model (See Definition 21.15.1). There is a protocol( in Algorithm 21.43) which solves the problem in Definition 21.15.2 with constant success probability. In addition, the communication complexity of the protocol is $s(\operatorname{poly}(k / \epsilon)+O(k n))$ words.

Proof. Correctness. The correctness is implied by Algorithm 21.2 and Algorithm 21.3 (Theorem 21.4.1.) Notice that $A_{1}=\sum_{i=1}^{s} A_{i, 1}, A_{2}=\sum_{i=1}^{s} A_{i, 2}, A_{3}=\sum_{i=1}^{s} A_{i, 3}$, which means that

$$
Y_{1}=T_{1} A_{1} S_{1}, Y_{2}=T_{2} A_{2} S_{2}, Y_{3}=T_{3} A_{3} S_{3}
$$

and

$$
C=A\left(T_{1}, T_{2}, T_{3}\right)
$$

According to line 23,

$$
X_{1}^{*}, X_{2}^{*}, X_{3}^{*}=\underset{X_{1}, X_{2}, X_{3}}{\arg \min }\left\|\sum_{j=1}^{k}\left(Y_{1} X_{1}\right)_{j} \otimes\left(Y_{2} X_{2}\right)_{j} \otimes\left(Y_{3} X_{3}\right)_{j}-C\right\|_{F}
$$

According to Lemma 21.4.3, we have

$$
\begin{aligned}
& \quad\left\|\sum_{j=1}^{k}\left(T_{1} A_{1} S_{1} X_{1}^{*}\right)_{j} \otimes\left(T_{2} A_{2} S_{2} X_{2}^{*}\right)_{j} \otimes\left(T_{3} A_{3} S_{3} X_{3}^{*}\right)_{j}-A\left(T_{1}, T_{2}, T_{3}\right)\right\|_{F}^{2} \\
& \leq(1+O(\epsilon))_{X_{1}, X_{2}, X_{3}} \min _{j=1}\left\|\sum_{j=1}^{k}\left(A_{1} S_{1} X_{1}\right)_{j} \otimes\left(A_{2} S_{2} X_{2}\right)_{j} \otimes\left(A_{3} Y_{3} X_{3}\right)_{j}-A\right\|_{F}^{2} \\
& \leq(1+O(\epsilon))_{U, V, W} \| \sum_{i=1}^{k} U_{i} \otimes V_{i} \otimes W_{i}-A
\end{aligned}
$$

where the last inequality follows by the proof of Theorem 21.4.1. By scaling a constant of $\epsilon$, we complete the proof of correctness.

Communication complexity. Since $S_{1}, S_{2}, S_{3}$ are $w_{1}$-wise independent, and $T_{1}, T_{2}, T_{3}$ are $w_{2}$-wise independent, the communication cost of sending random seeds in line 5 is $O\left(s\left(w_{1}+w_{2}\right)\right)$ words, where $w_{1}=O(k), w_{2}=O(1)$ (see [KVW14, CW13, Woo14b, KN14]). The communication cost in line 18 is $s \cdot \operatorname{poly}(k / \epsilon)$ words due to $T_{1} A_{i, 1} S_{1}, T_{2} A_{i, 2} S_{2}, T_{3} A_{i, 3} S_{3} \in$ $\mathbb{R}^{\text {poly }(k / \epsilon) \times O(k / \epsilon)}$ and $C_{i}=A_{i}\left(T_{1}, T_{2}, T_{3}\right) \in \mathbb{R}^{\text {poly }(k / \epsilon) \times \operatorname{poly}(k / \epsilon) \times \operatorname{poly}(k / \epsilon)}$.

Notice that, since $\forall i \in[s]$ each entry of $A_{i}$ has at most $O(\log (s n))$ bits, each entry of $Y_{1}, Y_{2}, Y_{3}, C$ has at most $O(\log (s n))$ bits. Due to Theorem 21.15.2, each entry of $X_{1}^{*}, X_{2}^{*}, X_{3}^{*}$ has at most $O(\log (s n))$ bits, and the sizes of $X_{1}^{*}, X_{2}^{*}, X_{3}^{*}$ are poly $(k / \epsilon)$ words. Thus the communication cost in line 24 is $s \cdot \operatorname{poly}(k / \epsilon)$ words.

Finally, since $\forall i \in[s], U_{i}^{*}, V_{i}^{*}, W_{i}^{*} \in \mathbb{R}^{n \times k}$, the communication here is at most $O(s k n)$ words. The total communication cost is $s(\operatorname{poly}(k / \epsilon)+O(k n))$ words.

Remark 21.15.1. If we slightly change the goal in Definition 21.15 .2 to the following: the coordinator does not need to output $U^{*}, V^{*}, W^{*}$, but each machine $i$ holds $U_{i}^{*}, V_{i}^{*}, W_{i}^{*}$ such that $U^{*}=\sum_{i=1}^{s} U_{i}^{*}, V^{*}=\sum_{i=1}^{s} V_{i}^{*}, W^{*}=\sum_{i=1}^{s} W_{i}^{*}$, then the protocol shown in Algorithm 21.43 does not have to do the line 28 . Thus the total communication cost is at most $s \cdot \operatorname{poly}(k / \epsilon)$ words in this setting.

Remark 21.15.2. Algorithm 21.43 needs exponential in poly $(k / \epsilon)$ running time since it solves a polynomial solver in line 23. Instead of solving line 23, we can solve the following opti-
mization problem:

$$
\alpha^{*}=\underset{\alpha \in \mathbb{R}^{s_{1} \times s_{2} \times s_{3}}}{\arg \min }\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} \alpha_{i, j, l} \cdot\left(Y_{1}\right)_{i} \otimes\left(Y_{2}\right)_{j} \otimes\left(Y_{3}\right)_{l}-C\right\|_{F} .
$$

Since it is actually a regression problem, it only takes polynomial running time to get $\alpha^{*}$. And according to Lemma 21.4.5,

$$
\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} \sum_{l=1}^{s_{3}} \alpha_{i, j, l}^{*} \cdot\left(Y_{1}\right)_{i} \otimes\left(Y_{2}\right)_{j} \otimes\left(Y_{3}\right)_{l}
$$

gives a rank- $O\left(k^{3} / \epsilon^{3}\right)$ bicriteria solution.
Further, similar to Theorem 21.4.8, we can solve

$$
\min _{U \in \mathbb{R}^{n \times s_{2} s_{3}}}\left\|\sum_{i=1}^{s_{1}} \sum_{j=1}^{s_{2}} U_{i+s_{1}(j-1)} \otimes\left(Y_{2}\right)_{i} \otimes\left(Y_{3}\right)_{j}-C\right\|_{F},
$$

where $C=\sum_{i} A_{i}\left(I, T_{2}, T_{3}\right)$. Thus, we can obtain a rank- $O\left(k^{2} / \epsilon^{2}\right)$ in polynomial time.
Remark 21.15.3. If we select sketching matrices $S_{1}, S_{2}, S_{3}, T_{1}, T_{2}, T_{3}$ to be random Cauchy matrices, then we are able to compute distributed entry-wise $\ell_{1}$ norm rank- $k$ tensor approximation (see Theorem 21.7.14). The communication cost is still $s(\operatorname{poly}(k / \epsilon)+O(k n))$ words. If we only require a bicriteria solution, then it only needs polynomial running time.

Using similar techniques as in the proof of Theorem 21.4.44, we can obtain:

Theorem 21.15.2. Let $\max _{i}\left\{t_{i}, d_{i}\right\} \leq n$. Given a $t_{1} \times t_{2} \times t_{3}$ tensor $A$ and three matrices: a $t_{1} \times d_{1}$ matrix $T_{1}$, a $t_{2} \times d_{2}$ matrix $T_{2}$, and a $t_{3} \times d_{3}$ matrix $T_{3}$. For any $\delta>0$, if there exists a solution to

$$
\min _{X_{1}, X_{2}, X_{3}}\left\|\sum_{i=1}^{k}\left(T_{1} X_{1}\right)_{i} \otimes\left(T_{2} X_{2}\right)_{i} \otimes\left(T_{3} X_{3}\right)_{i}-A\right\|_{F}^{2}:=\mathrm{OPT}
$$

and each entry of $X_{i}$ can be expressed using $O(\log n)$ bits, then there exists an algorithm that takes $\operatorname{poly}(\log n) \cdot 2^{O\left(d_{1} k+d_{2} k+d_{3} k\right)}$ time and outputs three matrices: $\widehat{X}_{1}, \widehat{X}_{2}$, and $\widehat{X}_{3}$ such that $\left\|\left(T_{1} \widehat{X}_{1}\right) \otimes\left(T_{2} \widehat{X}_{2}\right) \otimes\left(T_{3} \widehat{X}_{3}\right)-A\right\|_{F}^{2}=\mathrm{OPT}$.

```
Algorithm 21.42 Distributed Frobenius Norm Low Rank Approximation Protocol
    procedure DistributedFnormLowRankApproxProtocol \((A, \epsilon, k, s)\)
        \(A \in \mathbb{R}^{n \times n \times n}\) was arbitrarily partitioned into \(s\) matrices \(A_{1}, \cdots, A_{s} \in \mathbb{R}^{n \times n \times n}\) on \(s\)
    machines.
3: Coordinator
Machines \(i\)
Chooses a random seed.
Sends it to all machines.
\[
\begin{aligned}
------- & -> \\
& s_{i} \leftarrow O(k / \epsilon), \forall i \in[3] .
\end{aligned}
\]
\[
\text { Agree on } S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}, \forall i \in[3]
\]
\[
\text { which are } w_{1} \text {-wise independent random }
\]
\[
N\left(0,1 / s_{i}\right) \text { Gaussian matrices. }
\]
\[
t_{i} \leftarrow \operatorname{poly}(k / \epsilon), \forall i \in[3] .
\]
\[
\text { Agree on } T_{i} \in \mathbb{R}^{t_{i} \times n}, \forall i \in[3]
\]
\[
\text { which are } w_{2} \text {-wise independent random }
\]
sparse embedding matrices.
\[
\text { Compute } Y_{i, 1} \leftarrow T_{1} A_{i, 1} S_{1}
\]
\[
Y_{i, 2} \leftarrow T_{2} A_{i, 2} S_{2}, Y_{i, 3} \leftarrow T_{3} A_{i, 3} S_{3} .
\]
\[
\text { Send } Y_{i, 1}, Y_{i, 2}, Y_{i, 3} \text { to the coordinator. }
\]
Send \(C_{i} \leftarrow A_{i}\left(T_{1}, T_{2}, T_{3}\right)\) to the coordina-
tor.
```

Compute $Y_{1} \leftarrow \sum_{i=1}^{s} Y_{i, 1}, Y_{2} \leftarrow \sum_{i=1}^{s} Y_{i, 2}$,
$Y_{3} \leftarrow \sum_{i=1}^{s} Y_{i, 3}, C \leftarrow \sum_{i=1}^{s} C_{i}$.
Compute $X_{1}^{*}, X_{2}^{*}, X_{3}^{*}$ by solving
$\min _{X_{1}, X_{2}, X_{3}}\left\|\left(Y_{1} X_{1}\right) \otimes\left(Y_{2} X_{2}\right) \otimes\left(Y_{3} X_{3}\right)-C\right\|_{F}$
Send $X_{1}^{*}, X_{2}^{*}, X_{3}^{*}$ to machines.

$$
-------->
$$

Compute $U_{i}^{*} \leftarrow A_{i, 1} S_{1} X_{1}^{*}$,
$V_{i}^{*} \leftarrow A_{i, 2} S_{2} X_{2}^{*}, W_{i}^{*} \leftarrow A_{i, 3} S_{3} X_{3}^{*}$.
Send $U_{i}^{*}, V_{i}^{*}, W_{i}^{*}$ to the coordinator.
$<--------$
Compute $U^{*} \leftarrow \sum_{i=1}^{s} U_{i}^{*}$.
Compute $V^{*} \leftarrow \sum_{i=1}^{s} V_{i}^{*}$.
Compute $W^{*} \leftarrow \sum_{i=1}^{\bar{s}} W_{i}^{*}$.
return $U^{*}, V^{*}, W^{*}$.
end procedure

```
Algorithm 21.43 Distributed Frobenius Norm Low Rank Approximation Protocol
    procedure DistributedFnormLowRankApproxProtocol \((A, \epsilon, k, s)\)
        \(A \in \mathbb{R}^{n \times n \times n}\) was arbitrarily partitioned into \(s\) matrices \(A_{1}, \cdots, A_{s} \in \mathbb{R}^{n \times n \times n}\) on \(s\)
    machines.
```

4: Chooses a random seed.
5: $\quad$ Sends it to all machines.


```
                                    \(s_{i} \leftarrow O(k / \epsilon), \forall i \in[3]\).
                                    Agree on \(S_{i} \in \mathbb{R}^{n^{2} \times s_{i}}, \forall i \in[3]\)
                                    which are \(w_{1}\)-wise independent random
                                    \(N\left(0,1 / s_{i}\right)\) Gaussian matrices.
                                    \(t_{i} \leftarrow \operatorname{poly}(k / \epsilon), \forall i \in[3]\).
                    Agree on \(T_{i} \in \mathbb{R}^{t_{i} \times n}, \forall i \in[3]\)
                    which are \(w_{2}\)-wise independent random
                    sparse embedding matrices.
                    Compute \(Y_{i, 1} \leftarrow T_{1} A_{i, 1} S_{1}\),
                    \(Y_{i, 2} \leftarrow T_{2} A_{i, 2} S_{2}, Y_{i, 3} \leftarrow T_{3} A_{i, 3} S_{3}\).
                    Send \(Y_{i, 1}, Y_{i, 2}, Y_{i, 3}\) to the coordinator.
                    Send \(C_{i} \leftarrow A_{i}\left(T_{1}, T_{2}, T_{3}\right)\) to the coordina-
    tor.
            Compute \(Y_{1} \leftarrow \sum_{i=1}^{s} Y_{i, 1}, Y_{2} \leftarrow---\quad \sum_{i=1}^{s} Y_{i, 2}\),
            \(Y_{3} \leftarrow \sum_{i=1}^{s} Y_{i, 3}, C \leftarrow \sum_{i=1}^{s} C_{i}\).
            Compute \(X_{1}^{*}, X_{2}^{*}, X_{3}^{*}\) by solving
            \(\min _{X_{1}, X_{2}, X_{3}}\left\|\left(Y_{1} X_{1}\right) \otimes\left(Y_{2} X_{2}\right) \otimes\left(Y_{3} X_{3}\right)-C\right\|_{F}\)
            Send \(X_{1}^{*}, X_{2}^{*}, X_{3}^{*}\) to machines.
                                    \(-------->\)
                                    Compute \(U_{i}^{*} \leftarrow A_{i, 1} S_{1} X_{1}^{*}\),
                    \(V_{i}^{*} \leftarrow A_{i, 2} S_{2} X_{2}^{*}, W_{i}^{*} \leftarrow A_{i, 3} S_{3} X_{3}^{*}\).
                    Send \(U_{i}^{*}, V_{i}^{*}, W_{i}^{*}\) to the coordinator.
            Compute \(U^{*} \leftarrow \sum_{i=1}^{s} U_{i}^{*}\).
            Compute \(V^{*} \leftarrow \sum_{i=1}^{s} V_{i}^{*}\).
            Compute \(W^{*} \leftarrow \sum_{i=1}^{s=1} W_{i}^{*}\).
            return \(U^{*}, V^{*}, W^{*}\).
    end procedure
```


## Chapter 22

## Orthogonal Tensor Decomposition

A recent work [Wang, Tung, Smola, and Anandkumar, NIPS 2015] gives the fastest known algorithms for orthogonal tensor decomposition with provable guarantees. Their algorithm is based on computing sketches of the input tensor, which requires reading the entire input. We show in a number of cases one can achieve the same theoretical guarantees in sublinear time, i.e., even without reading most of the input tensor. Instead of using sketches to estimate inner products in tensor decomposition algorithms, we use importance sampling. To achieve sublinear time, we need to know the norms of tensor slices, and we show how to do this in a number of important cases. For symmetric tensors $A=\sum_{i=1}^{k} \lambda_{i} u_{i}^{\otimes p}$ with $\lambda_{i}>0$ for all $i$, we estimate such norms in sublinear time whenever $p$ is even. For the important case of $p=3$ and small values of $k$, we can also estimate such norms. For asymmetric tensors sublinear time is not possible in general, but we show if the tensor slice norms are just slightly below $\|A\|_{F}$ then sublinear time is again possible. One of the main strengths of our work is empirical - in a number of cases our algorithm is orders of magnitude faster than existing methods with the same accuracy.

### 22.1 Introduction

Tensors are a powerful tool for dealing with multi-modal and multi-relational data. In recommendation systems, often using more than two attributes can lead to better recommendations. This could occur, for example, in Groupon where one could look at users, activities, and time (season, time of day, weekday/weekend, etc.), as three attributes to base predictions on (see [Moi14b] for a discussion). Similar to low rank matrix approximation, we seek a tensor decomposition to succinctly store the tensor and to apply it quickly. A popular decomposition method is the canonical polyadic decomposition, i.e., the CANDECOMP/PARAFAC (CP) decomposition, where the tensor is decomposed into a sum of rank-1 components [Har70]. We refer the reader to [WTSA15], where applications of CP including data mining, computational neuroscience, and statistical learning for latent variable models are mentioned.

A natural question, given the emergence of large data sets, is whether such decompositions can be performed quickly. There are a number of works on this topic [Tso10, PTC13, CV14, KPHF12, $\mathrm{HNH}^{+} 13, \mathrm{BS} 15$, WLSH14]. Most related to ours are several recent works of Wang et al. [WTSA15] and Tung et al. [TWZS15], in which it is shown how to significantly speed up this decomposition for orthogonal tensor decomposition using the randomized technique of linear sketching [PP13]. In this work we also focus on orthogonal tensor decomposition. The idea in [WTSA15] is to create a succinct sketch of the input tensor, from which one can then perform implicit tensor decomposition by approximating inner products in existing decomposition methods.

Existing methods, like the power method, involve computing the inner product of a vector, which is now a rank-1 matrix, with another vector, which is now a slice of a
tensor. Such inner products can be approximated much faster by instead computing the inner product of the sketched vectors, which have significantly lower dimension. One can also replace the sketching with sampling to approximate inner products; we discuss some sampling schemes [Tso10, BS15] below and compare them to our work.

### 22.1.1 Our Contributions

We show in a number of important cases, one can achieve the same theoretical guarantees in the work of Wang et al. [WTSA15] (which was applied later by Tung et al. [TWZS15]), in sublinear time, that is, without reading most of the input tensor. While previous work needs to walk through the input at least once to create a sketch, we show one can instead perform importance sampling of the tensor based on the current iterate, together with reading a few entries of the tensor which help us learn the norms of tensor slices. We use a version of $\ell_{2}$-sampling for our importance sampling. One source of speedup in our work and in Wang et al. [WTSA15] comes from approximating inner products in iterations in the robust tensor power method (see below). To estimate $\langle u, v\rangle$ for $n$-dimensional vectors $u$ and $v$, their work computes sketches $S(u)$ and $S(v)$ and approximates $\langle u, v\rangle \approx\langle S(u), S(v)\rangle$. Instead, if one has $u$, one can sample coordinates $i$ proportional to $u_{i}^{2}$, which is known as $\ell_{2}$-sampling [MW10, CHW12]. One estimates $\langle u, v\rangle$ as $\frac{v_{i}\|u\|_{2}^{2}}{u_{i}}$, which is unbiased and has variance $O\left(\|u\|_{2}^{2}\|v\|_{2}^{2}\right)$. These guarantees are similar to those using sketching, though the constants are significantly smaller (see below), and unlike sketching, one does not need to read the entire tensor to perform such sampling.

Symmetric Tensors. As in [WTSA15], we focus on orthogonal tensor decomposition of symmetric tensors, though we explain the extension to the asymmetric case below. Symmetric tensors arise in engineering applications, for example, to represent the symmetric tensor field of stress, strain, and anisotropic conductivity. Another example is diffusion MRI in which one uses symmetric tensors to describe diffusion in the brain or other parts of the body. In spectral methods symmetric tensors are exactly those that come up in Latent

Dirichlet Allocation problems. Although one can symmetrize a tensor using simple matrix operations (see, e.g., $\left[\mathrm{AGH}^{+} 14\right]$ ), we cannot do this in sublinear time.

In orthogonal tensor decompostion of a symmetric matrix, there is an underlying $n \times n \cdots n$ tensor $A^{*}=\sum_{i=1}^{k} \lambda_{i} v_{i}^{\otimes p}$, and the input tensor is $A=A^{*}+E$, where $\|E\|_{2} \leq \epsilon$. We have $\lambda_{1}>\lambda_{2}>\cdots>\lambda_{k}>0$ and that $\left\{v_{i}\right\}_{i=1}^{k}$ is a set of orthonormal vectors. The goal is to reconstruct approximations $\widehat{v}_{i}$ to the vectors $v_{i}$, and approximations $\widehat{\lambda}_{i}$ to the $\lambda_{i}$. Our results naturally generalize to tensors with different lengths in different dimensions. For simplicity, we first focus on order $p=3$.

In the robust tensor power method $\left[\mathrm{AGH}^{+} 14\right]$, one generates a random initial vector $u$, and performs $T$ update steps $\widehat{u}=A(I, u, u) /\|A(I, u, u)\|_{2}$, where

$$
A(I, u, u)=\left[\sum_{j=1}^{n} \sum_{\ell=1}^{n} A_{1, j, \ell} u_{j} u_{\ell}, \sum_{j=1}^{n} \sum_{\ell=1}^{n} A_{2, j, \ell} u_{j} u_{\ell}, \cdots, \sum_{j=1}^{n} \sum_{\ell=1}^{n} A_{n, j, \ell} u_{j} u_{\ell}\right] .
$$

The matrices $A_{1, *, *}, \ldots, A_{n, *, *}$ are referred to as the slices. The vector $\widehat{u}$ typically converges to the top eigenvector in a small number of iterations, and one often chooses a small number $L$ of random initial vectors to boost confidence. Successive eigenvectors can be found by deflation. The algorithm and analysis immediately extend to higher order tensors.

We use $\ell_{2}$-sampling to estimate $A(I, u, u)$. To achieve the same guarantees as in [WTSA15], for typical settings of parameters (constant $k$ and several eigenvalue assumptions) naïvely one needs to take $O\left(n^{2}\right) \ell_{2}$-samples from $u$ for each slice in each iteration, resulting in $\Omega\left(n^{3}\right)$ time and destroying our sublinearity. We observe that if we additionally knew the squared norms $\left\|A_{1, *, *}\right\|_{F}^{2}, \ldots,\left\|A_{n, *, *}\right\|_{F}^{2}$, then we could take $O\left(n^{2}\right) \ell_{2}$-samples in total, where we take $\frac{\left\|A_{i, *, *}\right\|_{F}^{2}}{\|A\|_{F}^{2}} \cdot O\left(n^{2}\right) \ell_{2}$-samples from the $i$-th slice in expectation. Perhaps in some
applications such norms are known or cheap to compute in a single pass, but without further assumptions, how can one obtain such norms in sublinear time?

If $A$ is a symmetric tensor, then $A_{j, j, j}=\sum_{i=1}^{k} \lambda_{i} v_{i, j}^{3}+E_{j, j, j}$. Note that if there were no noise, then we could read off approximations to the slice norms, since $\left\|A_{j, *, *}\right\|_{F}^{2}=$ $\sum_{i=1}^{k} \lambda_{i}^{2} v_{i, j}^{2}$, and so $A_{j, j, j}^{2 / 3}$ is an approximation to $\left\|A_{j, *, *}\right\|_{F}^{2}$ up to factors depending on $k$ and the eigenvalues. However, there is indeed noise. To obtain non-trivial guarantees, the robust tensor power method assumes $\|E\|_{2}=O(1 / n)$, where

$$
\|E\|_{2}=\sup _{\|u\|_{2}=\|v\|_{2}=\|w\|_{2}=1} E(u, v, w)=\sup _{\|u\|_{2}=\|v\|_{2}=\|w\|_{2}=1} \sum_{i=1}^{n} \sum_{j=1}^{n} \sum_{k=1}^{n} E_{i, j, k} u_{i} v_{j} w_{k},
$$

which in particular implies $\left|E_{j, j, j}\right|=O(1 / n)$. This assumption comes from the $\Theta(1 / \sqrt{n})$ correlation of the random initial vector to $v_{1}$. This noise bound does not trivialize the problem; indeed, $E_{j, j, j}$ can be chosen adversarially subject to $\left|E_{j, j, j}\right|=O(1 / n)$, and if the $v_{i}$ were random unit vectors and the $\lambda_{i}$ and $k$ were constant, then $\sum_{i=1}^{k} \lambda_{i} v_{i, j}^{3}=O\left(1 / n^{3 / 2}\right)$, which is small enough to be completely masked by the noise $E_{j, j, j}$. Nevertheless, there is a lot of information about the slice norms. Indeed, suppose $k=1, \lambda_{1}=\Theta(1)$, and $\|A\|_{F}=1$. Then $A_{j, j, j}=\Theta\left(v_{1, j}^{3}\right)+E_{j, j, j}$, and one can show $\left\|A_{j, *, *}\right\|_{F}^{2}=\lambda_{1}^{2} v_{1, j}^{2} \pm O(1 / n)$. Again using that $\left|E_{j, j, j}\right|=O(1 / n)$, this implies $\left\|A_{j, *, *}\right\|_{F}^{2}=\omega\left(n^{-2 / 3}\right)$ if and only if $A_{j, j, j}=\omega(1 / n)$, and therefore one would notice this by reading $A_{j, j, j}$. There can only be $o\left(n^{2 / 3}\right)$ slices $j$ for which $\left\|A_{j, *, *}\right\|_{F}^{2}=\omega\left(n^{-2 / 3}\right)$, since $\|A\|_{F}^{2}=1$. Therefore, for each of them we can afford to take $O\left(n^{2}\right) \ell_{2}$-samples and still have an $O\left(n^{2+2 / 3}\right)=o\left(n^{3}\right)$ sublinear running time. The remaining slices all have $\left\|A_{j, *, *}\right\|_{F}^{2}=O\left(n^{-2 / 3}\right)$, and therefore if we also take $O\left(n^{1 / 3}\right) \ell_{2}$-samples from every slice, we will also estimate the contribution to $A(I, u, u)$ from these slices well. This is also a sublinear $O\left(n^{2+1 / 3}\right)$ number of samples.

While the previous paragraph illustrates the idea for $k=1$, for $k=2$ we need to read more than the $A_{j, j, j}$ entries to decide how many $\ell_{2}$-samples to take from a slice. The analysis is more complicated because of sign cancellations. Even for $k=2$ we could have $A_{j, j, j}=\lambda_{1} v_{1, j}^{3}+\lambda_{2} v_{2, j}^{3}+E_{j, j, j}$, and if $v_{1, j}=-v_{2, j}$ then we may not detect that $\left\|A_{j, *, *}\right\|_{F}^{2}$ is large. We fix this by also reading the entries $A_{i, j, j}, A_{j, i, j}$, and $A_{j, j, i}$ for every $i$ and $j$. This is still only $O\left(n^{2}\right)$ entries and so we are still sublinear time. Without additional assumptions, we only give a formal analysis of this for $k \in\{1,2\}$.

More importantly, if instead of third-order symmetric tensors we consider $p$-th order symmetric tensors for even $p$, we do not have such sign cancellations. In this case we do not have any restrictions on $k$ for estimating slice norms. One does need to show after deflation, the slice norms can still be estimated; this holds because the eigenvectors and eigenvalues are estimated sufficiently well.

We also give several per-iteration optimizations of our algorithm, based on careful implementations of generating a sorted list of random numbers and random permutations. We find empirically (see below) that we are much faster per iteration than previous sketching algorithms, in addition to not having to read the entire input tensor in a preprocessing step.

Asymmetric Tensors: For asymmetric tensors, e.g., 3rd-order tensors of the form $\sum_{i=1}^{k} \lambda_{i} u_{i} \otimes v_{i} \otimes w_{i}$, it is impossible to achieve sublinear time in general, since it is hard to distinguish $A=e_{i} \otimes e_{j} \otimes e_{k}$ for random $i, j, k \in\{1,2, \ldots, n\}$ from $A=0^{\otimes 3}$. We make a necessary and sufficient assumption that all the entries of the $u_{i}$ are less than $n^{-\gamma}$ for an arbitrarily small constant $\gamma>0$. In this case, all slice norms are $o\left(n^{-\gamma}\right)$ and by taking $O\left(n^{2-\gamma}\right)$ samples from each slice we achieve sublinear time. We can also apply such an assumption to symmetric tensors.

Empirical Results. One of the main strengths of our work is our empirical results. In each iteration we approximate $A(I, u, u)$ a total of $B$ times independently and take the median to increase our confidence. In the notation of [WTSA15], $B$ corresponds to the number of independent sketches used. While the median works empirically, there are some theoretical issues with it discussed in Remark 22.3.1. Also let $b$ be the total number of $\ell_{2}$-samples we take per iteration, which corresponds to the sketch size in the notation of [WTSA15]. We found that empirically we can set $B$ and $b$ to be much smaller than that in [WTSA15] and achieve the same error guarantees. One explanation for this is that the variance bound we obtain via importance sampling is a factor of $4^{3}=64$ smaller than in [WTSA15], and for $p$-th order tensors, a factor of $4^{p}$ smaller.

To give an idea of how much smaller we can set $b$ and $B$, to achieve roughly the same squared residual norm error on the synthetic data sets of dimension 1200 for finding a good rank-1 approximation, the algorithm of [WTSA15] would need to set parameters $b=2^{16}$ and $B=50$, whereas we can set $b=10 \times 1200$ and $B=5$. Our running time is 2.595 seconds and we have no preprocessing time, whereas the algorithm of [WTSA15] has a running time of 116.3 seconds and 55.34 seconds of preprocessing time. We refer the reader to Table 22.1 in Section 22.4. In total we are over 50 times faster.

We also demonstrate our algorithm in a real-world application using real datasets, even when the datasets are sparse. Namely, we consider a spectral algorithm for Latent Dirichlet Allocation $\left[\mathrm{AGH}^{+} 14, \mathrm{AFH}^{+} 12\right]$ which uses tensor decomposition as its core computational step. We show a significant speedup can be achieved on tensors occurring in applications such as LDA, and we refer the reader to Table 22.2 in Section 22.4. For example, on the wiki [WTSA15] dataset with a tensor dimension of 200, we run more than 5
times faster than the sketching-based method.

Previous Sampling Algorithms. Previous sampling-based schemes of [Tso10, BS15] do not achieve our guarantees, because [Tso10] uses uniform sampling, which does not work for tensors with spiky elements, while the non-uniform sampling in [BS15] requires touching all of the entries in the tensor and making two passes over it.

### 22.2 Notation.

Let $[n]$ denote $\{1,2, \ldots, n\}$. Let $\otimes$ denote the outer product, and $u^{\otimes 3}=u \otimes u \otimes u$. Let $A \in \mathbb{R}^{n^{p}}$, where $p$ is the order of tensor $A$ and $n$ is the dimension of tensor $A$. Let $\langle A, B\rangle$ denote the entry-wise inner product between two tensors $A, B \in \mathbb{R}^{n^{p}}$, e.g., $\langle A, B\rangle=$ $\sum_{i_{1}=1}^{n} \sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} A_{i_{1}, i_{2}, \cdots, i_{p}} \cdot B_{i_{1}, i_{2}, \cdots, i_{p}}$. For a tensor $A \in \mathbb{R}^{n^{p}}$,

$$
\|A\|_{F}=\left(\sum_{i_{1}=1}^{n} \sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} A_{i_{1}, \cdots, i_{p}}^{2}\right)^{\frac{1}{2}}
$$

For random variable $X$ let $\mathbb{E}[X]$ denote its expectation of $X$ and $\mathbb{V}[X]$ its variance (if these quantities exist).

For two functions $f, g$, we use the shorthand $f \lesssim g($ resp. $\gtrsim)$ to indicate that $f \leq C g$ (resp. $\geq$ ) for some absolute constant $C$.

### 22.3 Main Results

We explain the details of our main results in this section. First, we state the importance sampling lemmas for our tensor application. Second, we explain how to quickly produce a list of random tuples according to a certain distribution needed by our algorithm. Third, we combine the first and the second parts to get a fast way of approximating tensor contractions, which are used as subroutines in each iteration of the robust tensor power method. We then provide our main theoretical results, and how to estimate the slice norms needed by our main algorithm.

Importance sampling lemmas. Approximating an inner product is a simple application of importance sampling. Tensor contraction $A(u, v, w)$ can be regarded as the inner product between two $n^{3}$-dimensional vectors, and thus importance sampling can be applied. Lemma 22.3.1 suggests that we can take a few samples according to their importance, e.g., we can sample $A_{i, j, k} u_{i} v_{j} w_{k}$ with probability $\left|u_{i} v_{j} w_{k}\right|^{2} /\|u\|_{2}^{2}\|v\|_{2}^{2}\|w\|_{2}^{2}$. As long as the number of samples is large enough, it will approximate the true tensor contraction $\sum_{i} \sum_{j} \sum_{k} A_{i, j, k} u_{i} v_{j} w_{k}$ with small variance after a final rescaling.

Lemma 22.3.1. Suppose random variable $X=A_{i_{1}, \cdots, i_{p}} \prod_{j=1}^{p} u_{j, i_{j}} / \prod_{j=1}^{p} q_{j, i_{j}}$ with probability $\prod_{j=1}^{p} q_{j, i_{j}}$ where $q_{j, i_{j}}=\left|u_{j, i_{j}}\right|^{2} /\left\|u_{j}\right\|_{2}^{2}, \forall j \in[n], \forall i_{j} \in[n]$, and we take $L$ i.i.d. samples $X$, denote $X_{1}, X_{2}, \cdots, X_{L}$. Let $Y=\frac{1}{L} \sum_{\ell=1}^{L} X_{\ell}$. Then $1 . \mathbb{E}[Y]=\left\langle A, u_{1} \otimes u_{2} \otimes \cdots \otimes u_{p}\right\rangle$, and 2 . $\mathbb{V}[Y] \leq \frac{1}{L}\|A\|_{F}^{2} \cdot\left\|u_{1} \otimes u_{2} \otimes \cdots \otimes u_{p}\right\|_{F}^{2}$.

Similarly, we also have importance sampling for each slice $A_{i, *, *}$, i.e., "face" of $A$.
Lemma 22.3.2. For all $i_{1} \in[n]$, suppose random variable $X^{i_{1}}=A_{i_{1}, i_{2}, \cdots, i_{p}} \prod_{j=2}^{p} u_{j, i_{j}} / \prod_{j=2}^{p} q_{j, i_{j}}$ with probability $\prod_{j=2}^{p} q_{j, i_{j}}$ where $q_{j, i_{j}}=\left|u_{j, i_{j}}\right|^{2} /\left\|u_{j}\right\|_{2}^{2}, \forall j \in\{2,3, \cdots, n\}, \forall i_{j} \in[n]$ and we
take $L_{i_{1}}$ i.i.d. samples of $X^{i_{1}}$, say $X_{1}^{i_{1}}, X_{2}^{i_{1}}, \cdots, X_{L}^{i_{1}}$. Let $Y^{i_{1}}=\frac{1}{L_{i_{1}}} \sum_{\ell=1}^{L} X_{\ell}^{i}$. Then 1 . $\mathbb{E}\left[Y^{i}\right]=\left\langle A_{i_{1}, *, \cdots, *}, u_{2} \otimes \cdots \otimes u_{p}\right\rangle$, and $2 . \mathbb{V}\left[Y^{i_{1}}\right] \leq \frac{1}{L_{i_{1}}}\left\|A_{i, *, \cdots, *}\right\|_{F}^{2}\left\|u_{2} \otimes \cdots \otimes u_{p}\right\|_{F}^{2}$.

Generating importance samples in linear time. We need an efficient way to sample indices of a vector based on their importance. We view this problem as follows: imagine $[0,1]$ is divided into $z$ "bins" with different lengths corresponding to the probability of selecting each bin, where $z$ is the number of indices in a probability vector. We generate $m$ random numbers uniformly from $[0,1]$ and see which bin each random number belongs to. If a random number is in bin $i$, we sample the $i$-th index of a vector.

There are known algorithms [BP12, Wal77] to solve this problem in $O(z+m)$ time. We give an alternative algorithm GenRandTuples in Appendix 22.5. Our algorithm combines Bentley and Saxe's algorithm [BS80b] for efficiently generating $m$ sorted random numbers in $O(m)$ time, and Knuth's shuffling algorithm [Knu69] for generating a random permutation of $[m]$ in $O(m)$ time. We use the notation $\operatorname{CumProb}(v, w)$ and $\operatorname{CumProb}(u, v, w)$ for the algorithm creating the distributions on $\mathbb{R}^{n^{2}}$ and $\mathbb{R}^{n^{3}}$ of Lemma 22.3.2 and Lemma 22.3.1, respectively. We note that naïvely applying previous algorithms would require $z=O\left(n^{2}\right)$ and $z=O\left(n^{3}\right)$ time to form these two distributions, but we can take $O(m)$ samples from them implicitly in $O(n+m)$ time.

Fast approximate tensor contractions. We propose a fast way to approximately compute tensor contractions $A(I, v, w)$ and $A(u, v, w)$ with a sublinear number of samples of $A$, as shown in Alogrithm 22.3 and Algorithm 22.4. Naïvely computing tensor contractions using all of the entries of $A$ gives an exact answer but could take $n^{3}$ time. Also, to keep
our algorithm sublinear time, we never explicitly compute the deflated tensor; rather we represent it implicitly and sample from it.

The following theorem gives the error bounds of ApproxTIvw and ApproxTuvw (in Algorithm 22.3 and 22.4). Let $\widehat{b}_{i}$ be the number samples we take from slice $i \in[n]$ in ApproxTIVw, and let $\widehat{b}$ denote the total number of samples in our algorithm.

Theorem 22.3.3. For any $p \geq 3$, given tensor $A \in \mathbb{R}^{n^{p}}$, for any unit vector $u, v \in \mathbb{R}^{n}$, there is an algorithm that takes $\widehat{b}_{i}$ samples from $i$-th slice and it is able to output a value $\bar{A}(u, \cdots, u) \in \mathbb{R}$ and a vector $\bar{A}(I, u, \cdots, u) \in \mathbb{R}^{n}$ such that for any $b>0$ if $\widehat{b}_{i} \geq b\left\|A_{i, *, \cdots, *}\right\|_{F}^{2} /\|A\|_{F}^{2}$ then the following bounds hold:

$$
\begin{equation*}
\mathbb{E}\left[\|\bar{E}(I, u, \cdots, u)\|_{2}^{2}\right] \leq n\|A\|_{F}^{2} / b, \text { and } \mathbb{E}\left[|\bar{E}(u, \cdots, u)|^{2}\right] \leq\|A\|_{F}^{2} / b . \tag{22.1}
\end{equation*}
$$

and

$$
\begin{equation*}
\mathbb{E}\left[\left|v^{\top} \bar{E}(I, u, \cdots, u)\right|^{2}\right] \leq\|A\|_{F}^{2} / b \tag{22.2}
\end{equation*}
$$

where $\bar{E}:=\bar{A}-A$.

Eq. (22.2) can be obtained by observing that each random variable $[\bar{E}(I, u, \cdots, u)]_{i}$ is independent and so $\mathbb{E}\left[\left|v^{\top} \bar{E}(I, u, \cdots, u)\right|^{2}\right]=\sum_{i=1}^{n} v_{i}^{2} \frac{\left\|A_{i, *, \cdots, *}\right\|_{F}^{2}}{\widehat{b}_{i}} \leq\left(\sum_{i=1}^{n} v_{i}^{2}\right) \frac{\|A\|_{F}^{2}}{b}=\frac{\|A\|_{F}^{2}}{b}$. Remark 22.3.1. In [WTSA15], the coordinate-wise median of $B$ estimates to the $A(I, v, w)$ is used to boost the success probability. There appears to be a gap [Wan16] in their argument as it is unclear how to achieve (22.2) after taking a coordinate-wise median, which is (7) in Theorem 1 of [WTSA15]. To fix this, we instead pay a factor proportional to the number of iterations in Algorithm 22.5 in the sample complexity $\widehat{b}$. Since we have expectation bounds on the quantities in Theorem 22.3.3, we can apply a Markov bound and a union bound across


Figure 22.1: Sketching v.s. importance sampling. Running time with growing dimension
all iterations. This suffices for our main theorem concerning sublinear time below. One can obtain high probability bounds by running Algorithm 22.5 multiple times independently, and taking coordinate-wise medians of the output eigenvectors. Empirically, our algorithm works even if we take the median in each iteration, which is done in line 10 in Algorithm 22.3.

Replacing Theorem 1 in [WTSA15] by our Theorem 22.3.3, the rest of the analysis in [WTSA15] is unchanged. Our Algorithm 22.5 is the same as the sketching-based robust tensor power method in [WTSA15], except for lines 9, 11, 15, and 17, where the sketchingbased approximate tensor contraction is replaced by our importance sampling procedures ApproxTuvw and ApproxTIvw. Rather than use Theorem 2 of Wang et al. [WTSA15], the main theorem concerning the correctness of the robust tensor decomposition algorithm, we use a recent improvement of it by Wang and Anandkumar in Theorems 4.1 and 4.2 of [WA16], which states general guarantees for any algorithm satisfying per iteration noise guarantees. These theorems also remove many of the earlier eigenvalue assumptions in Theorem 2 of [WTSA15].

We combine techniques from [WTSA15], [WA16] and [AGH $\left.{ }^{+} 14\right]$ and extend the robust tensor power method analysis to any order $p \geq 3$.

Theorem 22.3.4. (Main, Abitrary order robust tensor power method) For any $p \geq 3, k \geq 1$, for any tensor $A=A^{*}+E \in \mathbb{R}^{n^{p}}$, where $A^{*}=\sum_{i=1}^{k} \lambda_{i} v_{i}^{\otimes p}$ with $\lambda_{i}>0$ and orthonormal basis vectors $\left\{v_{1}, \ldots, v_{k}\right\} \subseteq \mathbb{R}^{n}, n \geq k$. Let $\lambda_{1}, \lambda_{k}$ be the largest and smallest values in $\left\{\lambda_{i}\right\}_{i=1}^{k}$ and $\left\{\widehat{\lambda}_{i}, \widehat{v}_{i}\right\}_{i=1}^{k}$ be outputs of the robust tensor power method. For any sufficiently large constant $c_{0} \geq 100$, there exists a sufficiently small constant $c>0$, for any $\epsilon \in\left(0, c \lambda_{k} /\left(c_{0} p^{2} k n^{(p-2) / 2}\right)\right.$ if $E$ satisfies that

$$
\begin{equation*}
\|E\|_{2} \leq \epsilon /\left(c_{0} \sqrt{n}\right) \tag{22.3}
\end{equation*}
$$

and $T=\Omega\left(\log \left(\lambda_{1} n / \epsilon\right)\right), L=\Omega(k \log (k))$, then with probability at least $9 / 10$, there exists a permutation $\pi:[k] \rightarrow[k]$, such that $\forall i \in[k]$,

$$
\begin{equation*}
\left|\lambda_{i}-\widehat{\lambda}_{\pi(i)}\right| \leq \epsilon, \quad\left\|v_{i}-\widehat{v}_{\pi(i)}\right\|_{2} \leq \epsilon / \lambda_{i} \tag{22.4}
\end{equation*}
$$

Combining the previous theorem with our importance sampling analysis, we obtain:
Theorem 22.3.5 (Main, Sublinear time importance sampling robust tensor power method). Assume the notation of Theorem 22.3.4. For $p=3$, there exists some sufficiently small constant $\gamma>0$ such that $k \in\left[1, n^{\gamma}\right]$ and $\lambda_{k} \geq 1 / n^{\gamma}$, then there exists some sufficiently small constant $\alpha \in(0,1-10 \gamma)$, if $E$ satisfies (22.3), there exists an algorithm takes $O\left(n^{3-\alpha}\right)$ time and uses $O(n k)$ space such that with probability $9 / 10$ the output $\widehat{\lambda}_{i}, \widehat{v}_{i}$ satisfies (22.4).

Proof. For each $j \in[k]$, suppose we take $\widehat{b}^{(j)}=\sum_{i=1}^{n} \widehat{b}_{i}^{(j)}$ samples during the power iterations
for recovering $\widehat{\lambda}_{j}$ and $\widehat{v}_{j}$, the number of samples for slice $i$ is

$$
\widehat{b}_{i}^{(j)} \gtrsim b\left\|\left[A-\sum_{l=1}^{j-1} \widehat{\lambda}_{l} \widehat{v}_{l}^{\otimes 3}\right]_{i, *, *}\right\|_{F}^{2} /\left\|A-\sum_{l=1}^{j-1} \widehat{\lambda}_{l} \widehat{v}_{l}^{\otimes 3}\right\|_{F}^{2}
$$

where $b \gtrsim m \cdot n\|A\|_{F}^{2} / \epsilon^{2}$ and $m \geq k^{2} L T$. Then the output guarantees of Theorem 22.3.4 hold for Algorithm 22.5 with constant probability. Our total time is $O(m \widehat{b})$ and the space is $O(n k)$, where $\widehat{b}=\max _{j \in[k]} \widehat{b}^{(j)}$.

In Theorem 22.3.3, if we require $\widehat{b}_{i}=b\left\|A_{i, *, \cdots, *}\right\|_{F}^{2} /\|A\|_{F}^{2}$, we need to scan the entire tensor to compute $\left\|A_{i, *, \ldots, *}\right\|_{F}^{2}$, making our algorithm not sublinear. With the following mild assumption in Theorem 22.3.6, our algorithm is sublinear when sampling uniformly $\left(\widehat{b}_{i}=b / n\right)$ without computing $\left\|A_{i, *, \cdots, *}\right\|_{F}^{2}$ :

Theorem 22.3.6 (Bounded slice norm). There exists a sufficiently small constant $\gamma>0$ and a constant $\beta \in(0,1]$, there exists a constant $\alpha>0$ (that depends on $\gamma, \beta)$ such that for any order-p tensor $A=A^{*}+E \in \mathbb{R}^{n^{p}}$ with $\operatorname{rank}\left(A^{*}\right) \leq n^{\gamma}, p \leq n^{\gamma}, \lambda_{k} \geq 1 / n^{\gamma}$, $\left\|A_{i, *, \cdots, *}\right\|_{F}^{2} \leq \frac{1}{n^{\beta}}\|A\|_{F}^{2}$ for all $i \in[n]$, and $E$ satisfies (22.3), Algorithm 22.5 takes $O\left(n^{p-\alpha}\right)$ time.

The condition $\alpha \in(0,1]$ is a practical one. When $\alpha=1$, all tensor slices have equal Frobenius norm. The case $\alpha=0$ only occurs when $\left\|A_{i, *, \cdots, *}\right\|_{F}=\|A\|_{F}$; i.e., all except one slice is zero. This theorem can also be applied to asymmetric tensors, since the analysis in [WTSA15] can be extended to them.

For certain cases, we can remove the bounded slice norm assumption. The idea is to take a sublinear number of samples from the tensor to obtain upper bounds on all slice
norms. As outlined in Section 22.1, when $p$ is even, because we do not have sign cancellations we can show:

Theorem 22.3.7 (Even order, informal of Theorem 22.9.14). There is a constant $\alpha>0$ and a sufficiently small constant $\gamma>0$, such that, for any even order-p tensor $A=A^{*}+E \in \mathbb{R}^{n^{p}}$ with $\operatorname{rank}\left(A^{*}\right) \leq n^{\gamma}, p \leq n^{\gamma}$ and $\lambda_{k} \geq 1 / n^{\gamma}$ and $E$ satisfying (22.3), Algorithm 22.5 runs in $O\left(n^{p-\alpha}\right)$ time .

As outlined in Section 22.1, for $p=3$ and small $k$ we can take sign considerations into account:

Theorem 22.3.8 (Low rank). There is a constant $\alpha>0$ such that for any symmetric tensor $A=A^{*}+E \in \mathbb{R}^{n^{3}}$ with $E$ satisfying (22.3) and $\operatorname{rank}\left(A^{*}\right) \leq 2$ and $\lambda_{k} \geq 1 / n^{\gamma}$, Algorithm 22.5 runs in $O\left(n^{3-\alpha}\right)$ time.

### 22.4 Experiments

### 22.4.1 Experiment Setup and Datasets

Our implementation shares the same code base ${ }^{1}$ as the sketching-based robust tensor power method proposed in [WTSA15]. We ran our experiments on an i7-5820k CPU with 64 GB of memory in single-threaded mode. We ran two versions of our algorithm: the version with pre-scanning scans the full tensor to accurately measure per-slice Frobenius norms and make samples for each slice in proportion to its Frobenius norm in ApproxTIvw; the version without pre-scanning assumes that the Frobenius norm of each slice is bounded by $\frac{1}{n^{\alpha}}\|A\|_{F}^{2}, \alpha \in(0,1]$ and uses $b / n$ samples per slice, where $b$ is the total number of samples our algorithm makes, analogous to the sketch length $b$ in [WTSA15].

Synthetic datasets. We first generated an orthonormal basis $\left\{v_{i}\right\}_{i=1}^{k}$ and then computed the synthetic tensor as $A^{*}=\sum_{i=1}^{k} \lambda_{i} v_{i}{ }^{\otimes 3}$, with $\lambda_{1} \geq \cdots \geq \lambda_{k}$. Then we normalized $A^{*}$ such that $\left\|A^{*}\right\|_{F}=1$, and added a symmetric Gaussian noise tensor $E$ where $E_{i j l} \sim \mathcal{N}\left(0, \frac{\sigma}{n^{1.5}}\right)$ for $i \leq j \leq l$. Then $\sigma$ controls the noise-to-signal ratio and we kept it as 0.01 in all our synthetic tensors. For the eigenvalues $\lambda_{i}$, we generated three different decays: inverse decay $\lambda_{i}=\frac{1}{i}$, inverse square decay $\lambda_{i}=\frac{1}{i^{2}}$, and linear decay $\lambda_{i}=1-\frac{i-1}{k}$. We also set $k=100$ when generating tensors, since higher rank eigenvalues were almost indistinguishable from the added noise. To show the scalability of our algorithm, we generated tensors with different dimensions: $n=200,400,600,800,1000,1200$.

[^77]Real-life datasets. Latent Dirichlet Allocation [BNJ01] (LDA) is a powerful generative statistical model for topic modeling. A spectral method has been proposed to solve LDA models $\left[\mathrm{AGH}^{+} 14, \mathrm{AFH}^{+} 12\right]$ and the most critical step in spectral LDA is to decompose a symmetric $K \times K \times K$ tensor with orthogonal eigenvectors, where $K$ is the number of modeled topics. We followed the steps in [AGH ${ }^{+} 14$, TWZS15] and built a $K \times K \times K$ tensor $A_{\text {LDA }}$ for each dataset, and then ran our algorithms directly on $A_{\text {LDA }}$ to see how it works on those tensors in real applications. In our experiments we keep $K=200$. We used the two same datasets as the previous work [WTSA15]: Wiki and Enron, as well as four additional real-life datasets. We refer the reader to our GitHub repository ${ }^{2}$ for our code and full results.

[^78]
### 22.4.2 Results

We considered running time and the squared residual norm to evaluate the performance of our algorithms. Given a tensor $A \in \mathbb{R}^{n^{3}}$, let $\left\|A-\sum_{i=1}^{k} \lambda_{i} u_{i} \otimes v_{i} \otimes w_{i}\right\|_{F}^{2}$ denote the squared residual norm where $\left\{\left(\lambda_{1}, u_{1}, v_{1}, w_{1}\right), \cdots,\left(\lambda_{k}, u_{k}, v_{k}, w_{k}\right)\right\}$ are the eigenvalue/eigenvectors obtained by the robust power method. To reduce the experiment time we looked only for the first eigenvalue and eigenvector, but our algorithm is capable of finding any number of eigenvalues/eigenvectors. We list the pre-scanning time as preprocessing time in tables. It only depends on the tensor dimension $n$ and unlike the sketching based method, it does not depend on $b$. Pre-scanning time is very short, because it only requires one pass of sequential access to the tensor which is very efficient on hardware.

Sublinear time verification. Our theoretical result suggests the total number of samples $b_{\text {no-prescan }}$ for our algorithm without pre-scanning is $n^{1-\alpha}(\alpha \in(0,1])$ times larger than $b_{\text {prescan }}$ for our algorithm with pre-scanning. But in experiments we observe that when $b_{\text {no-prescan }}=$ $b_{\text {prescan }}$ both algorithms achieve very similar accuracy, indicating that in practice $\alpha \approx 1$.

Synthetic datasets. We ran our algorithm on a large number of synthetic tensors with different dimensions and different eigengaps. Table 22.1 shows results for a tensor with 1200 dimensions with 100 non-zero eigenvalues decaying as $\lambda_{i}=\frac{1}{i^{2}}$. To reach roughly the same residual norm, the running time of our algorithm is over 50 times faster than that of the sketching-based robust tensor power method, thanks to the fact that we usually need a relatively small $B$ and $b$ to get a good residual, and the hidden constant factor in running time of sampling is much smaller than that of sketching.

Our algorithm scales well on large tensors due to its sub-linear nature. In Figure 22.1(a), for the sketching-based method we kept $b=2^{16}, B=30$ for $n \leq 800$ and $B=50$ for $n>800$ (larger $n$ requires more sketches to observe a reasonable recovery). For our algorithm, we chose $b$ and $B$ such that for each $n$, our residual norm is on-par or better than the sketching-based method. Our algorithm needs much less time than the sketchingbased one over all dimensions. Another advantage of our algorithm is that there are zero or very minimal preprocessing steps. In Figure 22.1(b), we can see how the preprocessing time grows to prepare sketches when the dimension increases. For applications where only the first few eigenvectors are needed, the preprocessing time could be a large overhead.

Real-life datasets. Due to the small tensor dimension (200), our algorithm shows less speedup than the sketching-based method. But it is still $2 \sim 6$ times faster in each of the six real-life datasets, achieving the same squared residual norm. Table 22.2 reports results for one of the datasets in many different settings of $(b, B)$. Like in synthetic datasets, we also empirically observe that the constant $b$ in importance sampling is much smaller than the $b$ used in sketching to get the same error guarantee.

| Sketching based robust power method: $n=1200, \lambda_{i}=\frac{1}{i^{2}}$ |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  |  | Running time (s) |  |  | Preprocessing time (s) |  |  |
| $b \quad B$ | 10 | 30 | 50 | 10 | 30 | 50 | 10 | 30 | 50 |
| $2^{10}$ | 1.010 | 1.014 | 0.5437 | 0.6114 | 2.423 | 4.374 | 5.361 | 15.85 | 26.08 |
| $2^{12}$ | 1.020 | 0.2271 | 0.1549 | 1.344 | 4.563 | 8.022 | 5.978 | 17.23 | 28.31 |
| $2^{14}$ | 0.1513 | 0.1097 | 0.1003 | 4.928 | 15.51 | 27.87 | 8.788 | 24.72 | 40.4 |
| $2^{16}$ | 0.1065 | 0.09242 | 0.08936 | 22.28 | 69.7 | 116.3 | 13.76 | 34.74 | 55.34 |
| Importance sampling based robust power method (without prescanning): $n=1200, \lambda_{i}=\frac{1}{i^{2}}$ |  |  |  |  |  |  |  |  |  |
|  | Squared residual norm |  |  | Running time (s) |  |  | Preprocessing time (s) |  |  |
| $b \quad B$ | 10 | 30 | 50 | 10 | 30 | 50 | 10 | 30 | 50 |
| $5 n$ | 0.08684 | 0.08637 | 0.08639 | 2.595 | 8.3 | 15.46 | 0.0 | 0.0 | 0.0 |
| $10 n$ | 0.08784 | 0.08671 | 0.08627 | 4.42 | 13.68 | 25.84 | 0.0 | 0.0 | 0.0 |
| $20 n$ | 0.08704 | 0.08700 | 0.08618 | 8.02 | 24.51 | 46.37 | 0.0 | 0.0 | 0.0 |
| $30 n$ | 0.08697 | 0.08645 | 0.08625 | 11.63 | 35.35 | 66.71 | 0.0 | 0.0 | 0.0 |
| $40 n$ | 0.08653 | 0.08664 | 0.08611 | 15.19 | 46.12 | 87.24 | 0.0 | 0.0 | 0.0 |
| Importance sampling based robust power method (with prescanning): $n=1200, \lambda_{i}=\frac{1}{i^{2}}$ |  |  |  |  |  |  |  |  |  |
|  | Squared residual norm |  |  | Running time (s) |  |  | Preprocessing time (s) |  |  |
| $b \quad B$ | 10 | 30 | 50 | 10 | 30 | 50 | 10 | 30 | 50 |
| $5 n$ | 0.08657 | 0.08684 | 0.08636 | 3.1 | 10.47 | 18 | 2.234 | 2.236 | 2.234 |
| $10 n$ | 0.08741 | 0.08677 | 0.08668 | 5.427 | 17.43 | 30.26 | 2.232 | 2.233 | 2.233 |
| $20 n$ | 0.08648 | 0.08624 | 0.08634 | 9.843 | 31.42 | 54.49 | 2.226 | 2.226 | 2.226 |
| $30 n$ | 0.08635 | 0.08634 | 0.08615 | 14.33 | 45.4 | 63.85 | 2.226 | 2.224 | 2.227 |
| $40 n$ | 0.08622 | 0.08652 | 0.08619 | 18.68 | 59.32 | 82.83 | 2.225 | 2.225 | 2.225 |

Table 22.1: Synthetic tensor decomposition using the robust tensor power method. We use an order-3 normalized dense tensor with dimension $n=1200$ with $\sigma=0.01$ noise added. We run sketching-based and sampling-based methods to find the first eigenvalue and eigenvector by setting $L=50, T=30$ and varying $B$ and $b$.

### 22.5 Generate importance sampling indices

In the section we present the full algorithm to generate importance sampling indices.
We view this problem as follows: imagine $[0,1]$ is divided into $n$ "bins" with different lengths corresponding to the probability of selecting each bin. We generate $m$ random numbers uniformly from $[0,1]$ and see which bin each random number belongs to. If a random number is in bin $i$, we sample the $i$-th index of a vector. There are two straightforward ways to generate such indices, but neither of them gives us linear running time in $m$ and $n$. The first method is: (1) generate a list of random numbers (2) for each random number, choose

| Sketching based robust power method: dataset wiki, $\\|$ T $\\|_{F}^{2}=2.135 \mathrm{e}+07$ |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |
| $b \quad B$ | 10 | 30 | 10 | 30 | 10 | 30 |
| $2^{10}$ | $2.091 \mathrm{e}+07$ | $1.951 \mathrm{e}+07$ | 0.2346 | 0.8749 | 0.1727 | 0.2535 |
| $2^{11}$ | $1.971 \mathrm{e}+07$ | $1.938 \mathrm{e}+07$ | 0.4354 | 1.439 | 0.2408 | 0.3167 |
| $2^{12}$ | $1.947 \mathrm{e}+07$ | $1.930 \mathrm{e}+07$ | 1.035 | 2.912 | 0.4226 | 0.4275 |
| $2^{13}$ | $1.931 \mathrm{e}+07$ | $1.927 \mathrm{e}+07$ | 2.04 | 5.94 | 0.5783 | 0.6493 |
| $2^{14}$ | $1.928 \mathrm{e}+07$ | $1.926 \mathrm{e}+07$ | 4.577 | 13.93 | 1.045 | 1.121 |
| Importance sampling based robust power method (without prescanning): dataset wiki, $\\|$ T $\\|_{F}^{2}=2.135 \mathrm{e}+07$ |  |  |  |  |  |  |
|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |
| $b \quad B$ | 10 | 30 | 10 | 30 | 10 | 30 |
| $5 n$ | $1.931 \mathrm{e}+07$ | $1.928 \mathrm{e}+07$ | 0.3698 | 1.146 | 0.0 | 0.0 |
| $10 n$ | $1.931 \mathrm{e}+07$ | $1.929 \mathrm{e}+07$ | 0.5623 | 1.623 | 0.0 | 0.0 |
| $20 n$ | $1.935 \mathrm{e}+07$ | $1.926 \mathrm{e}+07$ | 0.9767 | 2.729 | 0.0 | 0.0 |
| $30 n$ | $1.929 \mathrm{e}+07$ | $1.926 \mathrm{e}+07$ | 1.286 | 3.699 | 0.0 | 0.0 |
| $40 n$ | $1.928 \mathrm{e}+07$ | $1.925 \mathrm{e}+07$ | 1.692 | 4.552 | 0.0 | 0.0 |
| Importance sampling based robust power method (with prescanning): dataset wiki, $\\|$ T $\\|_{F}^{2}=2.135 \mathrm{e}+07$ |  |  |  |  |  |  |
|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |
| $\square \quad B$ | 10 | 30 | 10 | 30 | 10 | 30 |
| $5 n$ | $1.931 \mathrm{e}+07$ | $1.930 \mathrm{e}+07$ | 0.4376 | 1.168 | 0.01038 | 0.01103 |
| $10 n$ | $1.928 \mathrm{e}+07$ | $1.930 \mathrm{e}+07$ | 0.6357 | 1.8 | 0.0104 | 0.01044 |
| $20 n$ | $1.931 \mathrm{e}+07$ | $1.927 \mathrm{e}+07$ | 1.083 | 2.962 | 0.01102 | 0.01042 |
| $30 n$ | $1.929 \mathrm{e}+07$ | $1.925 \mathrm{e}+07$ | 1.457 | 4.049 | 0.01102 | 0.01043 |
| $40 n$ | $1.929 \mathrm{e}+07$ | $1.925 \mathrm{e}+07$ | 1.905 | 5.246 | 0.01105 | 0.01105 |

Table 22.2: Tensor decomposition in LDA on the wiki dataset. The tensor is generated by spectral LDA with dimension $200 \times 200 \times 200$. It is symmetric but not normalized. We fix $L=50, T=30$ and vary $B$ and $b$.
the correct bin for it. Step (1) takes $O(m)$ time, but step (2) takes $O(m n)$ time in the worst case. The second method is: (1) generate a list of random numbers (2) sort these random numbers (3) sequentially for each random number, choose the correct bin for it. Step (1) takes $O(m)$ time, step (2) takes $O(m \log m)$ time, and step (3) takes $O(m+n)$ time as it merges two sorted lists.

In Algorithm 22.1, we instead present an $O(m+n)$ time algorithm which we call GENRandTuples for generating such indices by combining existing algorithms. We use Bentley and Saxe's algorithm GEnSortedRandN $(m)$ [BS80b] to efficiently generate $m$ sorted random numbers $(a)$ in $O(m)$ time. Then RandNToBins outputs a list of $m$ indices $(e)$ such
that for each random number in $a$, there is an index in $\widetilde{q}$. The running time of this step is $O(m+n)$. We use Knuth's shuffling algorithm GenRandPerm $(m)$ [Knu69] to generate a random permutation of $[m](c)$ in $O(m)$ time. Finally, GenRandTuples generates $m$ tuples with two (resp. three) entries given $n$ probabilities in array $\widetilde{q}$ for computing $A(I, v, w)$ $(\operatorname{resp} . A(u, v, w))$.

Claim 22.5.1. Function GEnSortedRandN $(m)$ uses Bentley and Saxe's algorithm [BS80b] and takes $O(m)$ time to generate a list of random numbers between 0 and 1 in increasing order.

Proof. For the correctness of this claim, we refer the reader to [BS80b].

Claim 22.5.2. Function RandNToBins $(A, m, Q, n)$ takes $O(m+n)$ time.

Proof. The input of RandNToBins is two arrays, $A$ and $Q$. The array $A$ is a list of sorted random numbers in $[0,1]$. The array $Q$ can be thought as a list of boundary values for bins, e.g., if some entry $x$ satisfies $Q[j]<x \leq Q[j+1]$, then we should assign $x$ to bin $j$. This procedure is equivalent to merging two sorted lists. Thus, it takes linear time in the summation of two lengths.

Claim 22.5.3. Function $\operatorname{GenRandPERM}(m)$ uses Knuth's shuffle algorithm [Knu69] to generate a random permutation of $[m]$ in $O(m)$ time.

Proof. For the correctness of this claim, we refer the reader to [Knu69] and [Wik16].

```
Algorithm 22.1 Generate importance sampling indices
    function GenRandTuples \((m, \widetilde{q})\)
    \(a[1 \ldots m] \leftarrow\) GEnSortedRandN \((m)\)
    \(e[1 \ldots m] \leftarrow \operatorname{RandNToBins}(a, m, \widetilde{q}, n)\)
    \(c[1 \ldots m] \leftarrow \operatorname{GEnRANDPERM}(m)\)
    \(\mathcal{L} \leftarrow \emptyset\)
    for \(i=1 \rightarrow m\) do
        \(\mathcal{L} \leftarrow \mathcal{L} \cup e[c[i]]\)
    end for
    return \(\mathcal{L}\)
    end function
    function RandNToBins \((A, m, Q, n)\)
    \(i \leftarrow 1, j \leftarrow 1\)
    while \(i \leq m\) do
        if \(A[i]<Q[j]\) then
            \(C[i] \leftarrow j, i \leftarrow i+1\)
        else
        \(j \leftarrow j+1\)
    end if
    end while
    return \(C[1 \ldots m]\)
    end function
    function GEnSortedRandN \((m)\) [BS80b]
    \(s \leftarrow 0\)
    for \(i=1 \rightarrow L\) do
    \(s \leftarrow s-\log\) (uniform \((0,1))\)
    \(A[i] \leftarrow s\)
    end for
    \(s \leftarrow s-\log (\operatorname{uniform}(0,1))\)
    for \(i=1 \rightarrow L\) do
        \(A[i] \leftarrow A[i] / s\)
    end for
    return \(A[1 \ldots m]\)
    end function
    function GenRandPERM \((m)\) [Knu69]
    for \(i=1 \rightarrow m\) do
        \(C[i] \leftarrow i\)
    end for
    for \(i=1 \rightarrow m\) do
        \(t \leftarrow \operatorname{uniform}(i), C[i] \leftarrow C[t], C[t] \leftarrow \$ 820\)
    end for
    return \(C[1 \ldots m]\)
    end function
```


### 22.6 Importance sampling lemmas

In this section we state some useful facts and give proofs for some lemmas that are not proved in the main text.

We shall make use of the following facts which follow from the definitions of a tensor. The following facts hold immediately by the definition of a tensor's $\ell_{2}$-norm and Frobenius norm.

Fact 22.6.1. For any $p \geq 3$, for any tensor $A \in \mathbb{R}^{n^{p}}$ and $p$ unit vectors $u_{1}, u_{2}, \cdots, u_{p} \in \mathbb{R}^{n}$,

1. $\left\|u_{1} \otimes u_{2} \otimes \cdots \otimes u_{p}\right\|_{F}^{2}=\prod_{j=1}^{p}\left\|u_{j}\right\|_{2}^{2}$.
2. $\left\langle A, u_{1} \otimes \cdots \otimes u_{p}\right\rangle=\sum_{i_{1}=1}^{n} \cdots \sum_{i_{p}=1}^{n} A_{i_{1}, \cdots, i_{p}} \prod_{j=1}^{p} u_{j, i_{j}}$.
3. $\left\langle A_{i, *, \cdots, *}, u_{2} \otimes \cdots \otimes u_{p}\right\rangle=\sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} A_{i, i_{2}, \cdots, i_{p}} \prod_{j=2}^{p} u_{j, i_{j}}$.

We provide the proofs for Lemma 22.3.1 and 22.3.2.
Lemma 22.3.1. Suppose random variable $X=A_{i_{1}, \cdots, i_{p}} \prod_{j=1}^{p} u_{j, i_{j}} / \prod_{j=1}^{p} q_{j, i_{j}}$ with probability $\prod_{j=1}^{p} q_{j, i_{j}}$ where $q_{j, i_{j}}=\left|u_{j, i_{j}}\right|^{2} /\left\|u_{j}\right\|_{2}^{2}, \forall j \in[n], \forall i_{j} \in[n]$, and we take $L$ i.i.d. samples $X$, denote $X_{1}, X_{2}, \cdots, X_{L}$. Let $Y=\frac{1}{L} \sum_{\ell=1}^{L} X_{\ell}$. Then $1 . \mathbb{E}[Y]=\left\langle A, u_{1} \otimes u_{2} \otimes \cdots \otimes u_{p}\right\rangle$, and 2 . $\mathbb{V}[Y] \leq \frac{1}{L}\|A\|_{F}^{2} \cdot\left\|u_{1} \otimes u_{2} \otimes \cdots \otimes u_{p}\right\|_{F}^{2}$.

Proof. Consider the expectation of $Y$,

$$
\mathbb{E}[Y]=\mathbb{E}\left[\frac{1}{L} \sum_{\ell=1}^{L} X_{\ell}\right]=\mathbb{E}\left[X_{\ell}\right]=\sum_{i_{1}=1}^{n} \cdots \sum_{i_{p}=1}^{n}\left(A_{i_{1}, \cdots, i_{p}} \prod_{j=1}^{p} \frac{1}{q_{j, i_{j}}} u_{j, i_{j}}\right) \cdot\left(\prod_{j=1}^{p} q_{j, i_{j}}\right)
$$

Furthermore, we can obtain,

$$
\mathbb{E}[Y]=\sum_{i_{1}=1}^{n} \cdots \sum_{i_{p}=1}^{n} A_{i_{1}, \cdots, i_{p}} \prod_{j=1}^{p} u_{j, i_{j}}=\left\langle A, u_{1} \otimes \cdots \otimes u_{p}\right\rangle .
$$

It remains to upper bound the variance of $Y$,

$$
\mathbb{V}[Y]=\mathbb{E}\left[Y^{2}\right]-(\mathbb{E}[Y])^{2}=\mathbb{E}\left[\frac{1}{L^{2}}\left(\sum_{\ell=1}^{L} X_{\ell}\right)^{2}\right]-\left(\mathbb{E}\left[\frac{1}{L} \sum_{\ell=1}^{L} X_{\ell}\right]\right)^{2}=\frac{1}{L^{2}}\left(\mathbb{E}\left[\left(\sum_{\ell=1}^{L} X_{\ell}\right)^{2}\right]-L^{2}\left(\mathbb{E}\left[X_{\ell}\right]\right)^{2}\right)
$$

Furthermore, we can obtain,

$$
\begin{aligned}
\mathbb{V}[Y] & =\frac{1}{L^{2}}\left(L \mathbb{E}\left[X_{\ell}^{2}\right]+L(L-1) \mathbb{E}\left[X_{\ell}\right]^{2}-L^{2} \mathbb{E}\left[X_{\ell}\right]^{2}\right) \\
& \leq \frac{1}{L} \mathbb{E}\left[X_{\ell}^{2}\right] \\
& =\frac{1}{L}\left(\sum_{i_{1}=1}^{n} \cdots \sum_{i_{p}=1}^{n} A_{i_{1}, \cdots, i_{p}}^{2}\left(\prod_{j=1}^{n} \frac{1}{q_{j, i_{j}}} u_{j, i_{j}}\right)^{2}\left(\prod_{j=1}^{p} q_{j, i_{j}}\right)\right) \\
& =\frac{1}{L} \sum_{i_{1}=1}^{n} \cdots \sum_{i_{p}=1}^{n} A_{i_{1}, \cdots, i_{p}}^{2} \prod_{j=1}^{p}\left\|u_{j}\right\|_{2}^{2} \\
& \leq \frac{1}{L}\|A\|_{F}^{2} \cdot\left\|u_{1} \otimes \cdots \otimes u_{p}\right\|_{F}^{2}
\end{aligned}
$$

Thus, we complete the proof.
Lemma 22.3.2. For all $i_{1} \in[n]$, suppose random variable $X^{i_{1}}=A_{i_{1}, i_{2}, \cdots, i_{p}} \prod_{j=2}^{p} u_{j, i_{j}} / \prod_{j=2}^{p} q_{j, i_{j}}$ with probability $\prod_{j=2}^{p} q_{j, i_{j}}$ where $q_{j, i_{j}}=\left|u_{j, i_{j}}\right|^{2} /\left\|u_{j}\right\|_{2}^{2}, \forall j \in\{2,3, \cdots, n\}, \forall i_{j} \in[n]$ and we take $L_{i_{1}}$ i.i.d. samples of $X^{i_{1}}$, say $X_{1}^{i_{1}}, X_{2}^{i_{1}}, \cdots, X_{L}^{i_{1}}$. Let $Y^{i_{1}}=\frac{1}{L_{i_{1}}} \sum_{\ell=1}^{L} X_{\ell}^{i}$. Then 1. $\mathbb{E}\left[Y^{i}\right]=\left\langle A_{i_{1}, *, \cdots, *}, u_{2} \otimes \cdots \otimes u_{p}\right\rangle$, and $2 . \mathbb{V}\left[Y^{i_{1}}\right] \leq \frac{1}{L_{i_{1}}}\left\|A_{i, *, \cdots, *}\right\|_{F}^{2}\left\|u_{2} \otimes \cdots \otimes u_{p}\right\|_{F}^{2}$.

Proof. Consider the expectation of $Y^{i_{1}}$,

$$
\mathbb{E}\left[Y^{i_{1}}\right]=\mathbb{E}\left[\frac{1}{L} \sum_{\ell=1}^{L} X_{\ell}^{i_{1}}\right]=\mathbb{E}\left[X^{i_{1}}\right]=\sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} A_{i, j, k}\left(\prod_{j=2}^{p} \frac{1}{q_{j, i_{j}}} u_{j, i_{j}}\right)\left(\prod_{j=2}^{p} q_{j, i_{j}}\right)
$$

Furthermore, we can obtain,

$$
\mathbb{E}\left[Y^{i_{1}}\right]=\sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} A_{i_{1}, i_{2}, \cdots, i_{i}} \prod_{j=2}^{p} u_{j, i_{j}}=\left\langle A_{i_{1}, *, \cdots, *,}, u_{2} \otimes \cdots \otimes u_{p}\right\rangle .
$$

Consider the variance of $Y^{i_{1}}$,

$$
\mathbb{V}\left[Y^{i_{1}}\right]=\mathbb{E}\left[\left(Y^{i_{1}}\right)^{2}\right]-\left(\mathbb{E}\left[Y^{i_{1}}\right]\right)^{2}=\frac{1}{L^{2}}\left(L \mathbb{E}\left[\left(X^{i_{1}}\right)^{2}\right]+L(L-1)\left(\mathbb{E}\left[X^{i_{1}}\right]\right)^{2}\right)-\left(\mathbb{E}\left[X^{i_{1}}\right]\right)^{2}
$$

Furthermore, we can obtain,

$$
\begin{aligned}
\mathbb{V}\left[Y^{i_{1}}\right] & =\frac{1}{L}\left(\mathbb{E}\left[\left(X^{i_{1}}\right)^{2}\right]-\left(\mathbb{E}\left[X^{i_{1}}\right]\right)^{2}\right) \\
& \leq \frac{1}{L} \mathbb{E}\left[\left(X^{i_{1}}\right)^{2}\right] \\
& =\frac{1}{L} \sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} A_{i_{1}, i_{2}, \cdots, i_{p}}^{2}\left(\prod_{j=2}^{p} \frac{1}{q_{j, i_{j}}} u_{j, i_{j}}\right)^{2}\left(\prod_{j=2}^{p} q_{j, i_{j}}\right) \\
& =\frac{1}{L} \sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} A_{i_{1}, i_{2}, \cdots, i_{p}}^{2} \prod_{j=2}^{n}\left\|u_{j}\right\|_{2}^{2} \\
& =\frac{1}{L}\left\|A_{i_{1}, *, \cdots, *}\right\|_{F}^{2}\left\|u_{2} \otimes \cdots \otimes u_{p}\right\|_{F}^{2} .
\end{aligned}
$$

Thus, we complete the proof.

### 22.7 Approximate Tensor Contractions

We need a slightly different version of Theorem 1 in [WTSA15]. Here we denote $\widehat{b}_{i}$ as the samples we take for tensor slice $i$.

Theorem 22.3.3. For any $p \geq 3$, given tensor $A \in \mathbb{R}^{n^{p}}$, for any unit vector $u, v \in \mathbb{R}^{n}$, there is an algorithm that takes $\widehat{b}_{i}$ samples from $i$-th slice and it is able to output a value $\bar{A}(u, \cdots, u) \in \mathbb{R}$ and a vector $\bar{A}(I, u, \cdots, u) \in \mathbb{R}^{n}$ such that for any $b>0$ if $\widehat{b}_{i} \geq b\left\|A_{i, *, \cdots, *}\right\|_{F}^{2} /\|A\|_{F}^{2}$ then the following bounds hold:

$$
\begin{equation*}
\mathbb{E}\left[\|\bar{E}(I, u, \cdots, u)\|_{2}^{2}\right] \leq n\|A\|_{F}^{2} / b, \text { and } \mathbb{E}\left[|\bar{E}(u, \cdots, u)|^{2}\right] \leq\|A\|_{F}^{2} / b \tag{22.1}
\end{equation*}
$$

and

$$
\begin{equation*}
\mathbb{E}\left[\left|v^{\top} \bar{E}(I, u, \cdots, u)\right|^{2}\right] \leq\|A\|_{F}^{2} / b . \tag{22.2}
\end{equation*}
$$

where $\bar{E}:=\bar{A}-A$.

## Proof. Part 1. of Equation 22.1

By combining Lemma (22.3.2) and $\widehat{b}_{i} \geq b\left\|A_{i, *, \cdots, *}\right\|_{F}^{2} /\|A\|_{F}^{2}$, we know that

$$
\mathbb{E}\left[\left|\bar{E}(I, u, \cdots, u)_{i}\right|^{2}\right]=\frac{\left\|A_{i, *, *}\right\|^{2}}{b\left\|A_{i, *, *}\right\|_{F}^{2} /\|A\|_{F}^{2}}=\|A\|_{F}^{2} / b .
$$

Part 2 of Equation (22.1) It follows by Lemma 22.3.1.

Equation (22.2) It follows by

$$
\begin{aligned}
\mathbb{E}\left[\left|v^{\top} \bar{E}\right|^{2}\right] & =\mathbb{E}\left[\sum_{i=1}^{n} v_{i}^{2} \bar{E}(I, u, \cdots, u)_{i}^{2}\right] \\
& \leq \sum_{i=1}^{n} v_{i}^{2}\left\|A_{i, *, \cdots, *}\right\|_{F}^{2} / \widehat{b}_{i} \\
& \leq \sum_{i=1}^{n} v_{i}^{2}\|A\|_{F}^{2} / b=\|A\|_{F}^{2} / b .
\end{aligned}
$$

Thus, we complete the proof.

Without loss of generality, we can assume that $\|A\|_{F}^{2}=1$. By choosing sufficiently large $b$, we can guarantee the error bounds of sketch noise $\bar{E}$ to be small.

Corollary 22.7.1. For any $m \geq 3$, if $b \geq m n c_{0}^{2} / \epsilon^{2}$, for any unit vector $u, v \in \mathbb{R}^{n}$, then with probability at least $1-3 / m$ such that,

1. $\|\bar{E}(I, u, \cdots, u)\|_{2} \leq \epsilon / c_{0}$.
2. $|\bar{E}(u, \cdots, u)| \leq \epsilon /\left(c_{0} \sqrt{n}\right)$.
3. $\left|v^{\top} \bar{E}(I, u, \cdots, u)\right| \leq \epsilon /\left(c_{0} \sqrt{n}\right)$.

Proof. Part 1. Using Markov's inequality $\operatorname{Pr}[x \geq a] \leq \mathbb{E}[x] / a$, we have

$$
\operatorname{Pr}\left[\|\bar{E}(I, u, \cdots, u)\|_{2}^{2} \geq\left(\epsilon / c_{0}\right)^{2}\right] \leq \frac{\mathbb{E}\left[\|\bar{E}(I, u, \cdots, u)\|_{2}^{2}\right]}{\left(\epsilon / c_{0}\right)^{2}} \leq \frac{n\|A\|_{F}^{2} / b}{\left(\epsilon / c_{0}\right)^{2}} \leq \frac{1}{m}
$$

where the last step follows by choosing $b \geq m n c_{0}^{2} / \epsilon^{2}$.
Part 2. Similarly, by Markov's inequality, we have

$$
\operatorname{Pr}\left[|\bar{E}(u, u, \cdots, u)|^{2} \geq\left(\epsilon /\left(c_{0} \sqrt{n}\right)\right)^{2}\right] \leq \frac{\mathbb{E}\left[\|\bar{E}(u, u, \cdots, u)\|_{2}^{2}\right]}{\left(\epsilon /\left(c_{0} \sqrt{n}\right)\right)^{2}} \leq \frac{\|A\|_{F}^{2} / b}{\left(\epsilon /\left(c_{0} \sqrt{n}\right)\right)^{2}} \leq \frac{1}{m},
$$

where the last step also follows by choosing $b \geq m n c_{0}^{2} / \epsilon^{2}$.

Part 3. This part also can be proved by Markov's inequality and it requires $b \geq$ $m n c_{0}^{2} / \epsilon^{2}$.

If we compare our Theorem 22.3.3 to the bounds obtained by sketching in [WTSA15], we see that our expectation bounds for $\bar{E}(I, u, \cdots, u)$ and $\bar{E}(u, \cdots, u)$ are the same. As long as the conditions for Theorem 22.3.3 holds, all analysis in [WTSA15, WA16] will follow. To guarantee the noise introduced by sampling satisfies the bounds in Theorem 22.3.4, we require $b \geq n\|A\|_{F}^{2} / \epsilon^{2}=O\left(n^{2} k\right)$.

However, as a condition of Theorem 22.3.3, we require that $\widehat{b}_{i} \gtrsim b\left\|A_{i, *, *}\right\|_{F}^{2} /\|A\|_{F}^{2}$. Because $\left\|A_{i, *, *}\right\|_{F}^{2}$ can be as large as $\|A\|_{F}^{2}$, if we made no assumptions, then $\widehat{b}_{i} \gtrsim b$; a summation over all $n$ slices is $n b=O\left(n^{3} k\right)$, making our algorithm not sublinear time.

We will discuss how to deal with the term $\left\|A_{i, *, *}\right\|_{F}^{2}$ and make our algorithm run in sublinear time in the following paragraphs.

### 22.7.1 Case 1: sampling with known per-slice Frobenius norm

If we do not uniformly take $\widehat{b}_{i} \gtrsim b$ samples per slice, but sample based the Frobenius norm of each slice, in other words, the number of samples for slice $i$ is exactly

$$
\widehat{b}_{i}=b \frac{\left\|A_{i, *, *}\right\|_{F}^{2}}{\|A\|_{F}^{2}}
$$

Then we will need only $\sum_{i=1}^{n} \widehat{b}_{i}=b=O\left(n^{2}\right)$ samples. However, to compute $\left\|A_{i, *, *}\right\|_{F}^{2}$, we have to pre-scan the tensor in $O\left(n^{3}\right)$ time.

Sometimes during the data collection process we can store the tensor's per-slice Frobenius norms as meta data and use them for free. But if we don't know the per-slice Frobenius norms, we have to compute them. This is not ideal because we need to scan the entire tensor before we start power iterations to compute per-slice Frobenius norms. Although in practice computing per-slice Frobenius norms only needs sequential access and can be quite efficient in memory or on storage device, it makes our algorithm depend on the full tensor and run in linear time. However, after making a weak assumption on the per-slice Frobenius norm which can be easily satisfied in practice, we can still obtain sublinear run time, as shown in Case 2.

### 22.7.2 Case 2: sampling without known per-slice Frobenius norm

Theorem 22.3.6 (Bounded slice norm). There exists a sufficiently small constant $\gamma>0$ and a constant $\beta \in(0,1]$, there exists a constant $\alpha>0$ (that depends on $\gamma, \beta)$ such that for any order-p tensor $A=A^{*}+E \in \mathbb{R}^{n^{p}}$ with $\operatorname{rank}\left(A^{*}\right) \leq n^{\gamma}, p \leq n^{\gamma}, \lambda_{k} \geq 1 / n^{\gamma}$, $\left\|A_{i, *, \cdots, *}\right\|_{F}^{2} \leq \frac{1}{n^{\beta}}\|A\|_{F}^{2}$ for all $i \in[n]$, and $E$ satisfies (22.3), Algorithm 22.5 takes $O\left(n^{p-\alpha}\right)$ time.

Proof. Because we assume $\left\|A_{i, *, *}\right\|_{F}^{2} \leq \frac{1}{n^{\alpha}}\|A\|_{F}^{2}$, we can set $\widehat{b}_{i}=\frac{b}{n^{\alpha}} \gtrsim b\left\|A_{i, *, *}\right\|_{F}^{2} /\|A\|_{F}^{2}$ for all $i$ such that the conditions for Theorem 22.3.3 hold. Thus, the number of samples we need is $\sum_{i=1}^{n} \widehat{b}_{i}=b n^{1-\alpha}=O\left(n^{3-\alpha}\right)$.

We should note that the condition $0<\alpha \leq 1$ is a very reasonable assumption. When $\alpha=1$, all tensor slices have the equal Frobenius norm. The case where $\alpha=0$ only occurs when $\left\|A_{i, *, *}\right\|_{F}=\|A\|_{F}$; in other words, all except one slices are zero. Except for this extreme case, our algorithm only needs to take $o\left(n^{3}\right)$ samples from the tensor.

Also, note that this proof does not make any assumption for the symmetry of the tensor. Thus, theorem 22.3.6 also applies to asymmetric tensors, as long as the analysis in [WTSA15] is extended to asymmetric case.

### 22.7.3 Case 3: improve the bounds in case 2

We are still interested in improve the $O\left(n^{3-\alpha}\right)$ running time bound of our algorithm. We want to go one step further and remove the assumption on per-slice Frobenius norm. The basic idea is to take a sublinear number of samples from the tensor to estimate per-slice Frobenius norm. With some properties of the tensor (symmetry, low-rank, orthogonal, etc), achieving a sub-linear runtime bound is possible, as we will discuss below.

It is worth noting that in our experiments we have seen that sampling with or without known Frobenius norm (with or without pre-scanning) does not make a big difference for most datasets. Because sampling with known per-slice Frobenius norm is the best case $\left(\sum_{i=1}^{n} \widehat{b}_{i}=O\left(n^{2}\right)\right.$ ), the algorithm described below with an improved bound on $\sum_{i=1}^{n} \widehat{b}_{i}=$ $\omega\left(n^{2}\right)$ is unlikely to outperform. Thus, we did not conduct experiments on this part, and focus on theoretical discussions here.

### 22.8 Estimating Slice Norms in Sublinear Time

We present the following facts before introducing our lemmas and theorems:

### 22.8.1 Definitions and Facts

Fact 22.8.1. For any symmetric matrix $A \in \mathbb{R}^{n \times n},\|A\|_{2}=\max _{\|u\|_{2}=1}\left|u^{\top} A u\right|$.

Proof. Because $A$ is symmetric, we choose $S$ such that $S^{\top} S=A$,

$$
\|A\|_{2}=\|S\|_{2}^{2}=\max _{\|u\|_{2}=1}\|S u\|_{2}^{2}=\max _{\|u\|_{2}=1}\left|u^{\top} S^{\top} S u\right|=\max _{\|u\|_{2}=1}\left|u^{\top} A u\right| .
$$

The following technical lemmas give necessary bounds on norms that we will use in proofs.

Lemma 22.8.2. For any $E \in \mathbb{R}^{n^{3}}$, if $\|E\|_{2}:=\max _{\|u\|_{2}=\|v\|_{2}=\|w\|_{2}=1}|E(u, v, w)| \leq 1 / n$, then for all $i \in[n]$,

1. $\max _{\|u\|_{2}=1}\|E(I, u, u)\|_{2} \leq\|E\|_{2}$,
2. $\left\|E_{i, *, *}\right\|_{2} \leq 1 / n$,
3. $\left\|E_{i, *, *}\right\|_{F} \leq 1 / \sqrt{n}$,
4. $\forall j \in[n],\left|E_{i, j, j}\right| \leq 1 / n$.

## Proof. Part 1.

$$
\begin{aligned}
\|E\|_{2}^{2} & =\max _{\|u\|_{2}=\|v\|_{2}=\|w\|_{2}=1}|E(u, v, w)|^{2} \\
& =\max _{\|u\|_{2}=\|v\|_{2}=\|w\|_{2}=1}\left(\sum_{i=1}^{n} u_{i} \cdot \sum_{j=1}^{n} \sum_{k=1}^{n} E_{i, j, k} v_{j} w_{k}\right)^{2} \\
& \geq \max _{\|v\|_{2}=\|w\|_{2}=1} \sum_{i=1}^{n} u_{i}^{2} \cdot \sum_{i=1}^{n}\left(\sum_{j=1}^{n} \sum_{k=1}^{n} E_{i, j, k} v_{j} w_{k}\right)^{2} \\
& =\max _{\|v\|_{2}=\|w\|_{2}=1} \sum_{i=1}^{n}\left(\sum_{j=1}^{n} \sum_{k=1}^{n} E_{i, j, k} v_{j} w_{k}\right)^{2} \\
& \geq \max _{\|v\|_{2}} \sum_{i=1}^{n}\left(\sum_{j=1}^{n} \sum_{k=1}^{n} E_{i, j, k} v_{j} v_{k}\right)^{2} \\
& =\max _{\|v\|_{2}}\|E(I, v, v)\|_{2}^{2}
\end{aligned}
$$

where the first inequality follows by setting $u_{i}=c \cdot \sum_{j=1}^{n} \sum_{k=1}^{n} E_{i, j, k} v_{j} v_{k}, \forall i \in[n]$ and $\sum_{i=1}^{n} u_{i}^{2}=1, c$ is a normalizing constant.

Part 2. By definition of $\|\cdot\|_{2}$ for a tensor $E \in \mathbb{R}^{n^{3}}$, we have

$$
\max _{\|u\|_{2}=1}\|E(I, u, u)\|_{2} \leq \max _{\|u\|_{2}=\|v\|_{2}=\|w\|_{2}=1}|E(u, v, w)|
$$

By definition of $E(I, u, u) \in \mathbb{R}^{n}$, we have

$$
E(I, u, u)_{i}=\sum_{j=1}^{n} \sum_{k=1}^{n} E_{i, j, k} u_{j} u_{k}=u^{\top} E_{i, *, *} u
$$

Thus, we can upper bound the $\|\cdot\|_{2}$ for matrix $E_{i, *, *} \in \mathbb{R}^{n^{2}}$ in the following sense,

$$
\begin{array}{rlrl}
1 / n & \geq\|E\|_{2} & \\
& \geq \max _{\|u\|_{2}=1}\|E(I, u, u)\|_{2} & \\
& =\max _{\|u\|_{2}=1}\left(\sum_{i=1}^{n}\left(u^{\top} E_{i, *, *} u\right)^{2}\right)^{1 / 2} & \\
& \geq \max _{\|u\|_{2}=1}\left|u^{\top} E_{i, *, *} u\right| & & \text { for any } i \in[n] \\
& =\left\|E_{i, *, *}\right\|_{2}, & & \text { by Fact } 22.8 .1
\end{array}
$$

which completes the proof of Part 2.
Part 3. Since $\left\|E_{i, *, *}\right\|_{2} \leq 1 / n$ and the rank of $n \times n$ matrix is at most $n$, we have

$$
\left\|E_{i, *, *}\right\|_{F}^{2}=\operatorname{Tr}\left(E_{i, *, *}^{\top} E_{i, *, *}\right)=\sum_{j=1}^{n} \lambda_{j} \leq n \cdot\left\|E_{i, *, *}\right\|_{2}^{2} \leq n \cdot(1 / n)^{2} \leq 1 / n
$$

where $\lambda_{j}$ are the eigenvalues of $E_{i, *, *}^{\top} E_{i, *, *}$ and the largest one is $\left\|E_{i, *, *}\right\|_{2}$, by the definition of $\|\cdot\|_{2}$ norm. The inequality implies $\left\|E_{i, *, *}\right\|_{F} \leq 1 / \sqrt{n}$ and finishes the proof of Part 3.

Part 4. For each slice $i \in[n]$, by $\left\|E_{i, *, *}\right\|_{2} \leq 1 / n$ and definition of $\|\cdot\|_{2}$ of a matrix, we have $\max _{\|u\|_{2}=1}\left|u^{\top} E_{i, *, *} u\right| \leq 1 / n$, choosing $u=e_{j}$ where $e_{j}$ is the vector that $j$ th position is 1 and all the other positions are 0 . Thus, we know for any $j \in[n],\left|E_{i, j, j}\right| \leq 1 / n$. Thus, we complete the proof of Part 4.

Fact 22.8.3. Given $p \geq 3$ and two vectors $u, v \in \mathbb{R}^{n}$ with $\|u\|_{2}=\|v\|_{2}=1$, then

1. $\left\|u^{\otimes p}-v^{\otimes p}\right\|_{F} \leq \sqrt{p}\|u-v\|_{2}$.
2. $\left|u_{i}-\widehat{u}_{i}\right| \leq\|u-v\|_{2}, \forall i \in[n]$.

$$
\text { 3. }\left\|u_{i} \cdot u^{\otimes p-1}-\widehat{u}_{i} \cdot \widehat{u}^{\otimes p-1}\right\|_{F} \leq \sqrt{2 p}\|u-v\|_{2}^{2}, \forall i \in[n] \text {. }
$$

Proof. We have $\|u\|_{2}=\|v\|_{2}=1$, and $\|u-v\|_{2}^{2}=\sum_{i=1}^{n} u_{i}^{2}+v_{i}^{2}-2 u_{i} v_{i}=\|u\|_{2}^{2}+\|v\|_{2}^{2}-2\langle u, v\rangle$, putting it all together, we can get $\langle u, v\rangle=1-\frac{1}{2}\|u-v\|_{2}^{2}$. For any $p \geq 3$, we consider $\left\|u^{\otimes p}-v^{\otimes p}\right\|_{F}^{2}$,

$$
\begin{aligned}
\left\|u^{\otimes p}-v^{\otimes p}\right\|_{F}^{2} & =\sum_{i_{1}=1}^{n} \sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n}\left(u_{i_{1}} u_{i_{2}} \cdots u_{i_{p}}-v_{i_{1}} v_{i_{2}} \cdots v_{i_{p}}\right)^{2} \\
& =\sum_{i_{1}=1}^{n} \sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} u_{i_{1}}^{2} u_{i_{2}}^{2} \cdots u_{i_{p}}^{2}-2 u_{i_{1}} u_{i_{2}} \cdots u_{i_{p}} v_{i_{1}} v_{i_{2}} \cdots v_{i_{p}}+v_{i_{1}}^{2} v_{i_{2}}^{2} \cdots v_{i_{p}}^{2} \\
& =\left(\|u\|_{2}^{2}\right)^{p}-2(\langle u, v\rangle)^{p}+\left(\|v\|_{2}^{2}\right)^{p} \\
& =2-2(\langle u, v\rangle)^{p} \\
& =2-2\left(1-\frac{1}{2}\|u-v\|_{2}^{2}\right)^{p} \\
& \leq p\|u-v\|_{2}^{2}
\end{aligned}
$$

We can upper bound $\left\|u_{i} \cdot u^{\otimes p-1}-\widehat{u}_{i} \cdot \widehat{u}^{\otimes p-1}\right\|_{F}$,

$$
\begin{aligned}
& \left\|u_{i} \cdot u^{\otimes p-1}-\widehat{u}_{i} \cdot \widehat{u}^{\otimes p-1}\right\|_{F} \\
\leq & \left\|u_{i} \cdot u^{\otimes p-1}-u_{i} \cdot \widehat{u}^{\otimes p-1}\right\|_{F}+\left\|u_{i} \cdot \widehat{u}^{\otimes p-1}-\widehat{u}_{i} \cdot \widehat{u}^{\otimes p-1}\right\|_{F} \\
= & \left|u_{i}\right| \cdot\left\|u^{\otimes p-1}-\widehat{u}^{\otimes p-1}\right\|_{F}+\left|u_{i}-\widehat{u}_{i}\right| \cdot\left\|\widehat{u}^{\otimes p-1}\right\|_{F} \\
\leq & 1 \cdot \sqrt{p-1}\|u-v\|_{2}+\|u-v\|_{2} \cdot 1 \\
\leq & \sqrt{2 p}\|u-v\|_{2} .
\end{aligned}
$$

Definition 22.8.1. For any $p \geq 3$, for any tensor $A \in \mathbb{R}^{n^{p}}$, we say vector $\tau \in \mathbb{R}^{n}$ is a "good"
estimate of $A$ if for any $i \in[n]$,

$$
\begin{equation*}
\left\|A_{i, *, \cdots, *}\right\|_{F}^{2} \leq \tau_{i} \tag{22.5}
\end{equation*}
$$

We say vector $\tau \in \mathbb{R}^{n}$ is an "efficient" estimate of $A$, if there exists some constant $\alpha>0$ such that,

$$
\begin{equation*}
\sum_{i=1}^{n} \min \left(n^{p-1} \tau_{i} /\|A\|_{F}^{2}, n^{p-1}\right)=O\left(n^{p-\alpha}\right) \tag{22.6}
\end{equation*}
$$

Remark 22.8.1. For $p=3$, Equation (22.6) implies $\sum_{i=1}^{n} \widehat{b}_{i}=o\left(n^{3}\right)$ as long as we choose $\widehat{b}_{i}=\min \left(n^{2} \tau_{i} /\|A\|_{F}^{2}, n^{2}\right)$. Note that in the most part of this section, we focus on order $p=3$ tensor.

### 22.8.2 Order $p=3$, rank $k=1$ tensor

We start by considering the simplest case, where the tensor $A^{*}$ is assumed to be rank-1, $A^{*}=u^{\otimes 3}$. We can observe the true tensor plus some noise, i.e., $A=A^{*}+E$. As long as the noise $E$ is sufficiently small, we can get a good bound on the tensor's per-slice Frobenius norm by just looking at its diagonals! The intuition is that Frobenious norm of slice $i\left(\left\|A_{i, *, *}\right\|_{F}^{2}\right)$ is in proportional to $u_{i}^{2}$, and the diagonals of the tensor $\left(\left|A_{i, i, i}\right|\right)$ is actually $\left|u_{i}\right|^{3}$, giving us an accurate estimate. We can choose per-slice samples $\widehat{b}_{i}$ accordingly. For those slices that potentially have $\left\|A_{i, *, *}\right\|_{F} \gtrsim 1 / n^{1 / 3}$, to estimate the norm of that slice $i$, we oversample them by choosing all the $n^{2}$ samples of that $n \times n$ matrix $A_{i, *, *}$. Since we can show there are only $o(n)$ slices that have large estimates for $\left\|A_{i, *, *}\right\|_{F}$, then the total number of samples is $o\left(n^{3}\right)$, which is sublinear in the size of tensor $A$.

Lemma 22.8.4 (rank-1 tensor). There exists some constant $\alpha>0$, for any tensor $A=$ $E+u^{\otimes 3} \in \mathbb{R}^{n^{3}}$ with $\|u\|_{2}=1,\|E\|_{2} \leq 1 / n$. There exists an algorithm that takes $O\left(n^{3-\alpha}\right)$ samples/time is able to output a vector $\tau \in \mathbb{R}^{n}$ which is a good and efficient estimate of $A$, i.e., $\tau, A$ satisfies (22.5) and (22.6).

Proof. For each slice $i \in[n]$, we have $\left\|E_{i, *, *}\right\|_{F} \leq 1 / \sqrt{n}$ (by part 3 of Lemma 22.8.2) and $\left\|A_{i, *, *}\right\|_{F}=\left\|E_{i, *, *}+u_{i} \cdot u u^{\top}\right\|_{F}$, which implies

$$
\left|u_{i}\right|-1 / \sqrt{n} \leq\left\|A_{i, *, *}\right\|_{F} \leq\left|u_{i}\right|+1 / \sqrt{n} .
$$

Also, observing the entries from main diagonal of tensor which is this form, $\left|A_{i, i, i}\right|=\mid E_{i, i, i}+$ $u_{i}^{3} \mid$, using $\left|E_{i, i, i}\right| \leq 1 / n$ (by part 4 of Lemma 22.8.2), we have

$$
\left|u_{i}\right|^{3}-1 / n \leq\left|A_{i, i, i}\right| \leq\left|u_{i}\right|^{3}+1 / n .
$$

If there is no noise, the diagonal term is exactly $u_{i}^{3}$, and we can get accurate $\left\|A_{i, *, *}\right\|_{F}=$ $u_{i}^{2}$. However, consider the potential noise term, we have to discuss the following two cases:

Part 1. $\left|A_{i, i, i}\right| \geq 2 / n$. In this case, $\left|u_{i}\right|^{3} \geq 1 / n$ and $\left\|A_{i, *, *}\right\|_{F} \geq\left|u_{i}\right|-1 / \sqrt{n} \geq$ $1 / n^{1 / 3}-1 / \sqrt{n} \geq 1 /\left(2 n^{1 / 3}\right)$. For such slice $i$, we need to take $\widehat{b}_{i}=O\left(n^{2}\right)$ examples to guarantee that the bounds on $\left|\left[\epsilon_{2, A}(u)\right]_{i}\right|$ and $\left\langle\omega, \epsilon_{2, A}(u)\right\rangle^{2}$ are enough.

Let $S$ denote a set of indices such that, for each $i \in S,\left|A_{i, i, i}\right| \geq 2 / n$. By triangle inequality, it means $\left|u_{i}\right|^{3} \geq 1 / n$ and then $\left|u_{i}\right| \geq 1 / n^{1 / 3}$. We can show that $|S|$ is at most $n^{2 / 3}$, because $u$ is a unit vector and there are at most $n^{2 / 3}$ coordinates $i$ could have $\left|u_{i}\right| \geq 1 / n^{1 / 3}$. Overall, the sample complexity for all these slices is $n^{3-1 / 3}$.

Part 2. $\left|A_{i, i, i}\right|<2 / n$. In this case, $\left\|A_{i, *, *}\right\|_{F} \leq\left|u_{i}\right|+1 / \sqrt{n} \leq 3 / n^{1 / 3}+1 / \sqrt{n} \leq 4 / n^{1 / 3}$. Thus, for such slice, we choose $\widehat{b}_{i}=O\left(b\left\|A_{i, *, *}\right\|_{F}^{2} /\|A\|_{F}^{2}\right)=O\left(b / n^{2 / 3}\right)$ samples is sufficient for the error bounds. The number of such slice can be $O(n)$, thus the total number of samples for this part is $O\left(n^{3-2 / 3}\right)$.

We set $\tau_{i}=1$ for part 1 and $\tau_{i}=O\left(1 / n^{2 / 3}\right)$ for part 2 .

Now we consider a further case with tensor rank $k=2$. Because there are two distinct eigenvectors, the way of estimating $\left\|A_{i, *, *}\right\|_{F}$ becomes trickier. Besides the diagonal terms of the tensor, we also need to look at the diagonal terms of each tensor slice, $A_{i, j, j}$. Given parameters $b, \alpha \in(0,1)$ and tensor $A \in \mathbb{R}^{n^{3}}$, we use the sampling procedure in Algorithm 22.2.

### 22.8.3 Order $p=3$, rank $k=2$ tensor

Lemma 22.8.5 (rank-2 tensor). There exists some constant $\alpha>0$, for any tensor $A=$ $E+u^{\otimes 3}+v^{\otimes 3} \in \mathbb{R}^{n^{3}}$ with $\|u\|_{2}=\|v\|_{2}=1, u \cdot v=0,\|E\|_{2} \leq 1 / n$. There exists an algorithm that takes $O\left(n^{3-\alpha}\right)$ samples/time is able to output a vector $\tau \in \mathbb{R}^{n}$ which is a good and efficient estimate of $A$ with probability $1-1 / \operatorname{poly}(n)$, i.e., $\tau, A$ satisfies (22.5) and (22.6).

Proof. We assume that $A=E+u \otimes u \otimes u+v \otimes v \otimes v$. Note that we ignore the eigenvalues $\lambda_{1}$ and $\lambda_{2}$ because it will only be constants before $u \otimes u \otimes u$ and $v \otimes v \otimes v$ and our proof is unaffected.

Because vector $u \in \mathbb{R}^{n}$ is orthogonal to vector $v \in \mathbb{R}^{n}$, then

$$
\begin{equation*}
u^{\top} \cdot v=\sum_{i=1}^{n} u_{i} \cdot v_{i}=0 \tag{22.7}
\end{equation*}
$$

Consider the absolute value on main diagonal of tensor,

$$
\left|A_{i, i, i}\right| \leq\left|E_{i, i, i}\right|+\left|u_{i}^{3}+v_{i}^{3}\right| \leq 1 / n+\left|u_{i}^{3}+v_{i}^{3}\right|
$$

where the first step follows by triangle inequality and the second step follows by Part 3 of Lemma 22.8.2 where $\left|E_{i, i, i}\right| \leq 1 / n$. Now we start to prove the 4 parts of our procedure.

Part 1. Following the same reasoning in Lemma 22.8.4, for $\left|A_{i, i, i}\right| \geq 3 / n$, either one of $u_{i}^{3}$ and $v_{i}^{3}$ is large or both $u_{i}^{3}$ and $v_{i}^{3}$ is large and have the same sign, both of these two situations imply that $\left\|A_{i, *, *}\right\|_{F}^{2}$ is large. Then we need to set $\tau_{i}=1$ and take $\widehat{b}_{i}=O\left(n^{2}\right)$ samples for these slice. Also, the number of such slice is $O\left(n^{2 / 3}\right)$ by using the similar argument in part 1 of Lemma 22.8.4. But unlike the rank-1 case, even if $\left|A_{i, i, i}\right|<3 / n$, we cannot conclude that $\left\|A_{i, *, *}\right\|_{F}^{2}$ is small (i.e., $\left\|A_{i, *, *}\right\|_{F}^{2} \lesssim 1 / n^{2 / 3}$ ), because $u_{i}^{3}$ and $v_{i}^{3}$ can have opposite sign and then cancel each other.

Part 2. For each $i \in[n]$, if $\left|A_{i, i, i}\right|<3 / n$, we take $O\left(b / n^{\alpha}\right)$ samples for slice $i$. By doing this, we can guarantee that the error bounds given by Equation (22.2) are sufficient for slice $i$ as long as $\left\|A_{i, *, *}\right\|_{F}^{2} \leq 5 / n^{\alpha}$.

Now we need to handle the slice $i$ where $\left\|A_{i, *, *}\right\|_{F}^{2}>5 / n^{\alpha}$ but $\left|A_{i, i, i}\right|<3 / n$ in the next part.

Part 3. Let us consider slice $i \in[n]$ such that

$$
\begin{equation*}
\left\|A_{i, *, *}\right\|_{F}^{2}=\left\|E_{i, *, *}+u_{i} \cdot u u^{\top}+v_{i} \cdot v v^{\top}\right\|_{F}^{2}>5 / n^{\alpha}, \text { and }\left|A_{i, i, i}\right|=\left|E_{i, i, i}+u_{i}^{3}+v_{i}^{3}\right|<3 / n . \tag{22.8}
\end{equation*}
$$

Using Part 3 and 4 of Lemma 22.8.2, we have

$$
\begin{equation*}
\left\|E_{i, *, *}\right\|_{F}^{2} \leq 1 / n, \text { and }\left|E_{i, i, i}\right| \leq 1 / n \tag{22.9}
\end{equation*}
$$

Combining Equation (22.8) and (22.9) with triangle inequality implies,

$$
u_{i}^{2}+v_{i}^{2}>4 / n^{\alpha}, \text { and }\left|u_{i}^{3}+v_{i}^{3}\right|<4 / n
$$

Using $u_{i}^{2}+v_{i}^{2}>4 / n^{\alpha}$ we can obtain the following result,

$$
\begin{array}{lll}
\Longrightarrow & \text { either } & u_{i}^{2}>2 / n^{\alpha}
\end{array} \text { or } v_{i}^{2}>2 / n^{\alpha} .
$$

Note that $\left|u_{i}^{3}+v_{i}^{3}\right|<4 / n, 3 \alpha / 2<1$ and $\operatorname{sign}\left(u_{i}\right) \neq \operatorname{sign}\left(v_{i}\right)$, without loss of generality, assume $\left|u_{i}\right|^{3}>2 / n^{3 \alpha / 2}$, then we have $\left|v_{i}\right|^{3}>\left|u_{i}\right|^{3}-\left|u_{i}^{3}+v_{i}^{3}\right|>2 / n^{3 \alpha / 2}-4 / n>1 / n^{3 \alpha / 2}$. Therefore, we can obtain that

$$
\text { both }\left|u_{i}\right|^{3}>1 / n^{3 \alpha / 2} \text { and }\left|v_{i}\right|^{3}>1 /\left(n^{3 \alpha / 2}\right) .
$$

Thus, $\left|u_{i} v_{i}\right|>1 / n^{\alpha}$. Recall that the $i$ is the index we choose at the beginning, and $u, v$ are unit vectors, i.e., $\|u\|_{2}=\|v\|_{2}=1$.

Define set $S \subseteq[n]$, such that for any $j \in S, \operatorname{sign}\left(u_{j}\right)=\operatorname{sign}\left(v_{j}\right)$. Define $\bar{S}=[n] \backslash S$. By Equation (22.7),

$$
0=\sum_{j \in S} u_{j} v_{j}+\sum_{j^{\prime} \in \bar{S}} u_{j}^{\prime} v_{j}^{\prime}
$$

It is obvious that the $i$ satisfying Equation (22.8) belongs to $\bar{S}$, then we can lower bound the contribution from set $\bar{S}$,

$$
\left|\sum_{j^{\prime} \in \bar{S}} u_{j^{\prime}} v_{j^{\prime}}\right| \geq\left|u_{i} v_{i}\right|+\left|\sum_{j^{\prime} \in \bar{S} \backslash\{i\}} u_{j^{\prime}} v_{j^{\prime}}\right|>1 / n^{\alpha}
$$

where the first step follows by triangle inequality and the second step follows by $\left|u_{i} v_{i}\right|>$ $1 / n^{\alpha}$. Using triangle inequality again, we can also lower bound the contribution from set $S$, $\left|\sum_{j \in S} v_{j} u_{j}\right|>1 / n^{\alpha}$.

For any $j \in S,\left|u_{j}^{3}+v_{j}^{3}\right|<4 / n$ and $\operatorname{sign}\left(u_{j}\right)=\operatorname{sign}\left(v_{j}\right)$. Then, $\left|u_{j}^{3}\right|<4 / n$ and $\left|v_{j}^{3}\right|<$ $4 / n$, which implies $\left|u_{j}\right|<2 / n^{1 / 3},\left|v_{j}\right|<2 / n^{1 / 3}$. Furthermore, we can obtain $\left|u_{j} \cdot v_{j}\right|<4 / n^{2 / 3}$, so we can lower bound the number of entries in set $S$,

$$
|S| \geq\left|\sum_{j \in S} u_{j} v_{j}\right| /\left(\max _{j \in S}\left|u_{j} v_{j}\right|\right)>\left(1 / n^{\alpha}\right) /\left(4 / n^{2 / 3}\right)=n^{2 / 3-\alpha} / 4
$$

If we take $\Theta\left(n^{1 / 3+\alpha} \log n\right)$ samples out of $[n]$ slices uniformly at random, then with probability $1-1 / \operatorname{poly}(n)$, there exists at least one index $j$ from set $S$. Since for all $j^{\prime} \in S$, $\operatorname{sign}\left(u_{j^{\prime}}\right)=\operatorname{sign}\left(v_{j^{\prime}}\right)$. Then $u_{j}$ and $v_{j}$ have the same sign. But set $S$ is not good enough, because there is no lower bound for $\left|u_{j} v_{j}\right|$, for all $j \in S$. To remove this issue, we can remove the small entries from set $S$ to obtain another set $F$.

Let $F \subseteq S$ denote the set of coordinates such that for any $j \in F,\left|u_{j} v_{j}\right| \geq 1 /\left(2 n^{1+\alpha}\right)$. We can upper bound the contribution from set $S \backslash F$,

$$
\left|\sum_{j \in S \backslash F} u_{j} v_{j}\right|=\sum_{j \in S \backslash F}\left|u_{j} v_{j}\right| \leq|S \backslash F| \cdot 1 /\left(2 n^{1+\alpha}\right) \leq n \cdot 1 /\left(2 n^{1+\alpha}\right)=1 /\left(2 n^{\alpha}\right)
$$

Thus using triangle inequality, we can lower bound the contribution from set $F$,

$$
\left|\sum_{j \in F} u_{j} v_{j}\right| \geq\left|\sum_{j \in S} u_{j} v_{j}\right|-\left|\sum_{j \in S \backslash F} u_{j} v_{j}\right| \geq 1 / n^{\alpha}-1 /\left(2 n^{\alpha}\right)=1 /\left(2 n^{\alpha}\right)
$$

which implies,

$$
|F| \geq\left|\sum_{j \in F} u_{j} v_{j}\right| /\left(\max _{j \in F}\left|u_{j} v_{j}\right|\right)>\left(1 / 2 n^{\alpha}\right) /\left(4 / n^{2 / 3}\right)=n^{2 / 3-\alpha} / 8
$$

Notice that $|F| \gtrsim|S|$, thus after taking $\Theta\left(n^{1 / 3+\alpha} \log n\right)$ samples over [n], we still can hope to see one $j$ from $F$. Now the idea is, we use index $j$ to find out the large entry on diagonal of slice $A_{j, *, *}$. By doing this, we can recover some index $l$ such that the norm of $l$-th slice is large, but $\left|A_{l, l, l}\right|$ looks small.

For each $j \in F$, let us look at the diagonal entry of slice $A_{j, *, *}$, consider entry $A_{j, l, l}$,

$$
\left|A_{j, l, l}\right|=\left|E_{j, l, l}+u_{j} u_{l}^{2}+v_{j} v_{l}^{2}\right| .
$$

Because both $u_{l}^{2}>2 / n^{\alpha}$ and $v_{l}^{2}>2 / n^{\alpha}$, and either $\left|u_{j}\right| \geq 1 /\left(2 n^{1 / 2+\alpha / 2}\right)$ or $\left|v_{j}\right| \geq 1 /\left(2 n^{1 / 2+\alpha / 2}\right)$, then we have

$$
\left|A_{j, l, l}\right| \geq\left|u_{j} u_{l}^{2}+v_{j} v_{l}^{2}\right|-\left|E_{j, l, l}\right| \geq 1 / n^{1 / 2+3 \alpha / 2}-1 / n \geq 1 /\left(2 n^{1 / 2+3 \alpha / 2}\right)
$$

We need to show that $\left\|A_{j, *, *}\right\|_{F}^{2}$ is small. Since $u_{j}$ and $v_{j}$ have the same sign and $\left|A_{j, j, j}\right| \leq 3 / n$, then $\left|u_{j}\right|^{3} \leq 4 / n$ which implies $\left|u_{j}\right|^{2} \leq 4 / n^{2 / 3}$. Thus $\left|u_{j}^{2}+v_{j}^{2}\right| \leq 8 / n^{2 / 3}$. Furthermore, we
have

$$
\left\|A_{j, *, *}\right\|_{F}^{2} \leq 2\left(\left\|E_{j, *, *}\right\|_{F}^{2}+\left|u_{j}^{2}+v_{j}^{2}\right|\right) \leq 2\left(1 / n+8 / n^{2 / 3}\right) \leq 18 / n^{2 / 3} .
$$

Thus,

$$
\left\{\# l \text { such that }\left|A_{j, l, l}\right|^{2} \geq 1 /\left(4 n^{1+3 \alpha}\right)\right\} \leq \frac{18 / n^{2 / 3}}{1 /\left(4 n^{1+3 \alpha}\right)}=72 n^{1 / 3+3 \alpha}
$$

which means number of $\left|A_{j, l, l}\right|^{2}$ is at least $\Theta\left(1 / n^{1+3 \alpha}\right)$ is at most $O\left(n^{1 / 3+3 \alpha}\right)$.

Part 4. We consider the $i$ such that

$$
\operatorname{sign}\left(u_{i}\right)=\operatorname{sign}\left(v_{i}\right) \text { and }\left|A_{i, i, i}\right|=\left|E_{i, i, i}+u_{i}^{3}+v_{i}^{3}\right| \geq 3 / n .
$$

Using $\left|E_{i, i, i}\right| \leq 1 / n$, then we have either $\left|u_{i}\right| \geq 1 / n^{1 / 3}$ or $\left|v_{i}\right| \geq 1 / n^{1 / 3}$. Let us consider the $j$ th diagonal entry of slice $i$ (which is matrix $A_{i, *, *}$ ),

$$
\left|A_{i, j, j}\right|=\left|E_{i, j, j}+u_{i} u_{j}^{2}+v_{i} v_{j}^{2}\right|
$$

We want to recover the $j$ such that

$$
\left\|A_{j, *, *}\right\|_{F}^{2} \geq 5 / n^{\alpha} .
$$

Then either $u_{j}^{2} \geq 1 / n^{\alpha}$ or $v_{j}^{2} \geq 1 / n^{\alpha}$. Then

$$
\left|A_{i, j, j}\right| \geq\left|u_{i} u_{j}^{2}+v_{i} v_{j}^{2}\right|-\left|E_{i, j, j}\right| \geq 1 / n^{1 / 3+\alpha}-1 / n \geq 1 /\left(2 n^{1 / 3+\alpha}\right)
$$

Thus, $\left|A_{i, j, j}\right|^{2} \geq 1 /\left(4 n^{2 / 3+2 \alpha}\right)$. Because $\|A\|_{F}^{2}=1$, the total number of such entries over entire tensor is at most $O\left(n^{2 / 3+2 \alpha}\right)$.

Sample complexity. Overall the sample complexity over four parts is dominated by $n^{1-\alpha}$ and $n^{2 / 3+4 \alpha}$. To minimize the sample complexity, we can set $1-\alpha=2 / 3+4 \alpha$ which implies $\alpha=1 / 15$. Thus, the sample is $O\left(n^{3-1 / 15}\right)$, which is sublinear in tensor size.

For the rank-2 case, we may also want to recover the second eigenvalue and eigenvector. Because the recovered first eigenvalue $\widehat{\lambda}_{1}$ and eigenvector $\widehat{v}_{1}$ contain noise, we need to show that under this noise, the oversampling procedure still works.

Lemma 22.8.6 (Deflation for rank-2 tensor). Given tensor $A=E+\left(u^{\otimes 3}-\widehat{u}^{\otimes 3}\right)+v^{\otimes 3} \in \mathbb{R}^{n^{3}}$ where $\|u\|_{2}=\|\widehat{u}\|_{2}=\|v\|_{2}=1,\|u-\widehat{u}\|_{2} \leq 1 / \sqrt{n}, u \cdot v=0$ and $\|E\| \leq 1 / n$. There exists an algorithm that takes $O\left(n^{3-\alpha}\right)$ samples/time is able to output a vector $\tau \in \mathbb{R}^{n}$ which is a good and efficient estimate of $A$ with probability $1-1 / \operatorname{poly}(n)$, i.e., $\tau, A$ satisfies (22.5) and (22.6).

Proof. We use $u_{i}$ to denote the $i$-th coordinate of vector $u \in \mathbb{R}^{n}$. Because $\|u-\widehat{u}\|_{2} \leq 1 / \sqrt{n}$, then for all $i \in[n],\left|u_{i}-\widehat{u}_{i}\right| \leq 1 / \sqrt{n}$. The proof sketch follows the proof of Lemma 22.8.4, we need to use the information from $\left|A_{i, i, i}\right|$ and $\left|A_{i, j, j}\right|$ to figure out if $i$-th slice has large Frobenius norm, for any $i, j \in[n]$, we have,

$$
\begin{aligned}
& \left|A_{i, i, i}\right|=\left|E_{i, i, i}+u_{i}^{3}-\widehat{u}_{i}^{3}+v_{i}^{3}\right|, \quad \quad \text { where }\left|E_{i, i, i}\right| \leq 1 / n,\left|u_{i}^{3}-\widehat{u}_{i}^{3}\right| \leq 4 / \sqrt{n} ; \\
& \left|A_{i, j, j}\right|=\left|E_{i, j, j}+u_{i} u_{j}^{2}-\widehat{u}_{i} \widehat{u}_{j}^{2}+v_{i} v_{j}^{2}\right|, \quad \text { where }\left|E_{i, j, j}\right| \leq 1 / n,\left|u_{i} u_{j}^{2}-\widehat{u}_{i} \widehat{u}_{j}^{2}\right| \leq 2 / \sqrt{n} .
\end{aligned}
$$

Consider the Frobenious norm of $i$-th slice, it can be written as

$$
\left\|A_{i, *, *}\right\|_{F}=\left\|E_{i, *, *}+\left(u_{i} \cdot u u^{\top}-\widehat{u}_{i} \cdot \widehat{u} \widehat{u}^{\top}\right)+v_{i} \cdot v v^{\top}\right\|_{F},
$$

where $\left\|E_{i, *, *}\right\|_{F} \leq 1 / \sqrt{n}$ (by part 3 of Lemma 22.8 .2 ) and $\left\|u_{i} \cdot u u^{\top}-\widehat{u}_{i} \cdot \widehat{u} \widehat{u}^{\top}\right\|_{F} \leq 3 / \sqrt{n}($ by part 4 of Fact 22.8.3).

For any $\beta \in(0,1)$, we use $R$ to denote a set of indices such that $\left\{i\left|\left|u_{i}-\widehat{u}_{i}\right|>\right.\right.$
$\left.1 / n^{1 / 2+\beta / 2}, i \in[n]\right\}$. Then we know that $|R| \leq n^{\beta}$, otherwise

$$
\|u-\widehat{u}\|_{2}^{2}=\sum_{i \in[n]}\left|u_{i}-\widehat{u}_{i}\right|^{2} \geq \sum_{i \in R}\left|u_{i}-\widehat{u}_{i}\right|^{2}>1 / n
$$

which is a contradiction to $\|u-\widehat{u}\|_{2} \leq 1 / \sqrt{n}$.
Part 1. Follow the similar reason in Lemma 22.8.4 and 22.8.5, for $\left|A_{i, i, i}\right| \geq 4 / n$, then at least one of $u_{i}^{3}-\widehat{u}_{i}^{3}$ and $v_{i}^{3}$ has large absolute value or both of them are large and have the same sign. Both of these two situations imply that $\left\|A_{i, *, *}\right\|_{F}^{2}$ is large. Then we need to set $\tau_{i}=1$ and take $\widehat{b}_{i}=O\left(n^{2}\right)$ samples for each of these slices. Also, the number of such slice could have

$$
\max \left(\left|u_{i}\right|,\left|v_{i}\right|,\left|\widehat{u}_{i}\right|\right)>1 / n^{1 / 3}
$$

is $O\left(n^{2 / 3}\right)$ by using the similar argument in part 1 of Lemma 22.8.4. Even if $\left|A_{i, i, i}\right|<4 / n$, we cannot conclude that $\left\|A_{i, *, *}\right\|_{F}^{2}$ is small, because $u_{i}^{3}-\widehat{u}_{i}^{3}$ and $v_{i}^{3}$ can have opposite sign and then cancel each other.

Part 2. For each $i \in[n]$, if $\left|A_{i, i, i}\right|<4 / n$, we take $O\left(b / n^{\alpha}\right)$ samples for slice $i$. By doing this, we can guarantee that the error bounds are sufficient for slice $i$ as long as $\left\|A_{i, *, *}\right\|_{F}^{2} \leq 5 / n^{\alpha}$.

Now we need to handle the slice where $\left\|A_{i, *, *}\right\|_{F}^{2}>5 / n^{\alpha}$ but $\| A_{i, i, i} \mid<4 / n$ in the next part.

## Part 3.

Let us consider slice $i \in[n]$ such that

$$
\begin{align*}
\left\|A_{i, *, *}\right\|_{F}^{2} & =\left\|E_{i, *, *}+u_{i} \cdot u u^{\top}-\widehat{u}_{i} \cdot \widehat{u} \widehat{u}^{\top}+v_{i} \cdot v v^{\top}\right\|_{F}^{2}>5 / n^{\alpha},  \tag{22.10}\\
\left|A_{i, i, i}\right| & =\left|E_{i, i, i}+u_{i}^{3}-\widehat{u}_{i}^{3}+v_{i}^{3}\right|<3 / n . \tag{22.11}
\end{align*}
$$

Using Part 3 and 4 of Lemma 22.8.2, we have

$$
\begin{equation*}
\left\|E_{i, *, *}\right\|_{F}^{2} \leq 1 / n, \text { and }\left|E_{i, i, i}\right| \leq 1 / n \tag{22.12}
\end{equation*}
$$

Combining Equation (22.10), (22.11), (22.12) and Fact 22.8 .3 with triangle inequality implies,

$$
\begin{aligned}
2 v_{i}^{2} & \geq 5 / n^{\alpha}-2\left\|E_{i, *, *}\right\|_{F}^{2}-2\left\|u_{i} \cdot u u^{\top}-\widehat{u}_{i} \cdot \widehat{u} \widehat{u}^{\top}\right\|_{F}^{2} \\
& \geq 5 / n^{\alpha}-2 / n-2 \cdot 9 / n \\
& \geq 4 / n^{\alpha}
\end{aligned}
$$

and

$$
\left|u_{i}^{3}-\widehat{u}_{i}^{3}+v_{i}^{3}\right|<6 / n .
$$

Combining $\alpha \in(0,2 / 3)$ with the above two equations, we can conclude that $\operatorname{sign}\left(u_{i}-\widehat{u}_{i}\right) \neq$ $\operatorname{sign}\left(v_{i}\right),\left|v_{i}\right|^{3}>1 / n^{3 \alpha / 2}$ and $\left|u_{i}^{3}-\widehat{u}_{i}^{3}\right| \geq 1 / n^{3 \alpha / 2}$. As long as we choose $\alpha \in(0,1 / 3)$, then it is impossible that

$$
1 / n^{3 \alpha / 2} \leq\left|u_{i}^{3}-\widehat{u}_{i}^{3}\right| \leq 4 / \sqrt{n}
$$

Thus, we know part 3 will never happen.

### 22.8.4 Even order tensor

In this section, we prove the even order case.

Lemma 22.8.7. For any $E \in \mathbb{R}^{n^{4}}$ and any $\beta>0$, if $\|E\|_{2}:=\max _{\|u\|_{2}=\|v\|_{2}=\|w\|_{2}=\|x\|_{2}}|E(u, v, w, x)| \leq$ $1 / n^{1+\beta}$, then for all $i \in[n]$,

1. $\max _{\|v\|_{2}=\|w\|_{2}=\|x\|_{2}=1}\|E(I, v, w, x)\|_{2} \leq\|E\|_{2}$,
2. $\left\|E_{i, *, *, *}\right\|_{2} \leq 1 / n^{1+\beta}$,
3. $\left\|E_{i, *, *, *}\right\|_{F} \leq 1 / \sqrt{n}$,
4. $\forall j \in[n],\left|E_{i, j, j, j}\right| \leq 1 / n^{1+\beta}$.

## Proof. Part 1.

$$
\begin{aligned}
\|E\|_{2}^{2} & =\max _{\|u\|_{2}=\|v\|_{2}=\|w\|_{2}=\|x\|_{2}=1}|E(u, v, w, x)|^{2} \\
& =\max _{\|u\|_{2}=\|v\|_{2}=\|w\|_{2}=\|x\|_{2}=1}\left(\sum_{i=1}^{n} u_{i} \cdot \sum_{j=1}^{n} \sum_{k=1}^{n} E_{i, j, k, t} v_{j} w_{k} x_{t}\right)^{2} \\
& \geq \max _{\|v\|_{2}=\|w\|_{2}=\|x\|_{2}=1} \sum_{i=1}^{n} u_{i}^{2} \cdot \sum_{i=1}^{n}\left(\sum_{j=1}^{n} \sum_{k=1}^{n} E_{i, j, k, t} v_{j} w_{k} x_{t}\right)^{2} \\
& =\max _{\|v\|_{2}=\|w\|_{2}=\|x\|_{2}=1} \sum_{i=1}^{n}\left(\sum_{j=1}^{n} \sum_{k=1}^{n} E_{i, j, k, t} v_{j} w_{k} x_{t}\right)^{2} \\
& =\max _{\|v\|_{2}=\|w\|_{2}=\|x\|_{2}=1}\|E(I, v, w, x)\|_{2}^{2},
\end{aligned}
$$

where the first inequality follows by setting $u_{i}=c \cdot \sum_{j=1}^{n} \sum_{k=1}^{n} \sum_{t=1}^{n} E_{i, j, k, t} v_{j} v_{k} v_{t}, \forall i \in[n]$ and $\sum_{i=1}^{n} u_{i}^{2}=1$.

Part 2. By definition of $\|\cdot\|_{2}$ for a tensor $E \in \mathbb{R}^{n^{4}}$, we have

$$
\max _{\|u\|_{2}=1}\|E(I, v, w, x)\|_{2} \leq \max _{\|u\|_{2}=\|v\|_{2}=\|w\|_{2}=\|x\|_{2}=1}|E(u, v, w, x)|
$$

By definition of $E(I, v, w, x) \in \mathbb{R}^{n}$, we have

$$
E(I, v, w, x)_{i}=\sum_{j=1}^{n} \sum_{k=1}^{n} \sum_{t=1}^{n} E_{i, j, k, t} v_{j} w_{k} x_{t}
$$

Thus, we can upper bound the $\|\cdot\|_{2}$ for matrix $E_{i, *, *, *} \in \mathbb{R}^{n^{3}}$ in the following sense,

$$
\begin{array}{rlr}
1 / n^{1+\beta} & \geq\|E\|_{2} \\
& \geq \max _{\|v\|_{2}=\|w\|_{2}=\|x\|_{2}=1}\|E(I, v, w, x)\|_{2} & \quad \text { by Part } 1 . \\
& =\max _{\|v\|_{2}=\|w\|_{2}=\|x\|_{2}=1}\left(\sum_{i=1}^{n}\left(\sum_{j=1}^{n} \sum_{k=1}^{n} \sum_{t=1}^{n} E_{i, j, k, t} v_{j} w_{k} x_{t}\right)^{2}\right)^{1 / 2} \\
& \geq \max _{\|v\|_{2}=\|w\|_{2}=\|x\|_{2}=1}\left|\sum_{j=1}^{n} \sum_{k=1}^{n} \sum_{t=1}^{n} E_{i, j, k, t} v_{j} w_{k} x_{t}\right| & \\
& =\left\|E_{i, *, *, *}\right\|_{2} \quad & \quad \text { by Fact 22.8.1 }
\end{array}
$$

which completes the proof of Part 2.
Part 3. By the proof of Lemma 22.8.2, $\forall i, j \in[n]$ we have

$$
\left\|E_{i, j, *, *}\right\|_{2} \leq\left\|E_{i, *, *, *}\right\|_{2} \leq 1 / n^{1+\beta}
$$

Since $\left\|E_{i, j, *, *}\right\|_{2} \leq 1 / n^{1+\beta}$ and the rank of $n \times n$ matrix is at most $n$, thus,

$$
\left\|E_{i, *, *, *}\right\|_{F}^{2}=\sum_{j=1}^{n}\left\|E_{i, j, *, *}\right\|_{F}^{2} \leq \sum_{j=1}^{n} n \cdot\left\|E_{i, j, *, *}\right\|_{2}^{2} \leq n \cdot n \cdot\left(1 / n^{1+\beta}\right)^{2} \leq 1 / n^{2 \beta}
$$

which implies $\left\|E_{i, *, *, *}\right\|_{F} \leq 1 / n^{\beta}$ and finishes the proof of Part 3.

Part 4. For each slice $i \in[n]$, by $\left\|E_{i, *, *, *}\right\|_{2} \leq 1 / n^{1+\beta}$ and definition of $\|\cdot\|_{2}$ of a matrix, we have

$$
\max _{\|v\|_{2}=\|w\|_{2}=\|x\|_{2}=1}\left|\sum_{j=1}^{n} \sum_{k=1}^{n} \sum_{t=1}^{n} E_{i, j, k, t} v_{j} w_{k} x_{t}\right| \leq 1 / n
$$

choosing $v=w=x=e_{j}$ where $e_{j}$ is the vector that $j$ th position is 1 and all the other positions are 0 . Thus, we know for any $j \in[n],\left|E_{i, j, j, j}\right| \leq 1 / n^{1+\beta}$. Thus, we complete the proof of Part 4.

Lemma 22.8.8 (Order $p=4$ ). For sufficiently small constant $\gamma>0$, for arbitrarily small constant $\beta>0$, for any $k \leq n^{\gamma}, t<k$, for any given tensor $A=E+\sum_{i=1}^{k} u_{i}^{\otimes 4}-\sum_{i=1}^{t} \widehat{u}_{i}^{\otimes 4} \in$ $\mathbb{R}^{n^{4}}$ with $\|E\|_{2} \leq 1 / n^{1+\beta}, u_{i} \cdot u_{j}=0, \forall i, j \in[k],\left\|\widehat{u}_{i}-u_{i}\right\|_{2} \leq 1 / \sqrt{n}, \forall i \in[t]$. There exists an algorithm that takes $O\left(n^{4-2 \beta}\right)$ samples/time is able to output a vector $\tau \in \mathbb{R}^{n}$ which is a good and efficient estimate of $A$, i.e., $\tau, A$ satisfies (22.5) and (22.6).

Proof. Let $A=E+\sum_{i=1}^{k} u_{i}^{\otimes 4}-\sum_{i=1}^{t} \widehat{u}_{i}^{\otimes 4}$. For any $j \in[n]$, we use $u_{i, j}$ to denote the $j$-th coordinate of vector $u_{i} \in \mathbb{R}^{n}$. We can observe the main diagonal of $A$,

$$
\left|A_{j, j, j, j}\right|=\left|E_{j, j, j, j}+\sum_{i=1}^{t}\left(u_{i, j}^{4}-\widehat{u}_{i, j}^{4}\right)+\sum_{i=t+1}^{k} u_{i, j}^{4}\right| .
$$

Since $\forall i \in[t],\left\|u_{i}-\widehat{u}_{i}\right\|_{2} \leq 1 / \sqrt{n}$, then $\forall j \in[n],\left|u_{i, j}-\widehat{u}_{i, j}\right| \leq 1 / \sqrt{n}$. Furthermore, we have

$$
\left|u_{i, j}^{4}-\widehat{u}_{i, j}^{4}\right|=\left|\left(u_{i, j}-\widehat{u}_{i, j}\right)\left(u_{i, j}+\widehat{u}_{i, j}\right)\left(u_{i, j}^{2}+\widehat{u}_{i, j}^{2}\right)\right| \leq 4 / \sqrt{n} .
$$

By part 4 of Lemma 22.8.7, we have $\left|E_{j, j, j, j}\right| \leq 1 / n$. Because there is no cancellation in $\left|A_{j, j, j, j}\right|$ and $1 / \sqrt{n} \gg 1 / n$, we can use the same threshold $1 / n$ as Lemma 22.8.4 again.

Part 1, $\left|A_{j, j, j, j}\right| \geq 1 / n$. In this case, for $i \in\{t+1, t+2, \cdots, k\}$ one of $\left|u_{i, j}\right|$ is large and $\left\|A_{j, *, *, *}\right\|_{F}^{2}$ is also large. Let $S$ denote a set of indices such that, for each $j \in S$, $\left|A_{j, j, j, j}\right| \geq 1 / n$. We can show an upper bound for $|S|$. For each $j$, there must exists one $i>t$ such that $\left|u_{i, j}\right| \geq 1 /\left(k n^{1 / 4}\right)=1 / n^{1 / 4-\gamma}$. There are at most $k$ different $i$, for each $i$ we know $u_{i}$ is a unit vector and then there are at most $n^{1 / 2-2 \gamma}$ coordinates $j$ of vector $u_{i}$ could have $\left|u_{i, j}\right| \geq 1 / n^{1 / 4-\gamma}$. Thus $|S| \leq k n^{1 / 2-2 \gamma} \leq n^{1 / 2-\gamma}$. For each $i \in S$, we take $O\left(n^{3}\right)$ samples from slice $j$. Then the total number of samples is $O\left(n^{3+1 / 2-\gamma}\right)$.

Part 2, $\left|A_{j, j, j, j}\right|<1 / n$. In this case, we know there is no $i>t$ such that $\left|u_{i, j}\right| \geq$ $2 / n^{1 / 4}$. We can upper bound $\left\|A_{j, *, *, *}\right\|_{F}$,

$$
\begin{aligned}
& \left\|A_{j, *, *, *}\right\|_{F} \\
\leq & \left\|E_{j, *, *, *}\right\|_{F}+\left\|\sum_{i=1}^{t} u_{i, j} \cdot u_{i}^{\otimes 3}-\widehat{u}_{i, j} \cdot \widehat{u}_{i}^{\otimes 3}\right\|_{F}+\left\|\sum_{i=t+1}^{k} u_{i, j} \cdot u_{i}^{\otimes 3}\right\|_{F} \\
\leq & \left\|E_{j, *, *, *}\right\|_{F}+\left\|\sum_{i=1}^{t} u_{i, j} \cdot u_{i}^{\otimes 3}-u_{i, j} \cdot \widehat{u}_{i}^{\otimes 3}\right\|_{F}+\left\|\sum_{i=1}^{t} u_{i, j} \cdot \widehat{u}_{i}^{\otimes 3}-\widehat{u}_{i, j} \cdot \widehat{u}_{i}^{\otimes 3}\right\|_{F}+\left\|\sum_{i=t+1}^{k} u_{i, j} \cdot u_{i}^{\otimes 3}\right\|_{F} \\
\leq & 1 / n^{\beta}+k \cdot 3 / \sqrt{n}+k \cdot 1 / \sqrt{n}+k \cdot 2 / n^{1 / 4} \\
\leq & 3 / n^{\min (\beta, 1 / 4-\gamma)}
\end{aligned}
$$

where the bound of the second term in the third step follows by Fact 22.8.3.
Thus, as long as $\beta<1 / 4-\gamma$, for each such slice $i$, we choose $\widehat{b}_{i}=O\left(b\left\|A_{j, *, *, *}\right\|_{F}^{2} /\|A\|_{F}^{2}\right)=$ $O\left(b / n^{2 \beta}\right)$ samples. Notice that $b=O\left(n^{3}\right)$ in order 4 tensor. The number of such $i$ can be $O(n)$, thus the total sample complexity is $n^{4-2 \beta}$.

We can extend Lemma 22.8.8 to general case where order $p$ is even,

Theorem 22.8.9 (Even order). For sufficiently small constant $\gamma>0$, there exists some constant $\alpha>0$ (depends on $\gamma$ ) for any $k \in\left[1, n^{\gamma}\right], p \in\left[3, n^{\gamma}\right] r<k$, for any given tensor $A=E+\sum_{i=1}^{k} \lambda_{i} u_{i}^{\otimes p}-\sum_{i=1}^{t} \widehat{\lambda}_{i} \widehat{u}_{i}^{\otimes p} \in \mathbb{R}^{n^{p}}$.

For any sufficiently large constant $c_{0} \geq 100$, there exists a sufficiently small constant $c>0$, for any $\epsilon \in\left(0, c \lambda_{k} /\left(c_{0} p^{2} k n^{(p-2) / 2}\right)\right.$ if $E$ satisfies that $\|E\|_{2} \leq \epsilon /\left(c_{0} \sqrt{n}\right)$ and $\forall i \in[t]$.

$$
\left|\widehat{\lambda}_{i}-\lambda_{i}\right| \leq \epsilon, \text { and }\left\|\widehat{u}_{i}-u_{i}\right\|_{2} \leq \epsilon / \lambda_{i}
$$

There exists an algorithm that takes $O\left(n^{p-\alpha}\right)$ samples/time is able to output a vector $\tau \in \mathbb{R}^{n}$ which is a good and efficient estimate of A, i.e., $\tau, A$ satisfies (22.5) and (22.6).

### 22.9 Robust Tensor Power Method Analysis for General Order $p \geq$ 3

In Section 22.9.1, we extend several useful facts from 3rd order and 4th order to arbitrary $p$-th order tensor when $p \geq 3$. In Section 22.9.2, we extend analysis of the original robust tensor power method $\left[\mathrm{AGH}^{+} 14\right]$ from 3rd to arbitrary order $p \geq 3$, some of our proofs also borrow the idea from a very recent work [WA16].

### 22.9.1 Useful facts

Fact 22.9.1. For any three unit vectors $u, v, w \in \mathbb{R}^{n}$, for any $0 \leq x \leq 1$, if $\langle u, w\rangle \geq 1-x$ and $\langle v, w\rangle=0$, then $\langle u, v\rangle \leq \sqrt{2 x-x^{2}}$.

Proof.

$$
\begin{aligned}
\langle u, v\rangle & =\cos \theta(u, v) \\
& \leq|\cos \theta(u, w) \cos \theta(v, w)|+|\sin \theta(u, w) \sin \theta(v, w)| \\
& \leq 0+\sqrt{1-(1-x)^{2}} \\
& =\sqrt{2 x-x^{2}} .
\end{aligned}
$$

Fact 22.9.2. For any $E \in \mathbb{R}^{n^{p}}, u, v \in \mathbb{R}^{n}$, then $\left|v^{\top} E(I, u, \cdots, u)\right|=|E(v, u, \cdots, u)|$.

Proof. It follows

$$
\begin{aligned}
\left|v^{\top} E(I, u, \cdots, u)\right| & =\left|\sum_{i_{1}=1}^{n} v_{i_{1}} \cdot\left(\sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} E_{i_{1}, i_{2}, \cdots, i_{p}} u_{i_{2}} \cdots u_{i_{p}}\right)\right| \\
& =\left|\sum_{i_{1}=1}^{n} \sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} E_{i_{1}, i_{2}, \cdots, i_{p}} v_{i_{1}} u_{i_{2}} \cdots u_{i_{p}}\right| \\
& =|E(v, u, \cdots, u)| .
\end{aligned}
$$

Fact 22.9.3. For any $p \geq 3$, given orthogonal tensor $A^{*}=\sum_{i=1}^{k} \lambda_{j} v_{j}^{\otimes p} \in \mathbb{R}^{n^{p}}$ and vector $u \in \mathbb{R}^{n}$, for any $j \in[k]$, we have $\left|v_{j}^{\top} A^{*}(I, u, \cdots, u)\right|=\lambda_{j}\left|v_{j}^{\top} u\right|^{p-1}$.

Proof. For any $j \in[k]$, we have

$$
\begin{aligned}
\left|v_{j}^{\top} A^{*}(I, u, \cdots, u)\right| & =\left|\sum_{i=1}^{n} v_{j, i} \sum_{\ell=1}^{k} \lambda_{\ell} v_{\ell, i} \sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} v_{\ell, i_{2}} \cdots v_{\ell, i_{p}} u_{i_{2}} \cdots u_{i_{p}}\right| \\
& =\left|\lambda_{j} \sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} v_{j, i_{2}} \cdots v_{j, i_{p}} u_{i_{2}} \cdots u_{i_{p}}\right| \\
& =\left|\lambda_{j}\right|\left|v_{j}^{\top} u\right|^{p-1} .
\end{aligned}
$$

Fact 22.9.4. For any tensor $E \in \mathbb{R}^{n^{p}}$, for any unit vectors $u, v \in \mathbb{R}^{n}$,

1. $|E(u, v, \cdots, v)| \leq\|E\|_{2}$.
2. $\|E(I, v, \ldots, v)\|_{2} \leq \sqrt{n}\|E\|_{2}$.

Proof. The part 1 trivially follows by definition of $\|E\|_{2}$. For part 2 , we define a unit vector $w$ to be $(1 / \sqrt{n}, \cdots, 1 / \sqrt{n})$,

$$
\begin{aligned}
& \|E(I, v, \ldots, v)\|_{2}^{2} \\
= & \sum_{i_{1}=1}^{n}\left(\sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} E_{i_{1}, i_{2}, \cdots, i_{p}} v_{i_{2}} \cdots v_{i_{p}}\right)^{2} \\
= & n \sum_{i_{1}=1}^{n}\left(\sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} E_{i_{1}, i_{2}, \cdots, i_{p}} w_{i_{1}} v_{i_{2}} \cdots v_{i_{p}}\right)^{2} \\
\leq & n\|E\|_{2}^{2}
\end{aligned}
$$

which implies $\|E(I, v, \cdots, v)\|_{2} \leq \sqrt{n}\|E\|_{2}$.

Fact 22.9.5. For any $p \geq 3$, for any unit vectors $x, y, u, v \in \mathbb{R}^{n}$, we have

1. $\left\|\left[x \otimes v^{\otimes(p-1)}\right](I, u, \cdots, u)-\left[y \otimes v^{\otimes(p-1)}\right](I, u, \cdots, u)\right\|_{2}=|\langle u, v\rangle|^{p-1} \cdot\|x-y\|_{2}$.
2. $\forall j \in\{0,1, \cdots, p-2\}$,

$$
\begin{aligned}
& \left\|\left[v^{\otimes(1+j)} \otimes x \otimes v^{\otimes(p-2-j)}\right](I, u, \cdots, u)-\left[v^{\otimes(1+j)} \otimes y \otimes v^{\otimes(p-2-j)}\right](I, u, \cdots, u)\right\|_{2} \\
\leq & |\langle u, v\rangle|^{p-2} \cdot\|x-y\|_{2} .
\end{aligned}
$$

Proof. Part 1. We consider the $i$-th coordinate of vector $\left[x \otimes v^{\otimes(p-1)}\right](I, u, \cdots, u) \in \mathbb{R}^{n}$, it can be written as

$$
x_{i} \sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} v_{i_{2}} \cdots v_{i_{p}} u_{i_{2}} \cdots u_{i_{p}}=x_{i} \sum_{i_{2}=1}^{n} v_{i_{2}} u_{i_{2}} \cdots \sum_{i_{p}=1}^{n} v_{i_{p}} u_{i_{p}}=x_{i}\langle v, u\rangle^{p-1} .
$$

Thus,

$$
\mathrm{LHS}^{2}=\sum_{i=1}^{n}\left(x_{i}\langle v, u\rangle^{p-1}-y_{i}\langle v, u\rangle^{p-1}\right)^{2}=\|x-y\|_{2}^{2} \cdot|\langle v, u\rangle|^{2(p-1)} .
$$

Part 2. We consider the $i$-th coordinate of vector $\left[v^{\otimes(1+j)} \otimes x \otimes v^{\otimes(p-2-j)}\right](I, u, \cdots, u) \in$ $\mathbb{R}^{n}$, it can be written as $v_{i}\langle x, u\rangle \cdot\langle v, u\rangle^{p-2}$. Thus,
$\mathrm{LHS}^{2}=\sum_{i=1}^{n}\left(v_{i}\langle x, u\rangle\langle v, u\rangle^{p-2}-v_{i}\langle y, u\rangle\langle v, u\rangle^{p-2}\right)^{2}=\langle x-y, u\rangle^{2}\langle v, u\rangle^{2(p-2)} \leq\|x-y\|_{2}^{2}\langle v, u\rangle^{2(p-2)}$.

Fact 22.9.6. For any two unit vectors $u$, $v$, if $\theta(u, v) \in(0, \pi / 2)$, then $\|u-v\|_{2} \leq \tan \theta(u, v)$.

Proof. Because $\theta(u, v) \in(0, \pi / 2)$, we have $\cos \theta(u, v) \in(0,1)$. We use $x$ to denote $\langle u, v\rangle$. We want to show $\|u-v\|_{2}^{2} \leq \tan ^{2} \theta(u, v)$, which is equivalent to

$$
\begin{aligned}
& 2-2 x \leq \sin ^{2} \theta(u, v) / \cos ^{2} \theta(u, v) \\
\Longleftrightarrow & 2-2 x \leq 1-x^{2} / x^{2} \\
\Longleftrightarrow & 0 \leq(1-x)\left(1+x-2 x^{2}\right) \\
\Longleftrightarrow & 0 \leq(1-x)^{2}(1+2 x)
\end{aligned}
$$

which always holds. Thus we complete the proof.

Fact 22.9.7. Given an orthonormal basis, $v_{1}, v_{2}, \cdots, v_{n}$. Let $V=\left(v_{2}, \cdots, v_{n}\right) \in \mathbb{R}^{n \times(n-1)}$ and $A^{*}=\sum_{i=1}^{k} \lambda_{i} v_{i}^{\otimes p}$. Then for any vector $u \in \mathbb{R}^{n}$, we have $\left\|V^{\top} A^{*}(I, u, \cdots, u)\right\|_{2}^{2}=$ $\sum_{j=2}^{k} \lambda_{j}^{2}\left|v_{j}^{\top} u\right|^{2(p-1)}$.

Proof.

$$
\begin{aligned}
& \left\|V^{\top} A^{*}(I, u, \cdots, u)\right\|_{2}^{2} \\
= & \sum_{j=2}^{n}\left|v_{j}^{\top} A^{*}(I, u, \cdots, u)\right|^{2} \\
= & \sum_{j=2}^{n}\left(\sum_{i=1}^{n} v_{j, i} A^{*}(I, u, \cdots, u)_{i}\right)^{2} \\
= & \sum_{j=2}^{n}\left(\sum_{i=1}^{n} v_{j, i} \sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n} A_{i, i_{2}, \cdots, i_{p}}^{*} u_{i_{2}} \cdots u_{i_{p}}\right)^{2} \\
= & \sum_{j=2}^{n}\left(\sum_{i=1}^{n} v_{j, i} \sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n}\left(\sum_{\ell=1}^{n} \lambda_{\ell} v_{\ell, i} v_{\ell, i_{2}} \cdots v_{\ell, i_{p}}\right) u_{i_{2}} \cdots u_{i_{p}}\right)^{2} \quad \text { by definition of } A^{*} \\
= & \sum_{j=2}^{n}\left(\sum_{\ell=1}^{k} \lambda_{\ell} \sum_{i=1}^{n} v_{j, i} v_{\ell, i} \sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n}\left(v_{\ell, i_{2}} \cdots v_{\ell, i_{p}}\right) u_{i_{2}} \cdots u_{i_{p}}\right)^{2} \\
= & \sum_{j=2}^{k}\left(\lambda_{j} \sum_{i_{2}=1}^{n} \cdots \sum_{i_{p}=1}^{n}\left(v_{j, i_{2}} \cdots v_{j, i_{p}}\right) u_{i_{2}} \cdots u_{i_{p}}\right)^{2} \\
= & \sum_{j=2}^{k} \lambda_{j}^{2}\left|v_{j}^{\top} u\right|^{2(p-1)}
\end{aligned}
$$

### 22.9.2 Convergence guarantee and deflation

Lemma 22.9 .8 (Lemma C. 1 and Corollary C. 1 in [WA16]). For any $t \in[k]$ and $\eta \in(0,1 / 2)$.
Let $\mathcal{U}$ denote $a$ set of random Gaussian vectors in $\mathbb{R}^{n}$. If $|\mathcal{U}|=\Omega(k \log (1 / \eta))$, then with probability at least $1-\eta$ there exists a vector $u \in \mathcal{U}$ such that

$$
\max _{j \in[k] \backslash\{t\}}\left|v_{j}^{\top} u\right| \leq \frac{1}{4}\left|v_{t}^{\top} u\right| \text { and }\left|v_{t}^{\top} u\right| \geq 1 / \sqrt{n} .
$$

We generalize Lemma C. 2 in [WA16] from $p=3$ to arbitrary $p \geq 3$,

Lemma 22.9.9. For sufficiently small constant $c>0$, for any constant $c_{0} \geq 1$, for any $p \geq 3$, for any tensor $A=A^{*}+\widetilde{E} \in \mathbb{R}^{n^{p}}$ with $A^{*}=\sum_{i=1}^{k} \lambda_{i} v_{i}^{\otimes p}$, let

$$
u_{t+1}=A\left(I, u_{t}, \cdots, u_{t}\right) /\left\|A\left(I, u_{t}, \cdots, u_{t}\right)\right\|_{2}
$$

For any $\epsilon \leq c \lambda_{1} /\left(c_{0} p^{2} k n^{(p-2) / 2}\right)$, and unit vector $u_{t} \in \mathbb{R}^{n}$, if

$$
\left\|\widetilde{E}\left(I, u_{t}, \cdots, u_{t}\right)\right\|_{2} \leq \begin{cases}4 p \epsilon & \text { if }\left|v_{1}^{\top} u_{t}\right| \leq 1-1 /\left(c_{0}^{2} p^{2} k^{2}\right) \\ 6 \epsilon / c_{0} & \text { otherwise }\end{cases}
$$

and

$$
\left|\widetilde{E}\left(v, u_{t}, \cdots, u_{t}\right)\right| \leq \begin{cases}4 \epsilon / \sqrt{n} & \text { if }\left|v_{1}^{\top} u_{t}\right| \leq 1-1 /\left(c_{0}^{2} p^{2} k^{2}\right) \\ 6 \epsilon /\left(c_{0} \sqrt{n}\right) & \text { otherwise }\end{cases}
$$

then for any $t \in[T]$,
1.

$$
\begin{aligned}
& \qquad \tan \theta\left(v_{1}, u_{t+1}\right) \leq \begin{cases}0.8 \tan \theta\left(v_{1}, u_{t}\right) & \text { if }\left|v_{1}^{\top} u_{t}\right| \leq 1-1 /\left(c_{0}^{2} p^{2} k^{2}\right) \\
0.8 \tan \theta\left(v_{1}, u_{t}\right)+18 \epsilon /\left(c_{0} \lambda_{1}\right) & \text { otherwise }\end{cases} \\
& \text { 2. } \max _{j \in[k] \backslash\{1\}} \lambda_{j}\left|v_{j}^{\top} u_{t}\right|^{p-2} \leq(1 / 4) \lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2} .
\end{aligned}
$$

3. For any $j \in\{2, \cdots, k\}$,

$$
\left|v_{j}^{\top} u_{t+1}\right| /\left|v_{1}^{\top} u_{t+1}\right| \leq \begin{cases}0.8\left|v_{j}^{\top} u_{t}\right| /\left|v_{1}^{\top} u_{t}\right| & \text { if }\left|v_{1}^{\top} u_{t}\right| \leq 1-1 /\left(c_{0}^{2} p^{2} k^{2}\right) \\ 0.8\left|v_{j}^{\top} u_{t}\right| /\left|v_{1}^{\top} u_{t}\right|+18 \epsilon /\left(c_{0} \lambda_{1} \sqrt{n}\right) & \text { otherwise }\end{cases}
$$

Proof. Part 1. Let $V \in \mathbb{R}^{n \times(n-1)}$ denote an orthonormal basis of the complement of $v_{1}$, i.e., $V$ is $\left(v_{2}, \cdots, v_{k}, \cdots, v_{n}\right)$ and let $\widetilde{E}_{u_{t}}=\widetilde{E}\left(I, u_{t}, \cdots, u_{t}\right) \in \mathbb{R}^{n}$. We can give an upper bound for $\tan \theta\left(v_{1}, u_{t+1}\right)$,

$$
\begin{array}{rlr} 
& \tan \theta\left(v_{1}, u_{t+1}\right) & \\
= & \tan \theta\left(v_{1}, A\left(I, u_{t}, \cdots, u_{t}\right) /\left\|A\left(I, u_{t}, \cdots, u_{t}\right)\right\|_{2}\right) & \text { by definition of } u_{t+1} \\
= & \tan \theta\left(v_{1}, A\left(I, u_{t}, \cdots, u_{t}\right)\right) & \text { by definition of angle } \\
=\tan \theta\left(v_{1}, A^{*}\left(I, u_{t}, \cdots, u_{t}\right)+\widetilde{E}\left(I, u_{t}, \cdots, u_{t}\right)\right) & \text { by } A=A^{*}+\widetilde{E} \in \mathbb{R}^{n^{p}} \\
=\tan \theta\left(v_{1}, A^{*}\left(I, u_{t}, \cdots, u_{t}\right)+\widetilde{E}_{u_{t}}\right) & \text { by } \widetilde{E}_{u_{t}}:=\widetilde{E}\left(I, u_{t}, \cdots, u_{t}\right) \in \mathbb{R}^{n} \\
= & \frac{\left\|V^{\top}\left[A^{*}\left(I, u_{t}, \cdots, u_{t}\right)+\widetilde{E}_{u_{t}}\right]\right\|_{2}}{\mid v_{1}^{\top}\left[A^{*}\left(I, u_{t}, \cdots, u_{t}\right)+\widetilde{E}_{u_{t}}\right]} & \text { by definition of } \tan \theta \\
\leq & \frac{\left\|V^{\top} A^{*}\left(I, u_{t}, \cdots, u_{t}\right)\right\|_{2}+\left\|V^{\top} \widetilde{E}_{u_{t}}\right\|_{2}}{\left|v_{1}^{\top} A^{*}\left(I, u_{t}, \cdots, u_{t}\right)\right|-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right|} & \text { by triangle inequality }
\end{array}
$$

Using Fact 22.9.7, we have

$$
\begin{align*}
\left\|V^{\top} A^{*}\left(I, u_{t}, \cdots, u_{t}\right)\right\|_{2}^{2} & =\sum_{j=2}^{k} \lambda_{j}^{2}\left|v_{j}^{\top} u_{t}\right|^{2(p-1)}  \tag{22.13}\\
& \leq\left(\max _{j \in[k] \backslash\{1\}}\left|\lambda_{j}\right|^{2}\left|v_{j}^{\top} u_{t}\right|^{2(p-2)}\right) \cdot\left(\sum_{j=2}^{k}\left|v_{j}^{\top} u_{t}\right|^{2}\right) \tag{22.14}
\end{align*}
$$

Putting it all together, we have

$$
\begin{aligned}
& \tan \theta\left(v_{1}, u_{t+1}\right) \\
& \leq \frac{\left\|V^{\top} A^{*}\left(I, u_{t}, \cdots, u_{t}\right)\right\|_{2}+\left\|V^{\top} \widetilde{E}_{u_{t}}\right\|_{2}}{\left|v_{1}^{\top} A^{*}\left(I, u_{t}, \cdots, u_{t}\right)\right|-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right|} \\
& \leq \tan \theta\left(v_{1}, u_{t}\right) \cdot \frac{\left(\left\|V^{\top} A^{*}\left(I, u_{t}, \cdots, u_{t}\right)\right\|_{2}+\left\|V^{\top} \widetilde{E}_{u_{t}}\right\|_{2}\right) /\left\|V^{\top} u_{t}\right\|_{2}}{\left|v_{1}^{\top} A^{*}\left(I, u_{t}, \cdots, u_{t}\right)\right| /\left|v_{1}^{\top} u_{t}\right|-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right| /\left|v_{1}^{\top} u_{t}\right|} \\
& \leq \tan \theta\left(v_{1}, u_{t}\right) \cdot \frac{\max _{j \in[k] \backslash\{1\}} \lambda_{j}\left|v_{j}^{\top} u_{t}\right|^{p-2}+\left\|V^{\top} \widetilde{E}_{u_{t}}\right\|_{2} /\left\|V^{\top} u_{t}\right\|_{2}}{\left|v_{1}^{\top} A^{*}\left(I, u_{t}, \cdots, u_{t}\right)\right| /\left|v_{1}^{\top} u_{t}\right|-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right| /\left|v_{1}^{\top} u_{t}\right|} \\
& \leq \tan \theta\left(v_{1}, u_{t}\right) \cdot \frac{\max _{j \in[k \backslash \backslash 1\}} \lambda_{j}\left|v_{j}^{\top} u_{t}\right|^{p-2}+\left\|V^{\top} \widetilde{E}_{u_{t}}\right\|_{2} /\left\|V^{\top} u_{t}\right\|_{2}}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2}-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right| /\left|v_{1}^{\top} u_{t}\right|} \\
& \leq \tan \theta\left(v_{1}, u_{t}\right) \cdot \frac{(1 / 4) \lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2}+\left\|V^{\top} \widetilde{E}_{u_{t}}\right\|_{2} /\left\|V^{\top} u_{t}\right\|_{2}}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2}-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right| /\left|v_{1}^{\top} u_{t}\right|} \\
& \leq \tan \theta\left(v_{1}, u_{t}\right) \cdot(1 / 4) \cdot \underbrace{\frac{1}{1-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right| /\left(\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}\right)}}_{A} \\
& +\underbrace{\frac{1}{1-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right| /\left(\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}\right)}}_{A} \cdot \underbrace{\frac{\left\|V^{\top} \widetilde{E}_{u_{t}}\right\|_{2}}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}}}_{B} \\
& =\tan \theta\left(v_{1}, u_{t}\right) \cdot(1 / 4) \cdot A+A \cdot B
\end{aligned}
$$

where the second step follows from $\tan \theta\left(v_{1}, u_{t}\right)=\frac{\left\|V^{\top} u_{t}\right\|_{2}}{\left|v_{1}^{\top} u_{t}\right|}$, the third step follows from Equation (22.13), the forth step follows from Fact 22.9.3,
the fifth step follows from Part 2 of Lemma 22.9.9.
We show
Claim 22.9.10. For any $t \in[T],\left|v_{1}^{\top} u_{0}\right| \leq\left|v_{1}^{\top} u_{t}\right|$.

Proof. By induction hypothesis, we assume there exists some sufficiently small constant $c>0$
such that $\tan \theta\left(v_{1}, u_{t}\right) \leq 0.8 \tan \theta\left(v_{1}, u_{t-1}\right)+c$. Thus, we can show

$$
\begin{array}{rlr}
\tan \theta\left(v_{1}, u_{t}\right) & \leq 0.8\left(0.8 \tan \theta\left(v_{1}, u_{t-1}\right)+c\right)+c \\
& \leq 0.8^{t} \tan \theta\left(v_{1}, u_{0}\right)+c \sum_{j=0}^{t-1} 0.8^{j} & \\
& \leq 0.8^{t} \tan \theta\left(v_{1}, u_{0}\right)+5 c & \\
& \leq \tan \theta\left(v_{1}, u_{0}\right) & \text { by } \tan \theta\left(v_{1}, u_{0}\right)=\omega(1)
\end{array}
$$

which implies $\theta\left(v_{1}, u_{t}\right) \leq \theta\left(v_{1}, u_{0}\right)$ and hence $\left|v_{1}^{\top} u_{t}\right|=\cos \theta\left(v_{1}, u_{t}\right) \geq \cos \theta\left(v_{1}, u_{0}\right)=\left|v_{1}^{\top} u_{0}\right|$.

Then we can upper bound $A$ and $B$

Claim 22.9.11. $A \leq 1.1$.

Proof. Note that $\left|\widetilde{E}\left(v_{j}, u_{t}, \cdots, u_{t}\right)\right|=\left|v_{j}^{\top} \widetilde{E}\left(I, u_{t}, \cdots, u_{t}\right)\right|=\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right|$. Because

$$
\begin{aligned}
\left|\widetilde{E}\left(v_{j}, u_{t}, \cdots, u_{t}\right)\right| & \leq 4 \epsilon / \sqrt{n} & & \\
& \leq 4 c \lambda_{1} / n^{(p-1) / 2} & & \text { by } \epsilon \leq c \lambda_{1} / n^{(p-2) / 2} \\
& \leq 4 c \lambda_{1}\left|v_{1}^{\top} u_{0}\right|^{p-1} & & \text { by }\left|v_{1}^{\top} u_{0}\right| \geq 1 / \sqrt{n} .
\end{aligned}
$$

Thus, if we choose sufficiently small $c<1 / 40$, using $\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right| \leq \lambda_{1}\left|v_{1}^{\top} u_{0}\right|^{p-1} / 10$ and $\left|v_{1}^{\top} u_{0}\right| \leq$ $\left|v_{1}^{\top} u_{t}\right|$, we can show

$$
\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right| \leq \lambda_{1}\left|v_{1}^{\top} u_{0}\right|^{p-1} / 10 \leq \lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1} / 10
$$

Thus, $A \leq \frac{1}{1-1 / 11}=1.1$

It remains to bound $B$. We give a proof by considering two cases separately, one is $\left|v_{1}^{\top} u_{t}\right| \leq 1-\frac{1}{c_{0}^{2} p^{2} k^{2}}$ and the other is $\left|v_{1}^{\top} u_{t}\right|>1-\frac{1}{c_{0}^{2} p^{2} k^{2}}$. If $\left|v_{1}^{\top} u_{t}\right| \leq\left(1-\frac{1}{c_{0}^{2} p^{2} k^{2}}\right)$, we have

$$
B=\frac{\left\|V^{\top} \widetilde{E}_{u_{t}}\right\|_{2}}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}}
$$

$$
=\frac{\sqrt{1-\left|v_{1}^{\top} u_{t}\right|^{2}}}{\left|v_{1}^{\top} u_{t}\right|} \cdot \frac{\left\|V^{\top} \widetilde{E}_{u_{t}}\right\|_{2}}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2} \sqrt{1-\left|v_{1}^{\top} u_{t}\right|^{2}}}
$$

$$
=\tan \theta\left(v_{1}, u_{t}\right) \cdot \frac{\left\|V^{\top} \widetilde{E}_{u_{t}}\right\|_{2}}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2} \sqrt{1-\left|v_{1}^{\top} u_{t}\right|^{2}}}
$$

$$
\leq \tan \theta\left(v_{1}, u_{t}\right) \cdot \frac{\left\|\widetilde{E}_{u_{t}}\right\|_{2}}{\lambda_{1}\left|v_{1}^{\top} u_{\tau}\right|^{p-2} \sqrt{1-\left|v_{1}^{\top} u_{t}\right|^{2}}} \quad \text { by }\left\|V^{\top} \widetilde{E}_{u_{t}}\right\|_{2} \leq\left\|\widetilde{E}_{u_{t}}\right\|_{2}
$$

$$
\leq \tan \theta\left(v_{1}, u_{t}\right) \cdot \frac{c_{0} p k\left\|\widetilde{E}_{u_{t}}\right\|_{2}}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2}} \quad \quad \text { by } 1 / \sqrt{1-\left|v_{1}^{\top} u_{t}\right|^{2}} \leq c_{0} p k
$$

We need to bound $\frac{c_{0} p k\left\|\widetilde{E}_{u_{t}}\right\|_{2}}{\lambda_{1}\left|v_{1}^{v} u_{t}\right|^{p-2}}$. On one side, we can show

$$
\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2} \geq \lambda_{1} /\left(n^{(p-2) / 2}\right)
$$

On the other side, by Part 1 of Lemma 22.9.12 and assumption on $\bar{E}$ and $E$, we have $\left\|\widetilde{E}_{u_{t}}\right\|_{2}$

$$
c_{0} p k\left\|\widetilde{E}_{u_{t}}\right\|_{2} \leq c_{0} p k \cdot 4 p \epsilon
$$

Thus, as long as we choose sufficiently small $\epsilon$ such that $\epsilon \leq \lambda_{1} /\left(n^{(p-2) / 2} \cdot 40 \cdot c_{0} p^{2} k\right)$, then $B \leq 0.1 \tan \theta\left(v_{1}, u_{t}\right)$.

If $\left|v_{1}^{\top} u_{t}\right|>1-\frac{1}{c_{0}^{2} k^{2} p^{2}}$, we have

$$
\begin{array}{rr}
B & =\frac{\left\|V^{\top} \widetilde{E}_{u_{t}}\right\|_{2}}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}} \\
& \leq \frac{\left\|V^{\top} \widetilde{E}_{u_{t}}\right\|_{2}}{\lambda_{1}\left(1-\frac{1}{c_{0}^{2} p^{2} k^{2}}\right)^{p-1}} \\
& \leq 3\left\|V^{\top} \widetilde{E}_{u_{t}}\right\|_{2} / \lambda_{1} \\
& \text { by }\left|v_{1}^{\top} u_{t}\right|>1-\frac{1}{c_{0}^{2} p^{2} k^{2}} \\
& \text { by } 1 /\left(1-\frac{1}{c_{0}^{2} p^{2} k^{2}}\right)^{p-1} \leq 3, \forall p \geq 3, k \geq 1, c_{0} \geq 1
\end{array}
$$

By Part 1 of Corollary 22.9.13, we have $\left\|\widehat{E}_{u_{t}}\right\|_{2} \leq 4 \epsilon / c_{0}$. By assumption on $E$ and $\bar{E}$, we have $\left\|E_{u_{t}}\right\|_{2} \leq \epsilon / c_{0}$ and $\left\|\bar{E}_{u_{t}}\right\|_{2} \leq \epsilon / c_{0}$. Thus, we complete the proof $B \leq 18 \epsilon /\left(c_{0} \lambda_{1}\right)$.

Part 2. For any $j \in[k] \backslash\{1\}$, we have lower bound $\frac{\left|v_{1}^{\top} u_{t+1}\right|}{\left|v_{j}^{\top} u_{t+1}\right|}$,

$$
\begin{align*}
& \frac{\left|v_{1}^{\top} u_{t+1}\right|}{\left|v_{j}^{\top} u_{t+1}\right|} \\
= & \frac{\mid v_{1}^{\top}\left[A^{*}\left(I, u_{t}, \cdots, u_{t}\right)+\widetilde{E}_{u_{t}}| |\right.}{\left|v_{j}^{\top}\right| A^{*}\left(I, u_{t}, \cdots, u_{t}\right)+\widetilde{E}_{u_{t}} t \mid} \\
\geq & \frac{\left|v_{1}^{\top} A^{*}\left(I, u_{t}, \cdots, u_{t}\right)\right|-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right|}{\left|v_{j}^{\top} A^{*}\left(I, u_{t}, \cdots, u_{t}\right)\right|+\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right|} \\
\geq & \text { by definition of } u_{t+1} \\
\geq & \\
\frac{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right|}{\lambda_{j}\left|v_{j}^{\top} u_{t}\right|^{p-1}+\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right|} & \text { by triangle inequality } \\
\geq & \frac{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}-\frac{1}{10} \lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}}{\lambda_{j}\left|v_{j}^{\top} u_{t}\right|^{p-1}+\frac{1}{10} \lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}}  \tag{22.15}\\
\geq \frac{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}-\frac{1}{10} \lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}}{} \quad & \text { by Fact } 22.9 .3 \\
\geq & \text { by }\left|v_{j}^{\top} \widetilde{E}_{u_{t} \mid}\right| \leq \frac{1}{10} \lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1} \\
= & \text { by Part 1. } \lambda_{j}\left|v_{j}^{\top} u_{t}\right|^{p-2} \leq \frac{1}{4} \lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2} \\
\frac{\frac{9}{10}\left|v_{1}^{\top} u_{t}\right|}{\frac{1}{4}\left|v_{j}^{\top} u_{t}\right|+\frac{1}{10}\left|v_{1}^{\top} u_{t}\right|} &
\end{align*}
$$

If $\left|v_{j}^{\top} u_{t}\right|<\left|v_{1}^{\top} u_{t}\right|$, then

$$
\frac{\lambda_{1}\left|v_{1}^{\top} u_{t+1}\right|^{p-2}}{\lambda_{j}\left|v_{j}^{\top} u_{t+1}\right|^{p-2}} \geq \frac{\lambda_{1}}{\lambda_{j}} 2^{p-2}
$$

where the last step follows by $p \geq 4$. Note that for $p=3$ we apply the better analysis like the proof of Lemma C. 2 in [WA16]. Thus, it works for any $p \geq 3$.

If $\left|v_{j}^{\top} u_{t}\right| \geq\left|v_{1}^{\top} u_{t}\right|$, then

$$
\frac{\lambda_{1}\left|v_{1}^{\top} u_{t+1}\right|^{p-2}}{\lambda_{j}\left|v_{j}^{\top} u_{t+1}\right|^{p-2}} \geq \frac{\lambda_{1}}{\lambda_{j}} \frac{\left|v_{1}^{\top} u_{t}\right|^{p-2}}{\left|v_{j}^{\top} u_{t}\right|^{p-2}} 2^{p-2} \geq 4 \cdot 2^{p-2}
$$

Part 3. Similarly as Equation (22.15) and Part 1, we can also upper bound $\frac{\left|v_{j}^{\top} u_{t+1}\right|}{\left|v_{1}^{\top} u_{t+1}\right|}$,

$$
\begin{aligned}
\frac{\left|v_{j}^{\top} u_{t+1}\right|}{\left|v_{1}^{\top} u_{t+1}\right|} & \leq \frac{\lambda_{j}\left|v_{j}^{\top} u_{t}\right|^{p-1}+\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right|}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right|} \\
& =\frac{\left|v_{j}^{\top} u_{t}\right|}{\left|v_{1}^{\top} u_{t}\right|} \cdot \frac{\lambda_{j}\left|v_{j}^{\top} u_{t}\right|^{p-2}+\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right| /\left|v_{j}^{\top} u_{t}\right|}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2}-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right| /\left|v_{1}^{\top} u_{t}\right|} \\
& =\frac{\left|v_{j}^{\top} u_{t}\right|}{\left|v_{1}^{\top} u_{t}\right|} \cdot \frac{\lambda_{j}\left|v_{j}^{\top} u_{t}\right|^{p-2}}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2}-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right| /\left|v_{1}^{\top} u_{t}\right|}+\frac{\left|v_{j}^{\top} \epsilon_{t}\right|}{\frac{1}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right|}} \\
& \leq \frac{\left|v_{j}^{\top} u_{t}\right|}{\left|v_{1}^{\top} u_{t}\right|} \cdot \frac{1}{4} \cdot \underbrace{\frac{1}{1-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right| /\left(\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}\right)}}_{A}+\underbrace{\frac{1-\left|v_{1}^{\top} \widetilde{E}_{u_{t}}\right| /\left(\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}\right)}{\underbrace{}_{B}} \cdot \underbrace{\frac{\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right|}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}}}_{B}}_{A}
\end{aligned}
$$

Similarly as before, we can show $A \leq 1.1$ and if $\left|v_{1}^{\top} u_{t}\right| \geq 1-\frac{1}{c_{0}^{2} p^{2} k^{2}}$. It remains to bound $B$. We give a proof by considering two cases separately, one is $\left|v_{1}^{\top} u_{t}\right| \leq 1-\frac{1}{c_{0}^{2} p^{2} k^{2}}$ and the other
is $\left|v_{1}^{\top} u_{t}\right|>1-\frac{1}{c_{0}^{2} p^{2} k^{2}}$. If $\left|v_{1}^{\top} u_{t}\right| \leq\left(1-\frac{1}{c_{0}^{2} p^{2} k^{2}}\right)$, we have

$$
\begin{aligned}
B & =\frac{\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right|}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}} \\
& =\frac{\sqrt{1-\left|v_{1}^{\top} u_{t}\right|^{2}}}{\left|v_{1}^{\top} u_{t}\right|} \cdot \frac{\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right|}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2} \sqrt{1-\left|v_{1}^{\top} u_{t}\right|^{2}}} \\
& =\tan \theta\left(v_{1}, u_{t}\right) \cdot \frac{\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right|}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2} \sqrt{1-\left|v_{1}^{\top} u_{t}\right|^{2}}}
\end{aligned}
$$

$$
\leq \tan \theta\left(v_{1}, u_{t}\right) \cdot \frac{c_{0} p k\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right|}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2}} \quad \quad \text { by } 1 / \sqrt{1-\left|v_{1}^{\top} u_{t}\right|^{2}} \leq c_{0} p k
$$

We need to bound $\frac{c_{0} p k\left\|\widetilde{E}_{u_{t}}\right\|_{2}}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2}}$. On one side, we can show

$$
\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-2} \geq \lambda_{1} /\left(n^{(p-2) / 2}\right)
$$

On the other side, by Part 2 of Lemma 22.9.12 and assumption on $\bar{E}$ and $E$, we have $\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right|$

$$
c_{0} p k\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right| \leq c_{0} p k \cdot 4 \epsilon / \sqrt{n} .
$$

Thus, as long as we choose sufficiently small $\epsilon$ such that $\epsilon \leq \lambda_{1} \sqrt{n} /\left(n^{(p-2) / 2} \cdot 40 \cdot c_{0} p k\right)$, then $B \leq 0.1 \tan \theta\left(v_{1}, u_{t}\right)$.

If $\left|v_{1}^{\top} u_{t}\right|>1-\frac{1}{c_{0}^{2} k^{2} p^{2}}$, we have

$$
\begin{array}{rr}
B & =\frac{\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right|}{\lambda_{1}\left|v_{1}^{\top} u_{t}\right|^{p-1}} \\
& \leq \frac{\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right|}{\lambda_{1}\left(1-\frac{1}{\left.c_{0}^{p^{2} k^{2}}\right)^{p-1}}\right.} \begin{array}{lr} 
& \text { by }\left|v_{1}^{\top} u_{t}\right|>1-\frac{1}{c_{0}^{2} p^{2} k^{2}} \\
& \leq 3\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right| / \lambda_{1} \\
& \leq 3\left|v_{j}^{\top} \widetilde{E}_{u_{t}}\right| / \lambda_{1}
\end{array} \quad \text { by } 1 /\left(1-\frac{1}{c_{0}^{2} p^{2} k^{2}}\right)^{p-1} \leq 3, \forall p \geq 3, k \geq 1, c_{0} \geq 1
\end{array}
$$

By Part 1 of Corollary 22.9.13, we have $\left|v_{j}^{\top} \widehat{E}_{u_{t}}\right| \leq 4 \epsilon /\left(c_{0} \sqrt{n}\right)$. By assumption on $E$ and $\bar{E}$, we have $\left|v_{j}^{\top} E_{u_{t}}\right| \leq \epsilon /\left(c_{0} \sqrt{n}\right)$ and $\left|v_{j}^{\top} \bar{E}_{u_{t}}\right| \leq \epsilon /\left(c_{0} \sqrt{n}\right)$. Thus, we complete the proof $B \leq 18 \epsilon /\left(c_{0} \lambda_{1} \sqrt{n}\right)$.

Definition 22.9.1. For any $\epsilon>0$, we say $\left\{\widehat{\lambda}_{i}, \widehat{v}_{i}\right\}_{i=1}^{k}$ is $\epsilon$-close to $\left\{\lambda_{i}, v_{i}\right\}_{i=1}^{k}$ if for all $i \in[k]$,

1. $\left|\widehat{\lambda}_{i}-\lambda_{i}\right| \leq \epsilon$.
2. $\left\|\widehat{v}_{i}-v_{i}\right\|_{2} \leq \tan \theta\left(\widehat{v}_{i}, v_{i}\right) \leq \min \left(\sqrt{2}, \epsilon /\left(\lambda_{i}\right)\right)$.
3. $\left|\widehat{v}_{i}^{\top} v_{j}\right| \leq \epsilon /\left(\sqrt{n} \lambda_{i}\right), \forall j \in[k] \backslash[i]$.

Lemma 22.9.12. Let $\widehat{E}_{i}=\lambda_{i} v_{i}^{\otimes p}-\widehat{\lambda}_{i} \widehat{v}_{i}^{\otimes p} \in \mathbb{R}^{n^{p}}, \forall i \in[k]$. For any $\epsilon>0$, if $\left\{\widehat{\lambda}_{i}, \widehat{v}_{i}\right\}_{i=1}^{k}$ is $\epsilon$-close to $\left\{\lambda_{i}, v_{i}\right\}_{i=1}^{k}$, then for all $r \in[k]$, for any unit vector $u \in \mathbb{R}^{n}$,

1. $\left\|\sum_{i=1}^{r} \widehat{E}_{i}(I, u, \cdots, u)\right\|_{2} \leq 2 p \epsilon \kappa^{1 / 2}+2 \phi \epsilon$.
2. $\forall j \in[k] \backslash[i],\left|\sum_{i=1}^{r} \widehat{E}_{j}\left(v_{j}, u, \cdots, u\right)\right| \leq(2 \kappa \epsilon+\phi \epsilon) / \sqrt{n}$.
where $\kappa=\sum_{i=1}^{r}\left|u^{\top} v_{i}\right|^{2}$ and $\phi=2 k\left(\epsilon / \lambda_{k}\right)^{p-1}$.
Proof. Part 1. For each $i \in[r]$, the error $\widehat{E}_{i}$ is

$$
\widehat{E}_{i}(I, u, \cdots, u)=\lambda_{i}\left(u^{\top} v_{i}\right)^{p-1} v_{i}-\widehat{\lambda}_{i}\left(u^{\top} \widehat{v}_{i}\right)^{p-1} \widehat{v}_{i}
$$

lives in $\operatorname{span}\left\{v_{i}, \widehat{v}_{i}\right\}$; this space is the same as $\operatorname{span}\left\{v_{i}, \widehat{v}_{i}^{\perp}\right\}$, where

$$
\widehat{v}_{i}^{\perp}=\widehat{v}_{i}-\left(v_{i}^{\top} \widehat{v}_{i}\right) v_{i}
$$

is the projection of $\widehat{v}_{i}$ onto the subspace orthogonal to $v_{i}$. Since $\left\|\widehat{v}_{i}-v\right\|_{2}^{2}=2\left(1-v_{i}^{\top} \widehat{v}_{i}\right)$, it follows that

$$
c_{i}=v_{i}^{\top} \widehat{v}_{i}=1-\left\|\widehat{v}_{i}-v_{i}\right\|_{2}^{2} / 2 \geq 0
$$

(the inequality follows from the assumption $\left\|\widehat{v}_{i}-v_{i}\right\|_{2} \leq \sqrt{2}$, which in turn implies $0 \leq c_{i} \leq$ 1). By the Pythagorean theorem and the above inequality for $c_{i}$,

$$
\left\|\widehat{v}_{i}^{\perp}\right\|_{2}^{2}=1-c_{i}^{2} \leq\left\|\widehat{v}_{i}-v_{i}\right\|_{2}^{2} .
$$

We can obtain this bound, for any $p \geq 3$,

$$
\left|1-c_{i}^{p}\right|=\left|1-\left(1-\left\|\widehat{v}_{i}-v_{i}\right\|_{2}^{2} / 2\right)^{p}\right| \leq \frac{p}{2}\left\|\widehat{v}_{i}-v_{i}\right\|_{2}^{2}
$$

We first define $a_{i}, b_{i} \in \mathbb{R}^{n}$ such that $a_{i}=u^{\top} v_{i}$ and $b_{i}=u^{\top}\left(\widehat{v}_{i}^{\perp} /\left\|\widehat{v}_{i}^{\perp}\right\|_{2}\right)$. We now rewrite $\widehat{E}_{i}(I, u, \cdots, u)$ in terms of the coordinate system defined by $v_{i}$ and $\widehat{v}_{i}^{\perp}$,

$$
\begin{aligned}
& \widehat{E}_{i}(I, u, \cdots, u) \\
= & \lambda_{i}\left(u^{\top} v_{i}\right)^{p-1} v_{i}-\widehat{\lambda}_{i}\left(u^{\top} \widehat{v}_{i}\right)^{p-1} \widehat{v}_{i} \\
= & \lambda_{i} a_{i}^{p-1} v_{i}-\widehat{\lambda}_{i}\left(a_{i} c_{i}+\left\|\widehat{v}_{i}^{\perp} b_{i}\right\|_{2}\right)^{p-1}\left(c_{i} v_{i}+\widehat{v}_{i}^{\perp}\right) \\
= & \underbrace{\left(\lambda_{i} a_{i}^{p-1} v_{i}-\widehat{\lambda}_{i}\left(a_{i} c_{i}+\left\|\widehat{v}_{i}^{\perp}\right\|_{2} b_{i}\right)^{p-1} c_{i}\right)}_{A_{i}} \cdot v_{i}-\underbrace{\widehat{\lambda}_{i}\left(a_{i} c_{i}+\left\|\widehat{v}_{i}^{\perp}\right\|_{2} b_{i}\right)^{p-1}\left\|\widehat{v}_{i}^{\perp}\right\|_{2}}_{B_{i}} \cdot \widehat{v}_{i}^{\perp} /\left\|\widehat{v}_{i}^{\perp}\right\|_{2} \\
= & A_{i} \cdot v_{i}-B_{i} \cdot\left(\widehat{v}_{i}^{\perp} /\left\|\widehat{v}_{i}^{\perp}\right\|_{2}\right)
\end{aligned}
$$

The overall error can be written as

$$
\begin{aligned}
& \left\|\sum_{i=1}^{r} \widehat{E}_{i}(I, u, \cdots, u)\right\|_{2}^{2} \\
= & \left\|\sum_{i=1}^{r} A_{i} v_{i}-\sum_{i=1}^{t} B_{i}\left(\widehat{v}_{i}^{\perp} /\left\|\widehat{v}_{i}^{\perp}\right\|_{2}\right)\right\|_{2}^{2} \\
\leq & 2\left\|\sum_{i=1}^{r} A_{i} v_{i}\right\|_{2}^{2}+2\left\|\sum_{i=1}^{r} B_{i}\left(\widehat{v}_{i}^{\perp} /\left\|\widehat{v}_{i}^{\perp}\right\|_{2}\right)\right\|_{2}^{2} \\
\leq & 2 \sum_{i=1}^{r} A_{i}^{2}+2\left(\sum_{i=1}^{r}\left|B_{i}\right|\right)^{2}
\end{aligned}
$$

We first try to bound $A_{i}$ by using $\left|\lambda_{i}-\widehat{\lambda}_{i}\right| \leq \epsilon,\left\|\widehat{v}_{i}-v_{i}\right\|_{2} \leq \epsilon / \lambda_{i},\left\|\widehat{v}_{i}^{\perp}\right\|_{2}^{2} \leq\left\|\widehat{v}_{i}-v_{i}\right\|_{2}^{2}$, $0 \leq c_{i} \leq 1$ and $\left|b_{i}\right| \leq 1$,

$$
\begin{aligned}
& \left|A_{i}\right| \\
= & \left|\lambda_{i} a_{i}^{p-1}-\widehat{\lambda}_{i}\left(a_{i} c_{i}+\left\|\widehat{v}_{i}^{\perp}\right\|_{2} b_{i}\right)^{p-1} c_{i}\right| \\
\leq & \left|\lambda_{i} a_{i}^{p-1}-\widehat{\lambda}_{i} c_{i}^{p} a_{i}^{p-1}\right|+\sum_{j=1}^{p-1} \widehat{\lambda}_{i} c_{i}\binom{(p-1)}{j}\left|a_{i} c_{i}\right|^{(p-1)-j}\left\|\widehat{v}_{i}^{\perp}\right\|_{2}^{j} \\
\leq & \left|\lambda_{i} a_{i}^{p-1}-\widehat{\lambda}_{i} a_{i}^{p-1}\right|+\left|\left(1-c_{i}^{p}\right) \widehat{\lambda}_{i} a_{i}^{p-1}\right|+\sum_{j=1}^{p-1} \widehat{\lambda}_{i} c_{i}\binom{(p-1)}{j}\left|a_{i} c_{i}\right|^{(p-1)-j}\left(\epsilon / \lambda_{i}\right)^{j} \\
\leq & \epsilon\left|a_{i}\right|^{p-1}+\frac{p}{2}\left(\epsilon / \lambda_{i}\right)^{2} \widehat{\lambda}_{i}\left|a_{i}\right|^{p-1}+\sum_{j=1}^{p-1} \widehat{\lambda}_{i}\binom{(p-1)}{j}\left|a_{i}\right|^{(p-1)-j}\left(\epsilon / \lambda_{i}\right)^{j}
\end{aligned}
$$

where the second term can be bounded as

$$
\begin{array}{rlr}
\frac{p}{2}\left(\epsilon / \lambda_{i}\right)^{2} \widehat{\lambda}_{i}\left|a_{i}\right|^{p-1} & \leq \frac{p}{2}\left(\epsilon / \lambda_{i}\right)^{2}\left(\left|\widehat{\lambda}_{i}-\lambda_{i}\right|+\left|\lambda_{i}\right|\right)\left|a_{i}\right|^{p-1} & \\
& \leq \frac{p}{2}\left(\epsilon / \lambda_{i}\right)^{2} \epsilon\left|a_{i}\right|^{p-1}+\frac{p}{2}\left(\epsilon / \lambda_{i}\right) \cdot \epsilon\left|a_{i}\right|^{p-1} & \\
& \leq \frac{1}{100 k} \epsilon\left|a_{i}\right|^{p-1} & \text { by } \epsilon \ll \lambda_{i} / p k
\end{array}
$$

The third term can be divided into two parts $j=1, \cdots,(p-1) / 2$ and $j=(p-1) / 2, \cdots,(p-$ 1). For the first part,

$$
\begin{aligned}
& \sum_{j=1}^{(p-1) / 2} \widehat{\lambda}_{i}\binom{(p-1)}{j}\left|a_{i}\right|^{(p-1)-j}\left(\epsilon / \lambda_{i}\right)^{j} \\
= & \widehat{\lambda}_{i}(p-1)\left|a_{i}\right|^{p-2} \epsilon / \lambda_{i}+\sum_{j=2}^{(p-1) / 2} \widehat{\lambda}_{i}(p-1)^{j}\left|a_{i}\right|^{(p-1) / 2}\left(\epsilon / \lambda_{i}\right)^{j} \\
\leq & (p-1) \epsilon\left|a_{i}\right|^{p-2}+\frac{1}{100 k} \epsilon\left|a_{i}\right|^{(p-1) / 2}
\end{aligned}
$$

Similarly, we can bound the second part,

$$
\begin{aligned}
& \sum_{j=(p-1)}^{(p-1) / 2} \widehat{\lambda}_{i}\binom{(p-1)}{j}\left|a_{i}\right|^{(p-1)-j}\left(\epsilon / \lambda_{i}\right)^{j} \\
= & \widehat{\lambda}_{i} \cdot 1 \cdot 1 \cdot\left(\epsilon / \lambda_{i}\right)^{p-1}+\sum_{j=(p-1)-1}^{(p-1) / 2} \widehat{\lambda}_{i}\binom{(p-1)}{j}\left|a_{i}\right|^{(p-1)-j}\left(\epsilon / \lambda_{i}\right)^{j} \\
\leq & \frac{1}{100 k} \epsilon+\frac{1}{100 k} \epsilon\left|a_{i}\right|
\end{aligned}
$$

Thus, putting it all together, we get

$$
\left|A_{i}\right| \leq \epsilon\left|a_{i}\right|+\frac{1}{10 k} \epsilon\left|a_{i}\right|+(p-1) \epsilon\left|a_{i}\right|+\frac{1}{100 k} \epsilon
$$

which implies that

$$
\left|A_{i}\right|^{2} \leq 2\left(\left(\epsilon\left|a_{i}\right|\right)^{2}+\left(\frac{1}{10 k} \epsilon\left|a_{i}\right|\right)^{2}+\left((p-1) \epsilon\left|a_{i}\right|\right)^{2}+\left(\frac{1}{100 k} \epsilon\right)^{2}\right)
$$

It remains to bound $B_{i}$,

$$
\begin{aligned}
\left|B_{i}\right| & =\left|\widehat{\lambda}_{i}\left\|\widehat{v}_{i}^{\perp}\right\|_{2}\left(a_{i} c_{i}+\left\|\widehat{v}_{i}^{\perp}\right\|_{2} b_{i}\right)^{p-1}\right| \\
& \leq \widehat{\lambda}_{i}\left\|\widehat{v}_{i}^{\perp}\right\|_{2} \sum_{j=0}^{p-1}\binom{(p-1)}{j}\left|a_{i} c_{i}\right|^{p-1-j}\left\|\widehat{v}_{i}^{\perp}\right\|_{2}^{j} \\
& \leq \widehat{\lambda}_{i}\left(\epsilon / \lambda_{i}\right) \sum_{j=0}^{p-1}\binom{(p-1)}{j}\left|a_{i}\right|^{p-1-j}\left(\epsilon / \lambda_{i}\right)^{j} \\
& =\widehat{\lambda}_{i}\left(\epsilon / \lambda_{i}\right)\left|a_{i}\right|^{p-1}+\widehat{\lambda}_{i}\left(\epsilon / \lambda_{i}\right) \sum_{j=1}^{p-1}\binom{(p-1)}{j}\left|a_{i}\right|^{p-1-j}\left(\epsilon / \lambda_{i}\right)^{j}
\end{aligned}
$$

where the first term can be bounded as

$$
\widehat{\lambda}_{i}\left(\epsilon / \lambda_{i}\right)\left|a_{i}\right|^{p-1} \leq\left(\epsilon+\lambda_{i}\right)\left(\epsilon / \lambda_{i}\right)\left|a_{i}\right|^{p-1} \leq \epsilon\left|a_{i}\right|^{p-1}+\frac{1}{100 k} \epsilon\left|a_{i}\right|^{p-1}
$$

The second term can be divided into two parts $j=1, \cdots,(p-1) / 2$ and $j=(p-1) / 2, \cdots,(p-$ 1). For the first part,

$$
\begin{aligned}
& \left(\epsilon / \lambda_{i}\right) \sum_{j=1}^{(p-1) / 2} \widehat{\lambda}_{i}\binom{(p-1)}{j}\left|a_{i}\right|^{(p-1)-j}\left(\epsilon / \lambda_{i}\right)^{j} \\
= & \left(\epsilon / \lambda_{i}\right) \widehat{\lambda}_{i}(p-1)\left|a_{i}\right|^{p-2} \epsilon / \lambda_{i}+\left(\epsilon / \lambda_{i}\right) \sum_{j=2}^{(p-1) / 2} \widehat{\lambda}_{i}(p-1)^{j}\left|a_{i}\right|^{(p-1) / 2}\left(\epsilon / \lambda_{i}\right)^{j} \\
\leq & \frac{1}{100 k} \epsilon\left|a_{i}\right|^{p-2}+\frac{1}{100 k} \epsilon\left|a_{i}\right|^{(p-1) / 2}
\end{aligned}
$$

by $\epsilon \ll \lambda_{i} /(k p)$
Similarly, we can bound the second part,

$$
\begin{aligned}
& \left(\epsilon / \lambda_{i}\right) \sum_{j=(p-1)}^{(p-1) / 2} \widehat{\lambda}_{i}\binom{(p-1)}{j}\left|a_{i}\right|^{(p-1)-j}\left(\epsilon / \lambda_{i}\right)^{j} \\
= & \left(\epsilon / \lambda_{i}\right) \widehat{\lambda}_{i} \cdot 1 \cdot 1 \cdot\left(\epsilon / \lambda_{i}\right)^{p-1}+\left(\epsilon / \lambda_{i}\right) \sum_{j=(p-1)-1}^{(p-1) / 2} \widehat{\lambda}_{i}\binom{(p-1)}{j}\left|a_{i}\right|^{(p-1)-j}\left(\epsilon / \lambda_{i}\right)^{j} \\
\leq & \frac{\phi}{k} \epsilon+\frac{1}{100 k} \epsilon\left|a_{i}\right|
\end{aligned} \quad \text { by } \epsilon \ll \lambda_{i} / k p \text { p } \quad l l
$$

where $\phi=2 k\left(\epsilon / \lambda_{k}\right)^{p-1}$. Putting it all together, we have

$$
\begin{aligned}
& \left|B_{i}\right| \leq \epsilon\left|a_{i}\right|^{2}+\frac{1}{100 k^{2}} \epsilon\left|a_{i}\right|+\frac{\phi}{k} \epsilon \\
\Longrightarrow & \sum_{i=1}^{r}\left|B_{i}\right| \leq \epsilon \kappa+\frac{1}{100 k^{2}} \epsilon \sum_{i=1}^{t}\left|a_{i}\right|+\phi \epsilon \\
\Longrightarrow & \left(\sum_{i=1}^{r}\left|B_{i}\right|\right)^{2} \leq 2\left((\epsilon \kappa)^{2}+\left(\frac{1}{100 k} \epsilon \sum_{i=1}^{t}\left|a_{i}\right|\right)^{2}+(\phi \epsilon)^{2}\right) \\
\Longrightarrow & \left(\sum_{i=1}^{r}\left|B_{i}\right|\right)^{2} \leq 2\left((\epsilon \kappa)^{2}+\left(\frac{1}{100 k} \epsilon\right)^{2} \kappa k+(\phi \epsilon)^{2}\right)
\end{aligned}
$$

Recall that $\kappa=\sum_{i=1}^{t}\left|a_{i}\right|^{2} \leq 1$. Overall, we have

$$
\begin{aligned}
& \left\|\sum_{i=1}^{r} \widehat{E}_{i}(I, u, \cdots, u)\right\|_{2}^{2} \\
= & 2 \sum_{i=1}^{r}\left|A_{i}\right|^{2}+2\left(\sum_{i=1}^{r}\left|B_{i}\right|\right)^{2} \\
\leq & 4\left(\epsilon^{2} \cdot \kappa+\left(\frac{1}{10 k} \epsilon\right)^{2} \kappa+(p-1)^{2} \epsilon^{2} \kappa+\left(\frac{1}{100 \sqrt{k}} \epsilon\right)^{2}\right) \\
+ & 4\left((\epsilon \kappa)^{2}+\left(\frac{1}{100 k} \epsilon\right)^{2} \kappa k+(\phi \epsilon)^{2}\right) \\
\leq & 4 p^{2} \epsilon^{2} \kappa+4 \phi^{2} \epsilon^{2}
\end{aligned}
$$

which gives the desired bound.

Part 2. For any $j \in[k] \backslash[i]$, we have

$$
\begin{aligned}
\left|\sum_{i=1}^{r} \widehat{E}_{i}\left(v_{j}, u, \cdots, u\right)\right| & \leq \sum_{i=1}^{r}\left|\widehat{E}_{i}\left(v_{j}, u, \cdots, u\right)\right| \\
& =\sum_{i=1}^{r} \widehat{\lambda}_{i}\left|\widehat{v}_{i}^{\top} u\right|^{p-1}\left|\widehat{v}_{i}^{\top} v_{j}\right| \\
& \leq \sum_{i=1}^{t} \widehat{\lambda}_{i}\left|\widehat{v}_{i}^{\top} u\right|^{p-1} \frac{\epsilon}{\sqrt{n} \lambda_{i}}
\end{aligned}
$$

where the first step follows by triangle inequality, the second step follows by $v_{j} \cdot v_{i}, \forall i \neq j$, and the last step follows by part 3 of definition of $\epsilon$-close. Let's try to bound $\sum_{i=1}^{r} \frac{\widehat{\lambda}_{i}}{\lambda_{i}}\left|\widehat{v}_{i}^{\top} u\right|^{p-1}$,

$$
\begin{aligned}
& \sum_{i=1}^{r} \frac{\widehat{\lambda}_{i}}{\lambda_{i}}\left|\widehat{v}_{i}^{\top} u\right|^{p-1} \\
\leq & \sum_{i=1}^{r} \frac{\widehat{\lambda}_{i}}{\lambda_{i}}\left(\left|v_{i}^{\top} u\right|+\left|\left(v_{i}-\widehat{v}_{i}\right)^{\top} u\right|\right)^{p-1} \\
\leq & \sum_{i=1}^{r} \frac{\widehat{\lambda}_{i}}{\lambda_{i}}\left(\left|v_{i}^{\top} u\right|+\left\|v_{i}-\widehat{v}_{i}\right\|_{2}\right)^{p-1} \\
\leq & \sum_{i=1}^{r}\left(1+\epsilon / \lambda_{i}\right)\left(\left|v_{i}^{\top} u\right|+\epsilon / \lambda_{i}\right)^{p-1} \\
\leq & 2 \sum_{i=1}^{r}\left|v_{i}^{\top} u\right|^{2}+2 k\left(\epsilon / \lambda_{k}\right)^{p-1} \\
= & 2 \kappa+\phi .
\end{aligned}
$$

which gives the desired bound $(2 \kappa+\phi) \epsilon / \sqrt{n}$.

Corollary 22.9.13. Let $\widehat{E}_{i}=\lambda_{i} v_{i}^{\otimes p}-\widehat{\lambda}_{i} \hat{v}_{i}^{\otimes p} \in \mathbb{R}^{n^{p}}, \forall i \in[k]$. For any $c_{0} \geq 1$, for any $\epsilon \leq \lambda_{k} /\left(2 c_{0} k\right)$, if $\left\{\widehat{\lambda}_{i}, \widehat{v}_{i}\right\}_{i=1}^{k}$ is $\epsilon$-close to $\left\{\lambda_{i}, v_{i}\right\}_{i=1}^{k}$, then for all $r \in[k]$, for any unit vector $u \in \mathbb{R}^{n},\left|u^{\top} v_{r+1}\right| \geq 1-\frac{1}{c_{0}^{2} p^{2} k}$, then

1. $\left\|\sum_{i=1}^{r} \widehat{E}_{i}(I, u, \cdots, u)\right\|_{2} \leq 2 p \epsilon \kappa^{1 / 2}+2 \phi \epsilon \leq 4 \epsilon / c_{0}$.
2. $\forall j \in[k] \backslash i],\left|\sum_{i=1}^{r} \widehat{E}_{j}\left(v_{j}, u, \cdots, u\right)\right| \leq(2 \kappa \epsilon+\phi \epsilon) / \sqrt{n} \leq 4 \epsilon /\left(c_{0} \sqrt{n}\right)$.
where $\kappa=\sum_{i=1}^{r}\left|u^{\top} v_{i}\right|^{2}$ and $\phi=2 k\left(\epsilon / \lambda_{k}\right)^{p-1}$.
Proof. Using Fact 22.9.1, we have for any $i \in[r]$,

$$
\kappa=\sum_{i=1}^{r}\left|u^{\top} v_{i}\right|^{2} \leq k \cdot 1 /\left(c_{0}^{2} p^{2} k\right) \leq 1 / c_{0}^{2} p^{2} .
$$

which implies $2 p \epsilon \sqrt{\kappa} \leq 2 / c_{0}$. We also can bound $\phi$,

$$
\phi=2 k\left(\epsilon / \lambda_{k}\right)^{p-1} \leq 2 k\left(1 /\left(2 c_{0} k\right)\right)^{2} \leq 1 / c_{0} .
$$

### 22.9.3 Main result

We first extend the original robust tensor power method analysis to general order $p \geq 3$, i.e., Theorem 22.3.4. Second, we combine Theorem 22.3 .4 with our importance sampling techniques, i.e., Corollary 22.7.1 and Lemma 22.8.8.

Theorem 22.3.4. (Main, Abitrary order robust tensor power method) For any $p \geq 3, k \geq 1$, for any tensor $A=A^{*}+E \in \mathbb{R}^{n^{p}}$, where $A^{*}=\sum_{i=1}^{k} \lambda_{i} v_{i}^{\otimes p}$ with $\lambda_{i}>0$ and orthonormal basis vectors $\left\{v_{1}, \ldots, v_{k}\right\} \subseteq \mathbb{R}^{n}, n \geq k$. Let $\lambda_{1}, \lambda_{k}$ be the largest and smallest values in $\left\{\lambda_{i}\right\}_{i=1}^{k}$ and $\left\{\widehat{\lambda}_{i}, \widehat{v}_{i}\right\}_{i=1}^{k}$ be outputs of the robust tensor power method. For any sufficiently large constant $c_{0} \geq 100$, there exists a sufficiently small constant $c>0$, for any $\epsilon \in\left(0, c \lambda_{k} /\left(c_{0} p^{2} k n^{(p-2) / 2}\right)\right.$ if $E$ satisfies that

$$
\begin{equation*}
\|E\|_{2} \leq \epsilon /\left(c_{0} \sqrt{n}\right) \tag{22.3}
\end{equation*}
$$

and $T=\Omega\left(\log \left(\lambda_{1} n / \epsilon\right)\right), L=\Omega(k \log (k))$, then with probability at least $9 / 10$, there exists a permutation $\pi:[k] \rightarrow[k]$, such that $\forall i \in[k]$,

$$
\begin{equation*}
\left|\lambda_{i}-\widehat{\lambda}_{\pi(i)}\right| \leq \epsilon, \quad\left\|v_{i}-\widehat{v}_{\pi(i)}\right\|_{2} \leq \epsilon / \lambda_{i} \tag{22.4}
\end{equation*}
$$

Proof. We use $E \in \mathbb{R}^{n^{p}}$ to denote the original noise, use $\widehat{E}_{i}=\lambda_{i} v_{i}^{\otimes p}-\widehat{\lambda}_{i} \widehat{v}_{i}^{\otimes p} \in \mathbb{R}^{n^{p}}$ to denote the deflation noise, use $\bar{E} \in \mathbb{R}^{n^{p}}$ to denote the sketch noise and use $\widetilde{E}$ to denote the "true" noise that contains both original noise, deflation noise and sketch noise. Thus, in $t+1$ step, we actually observe tensor $A^{*}+\widetilde{E}$, where $\widetilde{E}=E+\sum_{i=1}^{t} \widehat{E}_{i}+\bar{E}$. To prove Theorem 22.3.4, in fact we don't need to worry about sketch noise $\bar{E}$. Instead of assuming it is $\bar{E}=0$, we prove a strong theorem by assuming sketch noise $\bar{E}$ is bounded, i.e., $\|\bar{E}\|_{2} \leq \epsilon /\left(c_{0} \sqrt{n}\right)$.

Base case $i=1$. Consider the first step, we get $\widehat{\lambda}_{1} \in \mathbb{R}, \widehat{v}_{1} \in \mathbb{R}^{n}$, we first show $\left\|\widehat{v}_{1}-v_{1}\right\|_{2}$ is bounded(as Part 2 of Definition 22.9.1) and then show $\left|\widehat{\lambda}_{1}-\lambda_{1}\right|$ is bounded(as Part 1 of Definition 22.9.1). At the end, we show $\left|\widehat{v}_{i}^{\top} v_{j}\right|$ is bounded(as Part 3 of Definition 22.9.1).

Bounding $\left|\widehat{v}_{1}-v_{1}\right|$. So, applying Lemma 22.9.8, we have

$$
\tan \theta\left(u_{0}, v_{1}\right)=\sin \theta\left(u_{0}, v_{1}\right) / \cos \left(\theta\left(u_{0}, v_{1}\right)\right) \leq \sqrt{n}
$$

Let $t^{*}$ denote the situation where $\left|u_{t^{*}}^{\top} v_{1}\right|=1-\frac{1}{c_{0}^{2} p^{2} k^{2}}$. We know

$$
\left\|u_{t^{*}}-v_{1}\right\|_{2}^{2}=2-2\left|u_{t^{*}}^{\top} v_{1}\right|=2 /\left(c_{0}^{2} p^{2} k^{2}\right) .
$$

We can upper bound

$$
\left\|u_{t^{*}}-v_{1}\right\|_{2} \leq \tan \theta\left(u_{t^{*}}, v_{1}\right) \leq 0.8 \tan \theta\left(u_{t^{*}-1}, v_{1}\right) \leq \cdots \leq 0.8^{t^{*}} \tan \theta\left(u_{0}, v_{1}\right) \leq 0.8^{t^{*}} \sqrt{n}
$$

Thus, we need to set $t^{*}=\Omega\left(\log \left(n k p c_{0}\right)\right)=\Omega\left(\log c_{0} n\right)$.
Now, we need to analyze $\left\|u_{T}-v_{1}\right\|_{2}$,

$$
\left\|u_{T}-v_{1}\right\|_{2} \leq 0.8\left(\tan \theta\left(u_{T}, v_{1}\right)+18 \epsilon /\left(c_{0} \lambda_{1}\right)\right) \leq \cdots \leq 0.8^{T-t^{*}} \tan \left(u_{t^{*}}, v_{1}\right)+5 \cdot 18 \epsilon /\left(c_{0} \lambda_{1}\right)
$$

To guarantee $\left\|u_{T}-v_{1}\right\|_{2} \leq \epsilon / \lambda_{1}$, we need to choose $T-t^{*}=\Omega\left(n \lambda_{1} / \epsilon\right)$ and $c_{0} \geq 100$. Thus, we get the desired property in Part 2 of Definition 22.9.1.

Bounding $\left|\widehat{\lambda}_{1}-\lambda_{1}\right|$. It remains to bound $\left|\widehat{\lambda}_{1}-\lambda_{1}\right|$.

$$
\begin{array}{rlr}
\left|\widehat{\lambda}_{1}-\lambda_{1}\right| & =\left|\left[A^{*}+\widetilde{E}\right]\left(\widehat{v}_{1}, \cdots, \widehat{v}_{1}\right)-\lambda_{1}\right| \\
& \leq\left|\widetilde{E}\left(\widehat{v}_{1}, \cdots, \widehat{v}_{1}\right)\right|+\left|A^{*}\left(\widehat{v}_{1}, \cdots, \widehat{v}_{1}\right)-\lambda_{1}\right| & \text { by triangle inequality } \\
& =\left|\widetilde{E}\left(\widehat{v}_{1}, \cdots, \widehat{v}_{1}\right)\right|+\left|\left[\sum_{i=1}^{k} \lambda_{i} v_{i}^{\otimes p}\right]\left(\widehat{v}_{1}, \cdots, \widehat{v}_{1}\right)-\lambda_{1}\right| \quad \text { by } A^{*}=\sum_{i=1}^{k} \lambda_{i} v_{i}^{\otimes p} \\
& \leq \underbrace{\left|\widetilde{E}\left(\widehat{v}_{1}, \cdots, \widehat{v}_{1}\right)\right|}_{A}+\underbrace{\left.\left|\lambda_{1}\right| v_{1}^{\top} \widehat{v}_{1}\right|^{p}-\lambda_{1} \mid}_{B}+\underbrace{\sum_{j=2}^{k} \lambda_{j}\left|v_{j}^{\top} \widehat{v}_{1}\right|^{p}}_{C} . \quad \text { by triangle inequality }
\end{array}
$$

For the term $A$, we have

$$
\begin{array}{rlr}
A & =\left|\widetilde{E}\left(\widehat{v}_{1}, \cdots, \widehat{v}_{1}\right)\right| \\
& \leq\left|E\left(\widehat{v}_{1}, \cdots, \widehat{v}_{1}\right)\right|\left|+\left|\bar{E}\left(\widehat{v}_{1}, \cdots, \widehat{v}_{1}\right)\right|\right. & \\
& \leq\|E\|_{2}+\left|\bar{E}\left(\widehat{v}_{1}, \cdots, \widehat{v}_{1}\right)\right| & \\
& \leq \epsilon /\left(c_{0} \sqrt{n}\right)+\epsilon /\left(c_{0} \sqrt{n}\right) & \text { by } c_{0} \geq 100, n \geq 1
\end{array}
$$

It remains to upper bound the second term $B$ and the third term $C$.

$$
\begin{array}{rlrl}
B & =\lambda_{1}-\lambda_{1}\left(1-\frac{1}{2}\left\|v_{1}-\widehat{v}_{1}\right\|_{2}^{2}\right)^{p} & \text { by } v_{1}^{\top} \widehat{v}_{1}=1-\frac{1}{2}\left\|v_{1}-\widehat{v}_{1}\right\|_{2}^{2} \\
& \leq \lambda_{1} p \frac{1}{2}\left\|v_{1}-\widehat{v}_{1}\right\|_{2}^{2} & & \text { by }\left\|v_{1}-\widehat{v}_{1}\right\|_{2}^{2} \ll 1 \\
& \leq p \epsilon^{2} /\left(2 \lambda_{1}\right) & & \text { by }\left\|v_{1}-\widehat{v}_{1}\right\|_{2} \leq \epsilon / \lambda_{1} \\
& \leq \epsilon / 12 . & \text { by } p \epsilon /\left(2 \lambda_{1}\right) \leq 1 / 12
\end{array}
$$

For the term $C$,

$$
C=\sum_{j=2}^{k} \lambda_{j}\left|v_{j}^{\top} \widehat{v}_{1}\right|^{p} \leq \sum_{j=2}^{k} \lambda_{j}\left(\epsilon /\left(\sqrt{n} \lambda_{j}\right)\right)^{p}=\epsilon \sum_{j=2}^{k}\left(\epsilon /\left(\lambda_{j} \sqrt{n}\right)\right)^{p-1} \leq \epsilon / 4 .
$$

where the first inequality follows by Part 3 of Definition 22.9.1, and the last step follows by $\left(\epsilon / \lambda_{k}\right)^{p-1} \leq 1 /(4 k)$. It suffices to set $\epsilon<\frac{1}{4} k^{1 /(p-1)} \lambda_{k}$ Thus, putting it all together, we have

$$
\left|\widehat{\lambda}_{1}-\lambda_{1}\right| \leq A+B+C \leq \epsilon
$$

Bounding $\left|\widehat{v}_{1}^{\top} v_{j}\right|$. We consider any $j \in\{2, \cdots, k\}$. We choose $t^{*}$ to be smallest integer such that $\left|v_{1}^{\top} u_{t^{*}}\right| \geq 1-\frac{1}{c_{0}^{2} p^{2} k^{2}}$. This implies $\left|v_{j}^{\top} u_{t^{*}}\right| \leq \frac{1}{c_{0} p k}$. By Part 3 of Lemma 22.9.9, we have

$$
\left|v_{j}^{\top} u_{t^{*}}\right| /\left|v_{1}^{\top} u_{t^{*}}\right| \leq 0.8\left|v_{j}^{\top} u_{t^{*}-1}\right| /\left|v_{1}^{\top} u_{t^{*}-1}\right| \leq \cdots \leq 0.8^{t^{*}} \cdot u_{0} \leq 0.8^{t^{*}} \cdot 1 /(1 / \sqrt{n})
$$

Thus, we need to choose $t^{*}=\Omega\left(\log c_{0} n\right)$.
Let analyze $T>t^{*}$,

$$
\left|v_{j}^{\top} u_{T}\right| /\left|v_{1}^{\top} u_{T}\right| \leq 0.8^{T-t^{*}}\left|v_{j}^{\top} u_{t^{*}}\right| /\left|v_{1}^{\top} u_{t^{*}}\right|+5 \cdot 18 \epsilon /\left(c_{0} \lambda_{1} \sqrt{n}\right) .
$$

Thus, we need to choose $T=\Omega\left(\log \left(n \lambda_{1} / \epsilon\right)\right)$ and $c_{0} \geq 100$ to guarantee that $\left|v_{j}^{\top} u_{T}\right| \leq$ $\epsilon /\left(\lambda_{1} \sqrt{n}\right)$.

The case for $i=r+1$. Conditioned on all the first $t$ cases are holding, let's consider $t+1$ case, then the "true" noise $\widetilde{E}=E+\sum_{i=1}^{r} E_{i}+\bar{E} \in \mathbb{R}^{n^{p}}$. Similarly as the base case $t=1$, we first show how to bound $\left\|\widehat{v}_{r+1}-v_{r+1}\right\|_{2}$ (as Part 2 of Definition 22.9.1), and then we show how to bound $\left|\widehat{\lambda}_{r+1}-\lambda_{r+1}\right|$ (as Part 1 of Definition 22.9.1), at the end we show how to bound $\left|v_{r+1}^{\top} v_{j}\right|$ (as Part 3 of Definition 22.9.1).

Bounding $\left\|\widehat{v}_{r+1}-v_{r+1}\right\|_{2}$. The proof is identical to the Base case, the difference is, we need to choose $T=\Omega\left(\log \left(n \lambda_{t+1} / \epsilon\right)\right)$.

Bounding $\left|\widehat{\lambda}_{r+1}-\lambda_{r+1}\right|$. We redefine $A^{*}$ and $\widetilde{E}$ to be $A^{*}=\sum_{i=t+1}^{k} \lambda_{i} v_{i}^{\otimes p}$ and $\widetilde{E}=E+\bar{E}+\sum_{i=1}^{t} \widehat{E}_{i}$. Then rewrite $\left|\widehat{\lambda}_{t+1}-\lambda_{t+1}\right|$ as follows

$$
\begin{aligned}
& \left|\widehat{\lambda}_{r+1}-\lambda_{r+1}\right| \\
= & \left|\left[A^{*}+\widetilde{E}\right]\left(\widehat{v}_{r+1}, \cdots, \widehat{v}_{r+1}\right)-\lambda_{r+1}\right| \\
\leq & \left|\widetilde{E}\left(\widehat{v}_{r+1}, \cdots, \widehat{v}_{r+1}\right)\right|+\left|A^{*}\left(\widehat{v}_{r+1}, \cdots, \widehat{v}_{r+1}\right)-\lambda_{r+1}\right| \\
= & \left|\widetilde{E}\left(\widehat{v}_{r+1}, \cdots, \widehat{v}_{r+1}\right)\right|+\left|\left[\sum_{i=r+1}^{k} \lambda_{i} v_{i}^{\otimes p}\right]\left(\widehat{v}_{r+1}, \cdots, \widehat{v}_{r+1}\right)-\lambda_{r+1}\right| \quad \text { by triangle inequality } \\
\leq & \underbrace{\left|\widetilde{E}\left(\widehat{v}_{r+1}, \cdots, \widehat{v}_{r+1}\right)\right|}_{A}+\underbrace{\left.\left|\lambda_{r+1}\right| v_{r+1}^{\top} \widehat{v}_{r+1}\right|^{p}-\lambda_{r+1} \mid}_{B}+\underbrace{\sum_{j=r+2}^{k} \lambda_{j}\left|v_{j}^{\top} \widehat{v}_{r+1}\right|^{p}}_{C} . \quad \text { by } A^{*}=\sum_{i=r+1}^{k} \lambda_{i} v_{i}^{\otimes p}
\end{aligned}
$$

We need to analyze $A$,

$$
\begin{aligned}
A & =\left|\widetilde{E}\left(\widehat{v}_{r+1}, \cdots, \widehat{v}_{r+1}\right)\right| \\
& =\left|E\left(\widehat{v}_{r+1}, \cdots, \widehat{v}_{r+1}\right)\right|+\left|\bar{E}\left(\widehat{v}_{r+1}, \cdots, \widehat{v}_{r+1}\right)\right|+\left|\sum_{i=1}^{r} \widehat{E}_{i}\left(\widehat{v}_{r+1}, \cdots, \widehat{v}_{r+1}\right)\right| \\
& \leq \epsilon /\left(c_{0} \sqrt{n}\right)+\epsilon /\left(c_{0} \sqrt{n}\right)+4 \epsilon /\left(c_{0} \sqrt{n}\right) \\
& \leq \epsilon / 12
\end{aligned}
$$

$$
\text { by } c_{0} \geq 100, n \geq 1
$$

Bounding $\left|\widehat{v}_{r+1}^{\top} v_{j}\right|$. We consider any $j \in\{r+2, \cdots, k\}$. The proof is identical to the base.

Theorem 22.9.14 (Arbitrary even order sublinear time robust tensor power method). There exists some sufficiently small constant $\gamma \in(0,1)$ such that any $p \in\left[3, n^{\gamma}\right], k \in\left[1, n^{\gamma}\right]$. There exists some sufficiently small constant $\alpha \in(0,1-20 \gamma)$ such that for any tensor $A=A^{*}+E \in \mathbb{R}^{n^{p}}$, where $A^{*}=\sum_{i=1}^{k} \lambda_{i} v_{i}^{\otimes p}$ with $\lambda_{i}>0$ and orthonormal basis vectors
$\left\{v_{1}, \ldots, v_{k}\right\} \subseteq \mathbb{R}^{n}$. Let $\lambda_{1}, \lambda_{k}\left(\geq 1 / n^{\gamma}\right)$ be the largest and smallest values in $\left\{\lambda_{i}\right\}_{i=1}^{k}$ and $\left\{\widehat{\lambda}_{i}, \widehat{v}_{i}\right\}_{i=1}^{k}$ be outputs of the "importance sampling" robust tensor power method. For any sufficiently large constant $c_{0} \geq 100$, there exists a sufficiently small constant $c>0$, for any $\epsilon \in\left(0, c \lambda_{k} /\left(c_{0} p^{2} k n^{(p-2) / 2}\right)\right.$ if $E$ satisfies that $\|E\|_{2} \leq \epsilon /\left(c_{0} \sqrt{n}\right)$ and if $T=\Omega\left(\log \left(\lambda_{1} n / \epsilon\right)\right), L=$ $\Omega(k \log (k)), m=\Omega\left(T L k^{2}\right), b=\Omega\left(m n c_{0}^{2} / \epsilon^{2}\right)$ then our algorithm uses $O(n k)$ spaces, runs in $O\left(n^{p-\alpha}\right)$ time and then with probability at least $9 / 10$, there exists a permutation $\pi:[k] \rightarrow[k]$, such that $\forall i \in[k]$,

$$
\left|\lambda_{i}-\widehat{\lambda}_{\pi(i)}\right| \leq \epsilon, \quad\left\|v_{i}-\widehat{v}_{\pi(i)}\right\|_{2} \leq \epsilon / \lambda_{i}
$$

Proof. It follows by combining Theorem 22.3.4, Corollary 22.7.1 and Theorem 22.8.9 . The total number of iterations is $m=\Omega\left(T L k^{2}\right)$, thus we need to take the union bound over all the iterations. The total running time is $O(m b)=n^{p-\alpha^{\prime}}\left(\log \lambda_{1} n\right)^{2} / \lambda_{k}^{2}$, where $\alpha^{\prime}$ is some sufficiently small constant.

### 22.10 Algorithm

```
Algorithm 22.3 Subroutine for approximate tensor contraction \(A(I, v, w)\)
    function \(\operatorname{ApproxTIvw}\left(A, v, w, n, B,\left\{\widehat{b}_{i}\right\}\right)\)
    \(\widetilde{q}, \widetilde{r} \leftarrow \operatorname{CumProb}(v, w)\)
    for \(d=1 \rightarrow B\) do
        \(\mathcal{L} \leftarrow \operatorname{GenRandTuples}\left(\sum_{i=1}^{n} \widehat{b}_{i}, \widetilde{q}, \widetilde{r}\right)\)
        for \(i=1 \rightarrow n\) do
            \(s_{i}^{(d)} \leftarrow 0\)
            for \(\ell=1 \rightarrow \widehat{b}_{i}\) do
                \((j, k) \leftarrow \mathcal{L}_{(i-1) b+\ell}\)
                \(s_{i}^{(d)} \leftarrow s_{i}^{(d)}+\frac{1}{q_{j} r_{k}} A_{i, j, k} \cdot u_{j} \cdot u_{k}\)
            end for
        end for
    end for
    \(\widehat{A}(I, v, w)_{i} \leftarrow \underset{d \in[B]}{\operatorname{median}} s_{i}^{(d)} / \widehat{b}_{i}, \forall i \in[n]\)
    return \(\widehat{A}(I, v, w)\)
```

```
Algorithm 22.4 Subroutine for approximate tensor contraction \(A(u, v, w)\)
    function \(\operatorname{ApproxTuVW}(A, u, v, w, n, B, \widehat{b})\)
    \(\widetilde{p}, \widetilde{q}, \widetilde{r} \leftarrow \operatorname{CumProb}(u, v, w)\)
    for \(d=1 \rightarrow B\) do
        \(\mathcal{L} \leftarrow \operatorname{GenRandTuples}(\widehat{b}, \widetilde{p}, \widetilde{q}, \widetilde{r})\).
        \(s^{(d)} \leftarrow 0\)
        for \((i, j, k) \in \mathcal{L}\) do
            \(s^{(d)} \leftarrow s^{(d)}+\frac{1}{p_{i} q_{j} r_{k}} A_{i, j, k} \cdot u_{i} \cdot u_{j} \cdot u_{k}\)
        end for
        \(s^{(d)} \leftarrow s^{(d)} / \widehat{b}\)
    end for
    \(\widehat{A}(u, v, w) \leftarrow \underset{d \in[B]}{\operatorname{median}} s^{(d)}\)
    return \(\widehat{A}(u, v, w)\)
```

```
Algorithm 22.5 Our main algorithm
    function ImportanceSamplingRobustTensorPowerMethod \((A, n, B, b)\)
    if \(s_{i}\) are known, where \(\left\|A_{i, *, *}\right\|_{F}^{2} \lesssim s_{i}\) then
        \(\widehat{b}_{i} \leftarrow b \cdot s_{i} /\|A\|_{F}^{2}, \forall i \in[n]\)
    else
        \(\widehat{b}_{i} \leftarrow b / n, \forall i \in[n]\)
    end if
    \(\widehat{b}=\sum_{i=1}^{n} \widehat{b}_{i}\)
    for \(\ell=1 \rightarrow L\) do
        for \(t=1 \rightarrow T\) do
            \(u^{(\ell)} \leftarrow \operatorname{ApproxTIvW}\left(A, u^{(\ell)}, u^{(\ell)}, n, B,\left\{\widehat{b}_{i}\right\}\right)\)
            \(u^{(\ell)} \leftarrow u^{(\ell)} /\left\|u^{(\ell)}\right\|_{2}\)
        end for
        \(\lambda^{(\ell)} \leftarrow \operatorname{ApproxTuvW}\left(A, u^{(\ell)}, u^{(\ell)}, u^{(\ell)}, n, B, \widehat{b}\right)\)
    end for
    \(\ell^{*} \leftarrow \arg \max _{\ell \in[L]} \lambda^{(\ell)}\)
    \(u^{*} \leftarrow u^{\left(\ell^{*}\right)}\)
    for \(t=1 \rightarrow T\) do
        \(u^{*} \leftarrow \operatorname{ApproxTIvw}\left(A, u^{*}, u^{*}, n, B,\left\{\widehat{b}_{i}\right\}\right)\)
        \(u^{*} \leftarrow u^{*} /\left\|u^{*}\right\|_{2}\)
    end for
    \(\lambda^{*} \leftarrow \operatorname{ApproxTuvw}\left(A, u^{*}, u^{*}, u^{*}, n, B, \widehat{b}\right)\)
    return \(\lambda^{*}, u^{*}\)
```


### 22.11 Additional Experiment Results

The following sections gather all of our experiment results we collected to compare our sampling-based robust tensor power method to a sketching-based robust tensor power method. Our algorithm has the following parameters to set:

1. $T$ : Number of power iterations
2. $L$ : The number of starting vectors of the robust tensor power method
3. B: The number of sketches used in sketching, or the number of repetitions of sampling
4. $b$ : The size of the sketch, or the number of indices sampled

In all of our experiments, we fix $T=30, L=50$, and change $B$ and $b$ for each input tensor. We observe the squared residual Frobenius norm to evaluate the performance of each algorithm. We only compute the first eigenvalue and eigenvector (rank-1 recovery) of each tensor. This enables us to run a large collection of tensors within a reasonable time.

### 22.11.1 Synthetic tensors with inverse decaying eigenvalues

The tables shown in this section use synthetic tensors generated with eigenvalues decaying as $\lambda_{i}=\frac{1}{i}$. All tensors have rank $k=100$ and added noise with $\sigma=0.01$.

| Sketching based robust power method |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  |  | Running time (s) |  |  | Preprocessing time (s) |  |  |
| $b \quad B$ | 10 | 30 | 50 | 10 | 30 | 50 | 10 | 30 | 50 |
| $2^{10}$ | 1.011 | 0.9263 | 0.6358 | 0.6034 | 2.407 | 4.382 | 5.386 | 15.96 | 26.47 |
| $2^{12}$ | 1.01 | 0.4912 | 0.4451 | 1.342 | 4.603 | 8.012 | 5.972 | 17.34 | 28.72 |
| $2^{14}$ | 0.4394 | 0.4161 | 0.4085 | 4.921 | 15.64 | 27.93 | 8.877 | 24.81 | 40.87 |
| $2^{16}$ | 0.4089 | 0.4029 | 0.4006 | 22.48 | 69.67 | 115.2 | 13.79 | 34.86 | 55.65 |

Table 22.3: Sketching based robust power method: $n=\mathbf{1 2 0 0}$, inverse decay, $\|\mathbf{T}\|_{F}^{2}=1.01$


Table 22.4: Importance sampling based robust power method, without prescanning: $n=$ 1200, inverse decay, $\|\mathbf{T}\|_{F}^{2}=1.01$

| Importance sampling based robust power method (with prescanning) |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  |  | Running time (s) |  |  | Preprocessing time (s) |  |  |
| b B | 10 | 30 | 50 | 10 | 30 | 50 | 10 | 30 | 50 |
| $5 n$ | 0.408 | 0.4025 | 0.4005 | 3.621 | 11.9 | 20.49 | 2.234 | 2.235 | 2.241 |
| $10 n$ | 0.4051 | 0.4003 | 0.3997 | 6.302 | 20.56 | 35.24 | 2.225 | 2.223 | 2.226 |
| $20 n$ | 0.4057 | 0.3991 | 0.3993 | 11.67 | 37.89 | 64.52 | 2.225 | 2.226 | 2.225 |
| $30 n$ | 0.4027 | 0.3993 | 0.399 | 16.99 | 54.83 | 93.14 | 2.225 | 2.225 | 2.226 |
| $40 n$ | 0.4024 | 0.399 | 0.3988 | 22.23 | 72.13 | 122.2 | 2.223 | 2.225 | 2.225 |

Table 22.5: Importance sampling based robust power method with prescanning: $n=\mathbf{1 2 0 0}$, inverse decay, $\|\mathbf{T}\|_{F}^{2}=1.01$

```
Algorithm 22.2 Sampling Procedure for Rank-2
    procedure SAmplingRank2 \((A, n, b, \alpha)\)
                                    \(\triangleright\) Lemma 22.8.5
    \(\mathcal{S} \leftarrow S_{1} \leftarrow S_{2} \leftarrow S_{3} \leftarrow S_{4} \leftarrow \emptyset\)
    for \(i \in[n]\) do
            if \(\left|A_{i, i, i}\right| \geq 3 / n\) then
                \(S_{1} \leftarrow S_{1} \cup i\)
            else
                \(S_{2} \leftarrow S_{2} \cup i\)
            end if
        end for \(\quad \triangleright\left|S_{1}\right|=O\left(n^{1 / 3}\right)\)
        for \(i \in S_{1}\) do \(\quad \triangleright\) Part 1 .
            Choose \(b\) i.i.d. samples from slice \(i\), add to \(\mathcal{S}\)
        end for
                \(\triangleright\) \# samples \(=O\left(b n^{1 / 3}\right)\)
        for \(i \in[n]\) do \(\quad \triangleright\) Part 2.
            Choose \(O\left(b / n^{\alpha}\right)\) i.i.d. samples from slice \(i\), add to \(\mathcal{S}\)
        end for
        \(\triangleright \#\) samples \(=O\left(b n^{1-\alpha}\right)\)
        Randomly choose a set \(S_{3} \subseteq S_{2}\) such that \(\left|S_{3}\right|=O\left(n^{1 / 3+\alpha} \log n\right) \quad \triangleright\) Part 3 .
        for \(i \in S_{3}\) do
            for \(l \in[n]\) do
                if \(\left|A_{i, l, l}\right|^{2} \geq 1 /\left(4 n^{1+3 \alpha}\right)\) then \(\quad \triangleright \# l=O\left(n^{1 / 3+3 \alpha}\right)\)
                Choose \(b\) i.i.d. samples from slice \(l\), add to \(\mathcal{S}\)
                end if
            end for
        end for \(\quad \triangleright \#\) samples \(=O\left(b n^{2 / 3+4 \alpha}\right)\)
        for \(i \in[n]\) do \(\quad \triangleright\) Part 4 .
            for \(j \in[n]\) do
                if \(\left|A_{i, j, j}\right|^{2} \geq 1 /\left(4 n^{2 / 3+2 \alpha}\right)\) then
                \(S_{4} \leftarrow S_{4} \cup\{i, j\}\)
                end if
            end for
        end for
        for \(i \in S_{4}\) do
                        \(\triangleright\left|S_{4}\right|=O\left(n^{2 / 3+2 \alpha}\right)\)
            Choose \(b\) i.i.d samples from slice \(i\), add to \(\mathcal{S}\)
        end for
                                    \(\triangleright \#\) samples \(=O\left(b n^{2 / 3+2 \alpha}\right)\)
        return \(\mathcal{S}\)
    end procedure
```


### 22.11.2 Synthetic tensors with inverse square decaying eigenvalues

The tables shown in this section use synthetic tensors generated with eigenvalues decaying as $\lambda_{i}=\frac{1}{i^{2}}$. All tensors have rank $k=100$ and added noise with $\sigma=0.01$.

| Sketching based robust power method |  |  |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |  |  |  |  |  |
| $b \quad B$ | 10 | 30 | 50 |  | 10 | 30 | 50 |  | 10 | 30 | 50 |
| $2^{10}$ | 1.01 | 1.014 | 0.5437 |  | 0.6114 | 2.423 | 4.374 |  | 5.361 | 15.85 | 26.08 |
| $2^{12}$ | 1.02 | 0.2271 | 0.1549 |  | 1.344 | 4.563 | 8.022 | 5.978 | 17.23 | 28.31 |  |
| $2^{14}$ | 0.1513 | 0.1097 | 0.1003 |  | 4.928 | 15.51 | 27.87 | 8.788 | 24.72 | 40.4 |  |
| $2^{16}$ | 0.1065 | 0.09242 | 0.08936 |  | 22.28 | 69.7 | 116.3 | 13.76 | 34.74 | 55.34 |  |

Table 22.6: Sketching based robust power method: $n=1200$, inverse-square decay, $\|\mathbf{T}\|_{F}^{2}=1.01$


Table 22.7: Importance sampling based robust power method, without prescanning: $n=$ 1200, inverse-square decay, $\|\mathbf{T}\|_{F}^{2}=1.01$

| Importance sampling based robust power method (with prescanning) |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  |  | Running time (s) |  |  | Preprocessing time (s) |  |  |
| $b \quad B$ | 10 | 30 | 50 | 10 | 30 | 50 | 10 | 30 | 50 |
| $5 n$ | 0.08657 | 0.08684 | 0.08636 | 3.1 | 10.47 | 18 | 2.234 | 2.236 | 2.234 |
| $10 n$ | 0.08741 | 0.08677 | 0.08668 | 5.427 | 17.43 | 30.26 | 2.232 | 2.233 | 2.233 |
| $20 n$ | 0.08648 | 0.08624 | 0.08634 | 9.843 | 31.42 | 54.49 | 2.226 | 2.226 | 2.226 |
| $30 n$ | 0.08635 | 0.08634 | 0.08615 | 14.33 | 45.4 | 63.85 | 2.226 | 2.224 | 2.227 |
| $40 n$ | 0.08622 | 0.08652 | 0.08619 | 18.68 | 59.32 | 82.83 | 2.225 | 2.225 | 2.225 |

Table 22.8: Importance sampling based robust power method with prescanning: $n=\mathbf{1 2 0 0}$, inverse-square decay, $\|\mathbf{T}\|_{F}^{2}=1.01$

### 22.11.3 Synthetic tensors with linearly decaying eigenvalues

The tables shown in this section use synthetic tensors generated with eigenvalues decaying as $\lambda_{i}=1-\frac{i-1}{k}$. All tensors have rank $k=100$ and added noise $\sigma=0.01$. Due to the slowly decaying eigenvalues, we expect to see large residual values of rank-1 decomposition, and both algorithms have to work harder to get a better recovery. Our sampling based method works better especially when the dimension is large.

| Sketching based robust power method |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  |  | Running time (s) |  |  | Preprocessing time (s) |  |  |
| $\stackrel{B}{ }$ | 10 | 30 | 50 | 10 | 30 | 50 | 10 | 30 | 50 |
| $2^{10}$ | 1.01 | 1.01 | 1.01 | 0.6153 | 2.457 | 4.381 | 5.364 | 15.78 | 26.25 |
| $2^{12}$ | 1.01 | 1.01 | 1.01 | 1.356 | 4.704 | 7.994 | 5.975 | 17.24 | 28.48 |
| $2^{14}$ | 1.01 | 1.01 | 0.9829 | 4.942 | 15.84 | 27.82 | 8.851 | 24.62 | 40.5 |
| $2^{16}$ | 0.9835 | 0.9813 | 0.9811 | 22.66 | 72.01 | 115.1 | 13.75 | 35.13 | 55.59 |

Table 22.9: Sketching based robust power method: $n=\mathbf{1 2 0 0}$, linear decay, $\|\mathbf{T}\|_{F}^{2}=1.01$


Table 22.10: Importance sampling based robust power method, without prescanning: $n=$ 1200, linear decay, $\|\mathbf{T}\|_{F}^{2}=1.01$

| Importance sampling based robust power method (with prescanning) |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  |  | Running time (s) |  |  | Preprocessing time (s) |  |  |
| $b \quad B$ | 10 | 30 | 50 | 10 | 30 | 50 | 10 | 30 | 50 |
| $5 n$ | 1.01 | 0.991 | 0.9882 | 3.668 | 11.64 | 19.48 | 2.234 | 2.234 | 2.234 |
| $10 n$ | 1.009 | 0.9883 | 0.9851 | 6.512 | 20.14 | 33.45 | 2.225 | 2.225 | 2.226 |
| $20 n$ | 0.9918 | 0.9852 | 0.9828 | 11.9 | 36.43 | 61.01 | 2.225 | 2.225 | 2.226 |
| $30 n$ | 0.9873 | 0.9849 | 0.9826 | 17.14 | 52.89 | 87.36 | 2.226 | 2.226 | 2.226 |
| $40 n$ | 0.9882 | 0.9829 | 0.9822 | 22.38 | 69.01 | 114.6 | 2.226 | 2.227 | 2.226 |

Table 22.11: Importance sampling based robust power method with prescanning: $n=\mathbf{1 2 0 0}$, linear decay, $\|\mathbf{T}\|_{F}^{2}=1.01$
22.11.4 Results from real-life datasets

| Dataset | \# Documents | \# Vocabulary | \# Words | Short Description |
| ---: | :---: | :---: | :---: | :--- |
| Wiki | 114274 | 10000 | 58508288 | Wikipedia articles, preprocessed <br> by authors of [WTSA15] |
| Enron $^{3}$ | 186501 | 10000 | 16971591 | Enron Email Dataset, prepro- <br> cessed by authors of [WTSA15] |
| AP $^{4}$ | 2246 | 10473 | 435838 | Associated Press [BNJ01] <br> NIPS $^{5}$ 1500 |
| 12419 | 1900000 | NIPS full papers, preprocessed by <br> UCI |  |  |
| KOS $^{6}$ | 3430 | 6906 | 467714 | KOS blog, preprocessed by UCI <br> NSF $^{7}$ 449078 |
| 22170 | 4406190 | NSF research award abstracts <br> $1990-2003$ |  |  |

Table 22.12: List of six real-life datasets

We used the two same datasets as the previous work [WTSA15]: Wiki and Enron, as well as four additional real-life datasets, shown in Table 22.12. We use the program from the authors of [TWZS15] to build a $K \times K \times K$ tensor $A_{\text {LDA }}$ for each dataset. Due to limited time our experiments only use $K=200$. We expect to gain more over the sketching-based method on larger tensors.

Since these tensors comes from real-world applications and are not normalized, their Frobenius norm is usually very large.

[^79]| Sketching based robust power method: dataset Enron |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |  |
| $b$ | 10 | 30 |  | 10 | 30 |  | 10 |
| $2^{10}$ | $4.733 \mathrm{e}+07$ | $4.608 \mathrm{e}+07$ |  | 0.3667 | 0.7942 | 0.2429 | 0.2378 |
| $2^{11}$ | $4.677 \mathrm{e}+07$ | $4.58 \mathrm{e}+07$ |  | 0.4434 | 1.486 | 0.2468 | 0.3238 |
| $2^{12}$ | $4.581 \mathrm{e}+07$ | $4.559 \mathrm{e}+07$ |  | 0.995 | 3.243 | 0.3647 | 0.6283 |
| $2^{13}$ | $4.568 \mathrm{e}+07$ | $4.546 \mathrm{e}+07$ |  | 2.023 | 6.017 | 0.5745 | 0.6458 |
| $2^{14}$ | $4.556 \mathrm{e}+07$ | $4.542 \mathrm{e}+07$ |  | 4.561 | 13.97 | 0.9698 | 1.061 |

Table 22.13: Sketching based robust power method: dataset Enron, $\|\mathbf{T}\|_{F}^{2}=4.96 \mathrm{e}+07$
Importance sampling based robust power method (without prescanning): dataset Enron

|  | Squared residual norm |  | Running time (s) |  |  | Preprocessing time (s) |  |
| ---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $b \quad B$ | 10 | 30 |  | 10 | 30 |  | 10 |
| $5 n$ | $4.579 \mathrm{e}+07$ | $4.557 \mathrm{e}+07$ |  | 0.3784 | 1.139 |  | 0.0 |
| $10 n$ | $4.563 \mathrm{e}+07$ | $4.559 \mathrm{e}+07$ |  | 0.5689 | 1.603 |  | 0.0 |
| $20 n$ | $4.56 \mathrm{e}+07$ | $4.545 \mathrm{e}+07$ |  | 1.003 | 2.706 |  | 0.0 |
| $30 n$ | $4.555 \mathrm{e}+07$ | $4.543 \mathrm{e}+07$ |  | 1.362 | 3.717 |  | 0.0 |
| $40 n$ | $4.564 \mathrm{e}+07$ | $4.544 \mathrm{e}+07$ |  | 1.693 | 4.753 |  | 0.0 |

Table 22.14: Importance sampling based robust power method, without prescanning: dataset Enron, $\|\mathbf{T}\|_{F}^{2}=4.96 \mathrm{e}+07$

| Importance sampling based robust power method (with prescanning): dataset Enron |  |  |  |  |  |  |  |  |
| ---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |  |  |
| $b$ | 10 | 30 |  | 10 | 30 |  | 10 | 30 |
| $5 n$ | $4.605 \mathrm{e}+07$ | $4.545 \mathrm{e}+07$ |  | 0.4014 | 1.176 |  | 0.0105 | 0.01043 |
| $10 n$ | $4.573 \mathrm{e}+07$ | $4.55 \mathrm{e}+07$ |  | 0.6255 | 1.773 |  | 0.01036 | 0.01044 |
| $20 n$ | $4.561 \mathrm{e}+07$ | $4.545 \mathrm{e}+07$ |  | 1.318 | 2.923 |  | 0.01041 | 0.01039 |
| $30 n$ | $4.563 \mathrm{e}+07$ | $4.546 \mathrm{e}+07$ |  | 1.459 | 4.041 |  | 0.01101 | 0.01104 |
| $40 n$ | $4.56 \mathrm{e}+07$ | $4.544 \mathrm{e}+07$ |  | 1.841 | 5.092 |  | 0.01039 | 0.01102 |

Table 22.15: Importance sampling based robust power method with prescanning: dataset Enron, $\|\mathbf{T}\|_{F}^{2}=4.96 \mathrm{e}+07$

| Sketching based robust power method: dataset AP |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |  |
| $b \quad B$ | 10 | 30 |  | 10 | 30 |  | 10 |
| $2^{10}$ | $7.065 \mathrm{e}+06$ | $6.703 \mathrm{e}+06$ |  | 0.2316 | 0.8128 | 0.172 | 0.2472 |
| $2^{11}$ | $6.674 \mathrm{e}+06$ | $6.613 \mathrm{e}+06$ |  | 0.4404 | 1.417 | 0.2419 | 0.3072 |
| $2^{12}$ | $6.656 \mathrm{e}+06$ | $6.57 \mathrm{e}+06$ |  | 0.9689 | 2.936 | 0.3586 | 0.4257 |
| $2^{13}$ | $6.585 \mathrm{e}+06$ | $6.573 \mathrm{e}+06$ |  | 2.058 | 6.172 | 0.5826 | 0.6592 |
| $2^{14}$ | $6.57 \mathrm{e}+06$ | $6.555 \mathrm{e}+06$ |  | 4.558 | 13.99 | 0.9601 | 1.059 |

Table 22.16: Sketching based robust power method: dataset AP, $\|\mathbf{T}\|_{F}^{2}=6.905 \mathrm{e}+06$
Importance sampling based robust power method (without prescanning): dataset AP

|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |  |
| ---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $b \quad B$ | 10 | 30 |  | 10 | 30 |  | 10 |
| $5 n$ | $6.859 \mathrm{e}+06$ | $6.805 \mathrm{e}+06$ |  | 0.3762 | 1.548 |  | 0.0 |
| $10 n$ | $6.631 \mathrm{e}+06$ | $6.904 \mathrm{e}+06$ |  | 0.5469 | 1.577 |  | 0.0 |
| $20 n$ | $6.572 \mathrm{e}+06$ | $6.72 \mathrm{e}+06$ |  | 0.9535 | 2.587 |  | 0.0 |
| $30 n$ | $6.64 \mathrm{e}+06$ | $6.557 \mathrm{e}+06$ |  | 1.263 | 3.53 | 0.0 |  |
| $40 n$ | $6.565 \mathrm{e}+06$ | $6.564 \mathrm{e}+06$ |  | 1.613 | 4.478 |  | 0.0 |

Table 22.17: Importance sampling based robust power method, without prescanning: dataset $\mathbf{A P},\|\mathbf{T}\|_{F}^{2}=6.905 \mathrm{e}+06$

| Importance sampling based robust power method (with prescanning): dataset AP |  |  |  |  |  |  |  |
| ---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |  |
| $b$ | 10 | 30 |  | 10 | 30 |  | 10 |
| 30 |  |  |  |  |  |  |  |
| $5 n$ | $6.58 \mathrm{e}+06$ | $6.893 \mathrm{e}+06$ |  | 0.4007 | 1.233 | 0.01045 | 0.01103 |
| $10 n$ | $6.604 \mathrm{e}+06$ | $6.567 \mathrm{e}+06$ |  | 0.7035 | 2.011 |  | 0.01061 |
| $20 n$ | $6.803 \mathrm{e}+06$ | $6.563 \mathrm{e}+06$ |  | 1.166 | 3.246 |  | 0.01042 |
| $30 n$ | $6.58 \mathrm{e}+06$ | $6.573 \mathrm{e}+06$ |  | 1.588 | 4.815 | 0.01103 | 0.01061 |
| $40 n$ | $6.66 \mathrm{e}+06$ | $6.555 \mathrm{e}+06$ |  | 1.989 | 5.601 |  | 0.01084 |

Table 22.18: Importance sampling based robust power method with prescanning: dataset $\mathbf{A P},\|\mathbf{T}\|_{F}^{2}=6.905 \mathrm{e}+06$

| Sketching based robust power method: dataset NSF |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |  |
| $b$ | 10 | 30 |  | 10 | 30 |  | 10 |
| $2^{10}$ | $4.342 \mathrm{e}+06$ | $3.983 \mathrm{e}+06$ |  | 0.2329 | 1.118 | 0.1731 | 0.2938 |
| $2^{11}$ | $4.183 \mathrm{e}+06$ | $3.906 \mathrm{e}+06$ |  | 0.4415 | 1.434 | 0.2439 | 0.3061 |
| $2^{12}$ | $3.929 \mathrm{e}+06$ | $3.869 \mathrm{e}+06$ |  | 0.9758 | 2.952 | 0.3554 | 0.4268 |
| $2^{13}$ | $3.885 \mathrm{e}+06$ | $3.86 \mathrm{e}+06$ |  | 2.045 | 6.015 | 0.6591 | 0.6489 |
| $2^{14}$ | $3.865 \mathrm{e}+06$ | $3.854 \mathrm{e}+06$ |  | 4.591 | 13.89 |  | 1.081 |

Table 22.19: Sketching based robust power method: dataset NSF, $\|\mathbf{T}\|_{F}^{2}=4.765 \mathrm{e}+06$

| Importance sampling based robust power method (without prescanning): dataset NSF |  |  |  |  |  |  |  |  |
| ---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |  |  |
| $b$ | 10 | 30 |  | 10 | 30 |  | 10 | 30 |
| $5 n$ | $3.899 \mathrm{e}+06$ | $3.885 \mathrm{e}+06$ |  | 0.3701 | 1.166 | 0.0 | 0.0 |  |
| $10 n$ | $3.891 \mathrm{e}+06$ | $3.862 \mathrm{e}+06$ |  | 0.5678 | 1.698 | 0.0 | 0.0 |  |
| $20 n$ | $3.896 \mathrm{e}+06$ | $3.861 \mathrm{e}+06$ |  | 1.029 | 2.927 | 0.0 | 0.0 |  |
| $30 n$ | $3.889 \mathrm{e}+06$ | $3.856 \mathrm{e}+06$ |  | 1.415 | 3.941 | 0.0 | 0.0 |  |
| $40 n$ | $3.867 \mathrm{e}+06$ | $3.857 \mathrm{e}+06$ |  | 1.823 | 5.287 |  | 0.0 | 0.0 |

Table 22.20: Importance sampling based robust power method, without prescanning: dataset NSF, $\|\mathbf{T}\|_{F}^{2}=4.765 \mathrm{e}+06$

| Importance sampling based robust power method (with prescanning): dataset NSF |  |  |  |  |  |  |  |
| ---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |  |
| $b \quad B$ | 10 | 30 | 10 | 30 | 10 | 30 |  |
| $5 n$ | $3.938 \mathrm{e}+06$ | $3.902 \mathrm{e}+06$ | 0.5219 | 1.566 | 0.01039 | 0.01108 |  |
| $10 n$ | $3.898 \mathrm{e}+06$ | $3.867 \mathrm{e}+06$ | 0.8981 | 2.568 | 0.01104 | 0.01043 |  |
| $20 n$ | $3.947 \mathrm{e}+06$ | $3.864 \mathrm{e}+06$ |  | 1.561 | 4.546 | 0.01097 | 0.01057 |
| $30 n$ | $3.889 \mathrm{e}+06$ | $3.852 \mathrm{e}+06$ | 2.207 | 6.477 | 0.011 | 0.01107 |  |
| $40 n$ | $3.891 \mathrm{e}+06$ | $3.853 \mathrm{e}+06$ | 2.827 | 8.446 | 0.01103 | 0.01065 |  |

Table 22.21: Importance sampling based robust power method with prescanning: dataset NSF, $\|\mathbf{T}\|_{F}^{2}=4.765 \mathrm{e}+06$

| Sketching based robust power method: dataset NIPS |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |  |  |
| $b \quad B$ | 10 | 30 |  | 10 | 30 |  | 10 | 30 |
| $2^{10}$ | $5.649 \mathrm{e}+09$ | $5.432 \mathrm{e}+09$ |  | 0.2383 | 0.8686 | 0.1724 | 0.2625 |  |
| $2^{11}$ | $5.522 \mathrm{e}+09$ | $5.383 \mathrm{e}+09$ |  | 0.45 | 1.52 | 0.2417 | 0.3247 |  |
| $2^{12}$ | $5.394 \mathrm{e}+09$ | $5.366 \mathrm{e}+09$ |  | 1.264 | 3.304 | 0.5082 | 0.6358 |  |
| $2^{13}$ | $5.375 \mathrm{e}+09$ | $5.353 \mathrm{e}+09$ |  | 2.07 | 6.068 | 0.5766 | 0.7604 |  |
| $2^{14}$ | $5.36 \mathrm{e}+09$ | $5.35 \mathrm{e}+09$ |  | 4.591 | 13.84 | 0.9817 | 1.044 |  |

Table 22.22: Sketching based robust power method: dataset NIPS, $\|\mathbf{T}\|_{F}^{2}=6.058 \mathrm{e}+09$

| Importance sampling based robust power method (without prescanning): dataset NIPS |  |  |  |  |  |  |  |
| ---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |  |
| $b \quad B$ | 10 | 30 | 10 | 30 | 10 | 30 |  |
| $5 n$ | $5.648 \mathrm{e}+09$ | $5.403 \mathrm{e}+09$ |  | 0.3205 | 0.9938 | 0.0 | 0.0 |
| $10 n$ | $5.417 \mathrm{e}+09$ | $5.347 \mathrm{e}+09$ | 0.4739 | 1.753 | 0.0 | 0.0 |  |
| $20 n$ | $5.352 \mathrm{e}+09$ | $5.349 \mathrm{e}+09$ | 0.8345 | 2.464 | 0.0 | 0.0 |  |
| $30 n$ | $5.361 \mathrm{e}+09$ | $5.35 \mathrm{e}+09$ | 1.19 | 3.317 | 0.0 | 0.0 |  |
| $40 n$ | $5.352 \mathrm{e}+09$ | $5.349 \mathrm{e}+09$ | 1.532 | 4.153 | 0.0 | 0.0 |  |

Table 22.23: Importance sampling based robust power method, without prescanning: dataset NIPS, $\|\mathbf{T}\|_{F}^{2}=6.058 \mathrm{e}+09$

| Importance sampling based robust power method (with prescanning): dataset NIPS |  |  |  |  |  |  |  |
| ---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |  |
| $b$ | 10 | 30 |  | 10 | 30 |  | 10 |
| 30 |  |  |  |  |  |  |  |
| $5 n$ | $6.019 \mathrm{e}+09$ | $5.943 \mathrm{e}+09$ | 0.4818 | 1.401 | 0.01039 | 0.01057 |  |
| $10 n$ | $5.352 \mathrm{e}+09$ | $5.389 \mathrm{e}+09$ | 0.797 | 2.197 | 0.01063 | 0.01105 |  |
| $20 n$ | $5.387 \mathrm{e}+09$ | $5.349 \mathrm{e}+09$ | 1.335 | 3.849 | 0.01042 | 0.01039 |  |
| $30 n$ | $5.363 \mathrm{e}+09$ | $5.346 \mathrm{e}+09$ | 1.895 | 5.491 | 0.0104 | 0.01042 |  |
| $40 n$ | $5.352 \mathrm{e}+09$ | $5.382 \mathrm{e}+09$ | 2.523 | 7.019 | 0.01064 | 0.011 |  |

Table 22.24: Importance sampling based robust power method with prescanning: dataset NIPS, $\|\mathbf{T}\|_{F}^{2}=6.058 \mathrm{e}+09$

| Sketching based robust power method: dataset KOS |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |  |
| $b \quad B$ | 10 | 30 |  | 10 | 30 |  | 10 |
| $2^{10}$ | $5.015 \mathrm{e}+04$ | $4.807 \mathrm{e}+04$ |  | 0.2356 | 0.8581 | 0.1781 | 0.2507 |
| $2^{11}$ | $4.858 \mathrm{e}+04$ | $4.787 \mathrm{e}+04$ |  | 0.4482 | 1.417 | 0.2411 | 0.3072 |
| $2^{12}$ | $4.797 \mathrm{e}+04$ | $4.773 \mathrm{e}+04$ |  | 0.9652 | 2.947 | 0.3556 | 0.4265 |
| $2^{13}$ | $4.776 \mathrm{e}+04$ | $4.764 \mathrm{e}+04$ |  | 2.046 | 5.985 | 0.5722 | 0.6517 |
| $2^{14}$ | $4.771 \mathrm{e}+04$ | $4.761 \mathrm{e}+04$ |  | 4.453 | 13.87 | 0.968 | 1.052 |

Table 22.25: Sketching based robust power method: dataset KOS, $\|\mathbf{T}\|_{F}^{2}=5.016 \mathrm{e}+04$
Importance sampling based robust power method (without prescanning): dataset KOS

|  | Squared residual norm |  | Running time (s) |  | Preprocessing time (s) |  |  |
| ---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $b$ | 10 | 30 |  | 10 | 30 |  | 10 |
| $B$ | $4.783 \mathrm{e}+04$ | $4.766 \mathrm{e}+04$ |  | 0.3968 | 1.181 |  | 0.0 |
| $10 n$ | $4.817 \mathrm{e}+04$ | $4.767 \mathrm{e}+04$ |  | 0.6259 | 1.704 |  | 0.0 |
| $20 n$ | $4.775 \mathrm{e}+04$ | $4.766 \mathrm{e}+04$ |  | 1.039 | 2.79 |  | 0.0 |
| $30 n$ | $4.771 \mathrm{e}+04$ | $4.762 \mathrm{e}+04$ |  | 1.357 | 3.57 |  | 0.0 |
| $40 n$ | $4.767 \mathrm{e}+04$ | $4.762 \mathrm{e}+04$ |  | 1.76 | 4.892 |  | 0.0 |

Table 22.26: Importance sampling based robust power method, without prescanning: dataset KOS, $\|\mathbf{T}\|_{F}^{2}=5.016 \mathrm{e}+04$

| Importance sampling based robust power method (with prescanning): dataset KOS |  |  |  |  |  |  |  |  |
| ---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Squared residual norm |  | Running time (s) |  |  | Preprocessing time (s) |  |  |
| $b \quad B$ | 10 | 30 |  | 10 | 30 |  | 10 | 30 |
| $5 n$ | $4.773 \mathrm{e}+04$ | $4.766 \mathrm{e}+04$ | 0.4808 | 1.35 | 0.01061 | 0.01099 |  |  |
| $10 n$ | $4.792 \mathrm{e}+04$ | $4.765 \mathrm{e}+04$ |  | 0.7363 | 2.024 | 0.01108 | 0.01044 |  |
| $20 n$ | $4.795 \mathrm{e}+04$ | $4.766 \mathrm{e}+04$ |  | 1.336 | 3.308 | 0.01056 | 0.01104 |  |
| $30 n$ | $4.785 \mathrm{e}+04$ | $4.765 \mathrm{e}+04$ |  | 1.728 | 4.631 |  | 0.01039 | 0.0104 |
| $40 n$ | $4.767 \mathrm{e}+04$ | $4.764 \mathrm{e}+04$ |  | 2.295 | 5.773 |  | 0.01063 | 0.01102 |

Table 22.27: Importance sampling based robust power method with prescanning: dataset KOS, $\|\mathbf{T}\|_{F}^{2}=5.016 \mathrm{e}+04$

## Part V

## Open Problems

## Chapter 23

## Open Problems

In this chapter we give a long list of open problems. Some of these problems have appeared in other places and might be well-known to researchers in the community; nevertheless we include them since they are related to this thesis in some way.

To encourage people to solve these problems, we offer cash for each problem. The reward of each problem is not decided by the importance of the problem, but rather determined by the time when it is solved. To decide the order of time, we will compare the arXiv ID. For the $i$-th problem solved, we will pay $0.01 \times 2^{i-1}$ dollars.

### 23.1 Linear Programs

Linear programs are among most important optimization problems from both theoretical and practical view. In a recent result, [CLS19] show how to solve linear programs in the current matrix multiplication time. It would be interesting if this is still true if we are in a world where $n \times n$ matrix multiplication can be done in $O\left(n^{2}\right)$ time.

Open Problem 1. Is it possible to solve linear programs in the matrix multiplication time?

Another question is whether we can solve linear programs faster on sparse input. [LS14, LS15] show how to solve linear program in $O\left(\sqrt{d} \cdot \mathrm{nnz}(A)+d^{2.5}\right)$ time, where $\mathrm{nnz}(A)$ is the number of non-zeros in $A$ and $d$ is the number of constraints.

Open Problem 2. Is it possible to solve linear programs in the input sparsity time?

There are other well-known open problems in the theory of linear programming. One important question is whether there is a strongly polynomial time algorithm.

Open Problem 3. Does LP admit a strongly polynomial time algorithm?

We concluded open problems on linear programs with a problem originally proposed by Seve Smale [Sma98] in 1998.

Open Problem 4 (Problem 9 in [Sma98]). Is there a polynomial-time algorithm over the real numbers which decides the feasibility of the linear system of inequalities $A x \geq b$ ?

### 23.2 Matrix Multiplication

Let $\omega$ denote the exponent of matrix multiplication. Let $\alpha$ denote the dual exponent of the matrix multiplication.

Open Problem 5. Is it possible that $\omega \geq 2.5-\alpha / 2$ ?

For the current value of $\omega \sim 2.38$ and $\alpha \sim 0.31$, this is true. If $\omega=2$, then $\alpha=1$, it is also true. For the motivation of this inequality, we refer readers to [CLS19, LSZ19].

We also state a very well-known problem about matrix multiplication

Open Problem 6. What is the right answer of $\omega$ ?

### 23.3 Ordinary Differential Equations

The following open problem is conjectured by Yin Tat Lee.

Open Problem 7. [LSV18] gives an algorithm for sampling logconcave densities under $m$-strong convexity and $M$-smoothness. It takes $\kappa^{1.5}$ iterations (where $\kappa:=M / m$ ), $O(1)$ gradient evaluation per iteration, and the total time is $\kappa^{1.5} d$. Is that possible to give an algorithm that takes $\kappa$ iterations, and $O(1)$ gradient evaluation per iteration?

### 23.4 Matrix Factorization

We discuss several open problems about matrix/tensor factorization in this section.

### 23.4.1 Weighted Factorization

Recall the definition of weighted low rank approximation

Definition 23.4.1. Given a matrix $A \in \mathbb{R}^{n \times n}$ and a matrix $W \in \mathbb{R}^{n \times n}$. Let $W$ denote a rank- $r$ weight matrix. For integer $k \geq 1$ and approximation ratio $\alpha>1$, the goal is to find a rank- $k$ matrix $B \in \mathbb{R}^{n \times n}$ such that

$$
\|W \circ(B-A)\|_{F} \leq \alpha \cdot \min _{\text {rank }-k A^{\prime}}\left\|W \circ\left(A^{\prime}-A\right)\right\|_{F} .
$$

We assume the entries of $A$ and $W$ are integers in the range $\{-M,-M+1, \cdots, M\}$ for an integer $M \leq 2^{\text {poly }(n)}$, i.e., that the entries of $A$ and $W$ can be specified with poly $(n)$ bits.

Open Problem 8. For $k=O(1)$ and $\alpha=O(1)$, [RSW16] showed an $n^{O(r)}$ upper bound and a $2^{\Omega(r)}$ lower bound under Feige-ETH. Is it possible to close the gap?

The above question was originally asked by David P. Woodruff (see Open Problem 9 in [CFHW17]).

Open Problem 9. If $W$ has $r$ distinct columns (which is a stronger assumption than $W$ has rank-r), [RSW16] showed $2^{\widetilde{O}\left(r^{2}\right)}$ upper bound and $2^{\Omega(r)}$ lower bound under Feige-ETH, is it possible to close the gap?

### 23.4.2 Nonnegative Matrix Factorization

Nonnegative matrix factorization was theoretically studied by [AGKM12] and [Moi13]. The formal definition is

Definition 23.4.2 (NMF). Given an $n \times n$ matrix $A$ and a parameter $k$, the goal is to output two nonnegative $n \times k$ matrices $U$ and $V^{\top}$ such that $U V=A$.

Open Problem 10. There is an algorithm for NMF that takes $n^{O\left(k^{2}\right)}$ time due to [Moi13]. Under ETH, any algorithm that solves NMF takes $n^{\Omega(k)}$ time. The question is, can we close the gap?

### 23.4.3 Tensor factorization

Tensor factorization/decomposition has many applications in computer science. We give a formal problem definition.

Definition 23.4.3. Given a $n \times n \times n$ tensor $A$ and a parameter $k$, the goal is to output three $n \times k$ matrices $U, V, W$ such that

$$
\|U \otimes V \otimes W-A\|_{F} \leq O(1) \min _{\text {rank }-k}\left\|A_{k}-A\right\|_{F} .
$$

Open Problem 11. [SWZ19b] propose an algorithm that takes $\mathrm{nnz}(A)+n \cdot \operatorname{poly}(k)+2^{O\left(k^{2}\right)}$ time. From the other perspective, [SWZ19b] show that there is no algorithm that takes $2^{o(k)}$ time unless Feige-ETH is false. For the situation $k=\omega(\log n)$, can we close the gap between $2^{O\left(k^{2}\right)}$ and $2^{\Omega(k)}$ ?

In fact, all three open problems are equivalent in some sense. If someone can solve one of them, then probably also can improve the remaining ones. From the lower bound
side, all the current hardness results are based on ETH-type assumption. We don't know how to boost the current reduction. From the algorithmic side, all the current results are based on polynomial system solver, which has exponential dependence on the number of variables. For the target rank- $k$ matrices, we don't know how to write down a system that only involves $o\left(k^{2}\right)$ variables.

### 23.5 Continuous Fourier Transform

### 23.5.1 One dimensional

We say $x(t): \mathbb{R} \rightarrow \mathbb{C}$ is a $k$-Fourier-sparse signal if $x(t)$ can be written as $x(t)=$ $\sum_{j=1}^{k} v_{j} e^{-2 \pi \mathrm{i} f_{j} t}$ where $v_{j} \in \mathbb{C}$ and $f_{j} \in \mathbb{R}, \forall j \in[k]$. Let $\|x\|_{T}=\left(\frac{1}{T} \int_{0}^{T}|x(t)|^{2} \mathrm{~d} t\right)^{1 / 2}$.

Open Problem 12. [CKPS16] showed that: for any $k$-Fourier-sparse signal $x(t): \mathbb{R} \rightarrow \mathbb{C}$ and any duration $T>0$, we have

$$
\max _{t \in[0, T]}|x(t)|^{2} \leq O\left(k^{4} \log ^{3} k\right) \cdot\|x\|_{T}^{2}
$$

[CKPS16] also showed that $\Omega\left(k^{2}\right)$ is necessary. Can we close the gap?

### 23.5.2 High dimensional

We say $x(t): \mathbb{R}^{d} \rightarrow \mathbb{C}$ is a $d$-dimensional $k$-Fourier-sparse signal if $x(t)$ can be written as $x(t)=\sum_{j=1}^{k} v_{j} e^{-2 \pi \mathbf{i} f_{j}^{\top} t}$ where $v_{j} \in \mathbb{C}$ and $f_{j} \in \mathbb{R}^{d}, \forall j \in[k]$. Let $\|x\|_{T}=$ $\left(\frac{1}{T^{d}} \int_{[0, T]^{d}}|x(t)|^{2} \mathrm{~d} t\right)^{1 / 2}$.

Open Problem 13. [Son17] showed that: for any d-dimensional $k$-Fourier-sparse signal $x(t): \mathbb{R}^{d} \rightarrow \mathbb{C}$ and any duration $T>0$, we have

$$
\max _{t \in[0, T]^{d}}|x(t)|^{2} \leq k^{O(d)} \cdot\|x\|_{T}^{2}
$$

Is it possible to improve the upper bound $k^{O(d)}$, e.g. to $\operatorname{poly}(k, d)$ ? Is it possible to prove a lower bound of $k^{\Omega(d)}$ ?

We provide the details of [Son17] in the Chapter 31.
Currently, [PS15] and [CKPS16] are working for $d=1$ case.
Open Problem 14. Can we generalize [PS15] and [CKPS16] to arbitrary dimension d?

### 23.6 Discrete Fourier Transform

### 23.6.1 High dimensional

We introduce some notations for high dimensional Fourier transform. We assume that the universe size $n=p^{d}$ for any positive integer $p$. Let $\omega=e^{2 \pi \mathbf{i} / p}$ where $\mathbf{i}=\sqrt{-1}$. We define the normalized $d$-dimensional Fourier transform

$$
\widehat{x}_{f}=\frac{1}{\sqrt{n}} \sum_{t \in[p]^{d}} x_{t} \cdot \omega^{f^{\top} t}, \forall f \in[p]^{d}
$$

and the inverse Fourier transform is

$$
x_{t}=\frac{1}{\sqrt{n}} \sum_{f \in[p]^{d}} \widehat{x}_{f} \cdot \omega^{-f^{\top} t}, \forall t \in[p]^{d} .
$$

One important goal for high dimensional Fourier transform is to achieve dimensionfree results. We suggest some problems for any $d \geq 1$.

Open Problem 15. For the noiseless setting, [KVZ19] give an algorithm that uses

$$
O\left(k^{3} \log ^{2} \log ^{2} n\right)
$$

samples and $O\left(k^{3} \log ^{2} k \log ^{2} n\right)$ time. This is already dimension-free! If the running time has to be $\operatorname{poly}(k \log n)$, what is the best sample complexity we can achieve?

Open Problem 16. For the noisy setting, [NSW19a] give an algorithm that uses

$$
O(k \log k \log n)
$$

samples and $\widetilde{O}(n)$ running time. Is $O(k \log k \log n)$ the right answer for sample complexity? Open Problem 17. For the noisy setting, can we get $\operatorname{poly}(k \log n)$ samples and $\operatorname{poly}(k \log n)$ running times?

For $d=O(1)$, the state-of-the-art result [Kap16] takes $k \log n \cdot \log \log n$ samples and $k$ poly $(\log n)$ running times.

Open Problem 18. For the noisy setting, for $d=O(1)$, what is the right answer for sample complexity?

### 23.7 Discrete time continuous frequency Fourier transform

There is a long line of work doing Fourier transform in the discrete setting [GMS05, Iwe08, Iwe10, HIKP12a, HIKP12b, Iwe13, IKP14, IK14, Kap16, Kap17, CI17, NSW18, LN19, NSW19a]. We provide the formal definition here. For ease of discussion, we call it discretediscrete setting.

Definition 23.7.1 (discrete time discrete frequency Fourier transform, [HIKP12a]). Given a vector $\widehat{x} \in \mathbb{C}^{n}$, an integer $k$ and an approximation ratio $\alpha>1$. design an algorithm that takes samples from $x \in \mathbb{C}^{n}$, and outputs a $O(k)$-sparse vector $y \in \mathbb{C}^{n}$ such that

$$
\|y-\widehat{x}\|_{2} \leq \alpha \cdot \min _{k-\text { sparse } z}\|z-\widehat{x}\|_{2}
$$

In general, we hope the number of samples and running time are both $k \operatorname{poly}(\log n)$.

To the best of our knowledge, [PS15] is the first work that defines the robust sparse Fourier transform in the continuous setting. For ease of discussion, we call it continuouscontinuous setting.

Definition 23.7.2 (continuous time continuous frequency Fourier transform, [PS15]). Let $x$ be a function $x(t)=x^{*}(t)+g(t)$ where $x^{*}(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$ (i.e. $x^{*}(t)$ is a $k$-Fourier-sparse signal) and $f_{j} \in[-F, F]$. Let $T>0$. The goal is to design an algorithm that allows to take samples from $x(t)$ and outputs a $O(k)$-Fourier-sparse signal $y(t)$ such that

$$
\frac{1}{T} \int_{0}^{T}|y(t)-x(t)|^{2} \mathrm{~d} t \leq \alpha \cdot \frac{1}{T} \int_{0}^{T}|g(t)|^{2} \mathrm{~d} t
$$

In general, we hope the number of samples and running time are $k \operatorname{poly}(\log F T)$.

For a function $x(t)$, we define

$$
\|x(t)\|_{T}=\left(\frac{1}{T} \int_{0}^{T}|x(t)|^{2} \mathrm{~d} t\right)^{1 / 2}
$$

The following definition is suggested by Jerry Li.

Definition 23.7.3 (discrete time continuous frequency Fourier transform). Given $k$ tones $\left(v_{j}, f_{j}\right) \in \mathbb{C} \times[0,1]$, we can access $x(t)=\sum_{j=1}^{k} v_{j} e^{-2 \pi \mathrm{i} f_{j} t}+g(t), \forall t \in[n]$. The goal is to output a $\operatorname{poly}(k \log n)$-sparse function $y(t)$ such that

$$
\sum_{t=1}^{n}(y(t)-x(t))^{2} \leq O(1) \cdot \sum_{t=1}^{n} g(t)^{2}
$$

Some interesting questions are:

Open Problem 19. Can we generalize the continuous Fourier transform techniques in [PS15] to the setting where we are only allowed to take restricted samples (i.e., from integers), if we may assume a frequency gap? Are we still able to recover the continuous frequencies up to some precision, and then reconstruct the signal?

Open Problem 20. Can we generalize the continuous Fourier transform techniques in [CKPS16] to the setting where we are only allowed to take restricted samples and reconstruct the signal with finding the true continuous frequencies?

The discrete time continuous frequency Fourier transform problem naturally arises in the study of learning low-rank Toeplitz covariance matrices. This is a well-studied problem in signal processing, for instance, in relationship to direction-of-arrival estimation. Here, we are given samples $X_{1}, \ldots, X_{n}$ drawn from a $d$-dimensional Gaussian with mean 0 and
covariance $\Sigma$, which we assume to be Toeplitz, and the goal is to approximately recover $\Sigma$. Moreover, we are allowed to access individual coordinates of each $X_{i}$ one-by-one, and only pay a cost of 1 in the sample complexity for every coordinate we access.

Classical algorithms (such as ruler-based methods, or matrix pencil-based methods), either require a number of samples which is polynomial in the ambient dimension, or are not very tolerant to noise. In [ELMM19], it was demonstrated that by solving the discrete time continuous frequency Fourier transform problem, one can obtain an algorithm which is both robust to noise, and which requires a number of samples which depends only polylogarithmically in the ambient dimension, when $\Sigma$ is low rank. The aforementioned paper gives an algorithm for this problem which is sample-efficient, but has exponential running time. By instead using hashing-based techniques, one could hope to obtain a polynomial-time algorithm for this problem, which would in turn yield the first efficient algorithm for learning low-rank Toeplitz covariance matrices using only poly $\log (d)$ samples.

### 23.8 Applications of Fourier transform

### 23.8.1 Speeding up linear programs

It is known that Subsampled Randomized Hadamard Transform gives the subspace embedding.

Theorem 23.8.1 ([AC06, Sar06, Tro11b, DMM06a, DMMS11, DMIMW12, LDFU13]). Let $S=\frac{\sqrt{n}}{\sqrt{r}} P H_{n} D$, where $D$ is an $n \times n$ diagonal matrix with i.i.d. diagonal entries $D_{i, i}$ in which $D_{i, i}=1$ with probability $1 / 2$, and $D_{i, i}=-1$ with probability $1 / 2 . H_{n}$ refers to the Hadamard matrix of size $n$, which we assume is a power of 2 . Here the $(i, j)$-th entry of $H_{n}$ is given by $(-1)^{\langle i, j\rangle} / \sqrt{n}$ where $\langle i, j\rangle=\sum_{z=1}^{\log n} i_{z} \cdot j_{z}$ and $\left(i_{\log n}, \cdots, i_{1}\right)$ and $\left(j_{\log n}, \cdots, j_{1}\right)$ are the binary representations of $i$ and $j$ respectively. Then $r \times n$ matrix $P$ samples $r$ coordinates of an $n$-dimensional vector uniformly at random, where

$$
r=\Omega\left(\epsilon^{-2} \log d \cdot(\sqrt{d}+\sqrt{\log n})^{2}\right)
$$

Then with probability at least .99, for any fixed $n \times d$ matrix $U$ with orthonormal columns,

$$
\left\|I_{d}-U^{\top} S^{\top} S U\right\|_{2} \leq \epsilon
$$

Further, for any vector $x \in \mathbb{R}^{n}$, $S x$ can be computed in $O(n \log r)$ time.

An interesting question is
Open Problem 21. If we don't use diagonal matrix D, can we show subspace embedding guarantee presented in Theorem 23.8.1, or the guarantee presented in Lemma 3.9.1?

If the above conjecture holds, then we have more power to design fast algorithm for several fundamental problems, i.e. linear programs.

### 23.8.2 Proving convergence result for deep neural network

Open Problem 22. Is the frequency gap assumption in continuous Fourier transform ([PS15]) somehow equivalent to data-separable assumption in deep learning theory ([AZLS18, AZLS19, SY19])?

The above problem seems strange in some sense, since Fourier transform and deep learning are completely different problems. The motivation of the stating that problem is : 1) Fourier transform seems to be related to kernel regression with random Fourier feature, 2) a recent work shows that an infinite width one-hidden-layer neural network is somehow equivalent to Kernel regression [ADH $\left.{ }^{+} 19\right]$.

### 23.9 An Over-parameterization Theory of Neural Networks

Over-parameterization theory for deep neural networks becomes extremely popular over the last few years. There is a long line (still growing very quickly) of work proving that (stochastic) gradient descent algorithm is able to find the global minimum if the network is wide enough [LL18, DZPS19, AZLS18, AZLS19, DLL+19, ZCZG18, SY19].

Formally speaking, the existing results show that as long as the width $m$ is at least polynomial of number of input data $n$, then (S)GD-type algorithm can work in the following sense: first randomly pick a weight matrix to be the initialization point, update the weight matrix according to gradient direction over each iteration, and eventually find the global minimum.

Yin Tat Lee [Lee18] conjectured the following

Open Problem 23. If $m$ is nearly in $n$ (ignore the polylog factors), then we can prove convergence result for one-hidden layer neural network.

Next, we want to ask the open questions for multiple layer neural networks. Let $L$ denote the number of layers of the neural network, let $n$ denote the number of input data points, let $d$ denote the dimension of input data points, let $\delta$ denote the failure probability, and let $m$ denote the width of the neural network.

Open Problem 24. For a deep neural network with ReLU activation function, what is the smallest $m$ that suffices to show convergence?

Similarly, we are also curious about the right answer for RNNs,

Open Problem 25. For a recurrent neural network with ReLU activation function, what is the smallest $m$ that suffices to show convergence?
[AZLS18] proved a large polynomial dependence on all the parameters, we believe it is improvable. However, we are not sure about the right answer. Note that all the above three open problems are about training error.

### 23.10 Matrix Chernoff and Discrepancy

### 23.10.1 Matrix Chernoff

Open Problem 26. [KS18] show that $O\left(\epsilon^{-2} \log ^{2} n\right)$ spanning trees suffice to produce a $(1 \pm \epsilon)$-spectral graph sparsifier, and that $\Omega\left(\epsilon^{-2} \log n\right)$ spanning trees are necessary. Which one is tight? Can we prove matching upper and lower bound?

For more background on this problem, we refer the reader to [BSST13, KS18].
We consulted some experts ${ }^{1}$. They are not aware of any form of the following Chernofftype result for hyperbolic polynomials. They believed the following result is interesting.

Open Problem 27. Given vectors $x_{1}, \cdots, x_{n} \in \mathbb{R}^{d}$ and a polynomial $h$ with respect to $e$, let us define the hyperbolic norm $\|x\|_{h}=\max _{i \in[n]}\left|\lambda_{i}(x)\right|$ where $h(t e-x)=h(e) \cdot \prod_{i=1}^{n}\left(t-\lambda_{i}(x)\right)$, and suppose $\left\|x_{i}\right\|_{h} \leq 1$ for all $i \in[n]$. If we sample i.i.d. $\epsilon_{1}, \cdots, \epsilon_{n}$ uniformly from $\{-1,1\}$, then does

$$
\left\|\sum_{i=1}^{n} \epsilon_{i} x_{i}\right\|_{h} \leq O(\sqrt{n \log n})
$$

hold with constant probability?

### 23.10.2 Discrepancy

Open Problem 28 (Matrix Spencer). For any symmetric matrices $A_{1}, \cdots, A_{n} \in \mathbb{R}^{n \times n}$ with $\left\|A_{i}\right\| \leq 1$, do there exist signs $\epsilon_{1}, \cdots, \epsilon_{n} \in\{-1,1\}$ such that

$$
\left\|\sum_{i=1}^{n} \epsilon_{i} A_{i}\right\|=O(\sqrt{n})
$$

[^80]The above problem is called the Matrix Spencer conjecture, see e.g. [Mek04] and [RR19].
[MSS15b] resolved the Kadison-Singer conjecture, a long-standing open problem, but their proof is non-constructive. From the algorithmic perspective, a major open question is the following:

Open Problem 29. Given positive semidefinite matrices $A_{1}, \cdots, A_{m} \in \mathbb{R}^{n \times n}$ and $\epsilon>0$ with $\sum_{i=1}^{m} A_{i}=I_{n}$ and $\operatorname{tr}\left[A_{i}\right] \leq \epsilon$ for all $i \in[m]$, is there a polynomial time algorithm to find signs $x_{1}, \ldots, x_{m} \in\{-1,1\}$ such that $\left\|\sum_{i=1}^{m} x_{i} A_{i}\right\| \leq O(\sqrt{\epsilon})$ ?

In order to describe the next open problems, we provide some definitions:

Definition 23.10.1 ([Brä18]). Given a degree $d$ polynomial $h$ that is hyperbolic with respect to $e \in \mathbb{R}^{n}$, The open hyperbolicity cone of $h$ with respect to $e$ is the set

$$
\Gamma_{++}(h, e)=\left\{x \in \mathbb{R}^{n}: \lambda_{i}(x)>0 \forall i \in[n]\right\} .
$$

We also define the closed hyperbolicity cone

$$
\Gamma_{+}(h, e)=\left\{x \in \mathbb{R}^{n}: \lambda_{i}(x) \geq 0 \forall i \in[n]\right\},
$$

where $h(t e-x)=h(e) \cdot \prod_{i=1}^{n}\left(t-\lambda_{i}(x)\right)$.
The trace, rank and spectral radius (with respect to $e$ ) of $x \in \mathbb{R}^{n}$ are defined as for matrices

$$
\operatorname{tr}[x]=\sum_{i=1}^{n} \lambda_{i}(x), \quad \operatorname{rank}_{h}(x)=\left|\left\{i \in[n]: \lambda_{i}(x) \neq 0\right\}\right|, \quad \text { and } \quad\|x\|_{h}=\max _{i \in[n]} \lambda_{i}(x) .
$$

Recently, [Brä18] generalized Kadison-Singer to hyperbolic polynomials and also higher rank. Note that [MSS15b] can be viewed as the special case when $h$ is the determinant polynomial with rank-1 matrices. [KLS19] generalized [MSS15b] from a different perspective. It would be interesting to combine [KLS19] and [Brä18] to show the following result.

Open Problem 30. Let $h$ be a polynomial that is hyperbolic with respect to $e \in \mathbb{R}^{n}$. Let $v_{1}, \cdots, v_{m} \in \Gamma_{+}(h, e)$ with $\operatorname{rank}_{h}\left(v_{i}\right) \leq 1$, and

$$
\sigma^{2}=\left\|\sum_{i=1}^{m} \operatorname{tr}\left[v_{i}\right] v_{i}\right\|_{h}
$$

Then there exists a choice of signs $\epsilon_{i} \in\{-1,1\}$ and some constant $c>0$ such that

$$
\left\|\sum_{i=1}^{m} \epsilon_{i} v_{i}\right\|_{h} \leq c \cdot \sigma
$$

We also wonder if similar techniques can be used to sparsify hyperbolicity cones:

Open Problem 31. Given $\epsilon>0$, a polynomial $h$ that is hyperbolic with respect to $e \in \mathbb{R}^{n}$ and $v_{1}, \cdots, v_{m} \in \Gamma_{+}(h, e)$ with $v_{1}+\cdots+v_{m}=e$, is it always possible to find coefficients $s \in \mathbb{R}_{\geq 0}^{m}$ such that $|\operatorname{supp}(s)| \leq n / \epsilon^{2}$ and $\left\|e-\sum_{i=1}^{m} s_{i} v_{i}\right\|_{h}=O(\epsilon)$ ?

The above problem can be viewed as a generalization of [BSS12].

### 23.11 Edit Distance and Longest Common Subsequence

In this section we consider two of the most ubiquitous measures of similarity between a pair of strings: the longest common subsequence (LCS) and the edit distance. The LCS of two $n$-character strings $A$ and $B$ is simply their longest (not necessarily contiguous) common substring. Edit distance is the minimum number of character insertions, deletions, and substitutions required to transform $A$ and $B$. In fact, under a slightly more restricted definition that does not allow substitutions, the two measures are complements to each other and the problems of computing them are exactly equivalent.

Under plausible fine-grained complexity assumptions such as SETH, neither edit distance nor LCS can be computed much faster [AWW14, ABW15, BI15, BK15, AHWW16].

For edit distance, an $\widetilde{O}\left(n+\Delta^{2}\right)$-time algorithm of [LMS98] (where $\Delta$ is the true edit distance between the strings) implies a linear-time $\sqrt{n}$-approximation algorithm. The approximation factor has been significantly improved in a series of works to $O\left(n^{3 / 7}\right)$ [BYJKK04], to $O\left(n^{0.34}\right)$ [BES06], to $O\left(2^{\widetilde{O}(\sqrt{\log n})}\right.$ [AO09], and finally to polylogarithmic [AKO10]. A recent work of Boroujeni, Ehsani, Ghodsi, HajiAghayi and Seddighin [BEG+18] obtains a constant factor approximation quantum algorithm for edit distance that runs in truly subquadratic time.

We state three open problems related to edit distance. Two of the problems were originally posted by Aviad Rubinstein [Rub18].

Open Problem 32. Is there a constant-factor approximation to edit distance that runs in near-linear time?

Currently, the best approximation ratio is $(3+\epsilon)$ [RSSS19].

Open Problem 33. Is there a $(1+\epsilon)$-approximation algorithm for edit distance in truly subquadratic (or near-linear) time?

Note that there is a more concrete plan to make progress on approximation ratio of edit distance problem. For more details, we refer the readers to [RSSS19].

Note that the current subquadratic time constant approximation edit distance algorithm $\left[\mathrm{CDG}^{+} 18\right]$ is a randomized algorithm. It would be interesting to derandomize the algorithm.

Open Problem 34. Is there a deterministic subquadratic time constant approximation edit distance algorithm?

Now let us come to the LCS problem. The textbook dynamic programming algorithm for computing LCS (or edit distance) runs in $O\left(n^{2}\right)$ time, and a slightly faster $O\left(n^{2} / \log ^{2}(n)\right)$ time algorithm is known due to Masek and Paterson [MP80]. There is a trivial algorithm that gives 2-approximation for binary strings. Finding faster algorithms is a central and long standing open problem both in theory and in practice (e.g. Problem 35 of [Knu72]).

Rubinstein and Song [RS19b] proposed a truly subquadratic time ( $2-\epsilon$ )-approximation algorithm for binary strings. Let $|\Sigma|$ be the alphabetical size. Then one natural question is if we can generalize the algorithm for $|\Sigma|=3,4$ or some constants.

Open Problem 35. For any constant $|\Sigma|$, is there a truly subquadratic time algorithm that gives $(|\Sigma|-\epsilon)$-approximation for some constants $\epsilon>0$ ?

Another interesting question is about de-randomization.
Open Problem 36. Can we de-randomize the LCS algorithm in [RSSS19]?

### 23.12 Map-Reduce

Many modern parallel systems, such as MapReduce, Hadoop, and Spark, can be modeled well by the MPC model. The MPC model captures well coarse-grained computation on large data - data is distributed to processors, each of which has a sublinear (in the input data) amount of memory and we alternate between rounds of computation and rounds of communication, where each machine can communicate an amount of data as large as the size of its memory.

We consider one fundamental graph problem: connectivity. On an undirected graph with $n$ and $m$ edges, $O(\log n)$ round connectivity algorithm has been known for over 35 years. In 2018, $\left[\mathrm{ASS}^{+} 18\right]$ proposed a new algorithm that only needs $O\left(\log D \cdot \log \log _{m / n} n\right)$ where $D$ is the diameter of the graph. Very recently, Behnezhad, Dhulipala, Esfandiari, Lacki and Mirrokni $\left[\mathrm{BDE}^{+} 19\right]$ improved the number of rounds to $O\left(\log D+\log \log _{m / n} n\right)$. An interesting open question is

Open Problem 37. Is there an $O(\log D)$ rounds algorithm for graph connectivity?

The above problem was originally suggested by Peilin Zhong. In next a few paragraphs, we discuss several other open problems which are proposed by Xiaorui Sun.

Roughgarden, Vassilvitskii and Wang [RVW18] showed that if there is an $\omega(1)$-round MPC lower bound for graph connectivity, then $N C_{1} \neq P$. Since $N C_{1}$ vs $P$ is a well-known hard problem, it is unlikely to prove an $\omega(1)$-round unconditional lower bound for connectivity. Hence, we ask if an conditional lower bound is possible.

Open Problem 38. Is there an $\omega(1)$-round conditional lower bound for graph connectivity?

Furthermore, people actually conjectured that any algorithm solve graph connectivity problem requires $\Omega(\log n)$ rounds. This conjecture has been use to show other graph problems are unlikely to have constant round algorithms (Yaroslavtsev and Vadapalli [YV18], Ghaffari, Kuhn and Uitto [GKU19]).

Open Problem 39. Is there an $\Omega(\log n)$-round conditional lower bound for graph connectivity?

In a more broader setting, it will be interesting to understand the complexity of MPC. In particular, if all the polynomial-time solvable problem have an efficient MPC algorithm. Since linear programming is P -Complete, we ask

Open Problem 40. Is there an $O\left(n^{o(1)}\right)$ rounds algorithm for linear programming?

Dynamic programming in MPC has been studied recently. Im, Moseley and Sun [IMS17a] and Hajiaghayi, Seddighin and Sun [HSS19] showed that many classic dynamic programming problems have constant rounds approximation algorithms. These problems include weighted interval scheduling, longest increasing subsequence, edit distance and longest common subsequence. But it is unclear if these problems can be solved exactly in constant rounds.

Open Problem 41. Can dynamic programming problems be solved exactly in $O(1)$ rounds?

## Part VI

## More Algorithms for Fundamental Problems

## Chapter 24

## $\ell_{\infty}$ Regression

Sketching has emerged as a powerful technique for speeding up problems in numerical linear algebra, such as regression. In the overconstrained regression problem, one is given an $n \times d$ matrix $A$, with $n \gg d$, as well as an $n \times 1$ vector $b$, and one wants to find a vector $\widehat{x}$ so as to minimize the residual error $\|A x-b\|_{2}$. Using the sketch and solve paradigm, one first computes $S \cdot A$ and $S \cdot b$ for a randomly chosen matrix $S$, then outputs $x^{\prime}=(S A)^{\dagger} S b$ so as to minimize $\left\|S A x^{\prime}-S b\right\|_{2}$.

The sketch-and-solve paradigm gives a bound on $\left\|x^{\prime}-x^{*}\right\|_{2}$ when $A$ is well-conditioned. Our main result is that, when $S$ is the subsampled randomized Fourier/Hadamard transform, the error $x^{\prime}-x^{*}$ behaves as if it lies in a "random" direction within this bound: for any fixed direction $a \in \mathbb{R}^{d}$, we have with $1-d^{-c}$ probability that

$$
\begin{equation*}
\left\langle a, x^{\prime}-x^{*}\right\rangle \lesssim \frac{\|a\|_{2}\left\|x^{\prime}-x^{*}\right\|_{2}}{d^{\frac{1}{2}-\gamma}} \tag{24.1}
\end{equation*}
$$

where $c, \gamma>0$ are arbitrary constants. This implies $\left\|x^{\prime}-x^{*}\right\|_{\infty}$ is a factor $d^{\frac{1}{2}-\gamma}$ smaller than $\left\|x^{\prime}-x^{*}\right\|_{2}$. It also gives a better bound on the generalization of $x^{\prime}$ to new examples: if rows of $A$ correspond to examples and columns to features, then our result gives a better bound for the error introduced by sketch-and-solve when classifying fresh examples. We show that not all oblivious subspace embeddings $S$ satisfy these properties. In particular, we give
counterexamples showing that matrices based on Count-Sketch or leverage score sampling do not satisfy these properties.

We also provide lower bounds, both on how small $\left\|x^{\prime}-x^{*}\right\|_{2}$ can be, and for our new guarantee (24.1), showing that the subsampled randomized Fourier/Hadamard transform is nearly optimal. Our lower bound on $\left\|x^{\prime}-x^{*}\right\|_{2}$ shows that there is an $O(1 / \epsilon)$ separation in the dimension of the optimal oblivious subspace embedding required for outputting an $x^{\prime}$ for which $\left\|x^{\prime}-x^{*}\right\|_{2} \leq \epsilon\left\|A x^{*}-b\right\|_{2} \cdot\left\|A^{\dagger}\right\|_{2}$, compared to the dimension of the optimal oblivious subspace embedding required for outputting an $x^{\prime}$ for which $\left\|A x^{\prime}-b\right\|_{2} \leq(1+\epsilon)\left\|A x^{*}-b\right\|_{2}$, that is, the former problem requires dimension $\Omega\left(d / \epsilon^{2}\right)$ while the latter problem can be solved with dimension $O(d / \epsilon)$. This explains the reason known upper bounds on the dimensions of these two variants of regression have differed in prior work.

This part is based upon the following previous publication

- Eric Price, Zhao Song, David P. Woodruff

Fast Regression with an $\ell_{\infty}$ Guarantee.
ICALP 2017 [PSW17]

### 24.1 Introduction

Oblivious subspace embeddings (OSEs) were introduced by Sarlos [Sar06] to solve linear algebra problems more quickly than traditional methods. An OSE is a distribution of matrices $S \in \mathbb{R}^{m \times n}$ with $m \ll n$ such that, for any $d$-dimensional subspace $U \subset \mathbb{R}^{n}$, with "high" probability $S$ preserves the norm of every vector in the subspace. OSEs are a generalization of the classic Johnson-Lindenstrauss lemma from vectors to subspaces. Formally,
we require that with probability $1-\delta$,

$$
\|S x\|_{2}=(1 \pm \epsilon)\|x\|_{2}
$$

simultaneously for all $x \in U$, that is, $(1-\epsilon)\|x\|_{2} \leq\|S x\|_{2} \leq(1+\epsilon)\|x\|_{2}$.
A major application of OSEs is to regression. The regression problem is, given $b \in \mathbb{R}^{n}$ and $A \in \mathbb{R}^{n \times d}$ for $n \geq d$, to solve for

$$
\begin{equation*}
x^{*}=\underset{x \in \mathbb{R}^{d}}{\operatorname{argmin}}\|A x-b\|_{2} \tag{24.2}
\end{equation*}
$$

Because $A$ is a "tall" matrix with more rows than columns, the system is overdetermined and there is likely no solution to $A x=b$, but regression will find the closest point to $b$ in the space spanned by $A$. The classic answer to regression is to use the Moore-Penrose pseudoinverse: $x^{*}=A^{\dagger} b$ where

$$
A^{\dagger}=\left(A^{\top} A\right)^{-1} A^{\top}
$$

is the "pseudoinverse" of $A$ (assuming $A$ has full column rank, which we will typically do for simplicity). This classic solution takes $O\left(n d^{\omega-1}+d^{\omega}\right)$ time, where $\omega<2.373$ is the matrix multiplication constant [CW90, Wil12, LG14]: $n d^{\omega-1}$ time to compute $A^{\top} A$ and $d^{\omega}$ time to compute the inverse.

OSEs speed up the process by replacing (24.2) with

$$
x^{\prime}=\underset{x \in \mathbb{R}^{d}}{\operatorname{argmin}}\|S A x-S b\|_{2}
$$

for an OSE $S$ on $d+1$-dimensional spaces. This replaces the $n \times d$ regression problem with an $m \times d$ problem, which can be solved more quickly since $m \ll n$. Because $A x-b$ lies in
the $d+1$-dimensional space spanned by $b$ and the columns of $A$, with high probability $S$ preserves the norm of $S A x-S b$ to $1 \pm \epsilon$ for all $x$. Thus,

$$
\left\|A x^{\prime}-b\right\|_{2} \leq \frac{1+\epsilon}{1-\epsilon}\left\|A x^{*}-b\right\|_{2}
$$

That is, $S$ produces a solution $x^{\prime}$ which preserves the cost of the regression problem. The running time for this method depends on (1) the reduced dimension $m$ and (2) the time it takes to multiply $S$ by $A$. We can compute these for "standard" OSE types:

- If $S$ has i.i.d. Gaussian entries, then $m=O\left(d / \epsilon^{2}\right)$ is sufficient (and in fact, $m \geq d / \epsilon^{2}$ is required [NN14]). However, computing $S A$ takes $O(m n d)=O\left(n d^{2} / \epsilon^{2}\right)$ time, which is worse than solving the original regression problem (one can speed this up using fast matrix multiplication, though it is still worse than solving the original problem).
- If $S$ is a subsampled randomized Hadamard transform (SRHT) matrix with random sign flips (see Theorem 2.4 in [Woo14a] for a survey, and also see [AC06, Sar06, DMM06c, DMMS11, KW11, Tro11b, DMIMW12, LDFU13, AL13, Bou14, HR16, CNW16]), then $m$ increases to $\widetilde{O}\left(d / \epsilon^{2} \cdot \log n\right)$, where $\widetilde{O}(f)=f \operatorname{poly}(\log (f))$. But now, we can compute $S A$ using the fast Hadamard transform in $O(n d \log n)$ time. This makes the overall regression problem take $O\left(n d \log n+d^{\omega} / \epsilon^{2}\right)$ time.
- If $S$ is a random sparse matrix with random signs (the "Count-Sketch" matrix), then $m=d^{1+\gamma} / \epsilon^{2}$ suffices for $\gamma>0$ a decreasing function of the sparsity [CW13, MM13, NN13a, BDN15, Coh16a]. (The definition of a Count-Sketch matrix is, for any $s \geq 1$, $S_{i, j} \in\{0,-1 / \sqrt{s}, 1 / \sqrt{s}\}, \forall i \in[m], j \in[n]$ and the column sparsity of matrix $S$ is $s$. Independently in each column $s$ positions are chosen uniformly at random without
replacement, and each chosen position is set to $-1 / \sqrt{s}$ with probability $1 / 2$, and $+1 / \sqrt{s}$ with probability $1 / 2$.) Sparse OSEs can benefit from the sparsity of $A$, allowing for a running time of $\widetilde{O}(\mathrm{nnz}(A))+\widetilde{O}\left(d^{\omega} / \epsilon^{2}\right)$, where $n n z(A)$ denotes the number of non-zeros in $A$.

When $n$ is large, the latter two algorithms are substantially faster than the naïve $n d^{\omega-1}$ method.

### 24.1.1 Our Contributions

Despite the success of using subspace embeddings to speed up regression, often what practitioners are interested is not in preserving the cost of the regression problem, but rather in the generalization or prediction error provided by the vector $x^{\prime}$. Ideally, we would like for any future (unseen) example $a \in \mathbb{R}^{d}$, that $\left\langle a, x^{\prime}\right\rangle \approx\left\langle a, x^{*}\right\rangle$ with high probability.

Ultimately one may want to use $x^{\prime}$ to do classification, such as regularized least squares classification (RLSC) [RYP03], which has been found in cases to do as well as support vector machines but is much simpler [ZP04]. In this application, given a training set of examples with multiple (non-binary) labels identified with the rows of an $n \times d$ matrix $A$, one creates an $n \times r$ matrix $B$, each column indicating the presence or absence of one of the $r$ possible labels in each example. One then solves the multiple response regression problem $\min _{X}\|A X-B\|_{F}$, and uses $X$ to classify future examples. A commonly used method is for a future example $a$, to compute $\left\langle a, x_{1}\right\rangle, \ldots,\left\langle a, x_{r}\right\rangle$, where $x_{1}, \ldots, x_{r}$ are the columns of $X$. One then chooses the label $i$ for which $\left\langle a, x_{i}\right\rangle$ is maximum.

For this to work, we would like the inner products $\left\langle a, x_{1}^{\prime}\right\rangle, \ldots,\left\langle a, x_{r}^{\prime}\right\rangle$ to be close to
$\left\langle a, x_{1}^{*}\right\rangle, \ldots,\left\langle a, x_{r}^{*}\right\rangle$, where $X^{\prime}$ is the solution to $\min _{X}\|S A X-S B\|_{F}$ and $X^{*}$ is the solution to $\min _{X}\|A X-B\|_{F}$. For any $O(1)$-accurate OSE on $d+r$ dimensional spaces [Sar06], which also satisfies so-called approximate matrix multiplication with error $\left.\epsilon^{\prime}=\epsilon / \sqrt{( } d+r\right)$, we get that

$$
\begin{equation*}
\left\|x^{\prime}-x^{*}\right\|_{2} \leq O(\epsilon) \cdot\left\|A x^{*}-b\right\|_{2} \cdot\left\|A^{\dagger}\right\|_{2} \tag{24.3}
\end{equation*}
$$

where $\left\|A^{\dagger}\right\|$ is the spectral norm of $A^{\dagger}$, which equals the reciprocal of the smallest singular value of $A$. To obtain a generalization error bound for an unseen example $a$, one has

$$
\begin{equation*}
\left|\left\langle a, x^{*}\right\rangle-\left\langle a, x^{\prime}\right\rangle\right|=\left|\left\langle a, x^{*}-x^{\prime}\right\rangle\right| \leq\left\|x^{*}-x^{\prime}\right\|_{2}\|a\|_{2}=O(\epsilon)\|a\|_{2}\left\|A x^{*}-b\right\|_{2}\left\|A^{\dagger}\right\|_{2}, \tag{24.4}
\end{equation*}
$$

which could be tight if given only the guarantee in (24.3). However, if the difference vector $x^{\prime}-x^{*}$ were distributed in a uniformly random direction subject to (24.3), then one would expect an $\widetilde{O}(\sqrt{d})$ factor improvement in the bound. This is what our main theorem shows:

Theorem 24.1.1 (Main Theorem, informal). Suppose $n \leq \operatorname{poly}(d)$ and matrix $A \in \mathbb{R}^{n \times d}$ and vector $b \in \mathbb{R}^{n}$ are given. Let $S \in \mathbb{R}^{m \times n}$ be a subsampled randomized Hadamard transform matrix with $m=d^{1+\gamma} / \epsilon^{2}$ rows for an arbitrarily small constant $\gamma>0$. For $x^{\prime}=\operatorname{argmin}_{x \in \mathbb{R}^{d}}\|S A x-S b\|_{2}$ and $x^{*}=\operatorname{argmin}_{x \in \mathbb{R}^{d}}\|A x-b\|_{2}$, and any fixed $a \in \mathbb{R}^{d}$,

$$
\begin{equation*}
\left|\left\langle a, x^{*}\right\rangle-\left\langle a, x^{\prime}\right\rangle\right| \leq \frac{\epsilon}{\sqrt{d}}\|a\|_{2}\left\|A x^{*}-b\right\|_{2}\left\|A^{\dagger}\right\|_{2} \tag{24.5}
\end{equation*}
$$

with probability $1-1 / d^{C}$ for an arbitrarily large constant $C>0$. This implies that

$$
\begin{equation*}
\left\|x^{*}-x^{\prime}\right\|_{\infty} \leq \frac{\epsilon}{\sqrt{d}}\left\|A x^{*}-b\right\|_{2}\left\|A^{\dagger}\right\|_{2} \tag{24.6}
\end{equation*}
$$

with $1-1 / d^{C-1}$ probability.
If $n>\operatorname{poly}(d)$, then by first composing $S$ with a Count-Sketch OSE with poly $(d)$ rows, one can achieve the same guarantee.
(Here $\gamma$ is a constant going to zero as $n$ increases; see Theorem 24.3.3 for a formal statement of Theorem 24.1.1.)

Notice that Theorem 24.1.1 is considerably stronger than that of (24.4) provided by existing guarantees. Indeed, in order to achieve the guarantee (24.6) in Theorem 24.1.1, one would need to set $\epsilon^{\prime}=\epsilon / \sqrt{d}$ in existing OSEs, resulting in $\Omega\left(d^{2} / \epsilon^{2}\right)$ rows. In contrast, we achieve only $d^{1+\gamma} / \epsilon^{2}$ rows. We can improve the bound in Theorem 24.1.1 to $m=d / \epsilon^{2}$ if $S$ is a matrix of i.i.d. Gaussians; however, as noted, computing $S \cdot A$ is slower in this case.

Note that Theorem 24.1.1 also makes no distributional assumptions on the data, and thus the data could be heavy-tailed or even adversarially corrupted. This implies that our bound is still useful when the rows of $A$ are not sampled independently from a distribution with bounded variance.

The $\ell_{\infty}$ bound (24.6) of Theorem 24.1.1 is achieved by applying (24.5) to the standard basis vectors $a=e_{i}$ for each $i \in[d]$ and applying a union bound. This $\ell_{\infty}$ guarantee often has a more natural interpretation than the $\ell_{2}$ guarantee - if we think of the regression as attributing the observable as a sum of various factors, (24.6) says that the contribution of each factor is estimated well. One may also see our contribution as giving a way for estimating the pseudoinverse $A^{\dagger}$ entrywise. Namely, we get that $(S A)^{\dagger} S \approx A^{\dagger}$ in the sense that each entry is within additive $O\left(\epsilon \sqrt{\frac{\log d}{d}}\left\|A^{\dagger}\right\|_{2}\right)$. There is a lot of work on computing entries of inverses of a matrix, see, e.g., [ADL ${ }^{+} 12$, LAKD08].

Another benefit of the $\ell_{\infty}$ guarantee is when the regression vector $x^{*}$ is expected to be $k$-sparse (e.g. [Lee12]). In such cases, thresholding to the top $k$ entries will yield an $\ell_{2}$ guarantee a factor $\sqrt{\frac{k}{d}}$ better than (24.3).

One could ask if Theorem 24.1.1 also holds for sparse OSEs, such as the Count-Sketch. Surprisingly, we show that one cannot achieve the generalization error guarantee in Theorem 24.1.1 with high probability, say, $1-1 / d$, using such embeddings, despite the fact that such embeddings do approximate the cost of the regression problem up to a $1+\epsilon$ factor with high probability. This shows that the generalization error guarantee is achieved by some subspace embeddings but not all.

Theorem 24.1.2 (Not all subspace embeddings give the $\ell_{\infty}$ guarantee; informal version of Theorem 24.8.1). The Count-Sketch matrix with $d^{1.5}$ rows and sparsity $d^{25}$-which is an OSE with exponentially small failure probability-with constant probability will have a result $x^{\prime}$ that does not satisfy the $\ell_{\infty}$ guarantee (24.6).

We can show that Theorem 24.1.1 holds for $S$ based on the Count-Sketch OSE $T$ with $d^{O(C)} / \epsilon^{2}$ rows with $1-1 / d^{C}$ probability. We can thus compose the Count-Sketch OSE with the SRHT matrix and obtain an $O(\mathrm{nnz}(A))+\operatorname{poly}(d / \epsilon)$ time algorithm to compute $S \cdot T A$ achieving (24.6). We can also compute $R \cdot S \cdot T \cdot A$, where $R$ is a matrix of Gaussians, which is more efficient now that $S T A$ only has $d^{1+\gamma} / \epsilon^{2}$ rows; this will reduce the number of rows to $d / \epsilon^{2}$.

Another common method of dimensionality reduction for linear regression is leverage score sampling [DMIMW12, LMP13, PKB14, CMM15], which subsamples the rows of $A$ by choosing each row with probability proportional to its "leverage scores". With $O\left(d \log (d / \delta) / \epsilon^{2}\right)$ rows taken, the result $x^{\prime}$ will satisfy the $\ell_{2}$ bound (24.3) with probability $1-\delta$. However, it does not give a good $\ell_{\infty}$ bound:

Theorem 24.1.3 (Leverage score sampling does not give the $\ell_{\infty}$ guarantee; informal version
of Theorem 24.9.1). Leverage score sampling with $d^{1.5}$ rows-which satisfies the $\ell_{2}$ bound with exponentially small failure probability - with constant probability will have a result $x^{\prime}$ that does not satisfy the $\ell_{\infty}$ guarantee (24.6).

Finally, we show that the $d^{1+\gamma} / \epsilon^{2}$ rows that SRHT matrices use is roughly optimal:
Theorem 24.1.4 (Lower bounds for $\ell_{2}$ and $\ell_{\infty}$ guarantees; informal versions of of Theorem 24.5.1 and Corollary 24.5.3). Any sketching matrix distribution over $m \times n$ matrices that satisfies either the $\ell_{2}$ guarantee (24.3) or the $\ell_{\infty}$ guarantee (24.6) must have $m \gtrsim \min \left(n, d / \epsilon^{2}\right)$.

Notice that our result shows the necessity of the $1 / \epsilon$ separation between the results originally defined in Equation (3) and (4) of Theorem 12 of [Sar06]. If we want to output some vector $x^{\prime}$ such that $\left\|A x^{\prime}-b\right\|_{2} \leq(1+\epsilon)\left\|A x^{*}-b\right\|_{2}$, then it is known that $m=\Theta(d / \epsilon)$ is necessary and sufficient. However, if we want to output a vector $x^{\prime}$ such that $\left\|x^{\prime}-x^{*}\right\|_{2} \leq$ $\epsilon\left\|A x^{*}-b\right\|_{2} \cdot\left\|A^{\dagger}\right\|_{2}$, then we show that $m=\Theta\left(d / \epsilon^{2}\right)$ is necessary and sufficient.

Comparison to Gradient Descent While this work is primarily about sketching methods, one could instead apply iterative methods such as gradient descent, after appropriately preconditioning the matrix, see, e.g., [AMT10, ZF13, CW13]. That is, one can use an OSE with constant $\epsilon$ to construct a preconditioner for $A$ and then run conjugate gradient using the preconditioner. This gives an overall dependence of $\log (1 / \epsilon)$.

The main drawback of this approach is that one loses the ability to save on storage space or number of passes when $A$ appears in a stream, or to save on communication or rounds when $A$ is distributed. Given increasingly large data sets, such scenarios are now quite common, see, e.g., [CW90] for regression algorithms in the data stream model. In
situations where the entries of $A$ appear sequentially, for example, a row at a time, one does not need to store the full $n \times d$ matrix $A$ but only the $m \times d$ matrix $S A$.

Also, iterative methods can be less efficient when solving multiple response regression, where one wants to minimize $\|A X-B\|$ for a $d \times t$ matrix $X$ and an $n \times t$ matrix $B$. This is the case when $\epsilon$ is constant and $t$ is large, which can occur in some applications (though there are also other applications for which $\epsilon$ is very small). For example, conjugate gradient with a preconditioner will take $\widetilde{O}(n d t)$ time while using an OSE directly will take only $\widetilde{O}\left(n d+d^{2} t\right)$ time (since one effectively replaces $n$ with $O(d)$ after computing $S \cdot A$ ), separating $t$ from d. Multiple response regression, arises, for example, in the RLSC application above.

Proof Techniques Theorem 24.1.1. As noted in Theorem 24.1.2, there are some OSEs for which our generalization error bound does not hold. This hints that our analysis is nonstandard and cannot use generic properties of OSEs as a black box. Indeed, in our analysis, we have to consider matrix products of the form $S^{\top} S\left(U U^{\top} S^{\top} S\right)^{k}$ for our random sketching matrix $S$ and a fixed matrix $U$, where $k$ is a positive integer. We stress that it is the same matrix $S$ appearing multiple times in this expression, which considerably complicates the analysis, and does not allow us to appeal to standard results on approximate matrix product (see, e.g., [Woo14a] for a survey). The key idea is to recursively reduce $S^{\top} S\left(U U^{\top} S^{\top} S\right)^{k}$ using a property of $S$. We use properties that only hold for specifics OSEs $S$ : first, that each column of $S$ is unit vector; and second, that for all pairs $(i, j)$ and $i \neq j$, the inner product between $S_{i}$ and $S_{j}$ is at most $\frac{\sqrt{\log n}}{\sqrt{m}}$ with probability $1-1 / \operatorname{poly}(n)$.

Theorems 24.8.1 and 24.9.1. To show that Count-Sketch does not give the $\ell_{\infty}$ guar-
antee, we construct a matrix $A$ and vector $b$ as in Figure 24.1, which has optimal solution $x^{*}$ with all coordinates $1 / \sqrt{d}$. We then show, for our setting of parameters, that there likely exists an index $j \in[d]$ satisfying the following property: the $j$ th column of $S$ has disjoint support from the $k$ th column of $S$ for all $k \in[d+\alpha] \backslash\{j\}$ except for a single $k>d$, for which $S_{j}$ and $S_{k}$ share exactly one common entry in their support. In such cases we can compute $x_{j}^{\prime}$ explicitly, getting $\left|x_{j}^{\prime}-x_{j}^{*}\right|=\frac{1}{s \sqrt{\alpha}}$. By choosing suitable parameters in our construction, this gives that $\left\|x^{\prime}-x^{*}\right\|_{\infty} \gg \frac{1}{\sqrt{d}}$. The lower bound for leverage score sampling follows a similar construction.

Theorem 24.5.1 and Corollary 24.5.3. The lower bound proof for the $\ell_{2}$ guarantee uses Yao's minimax principle. We are allowed to fix an $m \times n$ sketching matrix $S$ and design a distribution over $\left[\begin{array}{ll}A & b\end{array}\right]$. We first write the sketching matrix $S=U \Sigma V^{\top}$ in its singular value decomposition (SVD). We choose the $d+1$ columns of the adjoined matrix $[A, b]$ to be random orthonormal vectors. Consider an $n \times n$ orthonormal matrix $R$ which contains the columns of $V$ as its first $m$ columns, and is completed on its remaining $n-m$ columns to an arbitrary orthonormal basis. Then $S \cdot[A, b]=V^{\top} R R^{\top} \cdot[A, b]=\left[U \Sigma I_{m}, 0\right] \cdot\left[R^{\top} A, R^{\top} b\right]$. Notice that $\left[R^{\top} A, R^{\top} b\right]$ is equal in distribution to $[A, b]$, since $R$ is fixed and $[A, b]$ is a random matrix with $d+1$ orthonormal columns. Therefore, $S \cdot[A, b]$ is equal in distribution to [UइG,Uइh] where [ $G, h$ ] corresponds to the first $m$ rows of an $n \times(d+1)$ uniformly random matrix with orthonormal columns.

A key idea is that if $n=\Omega\left(\max (m, d)^{2}\right)$, then by a result of Jiang [Jia06], any $m \times(d+1)$ submatrix of a random $n \times n$ orthonormal matrix has $o(1)$ total variation distance to a $d \times d$ matrix of i.i.d. $N(0,1 / n)$ random variables, and so any events that
would have occurred had $G$ and $h$ been independent i.i.d. Gaussians, occur with the same probability for our distribution up to an $1-o(1)$ factor, so we can assume $G$ and $h$ are independent i.i.d. Gaussians in the analysis.

The optimal solution $x^{\prime}$ in the sketch space equals $(S A)^{\dagger} S b$, and by using that $S A$ has the form $U \Sigma G$, one can manipulate $\left\|(S A)^{\dagger} S b\right\|$ to be of the form $\left\|\widetilde{\Sigma}^{\dagger}(\Sigma R)^{\dagger} \Sigma h\right\|_{2}$, where the SVD of $G$ is $R \widetilde{\Sigma} T$. We can upper bound $\|\widetilde{\Sigma}\|_{2}$ by $\sqrt{r / n}$, since it is just the maximum singular value of a Gaussian matrix, where $r$ is the rank of $S$, which allows us to lower bound $\left\|\widetilde{\Sigma}^{\dagger}(\Sigma R)^{\dagger} \Sigma h\right\|_{2}$ by $\sqrt{n / r}\left\|(\Sigma R)^{\dagger} \Sigma h\right\|_{2}$. Then, since $h$ is i.i.d. Gaussian, this quantity concentrates to $\frac{1}{\sqrt{r}}\left\|(\Sigma R)^{\dagger} \Sigma h\right\|$, since $\|C h\|^{2} \approx\|C\|_{F}^{2} / n$ for a vector $h$ of i.i.d. $N(0,1 / n)$ random variables. Finally, we can lower bound $\left\|(\Sigma R)^{\dagger} \Sigma\right\|_{F}^{2}$ by $\left\|(\Sigma R)^{\dagger} \Sigma R R^{\top}\right\|_{F}^{2}$ by the Pythagorean theorem, and now we have that $(\Sigma R)^{\dagger} \Sigma R$ is the identity, and so this expression is just equal to the rank of $\Sigma R$, which we prove is at least $d$. Noting that $x^{*}=0$ for our instance, putting these bounds together gives $\left\|x^{\prime}-x^{*}\right\| \geq \sqrt{d / r}$. The last ingredient is a way to ensure that the rank of $S$ is at least $d$. Here we choose another distribution on inputs $A$ and $b$ for which it is trivial to show the rank of $S$ is at least $d$ with large probability. We require $S$ be good on the mixture. Since $S$ is fixed and good on the mixture, it is good for both distributions individually, which implies we can assume $S$ has rank $d$ in our analysis of the first distribution above.

### 24.1.2 Notation

For a positive integer, let $[n]=\{1,2, \ldots, n\}$. For a vector $x \in \mathbb{R}^{n}$, define $\|x\|_{2}=$ $\left(\sum_{i=1}^{n} x_{i}^{2}\right)^{\frac{1}{2}}$ and $\|x\|_{\infty}=\max _{i \in[n]}\left|x_{i}\right|$. For a matrix $A \in \mathbb{R}^{m \times n}$, define $\|A\|_{2}=\sup _{x}\|A x\|_{2} /\|x\|_{2}$ to be the spectral norm of $A$ and $\|A\|_{F}=\left(\sum_{i, j} A_{i, j}^{2}\right)^{1 / 2}$ to be the Frobenius norm of $A$. We
use $A^{\dagger}$ to denote the Moore-Penrose pseudoinverse of $m \times n$ matrix $A$, which if $A=U \Sigma V^{\top}$ is its SVD (where $U \in \mathbb{R}^{m \times n}, \Sigma \in \mathbb{R}^{n \times n}$ and $V \in \mathbb{R}^{n \times}$ for $m \geq n$ ), is given by $A^{\dagger}=V \Sigma^{-1} U^{\top}$.

In addition to $O(\cdot)$ notation, for two functions $f, g$, we use the shorthand $f \lesssim g$ (resp. $\gtrsim)$ to indicate that $f \leq C g$ (resp. $\geq$ ) for an absolute constant $C$. We use $f \approx g$ to mean $c f \leq g \leq C f$ for constants $c, C$.

Definition 24.1.1 (Subspace Embedding). A ( $1 \pm \epsilon$ ) $\ell_{2}$-subspace embedding for the column space of an $n \times d$ matrix $A$ is a matrix $S$ for which for all $x \in \mathbb{R}^{d},\|S A x\|_{2}^{2}=(1 \pm \epsilon)\|A x\|_{2}^{2}$.

Definition 24.1.2 (Approximate Matrix Product). Let $0<\epsilon<1$ be a given approximation parameter. Given matrices $A$ and $B$, where $A$ and $B$ each have $n$ rows, the goal is to output a matrix $C$ so that $\left\|A^{\top} B-C\right\|_{F} \leq \epsilon\|A\|_{F}\|B\|_{F}$. Typically $C$ has the form $A^{\top} S^{\top} S B$, for a random matrix $S$ with a small number of rows. In particular, this guarantee holds for the subsampled randomized Hadamard transform $S$ with $O\left(\epsilon^{-2}\right)$ rows [DMMS11].

### 24.2 Warmup: Gaussians OSEs

We first show that if $S$ is a Gaussian random matrix, then it satisfies the generalization guarantee. This follows from the rotational invariance of the Gaussian distribution.

Theorem 24.2.1. Suppose $A \in \mathbb{R}^{n \times d}$ has full column rank. If the entries of $S \in \mathbb{R}^{m \times n}$ are i.i.d. $N(0,1 / m), m=O\left(d / \epsilon^{2}\right)$, then for any vectors $a, b$ and $x^{*}=A^{\dagger} b$, we have, with probability $1-1 / \operatorname{poly}(d)$,

$$
\left|a^{\top}(S A)^{\dagger} S b-a^{\top} x^{*}\right| \lesssim \frac{\epsilon \sqrt{\log d}}{\sqrt{d}}\|a\|_{2}\left\|b-A x^{*}\right\|_{2}\left\|A^{\dagger}\right\|_{2}
$$



Figure 24.1: Our construction of $A$ and $b$ for the proof that Count-Sketch does not obey the $\ell_{\infty}$ guarantee. $\alpha<d$.

Because $S A$ has full column rank with probability $1,(S A)^{\dagger} S A=I$. Therefore

$$
\left|a^{\top}(S A)^{\dagger} S b-a^{\top} x^{*}\right|=\left|a^{\top}(S A)^{\dagger} S\left(b-A x^{*}\right)\right|=\left|a^{\top}(S A)^{\dagger} S\left(b-A A^{\dagger} b\right)\right| .
$$

Thus it suffices to only consider vectors $b$ where $A^{\dagger} b=0$, or equivalently $U^{\top} b=0$. In such cases, $S U$ will be independent of $S b$, which will give the result. The proof is in Appendix 24.6.

### 24.3 SRHT Matrices

We first provide the definition of the subsampled randomized Hadamard transform (SRHT): let $S=\frac{1}{\sqrt{r n}} P H_{n} D$. Here, $D$ is an $n \times n$ diagonal matrix with i.i.d. diagonal entries $D_{i, i}$, for which $D_{i, i}$ in uniform on $\{-1,+1\}$. The matrix $H_{n}$ is the Hadamard matrix of size $n \times n$, and we assume $n$ is a power of 2 . Here, $H_{n}=\left[H_{n / 2}, H_{n / 2} ; H_{n / 2},-H_{n / 2}\right]$ and $H_{1}=[1]$. The $r \times n$ matrix $P$ samples $r$ coordinates of an $n$ dimensional vector uniformly at random.

For other subspace embeddings, we no longer have that $S U$ and $S b$ are independent. To analyze them, we start with a claim that allows us to relate the inverse of a matrix to a power series.

Claim 24.3.1. Let $S \in \mathbb{R}^{m \times n}, A \in \mathbb{R}^{n \times d}$ have $S V D A=U \Sigma V^{\top}$, and define $T \in \mathbb{R}^{d \times d}$ by

$$
T=I_{d}-U^{\top} S^{\top} S U
$$

Suppose $S A$ has linearly independent columns and $\|T\|_{2} \leq 1 / 2$. Then

$$
\begin{equation*}
(S A)^{\dagger} S=V \Sigma^{-1}\left(\sum_{k=0}^{\infty} T^{k}\right) U^{\top} S^{\top} S \tag{24.7}
\end{equation*}
$$

Proof.

$$
\begin{aligned}
(S A)^{\dagger} S & =\left(A^{\top} S^{\top} S A\right)^{-1} A^{\top} S^{\top} S \\
& =\left(V \Sigma U^{\top} S^{\top} S U \Sigma V^{\top}\right)^{-1} V \Sigma U^{\top} S^{\top} S \\
& =V \Sigma^{-1}\left(U^{\top} S^{\top} S U\right)^{-1} U^{\top} S^{\top} S \\
& =V \Sigma^{-1}\left(I_{d}-T\right)^{-1} U^{\top} S^{\top} S \\
& =V \Sigma^{-1}\left(\sum_{k=0}^{\infty} T^{k}\right) U^{\top} S^{\top} S
\end{aligned}
$$

where in the last equality, since $\|T\|_{2}<1$, the von Neumann series $\sum_{k=0}^{\infty} T^{k}$ converges to $\left(I_{d}-T\right)^{-1}$.

We then bound the $k$ th term of this sum:
Lemma 24.3.2. Let $S \in \mathbb{R}^{r \times n}$ be the subsampled randomized Hadamard transform, and let $a$ be a unit vector. Then with probability $1-1 / \operatorname{poly}(n)$, we have

$$
\begin{aligned}
& \left|a^{\top} S^{\top} S\left(U U^{\top} S^{\top} S\right)^{k} b\right| \\
= & O\left(\log ^{k} n\right) \cdot(O(d(\log n) / r)+1)^{\frac{k-1}{2}} \cdot\left(\sqrt{d}\|b\|_{2}(\log n) / r+\|b\|_{2}\left(\log ^{\frac{1}{2}} n\right) / r^{\frac{1}{2}}\right)
\end{aligned}
$$

Hence, for $r$ at least $d \log ^{2 k+2} n \log ^{2}(n / \epsilon) / \epsilon^{2}$, this is at most $O\left(\|b\|_{2} \epsilon / \sqrt{d}\right)$ with probability at least $1-1 / \operatorname{poly}(n)$.

We defer the proof of this lemma to the next section, and now show how the lemma lets us prove that SRHT matrices satisfy the generalization bound with high probability:

Theorem 24.3.3. Suppose $A \in \mathbb{R}^{n \times d}$ has full column rank with $\log n=d^{o(1)}$. Let $S \in \mathbb{R}^{m \times n}$ be a subsampled randomized Hadamard transform with $m=O\left(d^{1+\alpha} / \epsilon^{2}\right)$ for $\alpha=\Theta\left(\sqrt{\frac{\log \log n}{\log d}}\right)$. For any vectors $a, b$ and $x^{*}=A^{\dagger} b$, we have

$$
\left|a^{\top}(S A)^{\dagger} S b-a^{\top} x^{*}\right| \lesssim \frac{\epsilon}{\sqrt{d}}\|a\|_{2}\left\|b-A x^{*}\right\|_{2}\left\|\Sigma^{-1}\right\|_{2}
$$

with probability $1-1 / \operatorname{poly}(d)$.

Proof. Define $\Delta=\Theta\left(\frac{1}{\sqrt{m}}\right)\left(\log ^{c} d\right)\|a\|_{2}\left\|b-A x^{*}\right\|_{2}\left\|\Sigma^{-1}\right\|_{2}$. For a constant $c>0$, we have that $S$ is a $(1 \pm \gamma)$-subspace embedding (Definition 24.1.1) for $\gamma=\sqrt{\frac{d \log ^{c} n}{m}}$ with probability $1-1 / \operatorname{poly}(d)$ (see, e.g., Theorem 2.4 of [Woo14a] and references therein), so $\|S U x\|_{2}=$
$(1 \pm \gamma)\|U x\|_{2}$ for all $x$, which we condition on. Hence for $T=I_{d}-U^{\top} S^{\top} S U$, we have $\|T\|_{2} \leq(1+\gamma)^{2}-1 \lesssim \gamma$. In particular, $\|T\|_{2}<1 / 2$ and we can apply Claim 24.3.1.

As in Section 24.2, $S A$ has full column rank if $S$ is a subspace embedding, so $(S A)^{\dagger} S A=I$ and we may assume $x^{*}=0$ without loss of generality.

By the approximate matrix product (Definition 24.1.2), we have for some $c$ that

$$
\begin{equation*}
\left|a^{\top} V \Sigma^{-1} U^{\top} S^{\top} S b\right| \leq \frac{\log ^{c} d}{\sqrt{m}}\|a\|_{2}\|b\|_{2}\left\|\Sigma^{-1}\right\|_{2} \leq \Delta \tag{24.8}
\end{equation*}
$$

with $1-1 / \operatorname{poly}(d)$ probability. Suppose this event occurs, bounding the $k=0$ term of (24.7). Hence it suffices to show that the $k \geq 1$ terms of (24.7) are bounded by $\Delta$.

By approximate matrix product (Definition 24.1.2), we also have with $1-1 / d^{2}$ probability that

$$
\left\|U^{\top} S^{\top} S b\right\|_{F} \leq \frac{\log ^{c} d}{\sqrt{m}}\left\|U^{\top}\right\|_{F}\|b\|_{2} \leq \frac{\log ^{c} d \sqrt{d}}{\sqrt{m}}\|b\|_{2}
$$

Combining with $\|T\|_{2} \lesssim \gamma$ we have for any $k$ that

$$
\left|a^{\top} V \Sigma^{-1} T^{k} U^{\top} S^{\top} S b\right| \lesssim \gamma^{k}\left(\log ^{c} d\right) \frac{\sqrt{d}}{\sqrt{m}}\|a\|_{2}\left\|\Sigma^{-1}\right\|_{2}\|b\|_{2}
$$

Since this decays exponentially in $k$ at a rate of $\gamma<1 / 2$, the sum of all terms greater than $k$ is bounded by the $k$ th term. As long as

$$
\begin{equation*}
m \gtrsim \frac{1}{\epsilon^{2}} d^{1+\frac{1}{k}} \log ^{c} n \tag{24.9}
\end{equation*}
$$

we have $\gamma=\sqrt{\frac{d \log ^{c} n}{m}}<\epsilon d^{-1 /(2 k)} / \log ^{c} n$, so that

$$
\sum_{k^{\prime} \geq k}\left|a^{\top} V \Sigma^{-1} T^{k^{\prime}} U^{\top} S^{\top} S b\right| \lesssim \frac{\epsilon}{\sqrt{d}}\|a\|_{2}\left\|\Sigma^{-1}\right\|_{2}\|b\|_{2}
$$

On the other hand, by Lemma 24.3.2, increasing $m$ by a $C^{k}$ factor, we have for all $k$ that

$$
\left|a^{\top} V^{\top} \Sigma^{-1} U^{\top} S^{\top} S\left(U U^{\top} S^{\top} S\right)^{k} b\right| \lesssim \frac{1}{2^{k}} \frac{\epsilon}{\sqrt{d}}\|a\|_{2}\|b\|_{2}\left\|\Sigma^{-1}\right\|_{2}
$$

with probability at least $1-1 / \operatorname{poly}(d)$, as long as $m \gtrsim d \log ^{2 k+2} n \log ^{2}(d / \epsilon) / \epsilon^{2}$. Since the $T^{k}$ term can be expanded as a sum of $2^{k}$ terms of this form, we get that

$$
\sum_{k^{\prime}=1}^{k}\left|a^{\top} V \Sigma^{-1} T^{k} U^{\top} S^{\top} S b\right| \lesssim \frac{\epsilon}{\sqrt{d}}\|a\|_{2}\|b\|_{2}\left\|\Sigma^{-1}\right\|_{2}
$$

with probability at least $1-1 / \operatorname{poly}(d)$, as long as $m \gtrsim d(C \log n)^{2 k+2} \log ^{2}(d / \epsilon) / \epsilon^{2}$ for a sufficiently large constant $C$. Combining with (24.9), the result holds as long as

$$
m \gtrsim \frac{d \log ^{c} n}{\epsilon^{2}} \max \left((C \log n)^{2 k+2}, d^{\frac{1}{k}}\right)
$$

for any $k$. Setting $k=\Theta\left(\sqrt{\frac{\log d}{\log \log n}}\right)$ gives the result.

Combining Different Matrices. In some cases it can make sense to combine different matrices that satisfy the generalization bound.

Theorem 24.3.4. Let $A \in \mathbb{R}^{n \times d}$, and let $R \in \mathbb{R}^{m \times r}$ and $S \in \mathbb{R}^{r \times n}$ be drawn from distributions of matrices that are $\epsilon$-approximate OSEs and satisfy the generalization bound (24.6). Then $R S$ satisfies the generalization bound with a constant factor loss in failure probability and approximation factor.

We defer the details to Appendix 24.7.

### 24.4 Proof of Lemma 24.3.2

Proof. Each column $S_{i}$ of the subsampled randomized Hadamard transform has the same distribution as $\sigma_{i} S_{i}$, where $\sigma_{i}$ is a random sign. It also has $\left\langle S_{i}, S_{i}\right\rangle=1$ for all $i$ and $\left|\left\langle S_{i}, S_{j}\right\rangle\right| \lesssim \frac{\sqrt{\log (1 / \delta)}}{\sqrt{r}}$ with probability $1-\delta$, for any $\delta$ and $i \neq j$. See, e.g., [LDFU13].

By expanding the following product into a sum, and rearranging terms, we obtain

$$
\begin{aligned}
& a^{\top} S^{\top} S\left(U U^{\top} S^{\top} S\right)^{k} b \\
= & \sum_{i_{0}, j_{0}, i_{1}, j_{1}, \cdots, i_{k}, j_{k}} a_{i_{0}} b_{j_{k}} \sigma_{i_{0}} \sigma_{i_{1}} \cdots \sigma_{i_{k}} \sigma_{j_{0}} \sigma_{j_{1}} \cdots \sigma_{j_{k}} \\
& \cdot\left\langle S_{i_{0}}, S_{j_{0}}\right\rangle\left(U U^{\top}\right)_{j_{0}, i_{1}}\left\langle S_{i_{1}}, S_{j_{1}}\right\rangle \cdots\left(U U^{\top}\right)_{j_{k-1}, i_{k}}\left\langle S_{i_{k}}, S_{j_{k}}\right\rangle \\
= & \sum_{i_{0}, j_{k}} a_{i_{0}} b_{j_{k}} \sigma_{i_{0}} \sigma_{j_{k}} \sum_{j_{0}, i_{1}, j_{1}, \cdots, i_{k}} \sigma_{i_{1}} \cdots \sigma_{i_{k}} \sigma_{j_{0}} \sigma_{j_{1}} \cdots \sigma_{j_{k-1}} \\
& \cdot\left\langle S_{i_{0}}, S_{j_{0}}\right\rangle\left(U U^{\top}\right)_{j_{0}, i_{1}}\left\langle S_{i_{1}}, S_{j_{1}}\right\rangle \cdots\left(U U^{\top}\right)_{j_{k-1}, i_{k}}\left\langle S_{i_{k}}, S_{j_{k}}\right\rangle \\
= & \sum_{i_{0}, j_{k}} \sigma_{i_{0}} \sigma_{j_{k}} Z_{i_{0}, j_{k}}
\end{aligned}
$$

where $Z_{i_{0}, j_{k}}$ is defined to be

$$
Z_{i_{0}, j_{k}}=a_{i_{0}} b_{j_{k}} \sum_{\substack{i_{1}, \cdots i_{k} \\ j_{0}, \cdots j_{k-1}}} \prod_{c=1}^{k} \sigma_{i_{c}} \prod_{c=0}^{k-1} \sigma_{j_{c}} \cdot \prod_{c=0}^{k}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle \prod_{c=1}^{k}\left(U U^{\top}\right)_{i_{c-1}, j_{c}}
$$

Note that $Z_{i_{0}, j_{k}}$ is independent of $\sigma_{i_{0}}$ and $\sigma_{j_{k}}$. We observe that in the above expression if $i_{0}=j_{0}, i_{1}=j_{1}, \cdots, i_{k}=j_{k}$, then the sum over these indices equals $a^{\top}\left(U U^{\top}\right) \cdots\left(U U^{\top}\right) b=0$, since $\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle=1$ in this case for all $c$. Moreover, the sum over all indices conditioned on $i_{k}=j_{k}$ is equal to 0 . Indeed, in this case, the expression can be factored into the form $\zeta \cdot U^{\top} b$, for some random variable $\zeta$, but $U^{\top} b=0$.

Let $W$ be a matrix with $W_{i, j}=\sigma_{i} \sigma_{j} Z_{i, j}$. We need Khintchine's inequality:

Fact 24.4.1 (Khintchine's Inequality). Let $\sigma_{1}, \ldots, \sigma_{n}$ be i.i.d. sign random variables, and let $z_{1}, \ldots, z_{n}$ be real numbers. Then there are constants $C, C^{\prime}>0$ so that

$$
\operatorname{Pr}\left[\left|\sum_{i=1}^{n} z_{i} \sigma_{i}\right| \geq C t\|z\|_{2}\right] \leq e^{-C^{\prime} t^{2}}
$$

We note that Khintchine's inequality sometimes refers to bounds on the moment of $\left|\sum_{i} z_{i} \sigma_{i}\right|$, though the above inequality follows readily by applying a Markov bound to the high moments.

We apply Fact 24.4.1 to each column of $W$, so that if $W_{i}$ is the $i$-th column, we have by a union bound that with probability $1-1 / \operatorname{poly}(n),\left\|W_{i}\right\|_{2}=O\left(\left\|Z_{i}\right\|_{2} \sqrt{\log n}\right)$ simultaneously for all columns $i$. It follows that with the same probability, $\|W\|_{F}^{2}=O\left(\|Z\|_{F}^{2} \log n\right)$, that is, $\|W\|_{F}=O\left(\|Z\|_{F} \sqrt{\log n}\right)$. We condition on this event in the remainder.

Thus, it remains to bound $\|Z\|_{F}$. By squaring $Z_{i_{0}, j_{0}}$ and using that $\mathbf{E}\left[\sigma_{i} \sigma_{j}\right]=1$ if $i=j$ and 0 otherwise, we have,

$$
\begin{equation*}
\underset{\sigma}{\mathbf{E}}\left[Z_{i_{0}, j_{k}}^{2}\right]=a_{i_{0}}^{2} b_{j_{k}}^{2} \sum_{\substack{i_{1}, \cdots i_{k} \\ j_{0}, \cdots j_{k-1}}} \prod_{c=0}^{k}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k}\left(U U^{\top}\right)_{i_{c-1}, j_{c}}^{2} \tag{24.10}
\end{equation*}
$$

We defer to Appendix 24.10 the proof that

$$
\underset{S}{\mathbf{E}}\left[\|Z\|_{F}^{2}\right] \leq(O(d(\log n) / r)+1)^{k-1} \cdot\left(d\|b\|_{2}^{2}\left(\log ^{2} n\right) / r^{2}+\|b\|_{2}^{2}(\log n) / r\right)
$$

Note that we also have the bound:

$$
(O(d(\log n) / r)+1)^{k-1} \leq\left(e^{O(d(\log n) / r)}\right)^{k-1} \leq e^{O(k d(\log n) / r)} \leq O(1)
$$

for any $r=\Omega(k d \log n)$.

Having computed the expectation of $\|Z\|_{F}^{2}$, we now would like to show concentration. Consider a specific
$Z_{i_{0}, j_{k}}=a_{i_{0}} b_{j_{k}} \sum_{i_{k}} \sigma_{i_{k}}\left\langle S_{i_{k}}, S_{j_{k}}\right\rangle \cdots \sum_{j_{1}} \sigma_{j_{1}}\left(U U^{\top}\right)_{j_{1}, i_{2}} \sum_{i_{1}} \sigma_{i_{1}}\left\langle S_{i_{1}}, S_{j_{1}}\right\rangle \sum_{j_{0}} \sigma_{j_{0}}\left\langle S_{i_{0}}, S_{j_{0}}\right\rangle\left(U U^{\top}\right)_{j_{0}, i_{1}}$.
By Fact 24.4.1, for each fixing of $i_{1}$, with probability $1-1 / \operatorname{poly}(n)$, we have

$$
\begin{equation*}
\sum_{j_{0}} \sigma_{j_{0}}\left\langle S_{i_{0}}, S_{j_{0}}\right\rangle\left(U U^{\top}\right)_{j_{0}, i_{1}}=O(\sqrt{\log n})\left(\sum_{j_{0}}\left\langle S_{i_{0}}, S_{j_{0}}\right\rangle^{2}\left(U U^{\top}\right)_{j_{0}, i_{1}}^{2}\right)^{\frac{1}{2}} \tag{24.11}
\end{equation*}
$$

Now, we can apply Khintchine's inequality for each fixing of $j_{1}$, and combine this with (24.11). With probability $1-1 / \operatorname{poly}(n)$, again we have

$$
\begin{aligned}
& \sum_{i_{1}} \sigma_{i_{1}}\left\langle S_{i_{1}}, S_{j_{1}}\right\rangle \sum_{j_{0}} \sigma_{j_{0}}\left\langle S_{i_{0}}, S_{j_{0}}\right\rangle\left(U U^{\top}\right)_{j_{0}, i_{1}} \\
= & \sum_{i_{1}} \sigma_{i_{1}}\left\langle S_{i_{1}}, S_{j_{1}}\right\rangle O(\sqrt{\log n})\left(\sum_{j_{0}}\left\langle S_{i_{0}}, S_{j_{0}}\right\rangle^{2}\left(U U^{\top}\right)_{j_{0}, i_{1}}^{2}\right)^{\frac{1}{2}} \\
= & O(\log n)\left(\sum_{i_{1}}\left\langle S_{i_{1}}, S_{j_{1}}\right\rangle^{2} \sum_{j_{0}}\left\langle S_{i_{0}}, S_{j_{0}}\right\rangle^{2}\left(U U^{\top}\right)_{j_{0}, i_{1}}^{2}\right)^{\frac{1}{2}}
\end{aligned}
$$

Thus, we can apply Khintchine's inequality recursively over all the $2 k$ indexes $j_{0}, i_{1}, j_{1}, \cdots, j_{k-1}, i_{k}$, from which it follows that with probability $1-1 / \operatorname{poly}(n)$, for each such $i_{0}, j_{k}$, we have $Z_{i_{0}, j_{k}}^{2}=O\left(\log ^{k} n\right) \underset{S}{\mathbf{E}}\left[Z_{i_{0}, j_{k}}^{2}\right]$, using (24.18). We thus have with this probability, that $\|Z\|_{F}^{2}=$ $O\left(\log ^{k} n\right) \underset{S}{\mathbf{E}}\left[\|Z\|_{F}^{2}\right]$, completing the proof.

### 24.5 Lower bound for $\ell_{2}$ and $\ell_{\infty}$ guarantee

We prove a lower bound for the $\ell_{2}$ guarantee, which immediately implies a lower bound for the $\ell_{\infty}$ guarantee.

Definition 24.5.1. Given a matrix $A \in \mathbb{R}^{n \times d}$, vector $b \in \mathbb{R}^{n}$ and matrix $S \in \mathbb{R}^{r \times n}$, denote $x^{*}=A^{\dagger} b$. We say that an algorithm $\mathcal{A}(A, b, S)$ that outputs a vector $x^{\prime}=(S A)^{\dagger} S b$ "succeeds" if the following property holds: $\left\|x^{\prime}-x^{*}\right\|_{2} \lesssim \epsilon\|b\|_{2} \cdot\left\|A^{\dagger}\right\|_{2} \cdot\left\|A x^{*}-b\right\|_{2}$.

Theorem 24.5.1. Suppose $\Pi$ is a distribution over $\mathbb{R}^{m \times n}$ with the property that for any $A \in \mathbb{R}^{n \times d}$ and $b \in \mathbb{R}^{n}, \operatorname{Pr}_{S \sim \Pi}[\mathcal{A}(A, b, S)$ succeeds $] \geq 19 / 20$. Then $m \gtrsim \min \left(n, d / \epsilon^{2}\right)$.

Proof. The proof uses Yao's minimax principle. Let $\mathcal{D}$ be an arbitrary distribution over $\mathbb{R}^{n \times(d+1)}$, then $\underset{(A, b) \sim \mathcal{D}}{\mathbb{E}} \underset{S \sim \Pi}{\mathbb{E}}[\mathcal{A}(A, b, S)$ succeeds $] \geq 1-\delta$. Switching the order of probabilistic quantifiers, an averaging argument implies the existence of a fixed matrix $S_{0} \in \mathbb{R}^{m \times n}$ such that

$$
\underset{(A, b) \sim \mathcal{D}}{\mathbb{E}}\left[\mathcal{A}\left(A, b, S_{0}\right) \text { succeeds }\right] \geq 1-\delta
$$

Thus, we must construct a distribution $\mathcal{D}_{\text {hard }}$ such that

$$
\underset{(A, b) \sim \mathcal{D}_{\text {hard }}}{\mathbb{E}}\left[\mathcal{A}\left(A, b, S_{0}\right) \text { succeeds }\right] \geq 1-\delta,
$$

cannot hold for any $\Pi_{0} \in \mathbb{R}^{m \times n}$ which does not satisfy $m=\Omega\left(d / \epsilon^{2}\right)$. The proof can be split into three parts. First, we prove a useful property. Second, we prove a lower bound for the case $\operatorname{rank}(S) \geq d$. Third, we show why $\operatorname{rank}(S) \geq d$ is necessary.
(I) We show that $[S A, S b]$ are independent Gaussian, if both $[A, b]$ and $S$ are or-
thonormal matrices. We can rewrite $S A$ in the following sense,

$$
\begin{align*}
\underbrace{S}_{m \times n} \cdot \underbrace{A}_{n \times d} & =\underbrace{S}_{m \times n} \underbrace{R}_{n \times n} \underbrace{R^{\top}}_{n \times n} \underbrace{A}_{n \times d} \\
& =S\left[\begin{array}{ll}
S^{\top} & \bar{S}^{\top}
\end{array}\right]\left[\begin{array}{c}
S \\
\bar{S}
\end{array}\right] A \\
& =\left[\begin{array}{ll}
I_{m} & 0
\end{array}\right]\left[\begin{array}{c}
S \\
\bar{S}
\end{array}\right] A \\
& =\left[\begin{array}{ll}
I_{m} & 0
\end{array}\right] \underbrace{\widetilde{A}}_{n \times d} \\
& =\underbrace{\widetilde{A}_{m}}_{m \times d} \tag{24.12}
\end{align*}
$$

where $\bar{S}$ is the complement of the orthonormal basis $S, I_{m}$ is a $m \times m$ identity matrix, and $\widetilde{A}_{m}$ is the left $m \times d$ submatrix of $\widetilde{A}$. Thus, using [Jia06] as long as $m=o(\sqrt{n})$ (because of $\left.n=\Omega\left(d^{3}\right)\right)$ the total variation distance between $[S A, S b]$ and a random Gaussian matrix is small, i.e.,

$$
\begin{equation*}
D_{T V}([S A, S b], H) \leq 0.01 \tag{24.13}
\end{equation*}
$$

where each entry of $H$ is i.i.d. Gaussian $\mathcal{N}(0,1 / n)$.
(II) Here we prove the theorem in the case when $S$ has rank $r \geq d$ (we will prove this is necessary in part III. Writing $S=U \Sigma V^{\top}$ in its SVD, we have

$$
\begin{equation*}
\underbrace{S}_{m \times n} A=\underbrace{U}_{m \times r} \underbrace{\Sigma}_{r \times r} \underbrace{V^{\top}}_{r \times n} R R^{\top} A=U \Sigma G \tag{24.14}
\end{equation*}
$$

where $R=\left[\begin{array}{ll}V & \bar{V}\end{array}\right]$. By a similar argument in Equation (24.12), as long as $r=o(\sqrt{n})$ we have that $G$ also can be approximated by a Gaussian matrix, where each entry is sampled from i.i.d. $\mathcal{N}(0,1 / n)$. Similarly, $S b=U \Sigma h$, where $h$ also can be approximated by a Gaussian matrix, where each entry is sampled from i.i.d. $\mathcal{N}(0,1 / n)$.

Since $U$ has linearly independent columns, $(U \Sigma G)^{\dagger} U \Sigma h=(\Sigma G)^{\dagger} U^{\top} U \Sigma h=(\Sigma G)^{\dagger} \Sigma h$.
The $r \times d$ matrix $G$ has $S V D G=\underbrace{R}_{r \times d} \underbrace{\widetilde{\Sigma}}_{d \times d} \underbrace{T}_{d \times d}$, and applying the pseudo-inverse property again, we have

$$
\begin{aligned}
\left\|(S A)^{\dagger} S b\right\|_{2} & =\left\|(\Sigma G)^{\dagger} \Sigma h\right\|_{2} \\
& =\left\|(\Sigma R \widetilde{\Sigma} T)^{\dagger} \Sigma h\right\|_{2} \\
& =\left\|T^{\dagger}(\Sigma R \widetilde{\Sigma})^{\dagger} \Sigma h\right\|_{2} \\
& =\left\|(\Sigma R \widetilde{\Sigma})^{\dagger} \Sigma h\right\|_{2} \\
& =\left\|\widetilde{\Sigma}^{\dagger}(\Sigma R)^{\dagger} \Sigma h\right\|_{2},
\end{aligned}
$$

where the the first equality follows by Equation (24.14), the second equality follows by the $S V D$ of $G$, the third and fifth equality follow by properties of the pseudo-inverse ${ }^{1}$ when $T$ has orthonormal rows and $\widetilde{\Sigma}$ is a diagonal matrix, and the fourth equality follows since $\left\|T^{\dagger}\right\|_{2}=1$ and $T$ is an orthonormal basis.

Because each entry of $G=R \widetilde{\Sigma} T \in \mathbb{R}^{r \times d}$ is sampled from an i.i.d. Gaussian $\mathcal{N}(0,1)$, using the result of [Ver10] we can give an upper bound for the maximum singular value of $G:\|\widetilde{\Sigma}\| \lesssim \sqrt{\frac{r}{n}}$ with probability at least .99. Thus,

$$
\begin{aligned}
\left\|\widetilde{\Sigma}^{\dagger}(\Sigma R)^{\dagger} \Sigma h\right\|_{2} & \geq \sigma_{\min }\left(\widetilde{\Sigma}^{\dagger}\right) \cdot\left\|(\Sigma R)^{\dagger} \Sigma h\right\|_{2} \\
& =\frac{1}{\sigma_{\max }(\widetilde{\Sigma})}\left\|(\Sigma R)^{\dagger} \Sigma h\right\|_{2} \\
& \gtrsim \sqrt{n / r}\left\|(\Sigma R)^{\dagger} \Sigma h\right\|_{2} .
\end{aligned}
$$

[^81]Because $h$ is a random Gaussian vector which is independent of $(\Sigma R)^{\dagger} \Sigma$, by Claim 24.5.2, $\mathbb{E}_{h}\left[\left\|(\Sigma R)^{\dagger} \Sigma h\right\|_{2}^{2}\right]=\frac{1}{n} \cdot\left\|(\Sigma R)^{\dagger} \Sigma\right\|_{F}^{2}$, where each entry of $h$ is sampled from i.i.d. Gaussian $\mathcal{N}(0,1 / n)$. Then, using the Pythagorean Theorem,

$$
\begin{aligned}
\left\|(\Sigma R)^{\dagger} \Sigma\right\|_{F}^{2} & =\left\|(\Sigma R)^{\dagger} \Sigma R R^{\top}\right\|_{F}^{2}+\left\|(\Sigma R)^{\dagger} \Sigma\left(I-R R^{\top}\right)\right\|_{F}^{2} \\
& \geq\left\|(\Sigma R)^{\dagger} \Sigma R R^{\top}\right\|_{F}^{2} \\
& =\left\|(\Sigma R)^{\dagger} \Sigma R\right\|_{F}^{2} \\
& =\operatorname{rank}(\Sigma R) \\
& =\operatorname{rank}(S A) \\
& =d
\end{aligned}
$$

Thus, $\left\|x^{\prime}-x^{*}\right\|_{2} \gtrsim \sqrt{d / r} \geq \sqrt{d / m}=\epsilon$.
(III) Now we show that we can assume that $\operatorname{rank}(S) \geq d$.

We sample $A, b$ based on the following distribution $\mathcal{D}_{\text {hard }}$ : with probability $1 / 2, A, b$ are sampled from $\mathcal{D}_{1}$; with probability $1 / 2, A, b$ are sampled from $\mathcal{D}_{2}$. In distribution $\mathcal{D}_{1}$, $A$ is a random orthonormal basis and $d$ is always orthogonal to $A$. In distribution $\mathcal{D}_{2}, A$ is a $d \times d$ identity matrix in the top- $d$ rows and 0 s elsewhere, while $b$ is a random unit vector. Then, for any $(A, b)$ sampled from $\mathcal{D}_{1}, S$ needs to work with probability at least $9 / 10$. Also for any $(A, b)$ sampled from $\mathcal{D}_{2}, S$ needs to work with probability at least $9 / 10$. The latter two statements follow since overall $S$ succeeds on $\mathcal{D}_{\text {hard }}$ with probability at least 19/20.

Consider the case where $A, b$ are sampled from distribution $\mathcal{D}_{2}$. Then $x^{*}=b$ and $\mathrm{OPT}=0$. Then consider $x^{\prime}$ which is the optimal solution to $\min _{x}\|S A x-S b\|_{2}^{2}$, so $x^{\prime}=$ $(S A)^{\dagger} S b=\left(S_{L}\right)^{\dagger} S_{L} b$, where $S$ can be decomposed into two matrices $S_{L} \in \mathbb{R}^{r \times d}$ and $S_{R} \in$
$\mathbb{R}^{r \times(n-d)}, S=\left[\begin{array}{ll}S_{L} & S_{R}\end{array}\right]$. Plugging $x^{\prime}$ into the original regression problem, $\left\|A x^{\prime}-b\right\|_{2}^{2}=$ $\left\|A\left(S_{L}\right)^{\dagger} S_{L} b-b\right\|_{2}^{2}$, which is at most $(1+\epsilon) \mathrm{OPT}=0$. Thus $\operatorname{rank}\left(S_{L}\right)$ is $d$. Since $S_{L}$ is a submatrix of $S$, the rank of $S$ is also $d$.

It remains to define several tools which are used in the main proof of the lower bound.
Claim 24.5.2. For any matrix $A \in \mathbb{R}^{n \times d}$, if each entry of a vector $g \in \mathbb{R}^{d}$ is chosen from an i.i.d Gaussian $\mathcal{N}\left(0, \sigma^{2}\right)$, then $\underset{g}{\mathbb{E}}\left[\|A g\|_{2}^{2}\right]=\sigma^{2}\|A\|_{F}^{2}$.

Proof.

$$
\begin{aligned}
\underset{g}{\mathbb{E}}\left[\|A g\|_{2}^{2}\right] & =\underset{g}{\mathbb{E}}\left[\sum_{i=1}^{n}\left(\sum_{j=1}^{d} A_{i j} g_{j}\right)^{2}\right] \\
& =\underset{g}{\mathbb{E}}\left[\sum_{i=1}^{n}\left(\sum_{j=1}^{d} A_{i j}^{2} g_{j}^{2}+\sum_{j \neq j^{\prime}} A_{i j} A_{i j^{\prime}} g_{j} g_{j^{\prime}}\right)\right] \\
& =\sum_{i=1}^{n} \sum_{j=1}^{d} A_{i j}^{2} \sigma^{2} \\
& =\sigma^{2}\|A\|_{F}^{2} .
\end{aligned}
$$

Let $g_{1}, g_{2}, \cdots, g_{t}$ be i.i.d. $\mathcal{N}(0,1)$ random variables. The random variables $\sum_{i=1}^{t} g_{i}^{2}$ are $X^{2}$ with $t$ degree of freedom. Furthermore, the tail bounds are known (see Lemma A.1.6).

Definition 24.5.2. Given a matrix $A \in \mathbb{R}^{n \times d}$, vector $b \in \mathbb{R}^{n}$ and matrix $S \in \mathbb{R}^{r \times n}$, denote $x^{*}=A^{\dagger} b$. We say that an algorithm $\mathcal{B}(A, b, S)$ that outputs a vector $x^{\prime}=(S A)^{\dagger} S b$ "succeeds" if the following property holds:

$$
\left\|x^{\prime}-x^{*}\right\|_{\infty} \lesssim \frac{\epsilon}{\sqrt{d}}\|b\|_{2} \cdot\left\|A^{\dagger}\right\|_{2} \cdot\left\|A x^{*}-b\right\|_{2}
$$

Applying $\left\|x^{\prime}-x\right\|_{\infty} \geq \frac{1}{\sqrt{d}}\left\|x^{\prime}-x\right\|_{2}$ to Theorem 24.5.1, we obtain the $\ell_{\infty}$ lower bound as a corollary,

Corollary 24.5.3. Suppose $\Pi$ is a distribution over $\mathbb{R}^{m \times n}$ with the property that for any $A \in \mathbb{R}^{n \times d}$ and $b \in \mathbb{R}^{n}$,

$$
\operatorname{Pr}_{S \sim \Pi}[\mathcal{B}(A, b, S) \text { succeeds }] \geq 9 / 10
$$

Then $m \gtrsim \min \left(n, d / \epsilon^{2}\right)$.

### 24.6 Proof for Gaussian case

Lemma 24.6.1. If the entries of $S \in \mathbb{R}^{m \times n}$ are i.i.d. $\quad N(0,1 / m), m=O\left(d / \epsilon^{2}\right)$, and $U^{\top} b=0$, then

$$
\left|a^{\top}(S A)^{\dagger} S b\right| \lesssim \frac{\epsilon \sqrt{\log d}}{\sqrt{d}}\|a\|_{2}\|b\|_{2}\left\|\Sigma^{-1}\right\|_{2}
$$

for any vectors $a, b$ with probability $1-1 / \operatorname{poly}(d)$.

Proof. With probability 1, the matrix $S A$ has linearly independent columns, and so $(S A)^{\dagger}$ is

$$
\begin{aligned}
& =\left(A^{\top} S^{\top} S A\right)^{-1} A^{\top} S^{\top} \\
& =\left(V \Sigma U^{\top} S^{\top} S U \Sigma V^{\top}\right)^{-1} V \Sigma U^{\top} S^{\top} \\
& =V \Sigma^{-1}\left(U^{\top} S^{\top} S U\right)^{-1} \Sigma^{-1} V^{\top} V \Sigma U^{\top} S^{\top} \\
& =V \Sigma^{-1}\left(U^{\top} S^{\top} S U\right)^{-1} U^{\top} S^{\top} .
\end{aligned}
$$

Hence, we would like to bound

$$
X=a^{\top} V \Sigma^{-1}\left(U^{\top} S^{\top} S U\right)^{-1} U^{\top} S^{\top} S b .
$$

It is well-known (stated, for example, explicitly in Theorem 2.3 of [Woo14a]) that with probability $1-\exp (-d)$, the singular values of $S U$ are $(1 \pm \epsilon)$ for $m=O\left(d / \epsilon^{2}\right)$. We condition on this event. It follows that

$$
\begin{aligned}
& \left\|V \Sigma^{-1}\left(U^{\top} S^{\top} S U\right)^{-1} U^{\top} S\right\|_{2} \\
= & \left\|\Sigma^{-1}\left(U^{\top} S^{\top} S U\right)^{-1} U^{\top} S\right\|_{2} \\
\leq & \left\|\Sigma^{-1}\right\|_{2}\left\|\left(U^{\top} S^{\top} S U\right)^{-1}\right\|_{2}\left\|U^{\top} S\right\|_{2} \\
\leq & \left\|\Sigma^{-1}\right\|_{2} \cdot \frac{1}{1-\epsilon} \cdot(1+\epsilon) \\
= & O\left(\left\|\Sigma^{-1}\right\|_{2}\right)
\end{aligned}
$$

where the first equality uses that $V$ is a rotation, the first inequality follows by submultiplicativity, and the second inequality uses that the singular values of $S U$ are in the range $[1-\epsilon, 1+\epsilon]$. Hence, with probability $1-\exp (-d)$,

$$
\begin{equation*}
\left\|a^{\top} V \Sigma^{-1}\left(U^{\top} S^{\top} S U\right)^{-1} U^{\top} S^{\top}\right\|_{2}=O\left(\left\|\Sigma^{-1}\right\|_{2}\|a\|_{2}\right) \tag{24.15}
\end{equation*}
$$

The main observation is that since $U^{\top} b=0, S U$ is statistically independent from $S b$. Hence, $S b$ is distributed as $N\left(0,\|b\|_{2}^{2} I_{m}\right)$, conditioned on the vector $a^{\top} V \Sigma^{-1}\left(U^{\top} S^{\top} S U\right)^{-1} U^{\top} S^{\top}$. It follows that conditioned on the value of $a^{\top} V \Sigma^{-1}\left(U^{\top} S^{\top} S U\right)^{-1} U^{\top} S^{\top}, X$ is distributed as

$$
N\left(0,\|b\|_{2}^{2}\left\|a^{\top} V \Sigma^{-1}\left(U^{\top} S^{\top} S U\right)^{-1} U^{\top} S^{\top}\right\|_{2}^{2} / m\right)
$$

and so using (24.15), with probability $1-1 / \operatorname{poly}(d)$, we have

$$
|X|=O\left(\epsilon \sqrt{\log d}\|a\|_{2}\|b\|_{2}\left\|\Sigma^{-1}\right\|_{2} / \sqrt{d}\right)
$$

### 24.7 Combining Different Matrices

In some cases it can make sense to combine different matrices that satisfy the generalization bound.

Theorem 24.3.4. Let $A \in \mathbb{R}^{n \times d}$, and let $R \in \mathbb{R}^{m \times r}$ and $S \in \mathbb{R}^{r \times n}$ be drawn from distributions of matrices that are $\epsilon$-approximate OSEs and satisfy the generalization bound (24.6). Then $R S$ satisfies the generalization bound with a constant factor loss in failure probability and approximation factor.

Proof. For any vectors $a, b$, and $x^{*}=A^{\dagger} b$ we want to show

$$
\left|a^{\top}(R S A)^{\dagger} R S b-a^{\top} x^{*}\right| \lesssim \frac{\epsilon}{\sqrt{d}}\|a\|_{2}\left\|b-A x^{*}\right\|_{2}\left\|A^{\dagger}\right\|_{2}
$$

As before, it suffices to consider the $x^{*}=0$ case. We have with probability $1-\delta$ that

$$
\left|a^{\top}(S A)^{\dagger} S b\right| \lesssim \frac{\epsilon}{\sqrt{d}}\|a\|_{2}\|b\|_{2}\left\|A^{\dagger}\right\|_{2}
$$

suppose this happens. We also have by the properties of $R$, applied to $S A$ and $S b$, that

$$
\left|a^{\top}(R S A)^{\dagger} R S b-a^{\top}(S A)^{\dagger} S b\right| \lesssim \frac{\epsilon}{\sqrt{d}}\|a\|_{2}\|S b\|_{2}\left\|(S A)^{\dagger}\right\|_{2}
$$

Because $S$ is an OSE, we have $\|S b\|_{2} \leq(1+\epsilon)$ and $\left\|(S A)^{\dagger}\right\|_{2} \gtrsim(1-\epsilon)\left\|A^{\dagger}\right\|_{2}$. Therefore

$$
\left|a^{\top}(R S A)^{\dagger} R S b\right| \lesssim \frac{\epsilon}{\sqrt{d}}\|a\|_{2}\|b\|_{2}\left\|A^{\dagger}\right\|_{2}
$$

We describe a few of the applications of combining sketches.

### 24.7.1 Removing dependence on $n$ via Count-Sketch

One of the limitations of the previous section is that the choice of $k$ depends on $n$. To prove that theorem, we have to assume that $\log d>\log \log n$. Here, we show an approach to remove that assumption.

The main idea is instead of applying matrix $S \in \mathbb{R}^{m \times n}$ to matrix $A \in \mathbb{R}^{n \times d}$ directly, we pick two matrices $S \in \mathbb{R}^{m \times \operatorname{poly}(d)}$ and $C \in \mathbb{R}^{\text {poly }(d) \times n}$, e.g. $S$ is FastJL matrix and $C$ is Count-Sketch matrix with $s=1$. We first compute $C \cdot A$, then compute $S \cdot(C A)$. The benefit of these operations is $S$ only needs to multiply with a matrix $(C A)$ that has poly $(d)$ rows, thus the assumption we need is $\log d>\log \log (\operatorname{poly}(d))$ which is always true. The reason for choosing $C$ as a Count-Sketch matrix with $s=1$ is: $(1) \mathrm{nnz}(C A) \leq \mathrm{nnz}(A)(2)$ The running time is $O(\operatorname{poly}(d) \cdot d+\operatorname{nnz}(A))$.

### 24.7.2 Combining Gaussians and SRHT

By combining Gaussians with SRHT matrices, we can embed into the optimal dimension $O\left(d / \epsilon^{2}\right)$ with fast $\widetilde{O}\left(n d \log n+d^{\omega} / \epsilon^{4}\right)$ embedding time.

### 24.7.3 Combining all three

By taking Gaussians times SRHT times Count-Sketch, we can embed into the optimal dimension $O\left(d / \epsilon^{2}\right)$ with fast $O\left(\operatorname{nnz}(A)+d^{4} \operatorname{poly}\left(\frac{1}{\epsilon}, \log d\right)\right)$ embedding time.

### 24.8 Count-Sketch does not obey the $\ell_{\infty}$ guarantee

Here we demonstrate an $A$ and a $b$ such that Count-Sketch will not satisfy the $\ell_{\infty}$ guarantee with constant probability, so such matrices cannot satisfy the generalization guar-
antee (24.6) with high probability.

Theorem 24.8.1. Let $S \in \mathbb{R}^{m \times n}$ be drawn as a Count-Sketch matrix with $s$ nonzeros per column. There exists a matrix $A \in \mathbb{R}^{n \times d}$ and $b \in \mathbb{R}^{n}$ such that, if $s^{2} d \lesssim m \lesssim \sqrt{d^{3} s}$, then the "true" solution $x^{*}=A^{\dagger} b$ and the approximation $x^{\prime}=(S A)^{\dagger} S b$ have large $\ell_{\infty}$ distance with constant probability:

$$
\left\|x^{\prime}-x^{*}\right\|_{\infty} \gtrsim \sqrt{\frac{d}{m s}}\|b\|_{2}
$$

Plugging in $m=d^{1.5}$ and $s=d^{0.25}$ we find that

$$
\left\|x^{\prime}-x^{*}\right\|_{\infty} \gtrsim 1 / d^{3 / 8}\|b\|_{2} \gg 1 / \sqrt{d}\|b\|_{2},
$$

even though such a matrix is an OSE with probability exponential in $s$. Therefore there exists a constant $c$ for which this matrix does not satisfy the generalization guarantee (24.6) with $1-\frac{c}{d}$ probability.

Proof. We choose the matrix $A$ to be the identity on its top $d$ rows: $A=\left[\begin{array}{c}I_{d} \\ 0\end{array}\right]$. Choose some $\alpha \geq 1$, set the value of the first $d$ coordinates of vector $b$ to be $\frac{1}{\sqrt{d}}$ and set the value to be $1 / \sqrt{\alpha}$ for the next $\alpha$ coordinates, with the remaining entries all zero. Note that $\|b\|_{2}=\sqrt{2}$, $x^{*}=(1 / \sqrt{d}, \ldots, 1 / \sqrt{d})$, and $\left\|A x^{*}-b\right\|_{2}=1$.

Let $S_{k}$ denote the $k$ th column vector of matrix $S \in \mathbb{R}^{m \times n}$. We define two events, Event I, $\forall k^{\prime} \in[d]$ and $k^{\prime} \neq k$, we have $\operatorname{supp}\left(S_{k^{\prime}}\right) \cap \operatorname{supp}\left(S_{k}\right)=\emptyset$; Event II, $\exists$ a unique $k^{\prime} \in\{d+1, d+2, \cdots, d+\alpha\}$ such that $\left|\operatorname{supp}\left(S_{k^{\prime}}\right) \cap \operatorname{supp}\left(S_{k}\right)\right|=1$, and all other $k^{\prime}$ have $\operatorname{supp}\left(S_{k^{\prime}}\right) \cap \operatorname{supp}\left(S_{k}\right)=\emptyset$. Using Claim 24.8.2, with probability at least .99 there exists a $k$ for which both events hold.

Given the constructions of $A$ and $b$ described early, it is obvious that

$$
A x-b=\left[x_{1}-\frac{1}{\sqrt{d}}, \cdots, x_{d}-\frac{1}{\sqrt{d}},-\frac{1}{\sqrt{\alpha}}, \cdots,-\frac{1}{\sqrt{\alpha}}, 0, \cdots, 0\right]^{\top} .
$$

Conditioned on event I and II are holding, then denote $\operatorname{supp}\left(S_{j}\right)=\left\{i_{1}, i_{2}, \cdots, i_{s}\right\}$. Consider the terms involving $x_{j}$ in the quadratic form

$$
\min _{x}\|S A x-S b\|_{2}^{2}
$$

it can be written as $(s-1)\left(x_{j}-1 / \sqrt{d}\right)^{2}+\left(x_{j}-1 / \sqrt{d} \pm 1 / \sqrt{\alpha}\right)^{2}$. Hence the optimal $x^{\prime}$ will have $x_{j}^{\prime}=\frac{1}{\sqrt{d}} \pm \frac{1}{s \sqrt{\alpha}}$, which is different from the desired $1 / \sqrt{d}$ by $\frac{1}{s \sqrt{\alpha}}$. Plugging in our requirement of $\alpha \bar{\sim} m^{2} /\left(s^{3} d^{2}\right)$, we have

$$
\left\|x^{\prime}-x^{*}\right\|_{\infty} \geq \frac{1}{s \sqrt{\alpha}} \gtrsim c \sqrt{\frac{s d^{2}}{m^{2}}} \gtrsim \frac{1}{\sqrt{d}}
$$

where the last inequality follows by $m \lesssim \sqrt{s d^{3}}$. Thus, we get the result.

Claim 24.8.2. If $m=\Omega\left(s^{2} d\right), m=o\left(d^{2}\right), \alpha<d$, and $\alpha=O\left(\frac{m^{2}}{s^{3} d^{2}}\right)$, with probability at least .99 there exists a $k \in[d]$ for which both event I and II hold.

Proof. If $m=\Omega\left(s^{2} d\right)$, then for any $i$ in $\{1,2, \ldots, d\}$, let $X_{i}$ be an indicator that the entries of column i are disjoint from all $i^{\prime}$ in $[d] \backslash\{i\}$. Then $\mathbb{E}\left[X_{i}\right] \geq .9999$, so by Markov's inequality, with probability .99 , we have $.99 d$ columns having this property (indeed, the expected value of $d-X$ is at most $.0001 d$, so $\left.\operatorname{Pr}[d-X \geq .01 d] \leq \frac{\mathbb{E}[d-X]}{.01 d} \leq \frac{.0001 d}{.01 d}=.01\right)$. Define Event $E$ to be that $.99 d$ columns of first $d$ columns have the property that the entries of that column


Figure 24.2: Count Sketch matrix $S \in \mathbb{R}^{m \times n}$. Event I, for any $k^{\prime} \in[d]$ and $k^{\prime} \neq k$, $\operatorname{supp}\left(S_{k}\right) \cap \operatorname{supp}\left(S_{k^{\prime}}\right)=\emptyset$. Event II, there exists a unique $k^{\prime} \in\{d+1, d+2, \cdots, d+\alpha\}$ such that $S_{k}$ and $S_{k^{\prime}}$ intersect at exactly one location(row index).
are disjoint from all the other $d-1$ columns. Let $S$ be the set of these $.99 d$ columns. Let $N$ be the union of supports of columns in $S$.

Each column $i$ in $\{d+1, \ldots, d+\alpha\}$ chooses $s$ non-zero entries. Define event $F$ ( which is similar as event $E$ ) to be that $.99 \alpha$ columns of the next $\alpha$ columns have the property that the entries of that column are disjoint from all the other $\alpha-1$ columns. By the same argument, since $\alpha<d$, with probability . 99 , we have $.99 \alpha$ columns in $\{d+1, \ldots, d+\alpha\}$ being disjoint from other columns in $\{d+1, \ldots, d+\alpha\}$. Condition on event $F$ holding. Let $L$ be the multiset union of supports of all columns in $\{d+1, \ldots, d+\alpha\}$. Then $L$ has size $\alpha \cdot s$. Let
$M$ be the union of supports of all columns in $\{d+1, \ldots, d+\alpha\}$, that is, the set union rather than the multiset union. Note that $|M| \geq .99 \alpha \cdot s$ because of $.99 \alpha$ columns are disjoint from each other.

The intersection size $x$ of $N$ and $M$ is hyper-geometrically distributed with expectation

$$
\mathbb{E}[x]=\frac{s|S| \cdot|M|}{m}
$$

By a lower tail bound for the hypergeometric distribution ${ }^{2}$,

$$
\operatorname{Pr}[x \leq(p-t) n] \leq \exp \left(-2 t^{2} n\right)
$$

where $p=s \cdot|S| / m$ and $n=|M|$, so

$$
\operatorname{Pr}[x \leq \mathbb{E}[x]-t \cdot|M|] \leq \exp \left(-2 t^{2} \cdot|M|\right) \leq 0.01
$$

where the last inequality follows by setting $t=\Theta(1 / \sqrt{|M|})$. Thus, we get with probability .99, the intersection size is at least $\frac{s|S| \cdot|M|}{m}-\Theta(\sqrt{|M|})$.

Now let $W$ be the distinct elements in $L \backslash M$, so necessarily $|W| \leq .01 \alpha \cdot s$. By an upper tail bound for the hypergeometric distribution, the intersection size $y$ of $W$ and $N$ satisfies

$$
\operatorname{Pr}[y \geq(p+t) n] \leq \exp \left(-2 t^{2} n\right)
$$

where $p=s \cdot|S| / m$ and $n=|W|$, we again get

$$
\operatorname{Pr}[y \geq \mathbb{E}[y]+t \cdot|W|] \leq \exp \left(-2 t^{2} \cdot|W|\right)
$$

[^82]If $|W|=0$, then $y=0$. Otherwise, we can set $t=\Theta(1 / \sqrt{|W|})$ so that this probability is less than .01 , and we get with probability .99 , the intersection size $y$ is at most $s \cdot|S| \cdot|W| / m+$ $\Theta(\sqrt{|W|})$. Note that we have that $\Theta(\sqrt{|M|})$ and $\Theta(\sqrt{|W|})$ are bounded by $\Theta(\sqrt{s \cdot \alpha})$. Setting $\alpha=O\left(\frac{m^{2}}{s^{3} d^{2}}\right)$ suffices to ensure $y$ is at most (1.01)s $\cdot|S| \cdot|W| / m$, and earlier that $x$ is at least $.99 \cdot s \cdot|S| \cdot|M| / m$.

The probability one of the $|S|$ blocks in $N$ has two or more intersections with $M$ is less than $\binom{x}{2}$ times the probability two random distinct items in the intersection land in the block. This probability is

$$
\frac{\binom{x}{2} \cdot\binom{s}{2}}{\binom{s \cdot|S|}{2}}=\Theta\left(x^{2} /|S|^{2}\right)=\Theta\left(x^{2} / d^{2}\right)=\Theta\left(m^{2} /\left(d^{4} s^{2}\right)\right)
$$

So the expected number of such blocks is $\Theta\left(m^{2} s d /\left(d^{4} s^{2}\right)\right)=\Theta\left(m^{2} /\left(d^{3} s\right)\right)$ which is less than $(.99 \cdot s \cdot|S| \cdot|M|) /(2 m) \leq X / 2$ if $m=o\left(d^{2}\right)$, which we have. So, there are at least $x / 2$ blocks which have intersection size exactly 1 with $N$. Note that the number of intersections of the $|S|$ blocks with $W$ is at most $y$, which is at most (1.01) $s \cdot|S| \cdot|W| / m \leq(1.01) s \cdot|S| \cdot \frac{1}{99} \cdot|M| / m<$ $x / 2$, and therefore there exists a block, that is, a column among the first $d$ columns, which intersects $M$ in exactly one position and does not intersect $W$. This is our desired column. Thus, we complete the proof.

### 24.9 Leverage score sampling does not obey the $\ell_{\infty}$ guarantee

Not only does Count-Sketch fail, but so does leverage score sampling, which is a technique that takes a subsample of rows of $A$ with rescaling. In this section we show an $A$ and a $b$ such that leverage score sampling will not satisfy the $\ell_{\infty}$ guarantee. We start with a formal definition of leverage scores.

Definition 24.9.1 (Leverage Scores). Given an arbitrary $n \times d$ matrix $A$, with $n>d$, let $U$ denote the $n \times d$ matrix consisting of the $d$ left singular vectors of $A$, let $U_{(i)}$ denote the $i$-th row of the matrix $U$, so $U_{(i)}$ is a row vector. Then the leverage scores of the rows of $A$ are given by $l_{i}=\left\|U_{(i)}\right\|_{2}^{2}$, for $i \in[n]$.

The leverage score sampling matrix can be thought of as a square diagonal matrix $D \in \mathbb{R}^{n \times n}$ with diagonal entries chosen from some distribution. If $D_{i i}=0$, it means we do not choose the $i$-th row of matrix $A$., If $D_{i i}>0$, it means we choose that row of the matrix $A$ and also rescale that row. We show that the leverage score sampling matrix cannot achieve $\ell_{\infty}$ guarantee, nor can it achieve our notion of generalization error.

Theorem 24.9.1. Let $D \in \mathbb{R}^{n \times n}$ be a leverage score sampling matrix with $m$ nonzeros on the diagonal. There exists a matrix $A \in \mathbb{R}^{n \times d}$ and a vector $b \in \mathbb{R}^{n}$ such that, if $m \lesssim d \sqrt{d}$, then the "true" solution $x^{*}=A^{\dagger} b$ and the approximation $x^{\prime}=(D A)^{\dagger} D b$ have large $\ell_{\infty}$ distance with constant probability:

$$
\left\|x^{\prime}-x^{*}\right\|_{\infty} \gtrsim \frac{1}{\sqrt{d}}\|b\|_{2}
$$

Therefore there exists a constant $c$ for which this matrix does not satisfy the generalization guarantee (24.6) with $1-\frac{c}{d}$ probability.

Proof. We choose the matrix $A$ to be the identity on its top $d$ rows, and $L$ scaled identity matrices $\frac{1}{\sqrt{\alpha d}} I_{d}$ for the next $d L$ rows, where $L$ satisfies $\frac{1}{d}+\frac{1}{\alpha d} L=1$ (to normalize each column of $A$ ), which implies $L=\alpha(d-1)$. Choose some $\beta \in[1, d)$. Set the value of the first $d$ coordinates of vector $b$ to be $\frac{1}{\sqrt{d}}$ and set the value to be $\frac{1}{\sqrt{\beta}}$ for the next $\beta$ coordinates, with the remaining entries all zero. Note that $\|b\|_{2}=\sqrt{2}$.

First, we compute $\|A x-b\|_{2}^{2}$. Because $\beta$ is less than $d$, there are two kinds of $x_{j}$ : one involves the following term,

$$
\begin{equation*}
\left(\frac{1}{\sqrt{d}} x_{j}-\frac{1}{\sqrt{d}}\right)^{2}+(L-1)\left(\frac{1}{\sqrt{\alpha d}} x_{j}\right)^{2} \tag{24.16}
\end{equation*}
$$

where the optimal $x_{j}$ should be set to $1 / d$. The other involves the term:

$$
\begin{equation*}
\left(\frac{1}{\sqrt{d}} x_{j}-\frac{1}{\sqrt{d}}\right)^{2}+\left(\frac{1}{\sqrt{\alpha d}} x_{j}-\frac{1}{\sqrt{\beta}}\right)^{2}+(L-1)\left(\frac{1}{\sqrt{\alpha d}} x_{j}\right)^{2} \tag{24.17}
\end{equation*}
$$

where the optimal $x_{j}$ should be set to $1 / d+1 / \sqrt{\alpha \beta d}$. Because we are able to choose $\alpha, \beta$ such that $\alpha \beta \gtrsim d$, then

$$
x_{j}=1 / d+1 / \sqrt{\alpha \beta d} \lesssim 1 / d
$$

Second, we compute $\|D A x-D b\|_{2}^{2}$. With high probability, there exists a $j$ satisfying Equation (24.17), but after applying leverage score sampling, the middle term of Equation (24.17) is removed. Let $p_{1}=\frac{1}{d}$ denote the leverage score of each of the top $d$ rows of $A$, and let $p_{2}=\frac{1}{\alpha d}$ denote the leverage score of each of the next $L d$ rows of $A$. We need to discuss the cases $m>d$ and $m \leq d$ separately.

If $m>d$, then the following term involves $x_{j}$,

$$
\begin{aligned}
& \left(\frac{1}{\sqrt{p_{1}}} \frac{1}{\sqrt{d}} x_{j}-\frac{1}{\sqrt{p_{1}}} \frac{1}{\sqrt{d}}\right)^{2}+\frac{m-d}{d} \cdot\left(\frac{1}{\sqrt{p_{2}}} \frac{1}{\sqrt{\alpha d}} x_{j}\right)^{2} \\
= & \frac{1}{p_{1}}\left(\frac{1}{\sqrt{d}} x_{j}-\frac{1}{\sqrt{d}}\right)^{2}+\frac{m-d}{d} \cdot \frac{1}{p_{2}}\left(\frac{1}{\sqrt{\alpha d}} x_{j}\right)^{2} \\
= & d\left(\left(\frac{1}{\sqrt{d}} x_{j}-\frac{1}{\sqrt{d}}\right)^{2}+\frac{m-d}{d} \alpha\left(\frac{1}{\sqrt{\alpha d}} x_{j}\right)^{2}\right) .
\end{aligned}
$$

where the optimal $x_{j}$ should be set to

$$
\begin{array}{rlr}
x_{j} & =\frac{1 / d}{1 / d+(m-d) \alpha /\left(\alpha d^{2}\right)} \\
& =\frac{1}{1+(m-d) / d} & \\
& \gtrsim \frac{1}{(m-d) / d} & \text { by } m \ll d \sqrt{d}
\end{array}
$$

If $m \leq d$, then the term involving $x_{j}$ is $\left(\frac{1}{\sqrt{p_{1}}} \frac{1}{\sqrt{d}} x_{j}-\frac{1}{\sqrt{p_{1}}} \frac{1}{\sqrt{d}}\right)^{2}$ where the optimal $x_{j}$ should be set to be $1 \gg 1 / \sqrt{d}$.

Third, we need to compute $\left\|A x^{*}-b\right\|_{2}^{2}$ and $\sigma_{\min }(A)$. It is easy to see that $\sigma_{\min }(A)$ because $A$ is an orthonormal matrix. The upper bound for $\left\|A x^{*}-b\right\|_{2}^{2}=2$, and the lower bound is also a constant, which can be proved in the following way:

$$
\left\|A x^{*}-b\right\|_{2}^{2}=\sum_{j=1}^{\beta}(24.16)+\sum_{j=\beta+1}^{d}(24.17) \geq d\left(\frac{1}{\sqrt{d}} \frac{1}{d}-\frac{1}{\sqrt{d}}\right)^{2} \gtrsim d \cdot \frac{1}{d}=1
$$

### 24.10 Bounding $\mathbb{E}\left[\|Z\|_{F}^{2}\right]$

Before getting into the proof details, we define the key property of $S$ being used in the rest of the proofs.

Definition 24.10.1 (All Inner Product $\operatorname{Small}(A I P S)$ Property). For any matrix $S \in \mathbb{R}^{r \times n}$, if for all $i, j \in[n]$ with $i \neq j$ we have

$$
\left|\left\langle S_{i}, S_{j}\right\rangle\right|=O(\sqrt{\log n} / \sqrt{r})
$$

we say that $S$ satisfies the "AIPS" property.

Claim 24.10.1. If $S \in \mathbb{R}^{r \times n}$ is a subsampled Hadamard transform matrix, then the AIPS property holds with probability at least $1-1 / \operatorname{poly}(n)$.

Proof. From the structure of $S$, for any $i \neq j$, we have with probability $1-1 / \operatorname{poly}(n)$ such that $\left|\left\langle S_{i}, S_{j}\right\rangle\right|=O(\sqrt{\log n} / \sqrt{r})$. Applying a union bound over $O\left(n^{2}\right)$ pairs, we obtain that

$$
\operatorname{Pr}[\text { AIPS holds }] \geq 1-1 / \operatorname{poly}(n)
$$

The main idea for bounding $\mathbb{E}\left[\|Z\|_{F}^{2}\right]$ is to rewrite it as $\mathbb{E}\left[\|Z\|_{F}^{2}\right]=\mathbb{E}\left[\|Z\|_{F}^{2} \mid\right.$ AIPS holds $]+$ $\mathbb{E}\left[\|Z\|_{F}^{2} \mid\right.$ AIPS does not hold $]$. Because $\operatorname{Pr}[$ AIPS does not hold $]$ is at most $1 / \operatorname{poly}(n)$, the first term dominates the second term, which means we only need to pay attention to the first term. We repeatedly apply this idea until all the $S$ are removed.

We start by boundinbg $\mathbb{E}\left[\|Z\|_{F}^{2}\right]$ by squaring $Z_{i_{0}, j_{0}}$ and using that $\mathbb{E}\left[\sigma_{i} \sigma_{j}\right]=1$ if $i=j$ and 0 otherwise. Then, we obtain,

$$
\begin{equation*}
\underset{\sigma}{\mathbb{E}}\left[Z_{i_{0}, j_{k}}^{2}\right]=a_{i_{0}}^{2} b_{j_{k}}^{2} \sum_{i_{1}, \cdots i_{k}, j_{0}, \cdots j_{k-1}} \prod_{c=0}^{k}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k}\left(U U^{\top}\right)_{i_{c-1}, j_{c}}^{2} . \tag{24.18}
\end{equation*}
$$

We thus have,

$$
\sum_{i_{0}, j_{k}, j_{k} \neq i_{k}} a_{i_{0}}^{2}\left\langle S_{i_{0}}, S_{j_{0}}\right\rangle^{2} b_{j_{k}}^{2}\left\langle S_{i_{k}}, S_{j_{k}}\right\rangle^{2}=a_{j_{0}}^{2}\|b\|_{2}^{2} O((\log n) / r)+\|a\|_{2}^{2}\|b\|_{2}^{2} O\left(\left(\log ^{2} n\right) / r^{2}\right) \stackrel{\text { def }}{=} C_{j_{0}}
$$

where the first equality is from our conditioning, and the second equality is the definition
of $C_{j_{0}}$. Hence, $\underset{S}{\mathbb{E}}\left[\|Z\|_{F}^{2}\right]$ is

$$
\begin{aligned}
& =\sum_{i_{1}, \cdots i_{k}, j_{0}, \cdots j_{k-1}} \prod_{c=1}^{k-1}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} \cdot \sum_{i_{0}, j_{k}, j_{k} \neq i_{k}} a_{i_{0}}^{2}\left\langle S_{i_{0}}, S_{j_{0}}\right\rangle^{2} b_{j_{k}}^{2}\left\langle S_{i_{k}}, S_{j_{k}}\right\rangle^{2} \\
& =\sum_{i_{1}, \cdots i_{k}, j_{0}, \cdots j_{k-1}} \prod_{c=1}^{k-1}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} C_{j_{0}} \\
& =\sum_{i_{1}, \cdots i_{k}, j_{0}, \cdots j_{k-1}}\left\langle S_{i_{k-1}}, S_{j_{k-1}}\right\rangle^{2}\left(U U^{\top}\right)_{j_{k-1}, i_{k}}^{2} \cdot \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} C_{j_{0}},
\end{aligned}
$$

where the first equality follows from (24.18), the second equality by definition of $C_{j_{0}}$, and the final equality by factoring out $c=k-1$ from one product and $c=k-2$ from the other product.

The way to bound the term $\left\langle S_{i_{k-1}}, S_{j_{k-1}}\right\rangle$ is by separating the diagonal term where $i_{k-1}=j_{k-1}$ and the non-diagonal term where $i_{k-1} \neq j_{k-1}$. We now use the aforementioned property of $S$, namely, that $\left\langle S_{i_{k-1}}, S_{j_{k-1}}\right\rangle=1$, if $i_{k-1}=j_{k-1}$, while for $i_{k-1} \neq j_{k-1}$, we have with probability $1-1 / \operatorname{poly}(n)$ that $\left|\left\langle S_{i_{k-1}}, S_{j_{k-1}}\right\rangle\right|=O(\sqrt{\log n} / \sqrt{r})$ conditioned on AIPS holding.

Conditioned on AIPS holding, we can recursively reduce the number of terms in the
product:

$$
\begin{aligned}
& \|Z\|_{F}^{2} \\
= & \sum_{i_{1}, \cdots i_{k}, j_{0}, \cdots j_{k-1}, i_{k-1} \neq j_{k-1}} O((\log n) / r) \cdot\left(U U^{\top}\right)_{j_{k-1}, i_{k}}^{2} \cdot \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} C_{j_{0}} \\
+ & \sum_{i_{1}, \cdots i_{k}, j_{0}, \cdots j_{k-1}, i_{k-1}=j_{k-1}} 1 \cdot\left(U U^{\top}\right)_{j_{k-1}, i_{k}}^{2} \cdot \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} C_{j_{0}} \\
\leq & \sum_{i_{1}, \cdots i_{k}, j_{0}, \cdots j_{k-1}} O((\log n) / r) \cdot\left(U U^{\top}\right)_{j_{k-1}, i_{k}}^{2} \cdot \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} C_{j_{0}} \\
+ & \sum_{i_{1}, \cdots i_{k}, j_{0}, \cdots j_{k-1}, i_{k-1}=j_{k-1}} 1 \cdot\left(U U^{\top}\right)_{j_{k-1}, i_{k}}^{2} \cdot \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} C_{j_{0}},
\end{aligned}
$$

where the first equality follows from the property just mentioned, and the inequality follows by including back the tuples of indices for which $i_{k-1}=j_{k-1}$, using that each summand is non-negative.

Our next step will be to bound the term $\left(U U^{\top}\right)_{j_{k-1}, i_{k}}^{2}$. We have, $\|Z\|_{F}^{2}$ is

$$
\begin{aligned}
& \leq \sum_{i_{k}, j_{k-1}}\left(U U^{\top}\right)_{i_{k}, j_{k-1}}^{2} \sum_{\substack{i_{1}, \cdots, i_{k-1} \\
j_{0}, \cdots, j_{k-2}}} O((\log n) / r) \\
& \cdot \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} C_{j_{0}} \\
& +\sum_{\substack{i_{1}, \cdots, i_{k} \\
j_{0}, \cdots, k_{k-1} \\
i_{k},-1=j_{k-1}}} 1 \cdot\left(U U^{\top}\right)_{j_{k-1}, i_{k}}^{2} \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} C_{j_{0}} \\
& =\underbrace{O(d(\log n) / r) \sum_{\substack{i_{1}, \cdots, i_{k-1} \\
j_{0}, \cdots, j_{k-2}}} \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} C_{j_{0}}}_{B} \\
& +\underbrace{}_{\begin{array}{c}
\sum_{\substack{i_{1}, \cdots, i_{k} \\
j_{0}, \cdots, j_{k-1} \\
i_{k-1}=j_{k-1}}} 1 \cdot\left(U U^{\top}\right)_{j_{k-1}, i_{k}}^{2} \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} C_{j_{0}}
\end{array}}
\end{aligned}
$$

where the equality uses that $\sum_{i_{k}, j_{k-1}}\left(U U^{\top}\right)_{i_{k}, j_{k-1}}^{2}=\left\|U U^{\top}\right\|_{F}^{2}=d$. We first upper bound
term $B$ :

$$
\begin{aligned}
& =\sum_{\substack{i_{1}, \cdots, i_{k} \\
j_{0}, \cdots, j_{k-1} \\
i_{k-1}=j_{k-1}}} 1 \cdot\left(U U^{\top}\right)_{j_{k-1}, i_{k}}^{2} \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} C_{j_{0}} \\
& =\sum_{\substack{i_{1}, \cdots, i_{k-1} \\
j_{0}, j_{1}, \cdots, j_{k-1} \\
i_{k-1}=j_{k-1}}} C_{j_{0}} \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} \sum_{i_{k}}\left(U U^{\top}\right)_{j_{k-1}, i_{k}}^{2} \\
& =\sum_{\substack{i_{1}, \cdots, i_{k-1} \\
j_{0}, j_{1}, \cdots, j_{k-1} \\
i_{k-1}=j_{k-1}}} C_{j_{0}} \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2}\left|e_{j_{k-1}} U U^{\top}\right|^{2} \\
& \leq \sum_{\substack{i_{1}, \cdots, i_{k-1} \\
j_{0}, j_{1}, \cdots, j_{k-1} \\
i_{k-1}=j_{k-1}}} C_{j_{0}} \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} 1 \\
& =\sum_{\substack{i_{1}, \cdots, i_{k-1} \\
j_{0}, j_{1}, \cdots, j_{k-2}}} C_{j_{0}} \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2}, \\
& =\sum_{c=1},
\end{aligned}
$$

where the first equality is the definition of $B$, the second equality follows by separating out the index $i_{k}$, the third equality uses that $\sum_{i_{k}}\left(U U^{\top}\right)_{j_{k-1}, i_{k}}^{2}=\left\|e_{j_{k-1}} U U^{\top}\right\|_{2}^{2}$, that is, the squared norm of the $j_{k-1}$-th row of $U U^{\top}$, the inequality follows since all rows of a projection matrix $U U^{\top}$ have norm at most 1 , and the final equality uses that $j_{k-1}$ no longer appears in the expression.

We now merge our bounds for the terms $A$ and $B$ in the following way:

$$
\begin{aligned}
& \|Z\|_{F}^{2} \\
\leq & A+B \\
\leq & O(d(\log n) / r) \sum_{\substack{i_{1}, \ldots, i_{k-1} \\
j_{0}, \cdots, j_{k-2}}} \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} C_{j_{0}} \\
+ & \sum_{\substack{i_{1}, \cdots, i_{k-1} \\
j_{0}, j_{1}, \cdots, j_{k-2}}} C_{j_{0}} \prod_{c=1}^{k-2}\left\langle S_{i_{c},}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} \\
= & (O(d(\log n) / r)+1) \sum_{\substack{i_{1}, \cdots, i_{k-1} \\
j_{0}, \cdots, j_{k-2}}} \prod_{c=1}^{k-2}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} C_{j_{0}} \\
\leq & \cdots \\
\leq & (O(d(\log n) / r)+1)^{2} \sum_{\substack{i_{1}, \cdots, i_{k-2} \\
j_{0}, \cdots, j_{k-3}}} \prod_{c=1}^{k-3}\left\langle S_{i_{c}}, S_{j_{c}}\right\rangle^{2} \prod_{c=1}^{k-2}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} C_{j_{0}} \\
\leq & \cdots \\
\leq & (O(d(\log n) / r)+1)^{k-1} \sum_{i_{1}, j_{0}} \prod_{c=1}^{1}\left(U U^{\top}\right)_{j_{c-1}, i_{c}}^{2} C_{j_{0}} \\
\leq & (O(d(\log n) / r)+1)^{k-1}\left(d\|b\|_{2}^{2}\left(\log ^{2} n\right) / r^{2}+\|b\|_{2}^{2}(\log n) / r\right),
\end{aligned}
$$

where the first two inequalities and first equality are by definition of $A$ and $B$ above. The first inequality follows by induction, since at this point we have replaced $k$ with $k-1$, and can repeat the argument, incurring another multiplicative factor of $O(d(\log n) / r)+$ 1. Repeating the induction in this way we arrive at the last inequality. Finally, the last
inequality follows by plugging in the definition of $C_{j_{0}}$, using that $\sum_{i_{1}, j_{0}}\left(U U^{\top}\right)_{j_{0}, i_{1}}^{2}=d$, and

$$
\sum_{j_{0}, i_{1}}\left(U U^{\top}\right)_{j_{0}, i_{1}}^{2} a_{j_{0}}^{2}=\sum_{j_{0}} a_{j_{0}}^{2} \sum_{i_{1}}\left(U U^{\top}\right)_{j_{0}, i_{1}}^{2}=\sum_{j_{0}} a_{j_{0}}^{2}\left\|e_{j_{0}} U U^{\top}\right\|_{2}^{2} \leq 1,
$$

where the inequality follows since each row of $U U^{\top}$ has norm at most 1 , and $a$ is a unit vector. The final result is that

$$
\|Z\|_{F}^{2} \leq(O(d(\log n) / r)+1)^{k-1}\left(d\|b\|_{2}^{2}\left(\log ^{2} n\right) / r^{2}+\|b\|_{2}^{2}(\log n) / r\right) .
$$

## Chapter 25

## Symmetric Norm Regression

We provide efficient algorithms for overconstrained linear regression problems with size $n \times d$ when the loss function is a symmetric norm (a norm invariant under sign-flips and coordinate-permutations). An important class of symmetric norms are Orlicz norms, where for a function $G$ and a vector $y \in \mathbb{R}^{n}$, the corresponding Orlicz norm $\|y\|_{G}$ is defined as the unique value $\alpha$ such that $\sum_{i=1}^{n} G\left(\left|y_{i}\right| / \alpha\right)=1$. When the loss function is an Orlicz norm, our algorithm produces a $(1+\epsilon)$-approximate solution for an arbitrarily small constant $\epsilon>0$ in input-sparsity time, improving over the previously best-known algorithm which produces a $d \cdot$ polylog $n$-approximate solution. When the loss function is a general symmetric norm, our algorithm produces a $\sqrt{d} \cdot \operatorname{polylog} n \cdot \operatorname{mmc}(\ell)$-approximate solution in input-sparsity time, where $\operatorname{mmc}(\ell)$ is a quantity related to the symmetric norm under consideration. To the best of our knowledge, this is the first input-sparsity time algorithm with provable guarantees for the general class of symmetric norm regression problem. Our results shed light on resolving the universal sketching problem for linear regression, and the techniques might be of independent interest to numerical linear algebra problems more broadly.

This part is based upon the following previous publication

- Zhao Song, Ruosong Wang, Lin F. Yang, Hongyang Zhang, Peilin Zhong Efficient Symmetric Norm Regression via Linear Sketching.

Manuscript 2019 [SWY ${ }^{+}$19]

### 25.1 Introduction

Linear regression is a fundamental problem in machine learning. For a data matrix $A \in \mathbb{R}^{n \times d}$ and a response vector $b \in \mathbb{R}^{n}$ with $n \gg d$, the overconstrained linear regression problem can be formulated as solving the following optimization problem:

$$
\begin{equation*}
\min _{x \in \mathbb{R}^{d}} \mathcal{L}(A x-b) \tag{25.1}
\end{equation*}
$$

where $\mathcal{L}: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is a loss function. Via the technique of linear sketching, we have witnessed many remarkable speedups for linear regression for a wide range of loss functions. Such technique involves designing a sketching matrix $S \in \mathbb{R}^{r \times n}$, and showing that by solving a linear regression instance on the data matrix $S A$ and the response vector $S b$, which is usually much smaller in size, one can obtain an approximate solution to the original linear regression instance in (25.1). Sarlós showed in [Sar06] that by taking $S$ as a Fast JohnsonLindenstrauss Transform matrix [AC06], one can obtain (1+ $1+$-approximate solutions to the least square regression problem $\left(\mathcal{L}(y)=\|y\|_{2}^{2}\right)$ in $O(n d \log n+\operatorname{poly}(d / \epsilon))$ time. The running time was later improved to $O(\mathrm{nnz}(A)+\operatorname{poly}(d / \epsilon))$ [CW13, MM13, NN13a, LMP13, Coh16a]. Here $\mathrm{nnz}(A)$ is the number of non-zero entries in the data matrix $A$, which could be much smaller than $n d$ for sparse data matrices. This technique was later generalized to other loss functions. By now, we have $\widetilde{O}(\mathrm{nnz}(A)+\operatorname{poly}(d / \epsilon))^{1}$ time algorithms for $\ell_{p}$ norms $\left(\mathcal{L}(y)=\|y\|_{p}^{p}\right)\left[\mathrm{DDH}^{+} 09, \mathrm{MM} 13, \mathrm{WZ13}, \mathrm{CP} 15\right.$, WW19], the quantile loss function [YMM13], and $M$-estimators [CW15b, CW15a].

Despite we have successfully applied the technique of linear sketching to many different loss functions, ideally, it would be more desirable to design algorithms that work

[^83]Table 25.1: $M$-estimators

| HUBER | $\begin{cases}x^{2} / 2 & \|x\| \leq c \\ c(\|x\|-c / 2) & \|x\|>c\end{cases}$ |
| :--- | :---: |
| $\ell_{1}-\ell_{2}$ | $2\left(\sqrt{1+x^{2} / 2}-1\right)$ |
| "FAIR" | $c^{2}(\|x\| / c-\log (1+\|x\| / c))$ |

for a wide range of loss functions, instead of designing a new sketching algorithm for every specific loss function. Naturally, this leads to the following problem, which is the linear regression version of the universal sketching problem ${ }^{2}$ studied in streaming algorithms [BO10, BCWY16]. We note that similar problems are also asked and studied for various algorithmic tasks, including principal component analysis [SWZ18], sparse recovery [NSW19b], approximate nearest neighbor search $\left[\mathrm{ANN}^{+} 17, \mathrm{ANN}^{+} 18\right]$ and mean estimation with statistical queries [FGV17, LNRW19].

Question 25.1.1. Is that possible to design sketching algorithms for linear regression, that work for a wide range of loss functions?

Prior to our work, [CW15b, CW15a] studied this problem in terms of $M$-estimators, where the loss function employs the form $\mathcal{L}(y)=\sum_{i=1}^{n} G\left(y_{i}\right)$ for some function $G$. See Table 25.1 for a list a $M$-estimators. However, much less is known for the case where the loss function $\mathcal{L}(\cdot)$ is a norm, except for $\ell_{p}$ norms. Recently, Andoni et al. [ALS ${ }^{+} 18$ ] tackle Problem 25.1.1 for Orlicz norms, which can be seen as a scale-invariant version of $M$-estimators. For a function $G$ and a vector $y \in \mathbb{R}^{n}$ with $y \neq 0$, the corresponding Orlicz

[^84]norm $\|y\|_{G}$ is defined as the unique value $\alpha$ such that
\[

$$
\begin{equation*}
\sum_{i=1}^{n} G\left(\left|y_{i}\right| / \alpha\right)=1 \tag{25.2}
\end{equation*}
$$

\]

When $y=0$, we define $\|y\|_{G}$ to be 0 . Note that Orlicz norms include $\ell_{p}$ norms as special cases, by taking $G(z)=|z|^{p}$ for some $p \geq 1$. Under certain assumptions on the function $G$, $\left[A L S^{+} 18\right]$ obtains the first input-sparsity time algorithm for solving Orlicz norm regression. More precisely, in $\widetilde{O}(\mathrm{nnz}(A)+\operatorname{poly}(d \log n))$ time, their algorithm obtains a solution $\widehat{x} \in \mathbb{R}^{d}$ such that $\|A \widehat{x}-b\|_{G} \leq d \cdot \operatorname{poly} \log n \cdot \min _{x \in \mathbb{R}^{d}}\|A x-b\|_{G}$.

There are two natural problems left open by the work of $[A L S+18]$. First, the algorithm in $\left[\mathrm{ALS}^{+} 18\right]$ has approximation ratio as large as $d \cdot \operatorname{poly} \log n$. Although this result is interesting from a theoretical point of view, such a large approximation ratio is prohibited for machine learning applications in practice. Is it possible to obtain an algorithm that runs in $\widetilde{O}(\mathrm{nnz}(A)+\operatorname{poly}(d / \epsilon))$ time, with approximation ratio $1+\epsilon$, for arbitrarily small $\epsilon$, similar to the case of $\ell_{p}$ norms? Moreover, although Orlicz norm includes a wide range of norms, many other important norms, e.g., top- $k$ norms (the sum of absolute values of the leading $k$ coordinates of a vector), max-mix of $\ell_{p}$ norms (e.g. $\max \left\{\|x\|_{2}, c\|x\|_{1}\right\}$ for some $c>0$ ), and sum-mix of $\ell_{p}$ norms (e.g. $\|x\|_{2}+c\|x\|_{1}$ for some $c>0$ ), are not Orlicz norms. More complicated examples include the $k$-support norm [AFS12] and the box-norm [MPS14], which have found applications in sparse recovery. In light of Problem 25.1.1, it is natural to ask whether it is possible to apply the technique of linear sketching to a broader class of norms. In this paper, we obtain affirmative answers to both problems, and make progress towards finally resolving Problem 25.1.1.

Notations. For a matrix $A \in \mathbb{R}^{n \times d}$, we use $A_{i} \in \mathbb{R}^{d}$ to denote its $i$-th row, viewed as a column vector. For $n$ real numbers $x_{1}, x_{2}, \cdots, x_{n}$, we define $\operatorname{diag}\left(x_{1}, x_{2}, \cdots, x_{n}\right) \in \mathbb{R}^{n \times n}$ to be the diagonal matrix where the $i$-th diagonal entry is $x_{i}$. For a vector $x \in \mathbb{R}^{n}$ and $p \geq 1$, we use $\|x\|_{p}$ to denote its $\ell_{p}$ norm, and $\|x\|_{0}$ to denote its $\ell_{0}$ norm, i.e., the number of non-zero entries in $x$. For two vectors $x, y \in \mathbb{R}^{n}$, we use $\langle x, y\rangle$ to denote their inner product. For any $n \in \mathbb{N}_{+}$, we use $[n]$ to denote the set $\{1,2, \cdots, n\}$. For $0 \leq p \leq 1$, we define $\operatorname{Ber}(p)$ to be the Bernoulli distribution with parameter $p$. We use $\mathbb{S}^{n-1}$ to denote the unit $\ell_{2}$ sphere in $\mathbb{R}^{n}$, i.e., $\mathbb{S}^{n-1}=\left\{x \in \mathbb{R}^{n} \mid\|x\|_{2}=1\right\}$. We use $\mathbb{R}_{\geq 0}$ to denote the set of all non-negative real numbers, i.e., $\mathbb{R}_{\geq 0}=\{x \in \mathbb{R} \mid x \geq 0\}$.

### 25.1.1 Our Contributions

Algorithm for Orlicz Norms. Our first contribution is a unified algorithm which produces $(1+\epsilon)$-approximate solutions to the linear regression problem in (25.1), when the loss function $\mathcal{L}(\cdot)$ is an Orlicz norm. Before introducing our results, we first give our assumptions on the function $G$ appeared in (25.2).

Assumption 25.1.2. We assume the function $G: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ satisfies the following properties:

1. $G$ is a strictly increasing convex function on $[0, \infty)$;
2. $G(0)=0$, and for all $x \in \mathbb{R}, G(x)=G(-x)$;
3. There exists some $C_{G}>0$, such that for all $0<x<y, G(y) / G(x) \leq C_{G}(y / x)^{2}$.

The first two conditions in Assumption 25.1.2 are necessary to make sure the corresponding Orlicz norm $\|\cdot\|_{G}$ is indeed a norm, and the third condition requires the function $G$
to have at most quadratic growth, which can be satisfied by all $M$-estimators in Table 25.1 and is also required by prior work [ALS $\left.{ }^{+} 18\right]$. Notice that our assumptions are weaker than those in $\left[\mathrm{ALS}^{+} 18\right]$. In $\left[\mathrm{ALS}^{+} 18\right]$, it is further required that $G(x)$ is a linear function when $x>1$, and $G$ is twice differentiable on an interval $\left(0, \delta_{G}\right)$ for some $\delta_{G}>0$. Given our assumptions on $G$, our main theorem is summarized as follow.

Theorem 25.1.3. For a function $G$ that satisfies Assumption 25.1.2, there exists an algorithm that, on any input $A \in \mathbb{R}^{n \times d}$ and $b \in \mathbb{R}^{n}$, finds a vector $x^{*}$ in time $\widetilde{O}(\mathrm{nnz}(A)+$ $\operatorname{poly}(d / \epsilon))$, such that with probability at least $0.9,\left\|A x^{*}-b\right\|_{G} \leq(1+\epsilon) \min _{x \in \mathbb{R}^{d}}\|A x-b\|_{G}$.

To the best of our knowledge, this is the first input-sparsity time algorithm with $(1+\epsilon)$-approximation guarantee, that goes beyond $\ell_{p}$ norms, the quantile loss function, and $M$-estimators. See Table 25.2 for a more comprehensive comparison with previous results.

Algorithm for Symmetric Norms. We further study the case when the loss function $\mathcal{L}(\cdot)$ is a symmetric norm. Symmetric norm is a more general class of norms, which includes all norms that are invariant under sign-flips and coordinate-permutations. Formally, we define symmetric norms as follow.

Definition 25.1.1. A norm $\|\cdot\|_{\ell}$ is called a symmetric norm, if

$$
\left\|\left(y_{1}, y_{2}, \ldots, y_{n}\right)\right\|_{\ell}=\left\|\left(s_{1} y_{\sigma_{1}}, s_{2} y_{\sigma_{2}}, \ldots, s_{n} y_{\sigma_{n}}\right)\right\|_{\ell}
$$

for any permutation $\sigma$ and any assignment of $s_{i} \in\{-1,1\}$.

Symmetric norm includes $\ell_{p}$ norms and Orlicz norms as special cases. It also includes all examples provided in the introduction, i.e., top- $k$ norms, max-mix of $\ell_{p}$ norms,

Table 25.2: Comparison between input-sparsity time linear regression algorithms

| Reference | Loss Function | Approximation Ratio |
| :---: | :---: | :---: |
| $\left[\mathrm{DDH}^{+} 09\right.$, MM13, WZ13, CP15, WW19] | $\ell_{p}$ norms | $1+\epsilon$ |
| $[$ YMM13 $]$ | Quantile loss function | $1+\epsilon$ |
| $[$ CW15b, CW15a] | $M$-estimators | $1+\epsilon$ |
| $\left[\right.$ ALS $\left.^{+} 18\right]$ | Orlicz norms | $d \cdot \operatorname{polylog} n$ |
| Theorem 25.1.3 | Orlicz norms | $1+\epsilon$ |
| Theorem 25.1.4 | Symmetric norms | $\sqrt{d} \cdot \operatorname{polylog} n \cdot \operatorname{mmc}(\ell)$ |

sum-mix of $\ell_{p}$ norms, the $k$-support norm [AFS12] and the box-norm [MPS14], as special cases. Understanding this general set of loss functions can be seen as a preliminary step to resolve Problem 25.1.1. Our main result for symmetric norm regression is summarized in the following theorem.

Theorem 25.1.4. Given a symmetric norm $\|\cdot\|_{\ell}$, there exists an algorithm that, on any input $A \in \mathbb{R}^{n \times d}$ and $b \in \mathbb{R}^{n}$, finds a vector $x^{*}$ in time $\widetilde{O}(\operatorname{nnz}(A)+\operatorname{poly}(d))$, such that with probability at least $0.9,\left\|A x^{*}-b\right\|_{\ell} \leq \sqrt{d} \cdot \operatorname{polylog} n \cdot \operatorname{mmc}(\ell) \cdot \min _{x \in \mathbb{R}^{d}}\|A x-b\|_{\ell}$.

In the above theorem, $\operatorname{mmc}(\ell)$ is a characteristic of the symmetric norm $\|\cdot\|_{\ell}$, which has been proven to be essential in streaming algorithms for symmetric norms [ $\left.\mathrm{BBC}^{+} 17\right]$. See Definition 25.3.3 for the formal definition of $\operatorname{mmc}(\ell)$, and Section 25.3 for more details about $\operatorname{mmc}(\ell)$. In particular, for $\ell_{p}$ norms with $p \leq 2$, top- $k$ norms with $k \geq n /$ polylog $n$, max-mix of $\ell_{2}$ norm and $\ell_{1}$ norm $\left(\max \left\{\|x\|_{2}, c\|x\|_{1}\right\}\right.$ for some $\left.c>0\right)$, sum-mix of $\ell_{2}$ norm and $\ell_{1} \operatorname{norm}\left(\|x\|_{2}+c\|x\|_{1}\right.$ for some $\left.c>0\right)$, the $k$-support norm, and the box-norm, $\operatorname{mmc}(\ell)$ can all be upper bounded by polylog $n$, which implies our algorithm has approximation ratio $\sqrt{d} \cdot$ polylog $n$ for all these norms. This clearly demonstrates the generality of our algorithm.

### 25.1.2 Technical Overview

Similar to previous works on using linear sketching to speed up solving linear regression, our core technique is to provide efficient dimensionality reduction methods for Orlicz norms and general symmetric norms. In this section, we discuss the techniques behind our results.

Row Sampling Algorithm for Orlicz Norms. Compared to prior work on Orlicz norm regression $\left[\mathrm{ALS}^{+} 18\right]$ which is based on random projection ${ }^{3}$, our new algorithm is based on row sampling. For a given matrix $A \in \mathbb{R}^{n \times d}$, our goal is to output a sparse weight vector $w \in \mathbb{R}^{n}$ with at most poly $(d \log n / \epsilon)$ non-zero entries, such that with high probability, for all $x \in \mathbb{R}^{d}$,

$$
\begin{equation*}
(1-\epsilon)\|A x-b\|_{G} \leq\|A x-b\|_{G, w} \leq(1+\epsilon)\|A x-b\|_{G} . \tag{25.3}
\end{equation*}
$$

Here, for a weight vector $w \in \mathbb{R}^{n}$ and a vector $y \in \mathbb{R}^{n}$, the weighted Orlicz norm $\|y\|_{G, w}$ is defined as the unique value $\alpha$ such that $\sum_{i=1}^{n} w_{i} G\left(\left|y_{i}\right| / \alpha\right)=1$. See Definition 25.2.1 for the formal definition of weighted Orlicz norm. To obtain a $(1+\epsilon)$-approximate solution to Orlicz norm regression, by (25.3), it suffices to solve

$$
\begin{equation*}
\min _{x \in \mathbb{R}^{d}}\|A x-b\|_{G, w} . \tag{25.4}
\end{equation*}
$$

Since the vector $w \in \mathbb{R}^{n}$ has at most poly $(d \log n / \epsilon)$ non-zero entries, and we can ignore all rows of $A$ with zero weights, there are at most $\operatorname{poly}(d \log n / \epsilon)$ remaining rows in $A$ in the optimization problem in (25.4). Furthermore, as we show in Lemma 25.2.1, $\|\cdot\|_{G, w}$ is

[^85]a seminorm, which implies we can solve the optimization problem in (25.4) in poly $(d \log n / \epsilon)$ time, by simply solving a convex program with size poly $(d \log n / \epsilon)$. Thus, we focus on how to obtain the weight vector $w \in \mathbb{R}^{n}$ in the remaining part. Furthermore, by taking $\mathbf{A b}$ to be a matrix whose first $d$ columns are $A$ and last column is $b$, to satisfy (25.3), it suffices to find a weight vector $w$ such that for all $x \in \mathbb{R}^{d+1}$,
\[

$$
\begin{equation*}
(1-\epsilon)\|\mathbf{A} \mathbf{b} x\|_{G} \leq\|\mathbf{A} \mathbf{b} x\|_{G, w} \leq(1+\epsilon)\|\mathbf{A} \mathbf{b} x\|_{G} \tag{25.5}
\end{equation*}
$$

\]

Hence, we ignore the response vector $b$ in the remaining part of the discussion.
We obtain the weight vector $w$ via importance sampling. We compute a set of sampling probabilities $\left\{p_{i}\right\}_{i=1}^{n}$ for each row of the data matrix $A$, and sample the rows of $A$ according to these probabilities. The $i$-th entry of the weight vector $w$ is then set to be $w_{i}=1 / p_{i}$ with probability $p_{i}$ and $w_{i}=0$ with probability $1-p_{i}$. However, unlike $\ell_{p}$ norms, Orlicz norms are not "entry-wise" norms, and it is not even clear that such a sampling process gives an unbiased estimation. Our key insight here is that for a vector $A x$ with unit Orlicz norm, if for all $x \in \mathbb{R}^{d}$,

$$
\begin{equation*}
(1-\epsilon) \sum_{i=1}^{n} G\left((A x)_{i}\right) \leq \sum_{i=1}^{n} w_{i} G\left((A x)_{i}\right) \leq(1+\epsilon) \sum_{i=1}^{n} G\left((A x)_{i}\right) \tag{25.6}
\end{equation*}
$$

then (25.5) holds, which follows from the convexity of the function $G$. See Lemma 25.2.5 and its proof for more details. Therefore, it remains to develop a way to define and calculate $\left\{p_{i}\right\}_{i=1}^{n}$, such that the total number of sampled rows is small.

Our method for defining and computing sampling probabilities $p_{i}$ is inspired by row sampling algorithms for $\ell_{p}$ norms [ $\left.\mathrm{DDH}^{+} 09\right]$. Here, the key is to obtain an upper bound on the contribution of each entry to the summation $\sum_{i=1}^{n} G\left((A x)_{i}\right)$. Indeed, suppose for some
vector $u \in \mathbb{R}^{n}$ such that $G(A x)_{i} \leq u_{i}$ for all $x \in \mathbb{R}^{d}$ with $\|A x\|_{G}=1$, we can then sample each row of $A$ with sampling probability proportional to $u_{i}$. Now, by standard concentration inequalities and a net argument, (25.6) holds with high probability. It remains to upper bound the total number of sampled rows, which is proportional to $\sum_{i=1}^{n} u_{i}$.

We use the case of $\ell_{2}$ norm, i.e., $G(x)=x^{2}$, as an example to illustrate our main ideas for choosing the vector $u \in \mathbb{R}^{n}$. Suppose $U \in \mathbb{R}^{n \times d}$ is an orthonormal basis matrix of the column space of $A$, then the leverage score ${ }^{4}$ is defined to be the squared $\ell_{2}$ norm of each row of $U$. Indeed, leverage score gives an upper bound on the contribution of each row to $\|U x\|_{2}^{2}$, since by Cauchy-Schwarz inequality, for each row $U_{i}$ of $U$, we have $\left\langle U_{i}, x\right\rangle^{2} \leq\left\|U_{i}\right\|_{2}^{2}\|x\|_{2}^{2}=\left\|U_{i}\right\|_{2}^{2}\|U x\|_{2}^{2}$, and thus we can set $u_{i}=\left\|U_{i}\right\|_{2}^{2}$. It is also clear that $\sum_{i=1}^{n} u_{i}=d$.

For general Orlicz norms, leverage scores are no longer upper bounds on $G\left((U x)_{i}\right)$. Inspired by the role of orthonormal bases in the case of $\ell_{2}$ norm, we first define well-conditioned basis for general Orlicz norms as follow.

Definition 25.1.2. Let $\|\cdot\|_{G}$ be an Orlicz norm induced by a function $G$ which satisfies Assumption 25.1.2. We say $U \in \mathbb{R}^{n \times d}$ is a well-conditioned basis with condition number $\kappa_{G}=\kappa_{G}(U)$ if for all $x \in \mathbb{R}^{d},\|x\|_{2} \leq\|U x\|_{G} \leq \kappa_{G}\|x\|_{2}$.

Given this definition, when $\|U x\|_{G}=1$, by Cauchy-Schwarz inequality and monotonicity of $G$, we can show that $G\left((U x)_{i}\right) \leq G\left(\left\|U_{i}\right\|_{2}\|x\|_{2}\right) \leq G\left(\left\|U_{i}\right\|_{2}\|U x\|_{G}\right) \leq G\left(\left\|U_{i}\right\|_{2}\right)$. This also leads to our definition of Orlicz norm leverage scores.

[^86]Definition 25.1.3. Let $\|\cdot\|_{G}$ be an Orlicz norm induced by a function $G$ which satisfies Assumption 25.1.2. For a given matrix $A \in \mathbb{R}^{n \times d}$ and a well-conditioned basis $U$ of the column space of $A$, the Orlicz norm leverage score of the $i$-th row of $A$ is defined to be $G\left(\left\|U_{i}\right\|_{2}\right)$.

It remains to give an upper bound on the summation of Orlicz norm leverage scores of all rows. Unlike the $\ell_{2}$ norm, it is not immediately clear how to use the definition of well-conditioned basis to obtain such an upper bound for general Orlicz norms. To achieve this goal, we use a novel probabilistic argument. Suppose one takes $x$ to be a vector with i.i.d. Gaussian random variables. Then each entry of $U x$ has the same distribution as $\left\|U_{i}\right\|_{2} \cdot g_{i}$, where $\left\{g_{i}\right\}_{i=1}^{n}$ is a set of standard Gaussian random variables. Thus, with constant probability, $\sum_{i=1}^{n} G\left((U x)_{i}\right)$ is an upper bound on the summation of Orlicz norm leverage scores. Furthermore, by the growth condition of the function $G$, we have $\sum_{i=1}^{n} G\left((U x)_{i}\right) \leq$ $C_{G}\|U x\|_{G}^{2}$. Now by Definition 25.1.2, $\|U x\|_{G} \leq \kappa_{G}\|x\|_{2}$, and $\|x\|_{2} \leq O(\sqrt{d})$ with constant probability by tail inequalities of Gaussian random variables. This implies an upper bound on the summation of Orlicz norm leverage scores. See Lemma 25.2.2 and its proof for more details.

Our approach for constructing well-conditioned bases is inspired by [SW11]. In Lemma 25.2.3, we show that given a subspace embedding $\Pi$ which embeds the column space of $A$ with Orlicz norm $\|\cdot\|_{G}$ into the $\ell_{2}$ space with distortion $\kappa$, then one can construct a well-conditioned basis with condition number $\kappa_{G} \leq \kappa$. The running time is dominated by calculating $\Pi A$ and doing a QR-decomposition on $\Pi A$. To this end, we can use the oblivi-
ous subspace embedding for Orlicz norms in Corollary 25.3.4 ${ }^{5}$ to construct well-conditioned bases. The embedding in Corollary 25.3.4 has $O(d)$ rows and $\kappa=\operatorname{poly}(d \log n)$, and calculating $\Pi A$ can be done in $\widetilde{O}(\operatorname{nnz}(A)+\operatorname{poly}(d))$ time. Using such an embedding to construct the well-conditioned basis, our row sampling algorithm produces a vector $w$ that satisfies (25.6) with $\|w\|_{0} \leq \operatorname{poly}(d \log n / \epsilon)$ in time $\widetilde{O}(\operatorname{nnz}(A)+\operatorname{poly}(d))$.

Oblivious Subspace Embeddings for Symmetric Norms. To obtain a faster algorithm for linear regression when the loss function is a general symmetric norm, we show that there exists a distribution over embedding matrices, such that if $S$ is a random matrix drawn from that distribution, then for any $n \times d$ matrix $A$, with constant probability, for all $x \in \mathbb{R}^{d},\|A x\|_{\ell} \leq\|S A x\|_{2} \leq \operatorname{poly}(d \log n) \cdot \operatorname{mmc}(\ell) \cdot\|A x\|_{\ell}$. Moreover, the embedding matrix $S$ is sparse, and calculating $S A$ requires only $\widetilde{O}(\mathrm{nnz}(A)+\operatorname{poly}(d))$ time. Another favorable property of $S$ is that it is an oblivious subspace embeeding, meaning the distribution of $S$ does not depend on $A$. To achieve this goal, it is sufficient to construct a random diagonal matrix $D$ such that for any fixed vector $x \in \mathbb{R}^{n}$,

$$
\begin{equation*}
\operatorname{Pr}\left[\|D x\|_{2} \geq \Omega(1 / \operatorname{poly}(d \log n)) \cdot\|x\|_{\ell}\right] \geq 1-\exp (-\Omega(d \log n)) \tag{25.7}
\end{equation*}
$$

and

$$
\begin{equation*}
\operatorname{Pr}\left[\|D x\|_{2} \leq \operatorname{poly}(d \log n) \cdot \operatorname{mmc}(\ell) \cdot\|x\|_{\ell}\right] \geq 1-O(1 / d) \tag{25.8}
\end{equation*}
$$

Our construction is inspired by the sub-sampling technique in [IW05], which was used for sketching symmetric norms in data streams $\left[\mathrm{BBC}^{+} 17\right]$. Throughout the discussion, we

[^87]use $\xi^{(q)} \in \mathbb{R}^{n}$ to denote a vector with $q$ non-zero entries and each entry is $1 / \sqrt{q}$. Let us start with a special case where the vector $x \in \mathbb{R}^{n}$ has $s$ non-zero entries and each non-zero entry is 1 . It is easy to see $\|x\|_{\ell}=\sqrt{s}\left\|\xi^{(s)}\right\|_{\ell}$. Now consider a random diagonal matrix $D$ which corresponds to a sampling process, i.e., each diagonal entry is set to be 1 with probability $p$ and 0 with probability $1-p$. Our goal is to show that $\sqrt{1 / p}\left\|\xi^{(1 / p)}\right\|_{\ell} \cdot\|D x\|_{2}$ is a good estimator of $\|x\|_{\ell}$. If $p=\Theta(d \log n / s)$, then with probability at least $1-\exp (-\Omega(d \log n))$, $D x$ will contain at least one non-zero entry from $x$, in which case (25.7) is satisfied. However, we do not know $s$ in advance. Thus, we use $t=O(\log n)$ different matrices $D_{1}, D_{2}, \cdots, D_{t}$, where $D_{i}$ has sampling probability $1 / 2^{i}$. Clearly at least one such $D_{j}$ can establish (25.7). For the upper bound part, if $p$ is much smaller than $1 / s$, then $D x$ will never contain a non-zero entry from $x$. Otherwise, in expectation $D x$ will contain ps non-zero entries, in which case our estimation will be roughly $\sqrt{s}\left\|\xi^{(1 / p)}\right\|_{\ell}$, which can be upper bounded by $O\left(\log n \cdot \operatorname{mmc}(\ell) \cdot \sqrt{s}\left\|\xi^{(s)}\right\|_{\ell}\right)$. At this point, (25.8) follows from Markov's inequality. See Section 25.7.5 for the formal argument, and Section 25.3 for a detailed discussion on $\mathrm{mmc}(\ell)$.

To generalize the above argument to general vectors, for a vector $x \in \mathbb{R}^{n}$, we conceptually partition its entries into $\Theta(\log n)$ groups, where the $i$-th group contains entries with magnitude in $\left[2^{i}, 2^{i+1}\right)$. By averaging, at least one group of entries contributes at least $\Omega(1 / \log n)$ fraction to the value of $\|x\|_{\ell}$. To establish (25.7), we apply the lower bound part of the argument in the previous paragraph to this "contributing" group. To establish (25.8), we apply the upper bound part of the argument to all groups, which will only induce an additional $O(\log n)$ factor in the approximation ratio, by triangle inequality.

Since our oblivious subspace embedding embeds a given symmetric norm into the $\ell_{2}$ space, in order to obtain an approximate solution to symmetric norm regression, we only need
to solve a least square regression instance with much smaller size. This is another advantage of our subspace embedding, since the least square regression problem is a well-studied problem in optimization and numerical linear algebra, for which many efficient algorithms are known, both in theory and in practice.

### 25.2 Linear Regression for Orlicz Norms

In this section, we introduce our results for Orlicz norm regression.

Weighted Orlicz Norm. We first give the definition of weighted Orlicz norm.

Definition 25.2.1. For a function $G$ that satisfies Assumption 25.1.2 and a weight vector $w \in \mathbb{R}^{n}$ such that $w_{i} \geq 0$ for all $i \in[n]$, for a vector $x \in \mathbb{R}^{n}$, if $\sum_{i=1}^{n} w_{i} \cdot\left|x_{i}\right|=0$, then the weighted Orlicz norm $\|x\|_{G, w}$ is defined to be 0 . Otherwise, the weighted Orlicz norm $\|x\|_{G, w}$ is defined as the unique value $\alpha>0$ such that $\sum_{i=1}^{n} w_{i} G\left(\left|x_{i}\right| / \alpha\right)=1$.

When $w_{i}=1$ for all $i \in[n]$, we have $\|x\|_{G, w}=\|x\|_{G}$ where $\|x\|_{G}$ is the (unweighted) Orlicz norm. It is well known that $\|\cdot\|_{G}$ is a norm. We show in the following lemma that $\|\cdot\|_{G, w}$ is a seminorm.

Lemma 25.2.1. For a function $G$ that satisfies Assumption 25.1.2 and a weight vector $w \in \mathbb{R}^{n}$ such that $w_{i} \geq 0$ for all $i \in[n]$, for all $x, y \in \mathbb{R}^{n}$, we have (i) $\|x\|_{G, w} \geq 0$, (ii) $\|x+y\|_{G, w} \leq\|x\|_{G, w}+\|y\|_{G, w}$, and (iii) $\|a x\|_{G, w}=|a| \cdot\|x\|_{G, w}$ for all $a \in \mathbb{R}$.

Leverage Scores and Well-Conditioned Bases for Orlicz Norms. The following lemma establishes an upper bound on the summation of Orlicz norm leverage scores defined in Definition 25.1.3.

Lemma 25.2.2. Let $\|\cdot\|_{G}$ be an Orlicz norm induced by a function $G$ which satisfies Assumption 25.1.2. Let $U \in \mathbb{R}^{n \times d}$ be a well-conditioned basis with condition number $\kappa_{G}$ as in Definition 25.1.2. Then we have $\sum_{i=1}^{n} G\left(\left\|U_{i}\right\|_{2}\right) \leq O\left(C_{G} d \kappa_{G}^{2}\right)$,

Now we show that given a subspace embedding which embeds the column space of $A$ with Orlicz norm $\|\cdot\|_{G}$ into the $\ell_{2}$ space with distortion $\kappa$, then one can construct a well-conditioned basis with condition number $\kappa_{G} \leq \kappa$.

Lemma 25.2.3. Let $\|\cdot\|_{G}$ be an Orlicz norm induced by a function $G$ which satisfies Assumption 25.1.2. For a given matrix $A \in \mathbb{R}^{n \times d}$ and an embedding matrix $\Pi \in \mathbb{R}^{s \times n}$, suppose for all $x \in \mathbb{R}^{d},\|A x\|_{G} \leq\|\Pi A x\|_{2} \leq \kappa\|A x\|_{G}$. Let $Q \cdot R=\frac{1}{\kappa} \Pi A$ be a QR-decomposition of $\frac{1}{\kappa} \Pi A$. Then $A R^{-1}$ is a well-conditioned basis (see Definition 25.1.2) with $\kappa_{G}\left(A R^{-1}\right) \leq \kappa$.

The following lemma shows how to estimate Orlicz norm leverage scores given a change of basis matrix $R \in \mathbb{R}^{d \times d}$, in $\widetilde{O}(\operatorname{nnz}(A)+\operatorname{poly}(d))$ time.

Lemma 25.2.4. Let $\|\cdot\|_{G}$ be an Orlicz norm induced by a function $G$ which satisfies Assumption 25.1.2. For a given matrix $A \in \mathbb{R}^{n \times d}$ and $R \in \mathbb{R}^{d \times d}$, there exists an algorithm that outputs $\left\{u_{i}\right\}_{i=1}^{n}$ such that with probability at least 0.99 , $u_{i}=\Theta\left(G\left(\left\|\left(A R^{-1}\right)_{i}\right\|_{2}\right)\right)$ for all $1 \leq i \leq n$. The algorithm runs in $\widetilde{O}(\mathrm{nnz}(A)+\operatorname{poly}(d))$ time.

The Row Sampling Algorithm. Based on the notion of Orlicz norm leverage scores and well-conditioned bases, we design a row sampling algorithm for Orlicz norms.

Lemma 25.2.5. Let $\|\cdot\|_{G}$ be an Orlicz norm induced by a function $G$ which satisfies Assumption 25.1.2. Let $U \in \mathbb{R}^{n \times d}$ be a well-conditioned basis with condition number $\kappa_{G}=\kappa_{G}(U)$ as in Definition 25.1.2. For sufficiently small $\epsilon$ and $\delta$, and sufficiently large constant $C$, let $\left\{p_{i}\right\}_{i=1}^{n}$ be a set of sampling probabilities satisfying $p_{i} \geq \min \left\{1, C(\log (1 / \delta)+d \log (1 / \epsilon)) \epsilon^{-2} G\left(\left\|U_{i}\right\|_{2}\right)\right\}$. Let $w$ be a vector whose $i$-th entry is set to be $w_{i}=1 / p_{i}$ with probability $p_{i}$ and $w_{i}=0$ with
probability $1-p_{i}$, then with probability at least $1-\delta$, for all $x \in \mathbb{R}^{d}$, we have $(1-\epsilon)\|U x\|_{G} \leq$ $\|U x\|_{G, w} \leq(1+\epsilon)\|U x\|_{G}$.

Solving Linear Regression for Orlicz Norms. Now we combine all ingredients to give an algorithm for Orlicz norm regression. We use $\mathbf{A b} \in \mathbb{R}^{n \times(d+1)}$ to denote a matrix whose first $d$ columns are $A$ and the last column is $b$. The algorithm is described in Figure 25.1, and we prove its running time and correctness in Theorem 25.2.6. We assume we are given an embedding matrix $\Pi$, such that for all $x \in \mathbb{R}^{d+1},\|\mathbf{A b} x\|_{G} \leq\|\Pi \mathbf{A b} x\|_{2} \leq \kappa\|\mathbf{A b} x\|_{G}$. The construction of $\Pi$ and the value $\kappa$ will be given in Corollary 25.3.4. In Appendix 25.8.1, we use Theorem 25.2.6 and Corollary 25.3.4 to formally prove Theorem 25.1.3.

1. For the given embedding matrix $\Pi$, calculate $\Pi \mathbf{A b}$ and invoke $Q R$-decomposition on $\Pi \mathbf{A b} / \kappa$ to obtain $Q \cdot R=\Pi \mathbf{A b} / \kappa$.
2. Invoke Lemma 25.2.4 to obtain $\left\{u_{i}\right\}_{i=1}^{n}$ such that $u_{i}=\Theta\left(G\left(\left\|\left(\mathbf{A b} R^{-1}\right)_{i}\right\|_{2}\right)\right)$.
3. For a sufficiently large constant $C$, let $\left\{p_{i}\right\}_{i=1}^{n}$ be a set of sampling probabilities with $p_{i} \geq \min \left\{1, C \cdot d \cdot \epsilon^{-2} \log (1 / \epsilon) \cdot G\left(\left\|\left(\mathbf{A b} R^{-1}\right)_{i}\right\|_{2}\right)\right\}$, and $w$ be a vector whose $i$-th entry $w_{i}=1 / p_{i}$ with probability $p_{i}$ and $w_{i}=0$ with probability $1-p_{i}$.
4. Calculate $x^{*}=\operatorname{argmin}_{x \in \mathbb{R}^{d}}\|A x-b\|_{G, w}$. Return $x^{*}$.

Figure 25.1: Algorithm for Orlicz norm regression

Theorem 25.2.6. Let $\|\cdot\|_{G}$ be an Orlicz norm induced by a function $G$ which satisfies Assumption 25.1.2. Given an embedding matrix $\Pi$, such that for all $x \in \mathbb{R}^{d},\|\mathbf{A b} x\|_{G} \leq$ $\|\Pi \mathbf{A} \mathbf{b} x\|_{2} \leq \kappa\|\mathbf{A b} x\|_{G}$, with probability at least 0.9, the algorithm in Figure 25.1 outputs $x^{*} \in \mathbb{R}^{d}$ in time $\operatorname{poly}(d \kappa / \epsilon)+\mathcal{T}_{\mathrm{QR}}(\Pi \mathbf{A b})$, such that $\left\|A x^{*}-b\right\|_{G} \leq(1+\epsilon) \min _{x \in \mathbb{R}^{d}}\|A x-b\|_{G}$. Here, $\mathcal{T}_{\mathrm{QR}}(\Pi \mathbf{A b})$ is the running time for calculating $\Pi \mathbf{A b}$ and invoking $Q R$-decomposition
on ПAb.

### 25.3 Linear Regression for Symmetric Norms

In this section, we introduce SymSketch, a subspace embedding for symmetric norms.

Definition of SymSketch. We first formally define SymSketch. Due to space limitation, we give the definition of Gaussian embeddings, CountSketch embeddings and their compositions in Appendix 25.7.1.1.

Definition 25.3.1 (Symmetric Norm Sketch (SymSketch)). Let $t=\Theta(\log n)$. Let $\widetilde{D} \in$ $\mathbb{R}^{n(t+1) \times n}$ be a matrix defined as $\widetilde{D}=\left[\begin{array}{c}w_{0} D_{0} \\ w_{1} D_{1} \\ \vdots \\ w_{t} D_{t}\end{array}\right]$, where for each $i \in\{0,1, \ldots, t\}, D_{i}=$ $\operatorname{diag}\left(z_{i, 1}, z_{i, 2}, \ldots, z_{i, n}\right) \in \mathbb{R}^{n \times n}$ and $z_{i, j} \sim \operatorname{Ber}\left(1 / 2^{i}\right)$ for each $j \in[n]$. Moreover, $w_{i}=$ $\|(\underbrace{1,1, \ldots, 1}_{2^{i} \text { ones }}, 0, \ldots, 0)\|_{\ell}$. Let $\Pi \in \mathbb{R}^{O(d) \times n(t+1)}$ be a composition of Gaussian embedding and CountSketch embedding (Definition 25.7.4) with $\epsilon=0.1$, and $S=\Pi \widetilde{D}$. We say $S \in \mathbb{R}^{O(d) \times n}$ is a SymSketch.

Modulus of Concentration. Now we give the definition of $\mathrm{mmc}(\ell)$ for a symmetric norm.
Definition 25.3.2 ( $\left.\left[\mathrm{BBC}^{+} 17\right]\right)$. Let $X$ denote the uniform distribution over $\mathbb{S}^{n-1}$. The median of a symmetric norm $\|\cdot\|_{\ell}$ is the unique value $M_{\ell}$ such that $\operatorname{Pr}_{x \sim x}\left[\|x\|_{\ell} \geq M_{\ell}\right] \geq 1 / 2$ and $\operatorname{Pr}_{x \sim x}\left[\|x\|_{\ell} \leq M_{\ell}\right] \geq 1 / 2$.

Definition 25.3.3 ([ $\left.\left.\mathrm{BBC}^{+} 17\right]\right)$. For a given symmetric norm $\|\cdot\|_{\ell}$, we define the modulus of concentration to be $\operatorname{mc}(\ell)=\max _{x \in \mathbb{S}^{n-1}}\|x\|_{\ell} / M_{\ell}$, and define the maximum modulus of concentration to be $\operatorname{mmc}(\ell)=\max _{k \in[n]} \operatorname{mc}\left(\ell^{(k)}\right)$, where $\|\cdot\|_{\ell^{(k)}}$ is a norm on $\mathbb{R}^{k}$ which is defined to be $\left\|\left(x_{1}, x_{2}, \cdots, x_{k}\right)\right\|_{\ell^{(k)}}=\left\|\left(x_{1}, x_{2}, \cdots, x_{k}, 0, \cdots, 0\right)\right\|_{\ell}$.

It has been shown in $\left[\mathrm{BBC}^{+} 17\right]$ that $\operatorname{mmc}(\ell)=\Theta\left(n^{1 / 2-1 / p}\right)$ for $\ell_{p}$ norms when $p>2$, $\operatorname{mmc}(\ell)=\Theta(1)$ for $\ell_{p}$ norms when $p \leq 2, \operatorname{mmc}(\ell)=\widetilde{\Theta}(\sqrt{n / k})$ for top- $k$ norms, and $\operatorname{mmc}(\ell)=O(\log n)$ for the $k$-support norm [AFS12] and the box-norm [MPS14]. We show that $\operatorname{mmc}(\ell)$ is upper bounded by $O(1)$ for max-mix of $\ell_{2}$ norm and $\ell_{1}$ norm and sum-mix of $\ell_{2}$ norm and $\ell_{1}$ norm.

Lemma 25.3.1. For a real number $c>0$, let $\|x\|_{\ell_{a}}=\|x\|_{2}+c\|x\|_{1}$ and $\|x\|_{\ell_{b}}=\max \left\{\|x\|_{2}, c\|x\|_{1}\right\}$. We have $\operatorname{mmc}\left(\ell_{a}\right)=O(1)$ and $\operatorname{mmc}\left(\ell_{b}\right)=O(1)$.

Moreover, we show that for an Orlicz norm $\|\cdot\|_{G}$ induced by a function $G$ which satisfies Assumption 25.1.2, $\mathrm{mmc}(\ell)$ is upper bounded by $O\left(\sqrt{C_{G}} \log n\right)$.

Lemma 25.3.2. For an Orlicz norm $\|\cdot\|_{G}$ on $\mathbb{R}^{n}$ induced by a function $G$ which satisfies Assumption 25.1.2, $\operatorname{mmc}(\ell)$ is upper bounded by $O\left(\sqrt{C_{G}} \log n\right)$.

Subspace Embedding. The following theorem shows that SymSketch is a subspace embedding.

Theorem 25.3.3. Let $S \in \mathbb{R}^{O(d) \times n}$ be a SymSketch as defined in Definition 25.3.1. For a given matrix $A \in \mathbb{R}^{n \times d}$, with probability at least 0.9 , for all $x \in \mathbb{R}^{d}$,

$$
\Omega\left(\frac{1}{\sqrt{d} \log ^{3} n}\right) \cdot\|A x\|_{\ell} \leq\|S A x\|_{2} \leq O\left(d^{2} \operatorname{mmc}(\ell) \cdot \log ^{5 / 2} n\right) \cdot\|A x\|_{\ell}
$$

Furthermore, the running time of computing $S A$ is $\widetilde{O}(\mathrm{nnz}(A)+\operatorname{poly}(d))$.

Combine Theorem 25.3.3 with Lemma 25.3.2, we have the following corollary.

Corollary 25.3.4. Let $\|\cdot\|_{G}$ be an Orlicz norm induced by a function $G$ which satisfies Assumption 25.1.2. Let $S \in \mathbb{R}^{O(d) \times n}$ be a SymSketch as defined in Definition 25.3.1. For a given matrix $A \in \mathbb{R}^{n \times d}$, with probability at least 0.9 , for all $x \in \mathbb{R}^{d}$,

$$
\Omega\left(\frac{1}{\sqrt{d} \log ^{3} n}\right) \cdot\|A x\|_{\ell} \leq\|S A x\|_{2} \leq O\left(\sqrt{C_{G}} d^{2} \log ^{7 / 2} n\right) \cdot\|A x\|_{\ell} .
$$

Furthermore, the running time of computing $S A$ is $\widetilde{O}(\mathrm{nnz}(A)+\operatorname{poly}(d))$.

### 25.4 Conclusion

In this paper, we give efficient algorithms for solving the overconstrained linear regression problem, when the loss function is a symmetric norm. For the special case when the loss function is an Orlicz norm, our algorithm produces a $(1+\epsilon)$-approximate solution in $\widetilde{O}(\mathrm{nnz}(A)+\operatorname{poly}(d / \epsilon))$ time. When the loss function is a general symmetric norm, our algorithm produces a $\sqrt{d} \cdot \operatorname{poly} \log n \cdot \operatorname{mmc}(\ell)$-approximate solution in $\widetilde{O}(\operatorname{nnz}(A)+\operatorname{poly}(d))$ time.

In light of Problem 25.1.1, there are a few interesting problems that remain open. Is that possible to design an algorithm that produces $(1+\epsilon)$-approximate solutions to the linear regression problem, when the loss function is a general symmetric norm? Furthermore, is that possible to use the technique of linear sketching to speed up the overconstrained linear regression problem, when the loss function is a general norm? Answering these problems could lead to a better understanding of Problem 25.1.1.

### 25.5 Preliminaries

Notations. For a matrix $A \in \mathbb{R}^{n \times d}$, we use $A_{i}$ to denote its $i$-th row, $A^{i}$ to denote its $i$-th column, $\|A\|_{F}$ to denote the Frobenius norm of $A$, and $\|A\|_{2}$ to denote the spectral norm of $A$. For any $n^{\prime} \leq n$, we define $\xi^{\left(n^{\prime}\right)} \in \mathbb{R}^{n}$ to be a vector $\xi^{\left(n^{\prime}\right)}=\frac{1}{\sqrt{n^{\prime}}}(1,1, \cdots, 1,0,0, \cdots, 0)$.
$\epsilon$-nets. We use the standard upper bound on size of $\epsilon$-nets.

Definition 25.5.1. For a given set $\mathcal{S}$ and a norm $\|\cdot\|$, we say $\mathcal{N} \subseteq \mathcal{S}$ is a $\epsilon$-net of $\mathcal{S}$ if for any $s \in \mathcal{S}$, there exists some $\bar{s} \in \mathcal{N}$ such that $\|s-\bar{s}\| \leq \epsilon$.

Lemma 25.5.1 ([Woj96, II.E, 10]). Given a matrix $A \in \mathbb{R}^{n \times d}$ and a norm $\|\cdot\|$, let $\mathcal{S}$ be the unit $\|\cdot\|$-norm ball in the column space of $A$, i.e., $\mathcal{S}=\{A x \mid\|A x\|=1\}$. For $\epsilon \in(0,1)$, there exists an $\epsilon$-net $\mathcal{N}$ of $\mathcal{S}$ with size $|\mathcal{N}| \leq(1+1 / \epsilon)^{d}$.

### 25.6 Missing Proofs in Section 25.2

In this section, we give missing proofs in Section 25.2.
We first show that if a function $G$ satisfies Assumption 25.1.2, then $G$ has at least linear growth. We will use this fact in later proofs.

Lemma 25.6.1. Given a function $G$ that satisfies property $\mathcal{P}$, then for any $0<x \leq y$, $y / x \leq G(y) / G(x)$.

Proof. Due to the convexity of $G$ and $G(0)=0$, for any $y>x>0$, we have

$$
G(x) \leq G(y) x / y+G(0)(1-x / y)=G(y) x / y
$$

### 25.6.1 Proof of Lemma 25.2.1

Proof. The first condition is clear from the definition of $\|x\|_{G, w}$.
Now we prove the second condition. When $\|x+y\|_{G, w}=0$, the triangle inequality clearly holds since $\|x\|_{G, w} \geq 0$ and $\|y\|_{G, w} \geq 0$. When $\|x\|_{G, w}=0$ and $\|x+y\|_{G, w} \neq 0$, for any $\alpha>0$, we have

$$
\sum_{i=1}^{n} w_{i} G\left(\left|x_{i}+y_{i}\right| / \alpha\right)=\sum_{i \mid w_{i}>0} w_{i} G\left(\left|x_{i}+y_{i}\right| / \alpha\right)=\sum_{i \mid w_{i}>0} w_{i} G\left(\left|y_{i}\right| / \alpha\right)=\sum_{i=1}^{n} w_{i} G\left(\left|y_{i}\right| / \alpha\right),
$$

which implies $\|x+y\|_{G, w}=\|y\|_{G, w}$. Similarly, the second condition also holds if $\|y\|_{G, w}=0$ and $\|x+y\|_{G, w} \neq 0$. If $\|x+y\|_{G, w} \neq 0,\|x\|_{G, w} \neq 0$ and $\|y\|_{G, w} \neq 0$, by definition of $\|\cdot\|_{G, w}$, we have

$$
\sum_{i=1}^{n} w_{i} G\left(x_{i} /\|x\|_{G, w}\right)=1
$$

and

$$
\sum_{i=1}^{n} w_{i} G\left(y_{i} /\|y\|_{G, w}\right)=1
$$

Thus,

$$
\begin{aligned}
& \sum_{i=1}^{n} w_{i} G\left(\frac{x_{i}+y_{i}}{\|x\|_{G, w}+\|y\|_{G, w}}\right) \\
\leq & \sum_{i=1}^{n} w_{i} G\left(\frac{\left|x_{i}\right|+\left|y_{i}\right|}{\|x\|_{G, w}+\|y\|_{G, w}}\right) \\
\leq & \sum_{i=1}^{n} w_{i}\left(\frac{\|x\|_{G, w}}{\|x\|_{G, w}+\|y\|_{G, w}} \cdot G\left(\left|x_{i}\right| /\|x\|_{G, w}\right)+\frac{\|y\|_{G, w}}{\|x\|_{G, w}+\|y\|_{G, w}} \cdot G\left(\left|y_{i}\right| /\|y\|_{G, w}\right)\right) \\
= & 1
\end{aligned} \quad(G \text { is increasing })
$$

which implies $\|x+y\|_{G, w} \leq\|x\|_{G, w}+\|y\|_{G, w}$.
For the third condition, for any $a \in \mathbb{R}$ and $x \in \mathbb{R}^{n}$, if $\|x\|_{G, w}=0$ then $\|a x\|_{G, w}=0$.
If $a=0$, we have $\|a x\|_{G, w}=0$. Otherwise, we have

$$
\sum_{i=1}^{n} w_{i} G\left(x_{i} /\|x\|_{G, w}\right)=1
$$

which implies

$$
\sum_{i=1}^{n} w_{i} G\left(\frac{a \cdot x_{i}}{|a|\|x\|_{G, w}}\right)=1
$$

and thus $\|a x\|_{G, w}=|a|\|x\|_{G, w}$.

### 25.6.2 Proof of Lemma 25.2.2

Proof. Let $g \in \mathbb{R}^{d}$ be a vector whose entries are i.i.d. Gaussian random variables with zero mean and standard deviation $10^{2}$. We show that with probability at least 0.8 ,

$$
\sum_{i=1}^{n} G\left(\left\|U_{i}\right\|_{2}\right) \leq O\left(\sum_{i=1}^{n} G\left(\left\langle U_{i}, g\right\rangle\right)\right) \leq O\left(\max \left\{1, C_{G}\|U g\|_{G}^{2}\right\}\right) \leq O\left(C_{G} d \kappa_{G}^{2}\right)
$$

We divide our proofs into three parts.

Part I We will show that with probability at least 0.9 ,

$$
\sum_{i=1}^{n} G\left(\left\|U_{i}\right\|_{2}\right) \leq O\left(\sum_{i=1}^{n} G\left(\left\langle U_{i}, g\right\rangle\right)\right)
$$

For each $i \in[n],\left\langle U_{i}, g\right\rangle$ has the same distribution as $10^{2} \cdot\left\|U_{i}\right\|_{2} \cdot \mathcal{N}(0,1)$. For each $i \in[n]$, we let $B_{i}$ be the random variable such that

$$
B_{i}= \begin{cases}1 & \left|\left\langle U_{i}, g\right\rangle\right| \leq\left\|U_{i}\right\|_{2} \\ 0 & \text { otherwise }\end{cases}
$$

By tail inequalities of standard Gaussian random variables, $\operatorname{Pr}\left[B_{i}=1\right] \leq 0.01$. Thus,

$$
\mathbb{E}\left[B_{i} \cdot G\left(\left\|U_{i}\right\|_{2}\right)\right] \leq 0.01 \cdot G\left(\left\|U_{i}\right\|_{2}\right)
$$

which implies

$$
\mathbb{E}\left[\sum_{i=1}^{n} B_{i} \cdot G\left(\left\|U_{i}\right\|_{2}\right)\right] \leq 0.01 \cdot \sum_{i=1}^{n} G\left(\left\|U_{i}\right\|_{2}\right)
$$

By the monotonicity of $G$, since

$$
G\left(\left\langle U_{i}, g\right)\right\rangle \geq\left(1-B_{i}\right) G\left(\left\|U_{i}\right\|_{2}\right)
$$

we have

$$
\sum_{i=1}^{n} G\left(\left\langle U_{i}, g\right\rangle\right) \geq \sum_{i=1}^{n}\left(1-B_{i}\right) G\left(\left\|U_{i}\right\|_{2}\right)
$$

By Markov's inequality, with probability at least 0.9 , we have

$$
\sum_{i=1}^{n} B_{i} \cdot G\left(\left\|U_{i}\right\|_{2}\right) \leq 0.1 \sum_{i=1}^{n} G\left(\left\|U_{i}\right\|_{2}\right)
$$

which implies

$$
\sum_{i=1}^{n} G\left(\left\langle U_{i}, g\right\rangle\right) \geq 0.9 \sum_{i=1}^{n} G\left(\left\|U_{i}\right\|_{2}\right)
$$

Part II We will show that

$$
\sum_{i=1}^{n} G\left(\left\langle U_{i}, g\right\rangle\right) \leq \max \left\{1, C_{G} \cdot\|U g\|_{G}^{2}\right\}
$$

When $\|U g\|_{G} \leq 1$, by monotonicity of $G$, we must have

$$
\sum_{i=1}^{n} G\left(\left\langle U_{i}, g\right\rangle\right) \leq 1
$$

When $\|U g\|_{G} \geq 1$, we have

$$
\sum_{i=1}^{n} G\left(\left\langle U_{i}, g\right\rangle /\|U g\|_{G}\right)=1
$$

Since

$$
G\left(\left\langle U_{i}, g\right\rangle\right) \leq G\left(\left\langle U_{i}, g\right\rangle /\|U g\|_{G}\right) \cdot C_{G}\|U g\|_{G}^{2}
$$

and

$$
\sum_{i=1}^{n} G\left(\left\langle U_{i}, g\right\rangle /\|U g\|_{G}\right)=1
$$

we must have

$$
\sum_{i=1}^{n} G\left(\left\langle U_{i}, g\right\rangle\right) \leq \sum_{i=1}^{n} G\left(\left\langle U_{i}, g\right\rangle /\|U g\|_{G}\right) \cdot C_{G} \cdot\|U g\|_{G}^{2} \leq C_{G} \cdot\|U g\|_{G}^{2}
$$

Part III We will show that $\|U g\|_{G}^{2} \leq O\left(C_{G} d \kappa_{G}^{2}\right)$. By definition of a well-conditioned basis and tail inequalities of Gaussian random variables, with probability at least 0.9 , we have

$$
\|U g\|_{G} \leq \kappa_{G}\|g\|_{2} \leq O\left(\kappa_{G} \sqrt{d}\right)
$$

Thus, applying a union bound over three parts of the proof, we have with probability at least 0.8,

$$
\begin{equation*}
\sum_{i=1}^{n} G\left(\left\|U_{i}\right\|_{2}\right) \leq O\left(C_{G} d \kappa_{G}^{2}\right) \tag{25.9}
\end{equation*}
$$

However, the condition in (25.9) is deterministic. Thus, the condition in (25.9) always holds.

### 25.6.3 Proof of Lemma 25.2.3

Proof. Notice that for any $x \in \mathbb{R}^{d}$,

$$
\left\|A R^{-1} x\right\|_{G} \leq\left\|\Pi A R^{-1} x\right\|_{2}=\kappa\|Q x\|_{2}=\kappa\|x\|_{2}
$$

and

$$
\left\|A R^{-1} x\right\|_{G} \geq \frac{1}{\kappa}\left\|\Pi A R^{-1} x\right\|_{2}=\|Q x\|_{2}=\|x\|_{2} .
$$

### 25.6.4 Proof of Lemma 25.2.4

Proof. In Theorem 2.13 of [Woo14b], it has been shown how to calculate $\left\{l_{i}\right\}_{i=1}^{n}$ such that $l_{i}=\Theta\left(\left\|\left(A R^{-1}\right)_{i}\right\|_{2}\right)$ in $\widetilde{O}(\mathrm{nnz}(A)+\operatorname{poly}(d))$ time with probability at least 0.99 . We simply take $u_{i}=G\left(l_{i}\right)$. By Lemma 25.6.1 and the growth condition of $G$, we must have $u_{i}=$ $\Theta\left(G\left(\left\|\left(A R^{-1}\right)_{i}\right\|_{2}\right)\right)$.

### 25.6.5 Proof of Lemma 25.2.5

Proof. By homogeneity, we only need to prove that with probability $1-\delta$, for all $x$ which satisfies $\|U x\|_{G}=1$,

$$
(1-\epsilon)\|U x\|_{G} \leq\|U x\|_{G, w} \leq(1+\epsilon)\|U x\|_{G} .
$$

We first prove that for any fixed $x \in \mathbb{R}^{d}$ such that $\|U x\|_{G}=1$, with probability $1-\delta(1+4 / \epsilon)^{-d}$,

$$
(1-\epsilon / 4)\|U x\|_{G} \leq\|U x\|_{G, w} \leq(1+\epsilon / 4)\|U x\|_{G} .
$$

Let $x \in \mathbb{R}^{d}$ that satisfies $\|U x\|_{G}=1$ and $y=U x$. Let $Z_{i}$ be a random variable which denotes the value of $w_{i} G\left(y_{i}\right)$ and $Z=\sum_{i=1}^{n} Z_{i}$.

We will first show that if $Z \in[1-\epsilon / 4,1+\epsilon / 4]$, then $\|y\|_{G, w} \in[1-\epsilon / 4,1+\epsilon / 4]$. There are three cases:

1. If $\|y\|_{G, w}=1$, then $\|y\|_{G, w}$ is already in $[1-\epsilon / 4,1+\epsilon / 4]$.
2. If $\|y\|_{G, w}>1$, then by Lemma 25.6.1, we have

$$
\sum_{i=1}^{n} w_{i} G\left(y_{i}\right) \geq \sum_{i=1}^{n} w_{i}\|y\|_{G, w} \cdot G\left(y_{i} /\|y\|_{G, w}\right)
$$

Since

$$
\sum_{i=1}^{n} w_{i} \cdot G\left(y_{i} /\|y\|_{G, w}\right)=1
$$

we must have

$$
\|y\|_{G, w} \leq \sum_{i=1}^{n} w_{i} G\left(y_{i}\right)=Z \leq 1+\epsilon / 4
$$

3. If $\|y\|_{G, w}<1$, then by Lemma 25.6.1, we have

$$
1=\sum_{i=1}^{n} w_{i} G\left(y_{i} /\|y\|_{G, w}\right) \geq 1 /\|y\|_{G, w} \cdot \sum_{i=1}^{n} w_{i} G\left(y_{i}\right)
$$

which implies

$$
\|y\|_{G, w} \geq \sum_{i=1}^{n} w_{i} G\left(y_{i}\right)=Z \geq 1-\epsilon / 4
$$

Thus, it suffices to prove that

$$
\operatorname{Pr}[Z \in[1-\epsilon / 4,1+\epsilon / 4]] \geq 1-\delta(1+4 / \epsilon)^{-d}
$$

Consider the expectation of $Z$, we have

$$
\begin{aligned}
\mathbb{E}[Z] & =\sum_{i=1}^{n} \mathbb{E}\left[Z_{i}\right] \\
& =\sum_{i=1}^{n} \mathbb{E}\left[w_{i}\right] \cdot G\left((U x)_{i}\right) \\
& =\sum_{i=1}^{n} G\left((U x)_{i}\right) \\
& =1,
\end{aligned}
$$

where the last equality follows since $\|U x\|_{G}=1$.
Notice that $\left|Z_{i}-\mathbb{E}\left(Z_{i}\right)\right|$ is always upper bounded by

$$
\begin{aligned}
w_{i} G\left(y_{i}\right) & =w_{i} G\left((U x)_{i}\right) \\
& \leq w_{i} G\left(\left\|U_{i}\right\|_{2} \cdot\|x\|_{2}\right) \\
& \leq w_{i} G\left(\left\|U_{i}\right\|_{2}\right) \\
& \leq G\left(\left\|U_{i}\right\|_{2}\right) / p_{i} \\
& \leq \frac{\epsilon^{2}}{C(\log (1 / \delta)+d \log (1 / \epsilon))},
\end{aligned}
$$

where the first inequality follows from Cauchy-Schwarz inequality, the second inequality follows from the definition of well-conditioned basis in Definition 25.1.2 and monotonicity of $G$, the third inequality follows from definition of $w_{i}$ and the last inequality follows from the choice of $p_{i}$.

Consider the variance of $Z$, we have:

$$
\begin{aligned}
\mathbb{V}(Z) & =\sum_{i \mid p_{i}<1} \mathbb{V}\left(Z_{i}\right) \\
& \leq \sum_{i \mid p_{i}<1} \mathbb{E}\left(Z_{i}^{2}\right) \\
& =\sum_{i \mid p_{i}<1}\left(G\left((U x)_{i}\right)\right)^{2} / p_{i} \\
& \leq\left(\sum_{i \mid p_{i}<1} G\left((U x)_{i}\right)\right) \cdot \max _{i \mid p_{i}<1} G\left((U x)_{i}\right) / p_{i} \\
& \leq \frac{\epsilon^{2}}{C(\log (1 / \delta)+d \log (1 / \epsilon))},
\end{aligned}
$$

where the second inequality follows from Hölder's inequality and the last inequality follows from the upper bound of $G\left((U x)_{i}\right) / p_{i}$ and $\|U x\|_{G}=1$.

Thus, by Bernstein inequality, we have:

$$
\operatorname{Pr}(|Z-1|>\epsilon / 4) \leq(1+4 / \epsilon)^{-d} \delta .
$$

Thus, for a fixed $x$, with probability at least $1-(1+4 / \epsilon)^{-d} \delta$, we have

$$
(1-\epsilon / 4)\|U x\|_{G} \leq\|U x\|_{G, w} \leq(1+\epsilon / 4)\|U x\|_{G}
$$

Let $\mathcal{S}$ be the unit $\|\cdot\|_{G}$-norm ball in the column space of $U$, i.e., $\mathcal{S}=\left\{U x \mid\|U x\|_{G}=1\right\}$. According to Lemma 25.5.1, there exists an $\epsilon / 4$-net $\mathcal{N}$ of $\mathcal{S}$ with $|\mathcal{N}| \leq(1+4 / \epsilon)^{d}$. We use $\mathcal{E}$ to denote the event that for all $y \in \mathcal{N},\|y\|_{G, w} \in[1-\epsilon / 4,1+\epsilon / 4]$. By taking union bound over all vectors in $\mathcal{N}$, we have $\operatorname{Pr}[\mathcal{E}] \geq 1-\delta$.

Conditioned on $\mathcal{E}$, now we show that for all $y \in \mathcal{S},\|y\|_{G, w} \in[1-\epsilon, 1+\epsilon]$. Consider a
fixed vector $y \in \mathcal{S}$, since $\mathcal{N}$ is an $\epsilon / 4$-net of $\mathcal{S}$, we can choose a vector $u^{(1)} \in \mathcal{N}$ such that

$$
\left\|y-u^{(1)}\right\|_{G} \leq \epsilon / 4
$$

Thus, we have that

$$
\begin{aligned}
\|y\|_{G, w} & \leq\left\|u^{(1)}\right\|_{G, w}+\left\|y-u^{(1)}\right\|_{G, w} \\
& \leq(1+\epsilon / 4)+\left\|y-u^{(1)}\right\|_{G, w} .
\end{aligned}
$$

Let $\alpha^{(1)}=1 /\left\|y-u^{(1)}\right\|_{G}$. Then we have $\alpha^{(1)}\left(y-u^{(1)}\right) \in \mathcal{S}$. Thus, there exist $u^{(2)} \in \mathcal{N}$ such that

$$
\left\|u^{(2)}-\alpha^{(1)}\left(y-u^{(1)}\right)\right\|_{G} \leq \epsilon / 4 .
$$

It implies that

$$
\begin{aligned}
\left\|\left(y-u^{(1)}\right)-u^{(2)} / \alpha^{(1)}\right\|_{G} & \leq \epsilon /\left(4 \alpha^{(1)}\right) \\
& \leq(\epsilon / 4)^{2} .
\end{aligned}
$$

Thus,

$$
\begin{aligned}
\left\|y-u^{(1)}\right\|_{G, w} & \leq\left\|u^{(2)}\right\|_{G, w} / \alpha^{(1)}+\left\|y-u^{(1)}-u^{(2)} / \alpha^{(1)}\right\|_{G, w} \\
& \leq(1+\epsilon / 4) \epsilon / 4+\left\|y-u^{(1)}-u^{(2)} / \alpha^{(1)}\right\|_{G, w} .
\end{aligned}
$$

Let $\alpha^{(2)}=1 /\left\|y-u^{(1)}-u^{(2)} / \alpha^{(1)}\right\|_{G}$. Then we can repeat the above argument and get

$$
\begin{aligned}
\|y\|_{G, w} & \leq(1+\epsilon / 4)+(1+\epsilon / 4) \epsilon / 4+(1+\epsilon / 4)(\epsilon / 4)^{2}+\cdots \\
& =(1+\epsilon / 4) /(1-\epsilon / 4) \\
& \leq 1+\epsilon
\end{aligned}
$$

By applying the above upper bound on $\left\|\alpha^{(1)}\left(u^{(1)}-y\right)\right\|_{G, w}$, we can get

$$
\begin{aligned}
\|y\|_{G, w} & \geq\left\|u^{(1)}\right\|_{G, w}-\left\|u^{(1)}-y\right\|_{G, w} \\
& \geq(1-\epsilon / 4)-\left\|u^{(1)}-y\right\|_{G, w} \\
& \geq(1-\epsilon / 4)-\frac{1+\epsilon}{\alpha^{(1)}} \\
& \geq 1-\epsilon / 2-\epsilon^{2} / 4 \\
& \geq 1-\epsilon
\end{aligned}
$$

Thus, conditioned on $\mathcal{E}$, which holds with probability $1-\delta$, we have $\|y\|_{G, w} \in[1-\epsilon, 1+\epsilon]$ for all $y=U x$ with $\|y\|_{G}=1$.

### 25.6.6 Proof of Theorem 25.2.6

Proof. We first analyze the running time of the algorithm. In Step 1, we calculate חAb and invoke QR-decomposition on ПAb. In Step 2, we apply the algorithm in Lemma 25.2.4, which runs in $\widetilde{O}(\operatorname{nnz}(A)+\operatorname{poly}(d))$ time. Obtaining the weight vector $w \in \mathbb{R}^{n}$ in Step 3 requires $O(n)$ time.

Since for all $x \in \mathbb{R}^{d}$,
$\|\mathbf{A b} x\|_{G} \leq\|\Pi \mathbf{A b} x\|_{2} \leq \kappa\|\mathbf{A b} x\|_{G}$.
we have

$$
\begin{align*}
\mathbb{E}\left[\|w\|_{0}\right] & =\sum_{i=1}^{n} p_{i}=\sum_{i=1}^{n} O\left(d \log (1 / \epsilon) / \epsilon^{2} \cdot G\left(\left\|\left(\mathbf{A b} R^{-1}\right)_{i}\right\|_{2}\right)\right) \\
& \leq O\left(d \log (1 / \epsilon) / \epsilon^{2} \cdot C_{G} d\left(\kappa_{G}\left(\mathbf{A b} R^{-1}\right)\right)^{2}\right)  \tag{Lemma25.2.2}\\
& \leq O\left(C_{G} d^{2} \kappa^{2} \log (1 / \epsilon) / \epsilon^{2}\right) \tag{Lemma25.2.3}
\end{align*}
$$

By Markov's inequality, with constant probability we have $\|w\|_{0} \leq O\left(C_{G} d^{2} \kappa^{2} \log (1 / \epsilon) / \epsilon^{2}\right)$. Moreover, in order to solve $\min _{x}\|A x-b\|_{G, w}$, we can ignore all rows of $A$ with zero weights, and thus there are at most $O\left(C_{G} d^{2} \kappa^{2} \log (1 / \epsilon) / \epsilon^{2}\right)$ remaining rows in $A$. Furthermore, as we show in Lemma 25.2.1, $\|\cdot\|_{G, w}$ is a seminorm, which implies we can solve $\min _{x}\|A x-b\|_{G, w}$ in poly $\left(C_{G} d \kappa / \epsilon\right)$ time, by simply solve a convex program with size $O\left(C_{G} d^{2} \kappa^{2} \log (1 / \epsilon) / \epsilon^{2}\right)$.

Now we prove the correctness of the algorithm. The algorithm in Lemma 25.2.4 succeeds with constant probability. By Lemma 25.2.5, with constant probability, simultaneously for all $x \in \mathbb{R}^{d+1}$,

$$
(1-\epsilon / 3)\left\|\mathbf{A} \mathbf{b} R^{-1} x\right\|_{G} \leq\left\|\mathbf{A} \mathbf{b} R^{-1} x\right\|_{G, w} \leq(1+\epsilon / 3)\left\|\mathbf{A} \mathbf{b} R^{-1} x\right\|_{G} .
$$

Equivalently, with constant probability, simultaneously for all $x \in \mathbb{R}^{d+1}$,

$$
(1-\epsilon / 3)\|\mathbf{A} \mathbf{b} x\|_{G} \leq\|\mathbf{A} \mathbf{b} x\|_{G, w} \leq(1+\epsilon / 3)\|\mathbf{A} \mathbf{b} x\|_{G}
$$

Since $x^{*}=\operatorname{argmin}_{x}\|A x-b\|_{G, w}$, for all $x \in \mathbb{R}^{d}$, we have

$$
\begin{aligned}
\left\|A x^{*}-b\right\|_{G} & \leq 1 /(1-\epsilon / 3)\left\|A x^{*}-b\right\|_{G, w} \\
& \leq 1 /(1-\epsilon / 3)\|A x-b\|_{G, w} \\
& \leq(1+\epsilon / 3) /(1-\epsilon / 3)\|A x-b\|_{G} \\
& \leq(1+\epsilon)\|A x-b\|_{G}
\end{aligned}
$$

for sufficiently small $\epsilon$. Thus, $x^{*}$ is a $(1+\epsilon)$-approximate solution to $\min _{x}\|A x-b\|_{G}$.
Note that the failure probability of the algorithm can be reduced to an arbitrarily small constant by independent repetitions and taking the best solution found among all repetitions.

### 25.7 Missing Proofs in Section 25.3

In this section, we give missing proofs in Section 25.3.
Without loss of generality, throughout this section, for the symmetric norm $\|\cdot\|_{\ell}$ under consideration, we assume $\left\|\xi^{(1)}\right\|_{\ell}=1$.

### 25.7.1 Background

### 25.7.1.1 Known $\ell_{2}$ Oblivious Subspace Embeddings

In this section, we recall some known $\ell_{2}$ subspace embeddings.
Definition 25.7.1. We say $S \in \mathbb{R}^{t \times n}$ is an $\ell_{2}$ subspace embedding for the column space of $A \in \mathbb{R}^{n \times d}$, if for all $x \in \mathbb{R}^{d}$,

$$
(1-\epsilon)\|A x\|_{2} \leq\|S A x\|_{2} \leq(1+\epsilon)\|A x\|_{2}
$$

Definition 25.7.2. A CountSketch embedding is defined to be $\Pi=\Phi D \in \mathbb{R}^{m \times n}$ with $m=\Theta\left(d^{2} / \epsilon^{2}\right)$, where $D$ is an $n \times n$ random diagonal matrix with each diagonal entry independently chosen to be +1 or -1 with equal probability, and $\Phi \in\{0,1\}^{m \times n}$ is an $m \times n$ binary matrix with $\Phi_{h(i), i}=1$ and all remaining entries being 0 , where $h:[n] \rightarrow[m]$ is a random map such that for each $i \in[n], h(i)=j$ with probability $1 / m$ for each $j \in[m]$.

Theorem 25.7.1 ([CW13]). For a given matrix $A \in \mathbb{R}^{n \times d}$ and $\epsilon \in(0,1 / 2)$. Let $\Pi \in$ $\mathbb{R}^{\Theta\left(d^{2} / \epsilon^{2}\right) \times n}$ be a CountSketch embedding. With probability at least 0.9999 , $\Pi$ is an $\ell_{2}$ subspace embedding for the column space of $A$. Furthermore, $\Pi A$ can be computed in $O(\mathrm{nnz}(A))$ time. Definition 25.7.3. A Gaussian embedding $S$ is defined to be $\frac{1}{\sqrt{m}} \cdot G \in \mathbb{R}^{m \times n}$ with $m=$ $\Theta\left(d / \epsilon^{2}\right)$, where each entry of $G \in \mathbb{R}^{m \times n}$ is chosen independently from the standard Gaussian distribution.

Theorem 25.7.2 ([Woo14b]). For a given matrix $A \in \mathbb{R}^{n \times d}$ and $\epsilon \in(0,1 / 2)$. Let $S \in$ $\mathbb{R}^{\Theta\left(d / \epsilon^{2}\right) \times n}$ be a Gaussian embedding. With probability at least 0.9999, $S$ is an $\ell_{2}$ subspace embedding for the column space of $A$.

Definition 25.7.4. A composition of Gaussian embedding and CountSketch embedding is defined to be $S^{\prime}=S \Pi$, where $\Pi \in \mathbb{R}^{\Theta\left(d^{2} / \epsilon^{2}\right) \times n}$ is a CountSketch embedding and $S \in$ $\mathbb{R}^{\Theta\left(d / \epsilon^{2}\right) \times \Theta\left(d^{2} / \epsilon^{2}\right)}$ is a Gaussian embedding.

The following corollary directly follows from the above two theorems.
Corollary 25.7.3. For a given matrix $A \in \mathbb{R}^{n \times d}$ and $\epsilon \in(0,1 / 2)$. Let $S^{\prime} \in \mathbb{R}^{\Theta\left(d / \epsilon^{2}\right) \times n}$ be a composition of Gaussian embedding and CountSketch embedding. With probability at least 0.9998, $S^{\prime}$ is an $\ell_{2}$ subspace embedding for the column space of $A$. Furthermore, $S^{\prime} A$ can be computed in $O\left(\operatorname{nnz}(A)+d^{4} / \epsilon^{4}\right)$ time.

We remark that all $\ell_{2}$ subspace embeddings introduced in this section are oblivious, meaning that the distribution of the embedding matrix does not depend on the matrix $A$.

### 25.7.1.2 Properties of Symmetric Norms

General Properties. We first introduce several general properties of symmetric norms.

Lemma 25.7.4 (Lemma 2.1 in $\left.\left[\mathrm{BBC}^{+} 17\right]\right)$. For any symmetric norm $\|\cdot\|_{\ell}$ and $x, y \in \mathbb{R}^{n}$ such that for all $i \in[n]$ we have $\left|x_{i}\right| \leq\left|y_{i}\right|$, then $\|x\|_{\ell} \leq\|y\|_{\ell}$.

Lemma 25.7.5 (Fact 2.2 in $\left.\left[\mathrm{BBC}^{+} 17\right]\right)$. Suppose $\left\|\xi^{(1)}\right\|_{\ell}=1$, for any vector $x \in \mathbb{R}^{n}$,

$$
\|x\|_{\infty} \leq\|x\|_{\ell} \leq\|x\|_{1} .
$$

Lemma 25.7.6 (Lemma 3.12 in $\left.\left[\mathrm{BBC}^{+} 17\right]\right)$. Let $\|\cdot\|_{\ell}$ be a symmetric norm. Then

$$
\Omega\left(M_{\ell} / \sqrt{\log n}\right) \leq\left\|\xi^{(n)}\right\|_{\ell} \leq O\left(M_{\ell}\right)
$$

where $M_{\ell}$ is as defined in Definition 25.3.2.

Modulus of Approximation. We need the following quantity of a symmetric norm.
Definition 25.7.5. The maximum modulus of approximation of a symmetric norm $\|\cdot\|_{\ell}$ is defined as

$$
\operatorname{mma}(\ell, r)=\max _{1 \leq a \leq b \leq a r \leq n} \frac{M_{\ell^{(a r)}}}{M_{\ell^{(b)}}}
$$

where $\|\cdot\|_{\ell^{(k)}}$ is a norm on $\mathbb{R}^{k}$ which is defined to be

$$
\left\|\left(x_{1}, x_{2}, \cdots, x_{k}\right)\right\|_{\ell(k)}=\left\|\left(x_{1}, x_{2}, \cdots, x_{k}, 0, \cdots, 0\right)\right\|_{\ell},
$$

and $M_{\ell^{(k)}}$ is as defined in Definition 25.3.2.

Intuitively, mma $(\ell, r)$ characterizes how well the original symmetric norm can be approximated by a lower dimensional induced norm. We show in the following lemma that $\operatorname{mma}(\ell, r) \leq O(\sqrt{r \log n})$ for any symmetric norm.

Lemma 25.7.7. For any symmetric norm $\|\cdot\|_{\ell}$ and $r \in[n]$, $\mathrm{mma}(\ell, r) \leq O(\sqrt{r \log n})$.

Proof. By Lemma 25.7.6, for any $i \in[n], \Omega\left(M_{\ell^{(i)}} / \sqrt{\log n}\right) \leq\left\|\xi^{(i)}\right\|_{\ell} \leq O\left(M_{\ell^{(i)}}\right)$. Let ar $=$ $c_{1} b+c_{2}$, where $c_{1}, c_{2}$ are non-negative integers with $c_{1} \leq a r / b$ and $c_{2} \leq b$. Observe that we can rewrite $\xi^{(a r)}$ as

$$
\xi^{(a r)}=[\frac{\sqrt{b}}{\sqrt{a r}} \cdot(\underbrace{\xi^{(b)}, \xi^{(b)}, \xi^{(b)}, \ldots, \xi^{(b)}}_{c_{1} \text { times }}), \frac{\sqrt{c_{2}}}{\sqrt{a r}} \xi^{\left(c_{2}\right)}, 0, \ldots, 0] .
$$

Therefore, by triangle inequality, we have

$$
\begin{array}{rlr}
\left\|\xi^{(a r)}\right\|_{\ell} & \leq \frac{\sqrt{b}}{\sqrt{a r}} \cdot c_{1} \cdot\left\|\xi^{(b)}\right\|_{\ell}+\frac{\sqrt{c_{2}}}{\sqrt{a r}} \cdot\left\|\xi^{\left(c_{2}\right)}\right\|_{\ell} & \\
& \leq \sqrt{\frac{b}{a r}} \cdot \frac{a r}{b} \cdot\left\|\xi^{(b)}\right\|_{\ell}+\frac{\sqrt{b}}{\sqrt{a r}} \cdot\left\|\xi^{(b)}\right\|_{\ell} & \left(c_{1} \leq a r / b \text { and } c_{2} \leq b\right) \\
& \leq \sqrt{r} \cdot\left\|\xi^{(b)}\right\|_{\ell}+\left\|\xi^{(b)}\right\|_{\ell} & (a \leq b \leq a r) \\
& \leq 2 \sqrt{r} \cdot\left\|\xi^{(b)}\right\|_{\ell} & (r \geq 1)
\end{array}
$$

Now we apply Lemma 25.7.6 on both sides, which implies

$$
\frac{M_{\ell^{(a r)}}}{\sqrt{\log n}} \leq O\left(\sqrt{r} \cdot M_{\left.\ell^{(b)}\right)}\right)
$$

as desired.

Properties of SymSketch. Now we introduce several properties of SymSketch.
The following lemma shows that for a data matrix $A \in \mathbb{R}^{d}$, calculating $S A$ requires $\widetilde{O}(\mathrm{nnz}(A)+\operatorname{poly}(d))$ time for a SymSketch $S$.

Lemma 25.7.8. For a given matrix $A \in \mathbb{R}^{n \times d}$, let $S \in \mathbb{R}^{O(d) \times n}$ be a SymSketch as in Definition 25.3.1. $S A$ can be computed in $O(\mathrm{nnz}(A))+\operatorname{poly}(d)$ time in expectation, and in $O(\operatorname{nnz}(A) \log n)+\operatorname{poly}(d)$ time in the worst case.

Proof. Since $S$ is a SymSketch, $S=\Pi \widetilde{D}=\Pi \cdot\left[\begin{array}{c}w_{0} D_{0} \\ w_{1} D_{1} \\ \vdots \\ w_{t} D_{t}\end{array}\right]$, where $\Pi \in \mathbb{R}^{O(d) \times O(n \log n)}$. Since $D_{i}$ is a diagonal matrix, $\mathrm{nnz}\left(D_{i} A\right) \leq \mathrm{nnz}(A)$, and thus $\mathrm{nnz}(\widetilde{D} A) \leq(t+1) \cdot \mathrm{nnz}(A)=$ $O(\operatorname{nnz}(A) \log n)$, which implies $\widetilde{D} A$ can be computed in $(t+1) \cdot \operatorname{nnz}(A)=O(\operatorname{nnz}(A) \log n)$ time.

On the other hand, the expected number of nonzero entries of $D_{i} A$ is $2^{-i} \mathrm{nnz}(A)$. Thus, $\widetilde{D} A$ has $O(\mathrm{nnz}(A))$ nonzero entries in expectation, which implies $\widetilde{D} A$ can be computed in $O(\mathrm{nnz}(A))$ time.

Finally, notice that $\Pi$ is a composition of Gaussian embedding and CountSketch embedding, which implies $\Pi \widetilde{D} A$ can be computed in $n n z(\widetilde{D} A)+\operatorname{poly}(d)$ time.

The following lemma shows that with constant probability, for all $x \in \mathbb{R}^{n},\|S x\|_{2} \leq$ $\operatorname{poly}(n)\|x\|_{2}$.

Lemma 25.7.9. Let $S \in \mathbb{R}^{O(d) \times n}$ be a SymSketch as defined in Definition 25.3.1, then with probability at least $0.9999\|S\|_{2} \leq \operatorname{poly}(n)$.

Proof. Notice that $S=\Pi \widetilde{D}$, since $\|S\|_{2} \leq\|\Pi\|_{2} \cdot\|\widetilde{D}\|_{2}$, it suffices to bound $\|\Pi\|_{2}$ and $\|\widetilde{D}\|_{2}$. Since $\Pi$ is a composition of Gaussian embedding and CountSketch embedding (Definition 25.7.4), with probability at least $0.9999,\|\Pi\|_{2} \leq\|\Pi\|_{F} \leq \operatorname{poly}(n)$. Now consider $\widetilde{D}=\left[\begin{array}{c}w_{0} D_{0} \\ w_{1} D_{1} \\ \vdots \\ w_{t} D_{t}\end{array}\right]$. By Lemma 25.7.5, for all $j \in[t], w_{j} \leq \operatorname{poly}(n)$. Furthermore, $\left\|D_{j}\right\|_{2} \leq 1$ and $t=\Theta(\log n)$, which implies $\|\widetilde{D}\|_{2} \leq \operatorname{poly}(n)$.

Throughout this whole section we assume that for any non-zero vector $x \in \mathbb{R}^{n}$, we have $1 \leq\left|x_{j}\right| \leq \operatorname{poly}(n)$ for all $j \in[n]$. Notice that this assumption is without loss of generality, as shown in the following lemma.

Lemma 25.7.10. For any non-zero vector $x \in \mathbb{R}^{n}$, let $\bar{x} \in \mathbb{R}^{n}$ be a vector with $\bar{x}=\frac{\operatorname{poly}(n) \cdot x}{\|x\|_{\infty}}$, and $x^{\prime} \in \mathbb{R}^{n}$ where

$$
x_{i}^{\prime}=\left\{\begin{array}{ll}
\bar{x}_{i} & \text { if } \bar{x}_{i} \geq 1 \\
0 & \text { otherwise }
\end{array} .\right.
$$

For a symmetric norm $\|\cdot\|_{\ell}$, suppose $\|S\|_{2} \leq \operatorname{poly}(n)$ and

$$
\alpha\left\|x^{\prime}\right\|_{\ell} \leq\left\|S x^{\prime}\right\|_{2} \leq \beta\left\|x^{\prime}\right\|_{\ell}
$$

for some $\alpha, \beta \in[1 / \operatorname{poly}(n), \operatorname{poly}(n)]$, then

$$
\Omega(\alpha)\|x\|_{\ell} \leq\|S x\|_{2} \leq O(\beta)\|x\|_{\ell} .
$$

Proof. By triangle inequality and Lemma 25.7.4, we have

$$
\|\bar{x}\|_{\ell}-n \leq\left\|x^{\prime}\right\|_{\ell} \leq\|\bar{x}\|_{\ell} .
$$

By Lemma 25.7.5,

$$
\left\|x^{\prime}\right\|_{\ell} \geq\left\|x^{\prime}\right\|_{\infty}=\|\bar{x}\|_{\infty}=\operatorname{poly}(n)
$$

we have

$$
(1-1 / \operatorname{poly}(n))\|\bar{x}\|_{\ell} \leq\left\|x^{\prime}\right\|_{\ell} \leq\|\bar{x}\|_{\ell} .
$$

Notice that $\|S \bar{x}\|_{2}=\left\|S x^{\prime}+S\left(\bar{x}-x^{\prime}\right)\right\|_{2}$. By triangle inequality we have

$$
\left\|S x^{\prime}\right\|_{2}-\|S\|_{2}\left\|\bar{x}-x^{\prime}\right\|_{2} \leq\|S \bar{x}\|_{2} \leq\left\|S x^{\prime}\right\|_{2}+\|S\|_{2}\left\|\bar{x}-x^{\prime}\right\|_{2}
$$

By the given conditions, we have

$$
(1-1 / \operatorname{poly}(n))\left\|S x^{\prime}\right\|_{2} \leq\|S \bar{x}\|_{2} \leq(1+1 / \operatorname{poly}(n))\left\|S x^{\prime}\right\|_{2},
$$

which implies

$$
\Omega(\alpha)\|\bar{x}\|_{\ell} \leq\|S \bar{x}\|_{2} \leq O(\beta)\|\bar{x}\|_{\ell} .
$$

Since $\bar{x}=\frac{\operatorname{poly}(n) \cdot x}{\|x\|_{\infty}}$, we have

$$
\Omega(\alpha)\|x\|_{\ell} \leq\|S x\|_{2} \leq O(\beta)\|x\|_{\ell} .
$$

By Lemma 25.7.10, we can focus on those non-zero vectors $x \in \mathbb{R}^{n}$ such that $1 \leq$ $\left|x_{j}\right| \leq \operatorname{poly}(n)$ for all $j \in[n]$.

Definition 25.7.6. For a given vector $x \in \mathbb{R}^{n}$, suppose for all $j \in[n]$,

$$
1 \leq\left|x_{j}\right| \leq \operatorname{poly}(n)
$$

Let $g=\Theta(\log n)$. For each $i \in\{0,1, \ldots, g\}$, we define

$$
L_{i}(x)=\left\{j\left|2^{i} \leq\left|x_{j}\right|<2^{i+1}\right\} .\right.
$$

For each $i \in\{0,1, \ldots, g\}$, we define $V_{i}(x) \in \mathbb{R}^{n}$ to be the vector

$$
V_{i}(x)=(\underbrace{2^{i}, 2^{i}, \ldots, 2^{i}}_{\left|L_{i}(x)\right|}, 0, \ldots, 0)
$$

For each $i \in\{0,1, \ldots, g\}$, we say a level $i$ to be contributing if

$$
\left\|V_{i}(x)\right\|_{\ell} \geq \Omega(1 / g) \cdot\|x\|_{\ell}
$$

Lemma 25.7.11. Let $g=\Theta(\log n)$. For a given vector $x \in \mathbb{R}^{n}$ such that for all $j \in[n], 1 \leq$ $\left|x_{j}\right| \leq 2^{g}$, there exists at least one level $i \in\{0,1, \cdots, g\}$ which is contributing.

Proof. If none of $i \in\{0,1, \cdots, g\}$ is contributing, then $\|x\|_{\ell} \leq \sum_{i=0}^{g}\left\|V_{i}(x)\right\|_{\ell} \leq 1 /(2 g)$. $\sum_{i=0}^{g}\|x\|_{\ell} \leq \frac{1}{2}\|x\|_{\ell}$, which leads to a contradiction.

### 25.7.2 Proof of Lemma 25.3.1

Proof. Consider a fixed $n^{\prime} \in[n]$. By Lemma 25.7.6, we have

$$
M_{\ell_{a}^{\left(n^{\prime}\right)}}=\Omega\left(\left\|\xi^{\left(n^{\prime}\right)}\right\|_{\ell_{a}^{\left(n^{\prime}\right)}}\right)=\Omega\left(1+c \sqrt{n^{\prime}}\right)
$$

and

$$
M_{\ell_{b}^{\left(n^{\prime}\right)}}=\Omega\left(\left\|\xi^{\left(n^{\prime}\right)}\right\|_{\ell_{b}^{\left(n^{\prime}\right)}}\right)=\Omega\left(\max \left(1, c \sqrt{n^{\prime}}\right)\right) .
$$

It is also straightforward to verify that

$$
\max _{x \in \mathbb{S}^{\prime}-1}\|x\|_{\ell_{a}}=1+c \sqrt{n^{\prime}}
$$

and

$$
\max _{x \in \mathbb{S}^{\prime}-1}\|x\|_{\ell_{b}}=\max \left(1, c \sqrt{n^{\prime}}\right)
$$

Taking the ratio between $\max _{x \in \mathbb{S}^{n^{\prime}-1}}\|x\|_{\ell}$ and $M_{\ell\left(n^{\prime}\right)}$ for $\ell \in\left\{\ell_{a}, \ell_{b}\right\}$, we complete the proof.

### 25.7.3 Proof of Lemma 25.3.2

Proof. Let $\bar{G}(x)=\sqrt{x} \cdot G^{-1}(1 / x)$, where $G^{-1}(1 / x)$ is the unique value in $[0, \infty)$ such that $G\left(G^{-1}(1 / x)\right)=1 / x$. We first show that $\bar{G}(x)$ is an approximately decreasing function for $x \in(0, \infty)$. Let $m, n$ be two real numbers with $0<m \leq n$. We have $1 / n \leq 1 / m$, which implies $0<G^{-1}(1 / n) \leq G^{-1}(1 / m)$ by monotonicity of $G$. By the third condition in Assumption 25.1.2, we have

$$
\frac{G\left(G^{-1}(1 / m)\right)}{G\left(G^{-1}(1 / n)\right)} \leq C_{G} \cdot\left(\frac{G^{-1}(1 / m)}{G^{-1}(1 / n)}\right)^{2}
$$

which implies

$$
\sqrt{n} \cdot G^{-1}(1 / n) \leq \sqrt{C_{G}} \cdot \sqrt{m} \cdot G^{-1}(1 / m)
$$

Hence $\bar{G}(n) \leq \sqrt{C_{G}} \cdot \bar{G}(m)$.
We are now ready to prove the lemma. Recall that for the Orlicz norm $\|\cdot\|_{\ell}=\|\cdot\|_{G}$, we have

$$
\operatorname{mmc}(\ell)=\max _{n^{\prime} \in[n]} \operatorname{mc}\left(\ell^{\left(n^{\prime}\right)}\right)=\max _{n^{\prime} \in[n]} \frac{\max _{x \in \mathbb{S}^{n^{\prime}-1}}\|x\|_{\ell\left(n^{\prime}\right)}}{M_{\ell\left(n^{\prime}\right)}}
$$

By Lemma 25.7.6, we have,

$$
\Omega(1) \cdot\left\|\xi^{\left(n^{\prime}\right)}\right\|_{\ell\left(n^{\prime}\right)} \leq M_{\ell\left(n^{\prime}\right)} \leq O(\sqrt{\log n}) \cdot\left\|\xi^{\left(n^{\prime}\right)}\right\|_{\ell^{\left(n^{\prime}\right)}}
$$

Thus $\left\|\xi^{\left(n^{\prime}\right)}\right\|_{\ell\left(n^{\prime}\right)}$ provides an approximation to $M_{\ell\left(n^{\prime}\right)}$. By definition of $\|\cdot\|_{G}$, we have

$$
\left\|\xi^{\left(n^{\prime}\right)}\right\|_{\ell\left(n^{\prime}\right)}=\frac{1}{\sqrt{n^{\prime}} \cdot G^{-1}\left(1 / n^{\prime}\right)}=\frac{1}{\bar{G}\left(n^{\prime}\right)} .
$$

Hence

$$
\Omega(1) \leq M_{\ell\left(n^{\prime}\right)} \cdot \bar{G}\left(n^{\prime}\right) \leq O(\sqrt{\log n})
$$

Next, we compute $\max _{x \in \mathbb{S}^{n^{\prime}-1}}\|x\|_{\ell\left(n^{\prime}\right)}$. For an arbitrary unit vector $x \in \mathbb{S}^{n^{\prime}-1}$, we denote

$$
B_{j}=\left\{i \in[n]:\left|x_{i}\right| \in\left[1 / 2^{j}, 1 / 2^{j-1}\right)\right\}
$$

and $b_{j}=\left|B_{j}\right|$. For each $j$, let $x^{B_{j}} \in \mathbb{R}^{n}$ be the vector such that

$$
x_{i}^{B_{j}}= \begin{cases}x_{i} & \text { if } j \in B_{j} \\ 0 & \text { otherwise }\end{cases}
$$

Note that non-zero coordinates in $x^{B_{j}}$ have magnitude close to each other (within a factor
of 2 ), we thus have

$$
\begin{aligned}
\left\|x^{B_{j}}\right\|_{\ell} & =\left\|x^{B_{j}}\right\|_{2} \cdot\left\|\frac{x^{B_{j}}}{\left\|x^{B_{j}}\right\|_{2}}\right\|_{\ell} \\
& \leq \frac{\sqrt{b_{j}}}{2^{j-1}} \cdot\left\|\frac{x^{B_{j}}}{\left\|x^{B_{j}}\right\|_{2}}\right\|_{\ell} \\
& \leq \frac{\sqrt{b_{j}}}{2^{j-2}} \cdot\left\|\xi^{\left(b_{j}\right)}\right\|_{\ell^{\left(b_{j}\right)}} \\
& =\frac{\sqrt{b_{j}}}{2^{j-2}} \cdot \frac{1}{\bar{G}\left(b_{j}\right)} .
\end{aligned}
$$

Similarly,

$$
\begin{aligned}
\left\|x^{B_{j}}\right\|_{\ell} & \geq \frac{\sqrt{b_{j}}}{2^{j+2}} \cdot\left\|\xi^{\left(b_{j}\right)}\right\|_{\ell^{\left(b_{j}\right)}} \\
& =\frac{\sqrt{b_{j}}}{2^{j+2}} \cdot \frac{1}{\bar{G}\left(b_{j}\right)} \\
& \geq \frac{\sqrt{b_{j}}}{2^{j+2}} \cdot \frac{1}{\sqrt{C_{G}} \cdot \bar{G}(1)} .
\end{aligned}
$$

We claim there exists an constant $c>0$ such that

$$
\sum_{\substack{j>c \log n \\ b_{j}>0}}\left\|x^{B_{j}}\right\|_{\ell} \leq \sum_{j^{\prime} \leq c \log n}\left\|x^{B^{j^{\prime}}}\right\|_{\ell} .
$$

To show this, by Lemma 25.6.1, for any $b \geq 1$,

$$
b=\frac{G\left(G^{-1}(1)\right)}{G\left(G^{-1}(1 / b)\right)} \geq \frac{G^{-1}(1)}{G^{-1}(1 / b)}
$$

which implies

$$
G^{-1}(1 / b) \geq \frac{G^{-1}(1)}{b}
$$

Next, since $\|x\|_{2}=1$, there exists an $0 \leq \widetilde{j} \leq 4 \log n$ such that $b_{j} \geq 1$. Therefore,

$$
\left\|x^{B_{\tilde{j}}}\right\|_{\ell} \geq \frac{1}{2^{\tilde{j}+2}} \cdot \frac{1}{\sqrt{C_{G}} \cdot \bar{G}(1)}
$$

Thus, we have

$$
\begin{aligned}
\sum_{\substack{j>c \log n \\
b_{j}>0}}\left\|x^{B_{j}}\right\|_{\ell} & =\sum_{\substack{j>c \log n \\
b_{j}>0}}\left\|x^{B_{j}}\right\|_{2} \cdot\left\|\frac{x^{B_{j}}}{\left\|x^{B_{j}}\right\|_{2}}\right\|_{\ell} \\
& \leq \sum_{\substack{j>c \log n \\
b_{j}>0}} \frac{\sqrt{n}}{2^{j-2}} \cdot\left\|\xi^{\left(b_{j}\right)}\right\|_{\ell^{\left(b_{j}\right)}} \\
& \leq \sum_{\substack{j>c \log n \\
b_{j}>0}} \frac{\sqrt{n}}{2^{j-2}} \cdot \frac{1}{\sqrt{b_{j}} G^{-1}\left(1 / b_{j}\right)} \\
& \leq n \cdot \frac{\sqrt{n}}{2^{c \log n-2}} \cdot \frac{b_{j}}{\sqrt{b_{j}} G^{-1}(1)} \\
& \leq n \cdot 2^{\tilde{j}+2} \cdot \frac{\sqrt{n}}{2^{c \log n-2}} \cdot \sqrt{C_{G} n} \cdot\left\|x^{B_{\tilde{j}}}\right\|_{\ell} \\
& \leq\left\|x^{B_{\tilde{j}}}\right\|_{\ell} \\
& \leq \sum_{j^{\prime} \leq c \log n}\left\|x^{B^{j^{\prime}}}\right\|_{\ell}
\end{aligned}
$$

for some sufficiently large constant $c$.
Let

$$
j^{*}=\operatorname{argmax}_{j \leq c \log n}\left\|x^{B_{j}}\right\|_{\ell},
$$

we have

$$
\begin{aligned}
\left\|x^{B_{j^{*}}}\right\|_{\ell} \leq\|x\|_{\ell} & \leq \sum_{j \leq c \log n}\left\|x^{B_{j}}\right\|_{\ell}+\sum_{\substack{j>c \log n \\
b_{j}>0}}\left\|x^{B_{j}}\right\|_{\ell} \\
& \leq 2 \sum_{j \leq c \log n}\left\|x^{B_{j}}\right\|_{\ell} \\
& \leq O(\log n) \cdot\left\|x^{B_{j^{*}}}\right\|_{\ell} .
\end{aligned}
$$

Thus,

$$
\begin{aligned}
\max _{x \in \mathbb{S}^{\prime}-1}\|x\|_{\ell\left(n^{\prime}\right)} & \leq O\left(\log n^{\prime}\right) \max _{x \in \mathbb{S}^{\prime}-1}\left\|x^{B_{j^{*}}}\right\|_{\ell} \\
& \leq O\left(\log n^{\prime}\right) \max _{x \in \mathbb{S}^{\prime}-1}\left\|x^{B_{j^{*}}}\right\|_{2} \cdot\left\|\frac{x^{B_{j^{*}}}}{\left\|x^{B_{j^{*}}}\right\|_{2}}\right\|_{\ell} \\
& \leq O\left(\log n^{\prime}\right) \max _{x \in \mathbb{S}^{\prime}-1}\left\|\frac{x^{B_{j}}}{\left\|x^{B_{j}}\right\|_{2}}\right\|_{\ell} \\
& \leq O\left(\log n^{\prime}\right) \max _{b_{j *} \leq n^{\prime}}\left\|\xi^{\left.b_{j}\right)}\right\|_{\ell^{\left(b_{j}\right)}} \\
& \leq O\left(\log n^{\prime}\right) \max _{b_{j *} \leq n^{\prime}} \frac{1}{\bar{G}\left(b_{j^{*}}\right)} \\
& \leq \frac{O\left(\sqrt{C_{G}} \log n^{\prime}\right)}{\bar{G}\left(n^{\prime}\right)} .
\end{aligned}
$$

Thus, we have

$$
\begin{aligned}
\operatorname{mmc}(\ell) & =\max _{n^{\prime} \in[n]} \frac{\max _{x \in \mathbb{S}^{n^{\prime}-1}}\|x\|_{\ell\left(n^{\prime}\right)}}{M_{\ell\left(n^{\prime}\right)}} \\
& \leq O\left(\sqrt{C_{G}} \log n\right) .
\end{aligned}
$$

### 25.7.4 Contraction Bound of SymSketch

In this section we give the contraction bound of SymSketch. We first show that for a fixed vector $x \in \mathbb{R}^{n},\|\widetilde{D} x\|_{2} \geq 1 / \operatorname{poly}(d \log n) \cdot\|x\|_{\ell}$ with probability $1-2^{-\Theta(d \log n)}$.

Lemma 25.7.12. Let $\widetilde{D}$ be the matrix defined in Definition 25.3.1. For any fixed $x \in \mathbb{R}^{n}$, with probability $1-2^{-\Theta(d \log n)}$,

$$
\|\widetilde{D} x\|_{2} \geq 1 / \alpha_{0} \cdot\|x\|_{\ell}
$$

where $\alpha_{0}=O\left(\operatorname{mma}(\ell, d) \cdot \log ^{5 / 2} n\right)=O\left(\sqrt{d} \log ^{3} n\right)$.

Proof. The lemma follows from the following two claims. Recall that $t=\Theta(\log n)$.
Claim 25.7.13. For any fixed $x \in \mathbb{R}^{n}$. If there is a contributing level $i^{*} \in\{0,1,2, \ldots, g\}$ such that $\left|L_{i^{*}}(x)\right|=\Theta(d \log n) \cdot 2^{j}$ for some $j \in[t]$, then with probability at least $1-2^{-\Theta(d \log n)}$,

$$
\left\|w_{j} D_{j} x\right\|_{2} \geq \Omega\left(\frac{1}{\operatorname{mma}(\ell, d) \cdot \log ^{5 / 2} n}\right)\|x\|_{\ell}
$$

Proof. Let $y_{h}$ be a random variable such that

$$
y_{h}= \begin{cases}1 & \text { if the } h \text {-th diagonal entry of } D_{j} \text { is } 1 \\ 0 & \text { otherwise }\end{cases}
$$

Let $Y=\sum_{h \in L_{i^{*}}(x)} y_{h}$. By Chernoff bound, we have

$$
\operatorname{Pr}[Y \geq \Omega(d \log n)] \geq 1-2^{-\Theta(d \log n)}
$$

Conditioned on $Y \geq \Omega(d \log n)$, we have

$$
\begin{aligned}
\frac{\left\|w_{j} D_{j} x\right\|_{2}}{\|x\|_{\ell}} & =\frac{w_{j}\left\|D_{j} x\right\|_{2}}{\|x\|_{\ell}} \\
& \geq \frac{2^{i^{*}} w_{j} \sqrt{d \log n}}{\|x\|_{\ell}} \\
& \geq \frac{2^{i^{*}} \sqrt{2^{j}} M_{\ell^{\left(2^{j}\right)}} \sqrt{d \log n}}{2^{i^{*}+1} g M_{\ell^{\left(\left|L_{i^{*}}(x)\right|\right)} \sqrt{\log n} \sqrt{\left|L_{i^{*}}(x)\right|}}} \\
& \geq \Omega\left(1 / \log ^{3 / 2} n\right) \cdot \frac{M_{\left.\ell^{(2}\right)}}{M_{\ell^{\left(\left|L_{i^{*}}(x)\right|\right)}}} \\
& =\Omega\left(1 / \log ^{3 / 2} n\right) \cdot \frac{M_{\ell^{\left(2^{j}\right)}}}{M_{\ell^{\left(\left|L_{i^{*}}(x)\right| / \log n\right)}}} \cdot \frac{M_{\ell^{\left(\left|L_{i^{*}}(x)\right| \log n\right)}}}{M_{\ell^{\left(\left|L L_{i}(x)\right|\right)}}} \\
& \geq \Omega\left(\frac{1}{\operatorname{mma}(\ell, d) \cdot \operatorname{mma}(\ell, \log n) \cdot \log ^{3 / 2} n}\right) \\
& \geq \Omega\left(\frac{1}{\operatorname{mma}(\ell, d) \cdot \log ^{5 / 2} n}\right) .
\end{aligned}
$$

Here the first inequality follows from the fact that there are at least $\Omega(d \log n)$ coordinates sampled from $L_{i^{*}}(x)$. The second inequality follows from Lemma 25.7.6 and the fact that level $i^{*}$ is a contributing level. The third inequality follows from $\left|L_{i^{*}}(x)\right|=\Theta(d \log n) \cdot 2^{j}$ and $g=\Theta(\log n)$. The forth inequality follows from Definition 25.7.5. The last inequality follows from Lemma 25.7.7.

Claim 25.7.14. For any fixed $x \in \mathbb{R}^{n}$. If there is a contributing level $i^{*} \in\{0,1,2, \ldots, g\}$ such that $\left|L_{i^{*}}(x)\right|=O(d \log n)$, then we have

$$
\left\|w_{0} D_{0} x\right\|_{2} \geq \Omega\left(\frac{1}{\operatorname{mma}(\ell, d) \cdot \log ^{5 / 2} n}\right)\|x\|_{\ell}
$$

Proof.

$$
\begin{aligned}
\frac{\left\|w_{0} D_{0} x\right\|_{2}}{\|x\|_{\ell}} & =\frac{w_{0}\|x\|_{2}}{\|x\|_{\ell}} \\
& \geq \frac{2^{i^{*}} w_{0} \sqrt{\left|L_{i^{*}}(x)\right|}}{\|x\|_{\ell}} \\
& \geq \frac{2^{i^{*}} M_{\ell^{(1)}} \sqrt{\left|L_{i^{*}}(x)\right|}}{2^{i^{*}+1} g M_{\ell^{\left(\mid L_{i^{*}(x)| |}\right.} \sqrt{\log n} \sqrt{\left|L_{i^{*}}(x)\right|}}} \\
& \geq \Omega\left(1 / \log ^{3 / 2} n\right) \cdot \frac{M_{\ell^{(1)}}}{M_{\ell^{\left(\left|L_{i^{*}}(x)\right| / \log n\right)}}} \cdot \frac{M_{\ell^{\left(\mid L_{\left.L^{*}(x) \mid / \log n\right)}\right.}}}{M_{\ell^{\left(\left|L_{i^{*}}(x)\right|\right)}}} \\
& \geq \Omega\left(\frac{1}{\mathrm{mma}(\ell, d) \cdot \log ^{5 / 2} n}\right)
\end{aligned}
$$

The first inequality follows from the fact that we only consider the contribution of the coordinates in $L_{i^{*}}(x)$. The second inequality follows from Lemma 25.7.6 and the fact that level $i^{*}$ is a contributing level. The third inequality follows from $g=\Theta(\log n)$. The last inequality follows from Definition 25.7.5 and Lemma 25.7.7.

By Claim 25.7.13 and Claim 25.7.14, since any vector $x \in \mathbb{R}^{n}$ contains at least one contributing level, with probability at least $1-2^{-\Theta(d \log n)}$ we have $\|\widetilde{D} x\|_{2} \geq \Omega(1 /(\operatorname{mma}(\ell, d)$. $\left.\left.\log ^{5 / 2} n\right)\right) \cdot\|x\|_{\ell}$. We complete the proof by combining this with Lemma 25.7.7.

Now we show how to combine the contraction bound in Lemma 25.7.12 with a net argument to give a contraction bound for all vectors in a subspace.

Lemma 25.7.15. Let $S \in \mathbb{R}^{O(d) \times n}$ be a random matrix. For any $\alpha_{0}=\operatorname{poly}(n)$ and $A \in$ $\mathbb{R}^{n \times d}$, if

1. $\|S\|_{2} \leq \operatorname{poly}(n)$ holds with probability at least 0.999 ;
2. for any fixed $x \in \mathbb{R}^{n},\|S x\|_{2} \geq 1 / \alpha_{0} \cdot\|x\|_{\ell}$ holds with probability $1-e^{-C d \log n}$ for a sufficiently large constant $C$,
then with probability at least 0.998 , for all $y \in \mathbb{R}^{n}$ in the column space of $A$,

$$
\|S y\|_{2} \geq \Omega\left(1 / \alpha_{0}\right)\|y\|_{\ell}
$$

Proof. For the matrix $A \in \mathbb{R}^{n \times d}$, we define the set $\mathcal{B}=\left\{y \mid y=A x,\|y\|_{2}=1\right\}$. We define $\mathcal{N} \subset \mathbb{R}^{n}$ to be an $\epsilon$-net of $\mathcal{B}$ as in Definition 25.5.1. By Lemma 25.5.1, we have $|\mathcal{N}| \leq(1+1 / \epsilon)^{d}$, and for all $y \in \mathcal{B}$, there exists $z \in \mathcal{N}$ such that $\|y-z\|_{2} \leq \epsilon$. We take $\epsilon=1 / \operatorname{poly}(n)$ here .

Due to the second condition, since $|\mathcal{N}| \leq e^{O(d \log n)}$, by taking union bound over all vectors in $\mathcal{N}$, we know that with probability $1-e^{-\Theta(d \log n)}$, for all $z \in \mathcal{N},\|S z\|_{2} \geq 1 / \alpha_{0} \cdot\|z\|_{\ell}$.

Now, for any vector $y \in \mathcal{B}$, there exists $z \in \mathcal{N}$ such that $\|y-z\|_{2} \leq 1 / \operatorname{poly}(n)$, and we define $w=y-z$.

$$
\begin{array}{rlr}
\|S y\|_{2} & =\|S(z+w)\|_{2} & \\
& \geq\|S z\|_{2}-\|S w\|_{2} & \text { (triangle inequality) } \\
& \geq 1 / \alpha_{0} \cdot\|z\|_{\ell}-\|S w\|_{2} & \text { (by the second condition) } \\
& \geq 1 / \alpha_{0} \cdot\|z\|_{\ell}-\|S\|_{2}\|w\|_{2} & \left(\|S w\|_{2} \leq\|S\|_{2} \cdot\|w\|_{2}\right) \\
& \geq 1 / \alpha_{0} \cdot\|z\|_{\ell}-\operatorname{poly}(n) \cdot\|w\|_{2} & (y=z+w) \\
& \geq 1 / \alpha_{0} \cdot\|y-w\|_{\ell}-\operatorname{poly}(n) \cdot\|w\|_{2} & \text { (by the first condition) } \\
& \geq 1 / \alpha_{0} \cdot\|y\|_{\ell}-1 / \alpha_{0} \cdot\|w\|_{\ell}-\operatorname{poly}(n) \cdot\|w\|_{2} & \left(\|w\|_{2} \leq \epsilon\right) \\
& \geq 1 / \alpha_{0} \cdot\|y\|_{\ell}-1 / \alpha_{0} \cdot \sqrt{n}\|w\|_{2}-\operatorname{poly}(n) \cdot\|w\|_{2} & \\
& \geq 1 / \alpha_{0} \cdot\|y\|_{\ell}-\left(1 / \alpha_{0} \cdot \sqrt{n}+\operatorname{poly}(n)\right) \epsilon & \text { (Lemma 25.7.5) } \\
& \geq 0.5 / \alpha_{0} \cdot\|y\|_{\ell} . &
\end{array}
$$

Lemma 25.7.16. For a given matrix $A \in \mathbb{R}^{n \times d}$. Let $S \in \mathbb{R}^{O(d) \times n}$ be a SymSketch as defined in Definition 25.3.1. With probability at least 0.995, for all $x \in \mathbb{R}^{d},\|S A x\|_{2} \geq 1 / \alpha_{0} \cdot\|A x\|_{\ell}$ where $\alpha_{0}=O\left(\sqrt{d} \log ^{3} n\right)$.

Proof. By Lemma 25.7.12 and Lemma 25.7.9, the two conditions in Lemma 25.7.15 are satisfied. By Lemma 25.7.15, with probability at least 0.998 , for all $x \in \mathbb{R}^{d},\|\widetilde{D} A x\|_{2} \geq$ $\Omega\left(1 / \alpha_{0}\right)\|A x\|_{\ell}$. Since $\Pi \in \mathbb{R}^{O(d) \times n(t+1)}$ is a composition of Gaussian embedding and CountSketch embedding with $\epsilon=0.1$, by Corollary 25.7.3, with probability at least 0.999 , for all
$x \in \mathbb{R}^{d},\|\Pi \widetilde{D} A x\|_{2} \geq \Omega\left(\|\widetilde{D} A x\|_{2}\right)$. By a union bound, we know that with probability at least 0.995 , for all $x \in \mathbb{R}^{d},\|S A x\|_{2} \geq \Omega\left(1 / \alpha_{0}\right)\|A x\|_{\ell}$.

### 25.7.5 Dilation Bound of SymSketch

In this section we give the dilation bound of SymSketch. We first show that for any fixed $x \in \mathbb{R}^{n}$, with high probability, $\|\widetilde{D} x\|_{2} \leq \operatorname{poly}(d \log n) \cdot \operatorname{mac}(\ell) \cdot\|x\|_{\ell}$.

Lemma 25.7.17. Let $\widetilde{D}$ be the matrix defined in Definition 25.3.1. For any fixed vector $x \in \mathbb{R}^{n}$, with probability $1-\delta,\|\widetilde{D} x\|_{2} \leq \alpha_{1} / \delta \cdot\|x\|_{\ell}$, where $\alpha_{1}=O\left(\operatorname{mmc}(\ell) \log ^{5 / 2} n\right)$.

Proof. Consider a fixed vector $x \in \mathbb{R}^{n}$. Recall that $t=\Theta(\log n)$. Let $c>0$ be a fixed constant. We define the $j$-heavy level set $H_{j}$ as

$$
H_{j}=\left\{i| | L_{i}(x) \left\lvert\, \geq c \frac{\delta 2^{j}}{\log ^{2} n}\right., 0 \leq i \leq g\right\}
$$

Let $\bar{H}_{j}$ be the $j$-light level set, i.e., $\bar{H}_{j}=\{0,1, \ldots, g\} \backslash H_{j}$. Notice that

$$
\sum_{i \in \bar{H}_{j}}\left|L_{i}(x)\right| \cdot 2^{-j} \leq g \cdot c \delta 2^{j} / \log ^{2} n \cdot 2^{-j} \leq O(\delta / \log n)
$$

By Markov's inequality, with probability at least $1-\delta /(2 t)$, no element from a $j$-light level is sampled by $D_{j}$, i.e., for all $i \in \bar{H}_{j}, k \in L_{i}(x)$, the $k$-th diagonal entry of $D_{j}$ is 0 . By taking union bound over all $j \in[t]$, with probability at least $1-\delta / 2$, for all $j \in[t]$, no element from a $j$-light level is sampled by $D_{j}$. Let $\zeta$ denote this event. We condition on this event in the remaining part of the proof. In the following analysis, we show an upper bound of $\left\|w_{j} D_{j} x\right\|_{2}^{2}$ for each $j \in[t]$. Let $H_{j}$ be the set of $j$-heavy levels.

Consider a fixed $j \in[t]$, we have

$$
\begin{aligned}
\underset{D_{j}}{\mathbb{E}}\left[\left\|w_{j} D_{j} x\right\|_{2}^{2} \mid \zeta\right] & =w_{j}^{2} \underset{D_{j}}{\mathbb{E}}\left[\left\|D_{j} x\right\|_{2}^{2} \mid \zeta\right] \\
& =w_{j}^{2} \underset{D_{j}}{\mathbb{E}}\left[\sum_{h=1}^{n}\left(D_{j}(h, h)\right)^{2} x_{h}^{2} \mid \zeta\right] \\
& =w_{j}^{2} \underset{D_{j}}{\mathbb{E}}\left[\sum_{i=0}^{g} \sum_{h \in L_{i}(x)}\left(D_{j}(h, h)\right)^{2} x_{h}^{2} \mid \zeta\right] \\
& =w_{j}^{2} \frac{1}{2^{j}} \sum_{i \in H_{j}} \sum_{h \in L_{i}(x)} x_{h}^{2} \\
& \leq w_{j}^{2} \frac{1}{2^{j}} \sum_{i \in H_{j}}\left|L_{i}(x)\right| \cdot\left(2^{i+1}\right)^{2} .
\end{aligned}
$$

Claim 25.7.18. $w_{j}^{2} 2^{-j} \leq O\left(\left(M_{\left.\ell^{(2 j}\right)}\right)^{2}\right)$.

Proof.

$$
\begin{aligned}
w_{j}^{2} 2^{-j} & =\left(\|(1,1, \ldots, 1,0, \ldots, 0)\|_{\ell}\right)^{2} \cdot 2^{-j} \\
& =\left(\left\|\frac{1}{\sqrt{2^{j}}}(1,1, \ldots, 1,0, \ldots, 0)\right\|_{\ell} \sqrt{2^{j}}\right)^{2} \cdot 2^{-j} \\
& =\left(\left\|\xi^{\left(2^{j}\right)}\right\|_{\ell}\right)^{2} \cdot 2^{j} \cdot 2^{-j} \\
& =\left(\left\|\xi^{\left(2^{j}\right)}\right\|_{\ell}\right)^{2} \\
& \leq O\left(M_{\ell^{\left(2^{j}\right)}}\right)
\end{aligned}
$$

where the third step follows from the definition of $\xi^{\left(2^{j}\right)}$, and the last step follows from Lemma 25.7.6.

Using the above claim, we have

$$
\begin{aligned}
& w_{j}^{2} \sum_{i \in H_{j}}\left|L_{i}(x)\right| \cdot 2^{2 i-j} \\
& =\sum_{i \in H_{j}} \frac{w_{j}^{2} 2^{-j}}{\left(M_{\ell\left(\left|L_{i}(x)\right|\right)}\right)^{2}} \cdot\left|L_{i}(x)\right| \cdot\left(M_{\ell^{\left(\left|L_{i}(x)\right|\right)}}\right)^{2} \cdot 2^{2 i} \\
& \leq O\left(\sum_{i \in H_{j}}\left(\frac{M_{\ell(2 j)}}{M_{\ell\left(\left|L_{i}(x)\right|\right)}}\right)^{2} \cdot\left|L_{i}(x)\right| \cdot\left(M_{\ell\left(\left|L_{i}(x)\right| \mid\right.}\right)^{2} \cdot 2^{2 i}\right) \\
& \leq O\left(\sum_{i \in H_{j}}\left(\frac{M_{\ell^{\left(2^{j}\right)}}}{M_{\ell^{\left(\left|L_{i}(x)\right|\right)}}}\right)^{2} w_{\log \left|L_{i}(x)\right|}^{2} \log n \cdot 2^{2 i}\right) \\
& =\log n \cdot \underbrace{\sum_{i \in H_{j},\left|L_{i}(x)\right| \leq 2^{j}}\left(\frac{M_{\ell^{\left(2^{j}\right)}}}{M_{\ell\left(\left|L_{i}(x)\right|\right)}}\right)^{2} \cdot w_{\log \left|L_{i}(x)\right|}^{2} \cdot 2^{2 i}}_{\diamond} \\
& +\log n \cdot \underbrace{\sum_{i \in H_{j},\left|L_{i}(x)\right|>2^{j}}\left(\frac{M_{\ell\left(2^{j}\right)}}{M_{\ell\left(\left|L_{i}(x)\right|\right)}}\right)^{2} \cdot w_{\log \left|L_{i}(x)\right|}^{2} \cdot 2^{2 i}}_{0},
\end{aligned}
$$

where the second step follows from $w_{j}^{2} 2^{-j} \leq O\left(\left(M_{\ell^{\left(2^{j}\right)}}\right)^{2}\right)$ (Claim 25.7.18), and the third step follows from $\left|L_{i}(x)\right| \cdot\left(M_{\ell\left(\left|L_{i}(x)\right|\right)}\right)^{2} \leq O\left(w_{\log \left|L_{i}(x)\right|}^{2} \log n\right)$ (Lemma 25.7.6). It remains to upper bound $\diamond$ and $\odot$.

To given an upper bound for $\diamond$, we have

$$
\begin{aligned}
\diamond & \leq O\left(\sum_{i \in H_{j},\left|L_{i}(x)\right| \leq 2^{j}} \mathrm{mma}^{2}\left(\ell, \log ^{2} n / \delta\right) \cdot w_{\log \left|L_{i}(x)\right|}^{2} \cdot 2^{2 i}\right) \\
& \leq O\left(\mathrm{mma}^{2}\left(\ell, \log ^{2} n / \delta\right)\left(\sum_{i=0}^{g} w_{\log \left|L_{i}(x)\right|} \cdot 2^{i}\right)^{2}\right) \\
& \leq O\left(\mathrm{mma}^{2}\left(\ell, \log ^{2} n / \delta\right)\right)\|x\|_{\ell}^{2} \\
& \leq O\left(\log ^{3} n / \delta\right)\|x\|_{\ell}^{2},
\end{aligned}
$$

where the first step follows from the definition of mma, the second step follows from Minkowski inequality, the third step follows from the definition of $L_{i}(x), w_{\log \left|L_{i}(x)\right|}$ and triangle inequality, the last step follows from Lemma 25.7.7.

To give an upper bound for $\odot$, we have

$$
\begin{aligned}
\bigcirc & \leq O\left(\log n \cdot \sum_{i \in H_{j},\left|L_{i}(x)\right|>2^{j}} \operatorname{mmc}^{2}(\ell) w_{\log \left|L_{i}(x)\right|}^{2} \cdot 2^{2 i}\right) \\
& \leq O\left(\log n \cdot \mathrm{mmc}^{2}(\ell) \cdot\left(\sum_{i=0}^{g} w_{\log \left|L_{i}(x)\right|} \cdot 2^{i}\right)^{2}\right) \\
& \leq O\left(\log n \cdot \mathrm{mmc}^{2}(\ell) \cdot\|x\|_{\ell}^{2}\right),
\end{aligned}
$$

where the first step follows from $\left(M_{\left.\ell^{(2}\right)} / M_{\ell^{\left(\left|L_{i}(x)\right|\right)}}\right)^{2} \leq O\left(\log n \cdot \operatorname{mmc}^{2}(\ell)\right)$ (Lemma 3.14 in $\left.\left[\mathrm{BBC}^{+} 17\right]\right)$.

Putting it all together, we have

$$
\underset{D_{j}}{\mathbb{E}}\left[\left\|w_{j} D_{j} x\right\|_{2}^{2} \mid \zeta\right] \leq \log n \cdot(\diamond+\diamond) \leq O\left(\log ^{4} n / \delta+\log ^{2} n \cdot \operatorname{mmc}^{2}(\ell)\right)\|x\|_{\ell}^{2}
$$

Thus,

$$
\underset{\widetilde{D}}{\mathbb{E}}\left[\|\widetilde{D} x\|_{2}^{2} \mid \zeta\right] \leq \sum_{j=0}^{t} \underset{D_{j}}{\mathbb{E}}\left[\left\|w_{j} D_{j} x\right\|_{2}^{2} \mid \zeta\right] \leq O\left(\log ^{5} n / \delta+\log ^{3} n \cdot \operatorname{mmc}^{2}(\ell)\right)\|x\|_{\ell}^{2}
$$

By Markov's inequality, conditioned on $\zeta$, with probability at least $1-\delta / 2$,

$$
\|\widetilde{D} x\|_{2}^{2} \leq O\left(\log ^{5} n / \delta+\log ^{3} n \cdot \operatorname{mmc}^{2}(\ell)\right)\|x\|_{\ell}^{2} / \delta
$$

Since $\zeta$ holds with probability at least $1-\delta / 2$, with probability at least $1-\delta$, we have

$$
\|\widetilde{D} x\|_{2} \leq O\left(\log ^{5 / 2} n / \delta \cdot \operatorname{mmc}(\ell)\right) \cdot\|x\|_{\ell}
$$

Now we show how to use the dilation bound for a fixed vector in Lemma 25.7.17 to prove a dilation bound for all vectors in a subspace. We need the following existential result in our proof.

Lemma 25.7.19 ([Aue30]). Given a matrix $A \in \mathbb{R}^{n \times m}$ and a norm $\|\cdot\|$, there exists a basis matrix $U \in \mathbb{R}^{n \times d}$ of the column space of $A$, such that

$$
\sum_{i=1}^{d}\left\|U^{i}\right\| \leq d
$$

and for all $x \in \mathbb{R}^{d}$,

$$
\|x\|_{\infty} \leq\|U x\| .
$$

Lemma 25.7.20. Given a matrix $A \in \mathbb{R}^{n \times d}$. Let $S \in \mathbb{R}^{O(d) \times n}$ be a SymSketch as defined in Definition 25.3.1. With probability at least 0.99 , for all $x \in \mathbb{R}^{d}$,

$$
\|S A x\|_{2} \leq O\left(\alpha_{1} d^{2}\right)\|A x\|_{\ell}
$$

where $\alpha_{1}=O\left(\operatorname{mmc}(\ell) \cdot \log ^{5 / 2} n\right)$.

Proof. Recall that $S=\Pi \widetilde{D}$. Let $U$ be a basis matrix of the column space of $A$ as in Lemma 25.7.19. By Lemma 25.7.17, for a fixed $i \in[d]$, with probability at least $1-1 /(100 d)$, $\left\|\widetilde{D} U^{i}\right\|_{2} \leq O\left(\alpha_{1} d\right)\left\|U^{i}\right\|_{\ell}$. By taking a union bound over $i \in[d]$, with probability at least 0.999 , for all $i \in[d],\left\|\widetilde{D} U^{i}\right\|_{2} \leq \alpha_{1} d\left\|U^{i}\right\|_{\ell}$. Thus, for any $x \in \mathbb{R}^{d}$,

$$
\begin{aligned}
\|\widetilde{D} U x\|_{2} & \leq \sum_{i=1}^{d}\left|x_{i}\right| \cdot\left\|\widetilde{D} U^{i}\right\|_{2} \\
& \leq\|x\|_{\infty} \cdot \sum_{i=1}^{d}\left\|\widetilde{D} U^{i}\right\|_{2} \\
& \leq\|U x\|_{\ell} \cdot \sum_{i=1}^{d}\left\|\widetilde{D} U^{i}\right\|_{2} \\
& \leq O\left(\alpha_{1} d\right) \cdot\|U x\|_{\ell} \cdot \sum_{i=1}^{d}\left\|U^{i}\right\|_{\ell} \\
& \leq O\left(\alpha_{1} d^{2}\right) \cdot\|U x\|_{\ell}
\end{aligned}
$$

where the first step follows from triangle inequality, the second step follows from $\left|x_{i}\right| \leq\|x\|_{\infty}$ for all $i \in[d]$, the third step follows from $\|x\|_{\infty} \leq\|U x\|_{\ell}$, the fourth step follows from $\left\|\widetilde{D} U^{i}\right\|_{2} \leq O\left(\alpha_{1} d\right)\left\|U^{i}\right\|_{\ell}$, the last step follows from $\sum_{i=1}^{d}\left\|U^{i}\right\|_{\ell} \leq d$.

By Corollary 25.7.3, with probability at least $0.999, \Pi$ is an $\ell_{2}$ subspace embedding with $\epsilon=0.1$ for the column space of $\widetilde{D} U$. Thus, with probability at least 0.99 , for all $x \in \mathbb{R}^{d}$, $\|S A x\|_{2} \leq O\left(\alpha_{1} d^{2}\right)\|A x\|_{\ell}$.

### 25.7.6 Proof of Theorem 25.3.3

Proof. It directly follows from Lemma 25.7.16, Lemma 25.7.20 and Lemma 25.7.8.

### 25.8 Missing Proofs of Main Theorems

### 25.8.1 Proof of Theorem 25.1.3

Let $S \in \mathbb{R}^{O(d) \times n}$ be a SymSketch as defined in Definition 25.3.1, and $\Pi=O\left(\sqrt{d} \log ^{3} n\right)$. $S$. By Corollary 25.3.4, for a given matrix $A \in \mathbb{R}^{n \times d}$, with probability at least 0.9 , for all $x \in \mathbb{R}^{d}$,

$$
\|A x\|_{G} \leq\|\Pi A x\|_{2} \leq \kappa\|A x\|_{G},
$$

where $\kappa=O\left(\sqrt{C_{G}} d^{5 / 2} \log ^{13 / 2} n\right)$. We prove Theorem 25.1.3 by combining Theorem 25.2.6 with the embedding matrix $\Pi$ constructed above.

### 25.8.2 Proof of Theorem 25.1.4

Let $S \in \mathbb{R}^{O(d) \times n}$ be a SymSketch as defined in Definition 25.3.1. For a given data matrix $A \in \mathbb{R}^{n \times d}$ and response vector $b \in \mathbb{R}^{n}$, we calculate $x^{*}=\operatorname{argmin}_{x}\|S A x-S b\|_{2}$ and return $x^{*}$. The algorithm runs in $O(\operatorname{nnz}(A)+\operatorname{poly}(d))$ time, since by Lemma 25.7.8, the expected running time for calculating $S A$ is $O(\operatorname{nnz}(A)+\operatorname{poly}(d))$, and $x^{*}=(S A)^{+} S b$ can be calculated in poly $(d)$ time.

To see the correctness, let $\bar{x}=\operatorname{argmin}_{x}\|A x-b\|_{\ell}$. With probability at least 0.99 , we have

$$
\begin{aligned}
\left\|A x^{*}-b\right\|_{\ell} & \leq O\left(\sqrt{d} \log ^{3} n\right)\left\|S A x^{*}-S b\right\|_{2} \\
& \leq O\left(\sqrt{d} \log ^{3} n\right)\|S A \bar{x}-S b\|_{2} \\
& \leq O\left(\sqrt{d} \log ^{3} n\right)\|\widetilde{D} A \bar{x}-\widetilde{D} b\|_{\ell} \\
& =O\left(\sqrt{d} \log ^{11 / 2} n\right) \cdot \operatorname{mmc}(\ell) \cdot\|A \bar{x}-b\|_{\ell} .
\end{aligned}
$$

The first step follows by applying Lemma 25.7 .16 on $\mathbf{A b}$, where we use $\mathbf{A b} \in \mathbb{R}^{n \times(d+1)}$ to denote a matrix whose first $d$ columns are $A$ and the last column is $b$. The second step follows from the fact that $x^{*}=\operatorname{argmin}_{x}\|S A x-S b\|_{2}$. The third step follows by Definition 25.3.1 and Corollary 25.7.3. The last step follows by applying Lemma 25.7.17 on $A \bar{x}-b$.

## Chapter 26

## Kronecker Product Regression

We study Kronecker product regression, in which the design matrix to a regression problem is a Kronecker product of two or more matrices. Formally, given $A_{i} \in \mathbb{R}^{n_{i} \times d_{i}}$ for $i=1,2, \ldots, q$ where $n_{i} \gg d_{i}$ for each $i$, and $b \in \mathbb{R}^{n_{1} n_{2} \cdots n_{q}}$, let $\mathcal{A}=A_{i} \otimes A_{2} \otimes \cdots \otimes A_{q}$. Then for $p \in[1,2]$, the goal is to find $x \in \mathbb{R}^{d_{1} \cdots d_{q}}$ that approximately minimizes $\|\mathcal{A} x-b\|_{p}$. Recently, Diao, Song, Sun, and Woodruff (AISTATS, 2018) gave an algorithm which is faster than forming the Kronecker product $\mathcal{A} \in \mathbb{R}^{n_{1} \cdots n_{q} \times d_{1} \cdots d_{q}}$. Specifically, for $p=2$ they achieve a running time of $O\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+\mathrm{nnz}(b)\right)$, where $\mathrm{nnz}\left(A_{i}\right)$ is the number of non-zero entries in $A_{i}$. Note that $n n z(b)$ can be as large as $\Theta\left(n_{1} \cdots n_{q}\right)$. For $p=1, q=2$ and $n_{1}=n_{2}$, they achieve a worse bound of $O\left(n_{1}^{3 / 2} \operatorname{poly}\left(d_{1} d_{2}\right)+\mathrm{nnz}(b)\right)$. In this work, we provide significantly faster algorithms. For $p=2$, our running time is $O\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)\right)$, which has no dependence on $\mathrm{nnz}(b)$. For $p<2$, our running time is $O\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+\mathrm{nnz}(b)\right)$, which matches the prior best running time for $p=2$. We also consider the related all-pairs regression problem, where given $A \in \mathbb{R}^{n \times d}, b \in \mathbb{R}^{n}$, we want to solve $\min _{x \in \mathbb{R}^{d}}\|\bar{A} x-\bar{b}\|_{p}$, where $\bar{A} \in \mathbb{R}^{n^{2} \times d}, \bar{b} \in \mathbb{R}^{n^{2}}$ consist of all pairwise differences of the rows of $A, b$. We give an $O(\mathrm{nnz}(A))$ time algorithm for $p \in[1,2]$, improving the $\Omega\left(n^{2}\right)$ time required to form $\bar{A}$. Finally, we initiate the study of Kronecker product low rank approximation and low-trank approximation. For input $\mathcal{A}$ as above, we give an $O\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)\right)$ time algorithm, which is much faster than computing $\mathcal{A}$.

This part is based upon the following previous publication

- Huaian Diao, Rajesh Jayaram, Zhao Song, Wen Sun, David P. Woodruff Optimal Sketching for Kronecker Product Regression and Low Rank Approximation. Manuscript 2019 [DJS ${ }^{+}$19]


### 26.1 Introduction

In the $q$-th order Kronecker product regression problem, one is given $A_{i} \in \mathbb{R}^{n_{i} \times d_{i}}$ for each $i \in\{1,2, \ldots, q\}$, together with $b \in \mathbb{R}^{n_{1} n_{2} \cdots n_{q}}$, and the goal is to obtain a solution to the optimization problem:

$$
\min _{x \in \mathbb{R}^{d_{1} d_{2} \cdots d_{q}}}\left\|\left(A_{1} \otimes A_{2} \cdots \otimes A_{q}\right) x-b\right\|_{p}
$$

where $p \in[1,2]$, and for a vector $x \in \mathbb{R}^{n}$ the $\ell_{p}$ norm is defined by $\|x\|_{p}=\left(\sum_{i=1}^{n}\left|x_{i}\right|^{p}\right)^{1 / p}$. For $p=2$, this is least squares regression, and for $p=1$ this is least absolute deviation regression.

Kronecker product regression is a special case of ordinary regression, in which the design matrix is highly structured. Namely, it is the Kronecker product of two or more matrices. Kronecker product matrices naturally arise in applications such as spline regression, signal processing, and multivariate data fitting. We refer the reader to [VL92, VLP93, GVL13] for further background and applications of Kronecker product regression. As discussed in [DSSW18], Kronecker product regression also arises in structured blind deconvolution problems [OY05], and the bivariate problem of surface fitting and multidimensional density smoothing [EM06].

A recent work of Diao, Song, Sun, and Woodruff [DSSW18] utilizes sketching techniques to output an $x \in \mathbb{R}^{d_{1} d_{2} \cdots d_{q}}$ with objective function at most $(1+\epsilon)$-times larger than optimal, for both least squares and least absolute deviation Kronecker product regression. Importantly, their time complexity is faster than the time needed to explicitly compute $A_{1} \otimes \cdots \otimes A_{q}$. We note that sketching itself is a powerful tool for compressing extremely high dimensional data, and has been used in a number of tensor related problems, e.g., [SWZ16, LHW17, DSSW18, SWZ19b].

For least squares regression, [DSSW18] achieve $O\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+\mathrm{nnz}(b)+\operatorname{poly}(d / \epsilon)\right)$ time, where $\mathrm{nnz}(C)$ for a matrix $C$ denotes the number of non-zero entries of $C$. Note that the focus is on the over-constrained regression setting, when $n_{i} \gg d_{i}$ for each $i$, and so the goal is to have a small running time dependence on the $n_{i}$. We remark that over-constrained regression has been the focus of a large body of work over the past decade, which primarily attempts to design fast regression algorithms in the big data (large sample size) regime, see, e.g., [Mah11b, Woo14b] for surveys.

Observe that explicitly computing $A_{1} \otimes \cdots \otimes A_{q}$ would take $\prod_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)$ time, which can be as large as $\prod_{i=1}^{q} n_{i} d_{i}$, and so the results of [DSSW18] offer a large computational advantage. Unfortunately, since $b \in \mathbb{R}^{n_{1} n_{2} \cdots n_{q}}$, we can have $\operatorname{nnz}(b)=\prod_{i=1}^{q} n_{i}$, and therefore $n n z(b)$ is likely to be the dominant term in the running time. This leaves open the question of whether it is possible to solve this problem in time sub-linear in nnz(b), with a dominant term of $O\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)\right)$.

For least absolute deviation regression, the bounds of [DSSW18] achieved are still an improvement over computing $A_{1} \otimes \cdots \otimes A_{q}$, though worse than the bounds for least squares regression. The authors focus on $q=2$ and the special case $n=n_{1}=n_{2}$. Here, they obtain
a running time of $O\left(n^{3 / 2} \operatorname{poly}\left(d_{1} d_{2} / \epsilon\right)+\mathrm{nnz}(b)\right)^{1}$. This leaves open the question of whether an input-sparsity $O\left(\mathrm{nnz}\left(A_{1}\right)+\mathrm{nnz}\left(A_{2}\right)+\mathrm{nnz}(b)+\operatorname{poly}\left(d_{1} d_{2} / \epsilon\right)\right)$ time algorithm exists.

In this work, we also study the related all-pairs regression problem. Given $A \in$ $\mathbb{R}^{n \times d}, b \in \mathbb{R}^{n}$, the goal is to approximately solve the $\ell_{p}$ regression problem $\min _{x}\|\bar{A} x-\bar{b}\|_{p}$, where $\bar{A} \in \mathbb{R}^{n^{2} \times d}$ is the matrix formed by taking all pairwise differences of the rows of $A$ (and $\bar{b}$ is defined similarly). For $p=1$, this is known as the rank regression estimator, which has a long history in statistics. It is closely related to the renowned Wilconxon rank test [WL09], and enjoys the desirable property of being robust with substantial efficiency gain with respect to heavy-tailed random errors, while maintaining high efficiency for Gaussian errors [WKL09, WL09, WPB+ ${ }^{+}$18, Wan19a]. In many ways, it has properties more desirable in practice than that of the Huber M-estimator [WPB ${ }^{+} 18$, Wan19b]. Recently, the all-pairs loss function was also used by $\left[\mathrm{WPB}^{+} 18\right]$ as an alternative approach to overcoming the challenges of tuning parameter selection for the Lasso algorithm. However, the rank regression estimator is computationally intensive to compute, even for moderately sized data, since the standard procedure (for $p=1$ ) is to solve a linear program with $O\left(n^{2}\right)$ constraints. In this work, we demonstrate the first highly efficient algorithm for this estimator.

Finally, in addition to regression, we extend our techniques to the Low Rank Approximation (LRA) problem. Here, given a large data matrix $A$, the goal is to find a low rank matrix $B$ which well-approximates $A$. LRA is useful in numerous applications, such as

[^88]compressing massive datasets to their primary components for storage, denoising, and fast matrix-vector products. Thus, designing fast algorithms for approximate LRA has become a large and highly active area of research; see [Woo14b] for a survey. For an incomplete list of recent work using sketching techniques for LRA, see [CW13, MM13, NN13a, BW14, CW15b, CW15a, RSW16, BWZ16, SWZ17, MW17, CGK+17c, LHW17, SWZ18, BW18, SWZ19b, $\mathrm{BBB}^{+} 19 \mathrm{~b}$, IVWW19] and the references therein.

Motivated by the importance of LRA, we initiate the study of low-rank approximation of Kronecker product matrices. Given $q$ matrices $A_{1}, \cdots A_{q}$ where $A_{i} \in \mathbb{R}^{n_{i} \times d_{i}}, n_{i} \gg d_{i}, A=$ $\otimes_{i=1}^{q} A_{i}$, the goal is to output a rank- $k$ matrix $B \in \mathbb{R}^{n \times d}$ such that $\|B-A\|_{F}^{2} \leq(1+\epsilon) \mathrm{OPT}_{k}$, where $\mathrm{OPT}_{k}$ is the cost of the best rank- $k$ approximation, $n=n_{1} \cdots n_{q}$, and $d=d_{1} \cdots d_{q}$. Here $\|A\|_{F}^{2}=\sum_{i, j} A_{i, j}^{2}$. The fastest general purpose algorithms for this problem run in time $O(\operatorname{nnz}(A)+\operatorname{poly}(d k / \epsilon))$ [CW13]. However, as in regression, if $A=\otimes_{i=1}^{q} A_{i}$, we have $\mathrm{nnz}(A)=\prod_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)$, which grows very quickly. Instead, one might also hope to obtain a running time of $O\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+\operatorname{poly}(d k / \epsilon)\right)$.

### 26.1.1 Our Contributions

Our main contribution is an input sparsity time $(1+\epsilon)$-approximation algorithm to Kronecker product regression for every $p \in[1,2]$, and $q \geq 2$. Given $A_{i} \in \mathbb{R}^{n_{i} \times d_{i}}, i=1, \ldots, q$, and $b \in \mathbb{R}^{n}$ where $n=\prod_{i=1}^{q} n_{i}$, together with accuracy parameter $\epsilon \in(0,1 / 2)$ and failure probability $\delta>0$, the goal is to output a vector $x^{\prime} \in \mathbb{R}^{d}$ where $d=\prod_{i=1}^{q} d_{i}$ such that $\left\|\left(A_{1} \otimes \cdots \otimes A_{q}\right) x^{\prime}-b\right\|_{p} \leq(1+\epsilon) \min _{x}\left\|\left(A_{1} \otimes \cdots \otimes A_{q}\right) x-b\right\|_{p}$ holds with probability at least $1-\delta$. For $p=2$, our algorithm runs in $\left.\widetilde{O}\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)\right)+\operatorname{poly}\left(d \delta^{-1} / \epsilon\right)\right)$ time $^{2}$. Notice

[^89]that this is sub-linear in the input size, since it does not depend on nnz(b). For $p<2$, the running time is $\widetilde{O}\left(\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+\mathrm{nnz}(b)+\operatorname{poly}(d / \epsilon)\right) \log (1 / \delta)\right)$.

Observe that in both cases, this running time is significantly faster than the time to write down $A_{1} \otimes \cdots \otimes A_{q}$. For $p=2$, up to logarithmic factors, the running time is the same as the time required to simply read each of the $A_{i}$. Moreover, in the setting $p<2, q=2$ and $n_{1}=n_{2}$ considered in [DSSW18], our algorithm offers a substantial improvement over their running time of $O\left(n^{3 / 2} \operatorname{poly}\left(d_{1} d_{2} / \epsilon\right)\right)$.

We empirically evaluate our Kronecker product regression algorithm on exactly the same datasets as those used in [DSSW18]. For $p \in\{1,2\}$, the accuracy of our algorithm is nearly the same as that of [DSSW18] while the running time is significantly faster.

For the all-pairs (or rank) regression problem, we first note that for $A \in \mathbb{R}^{n \times d}$, one can rewrite $\bar{A} \in \mathbb{R}^{n^{2} \times d}$ as the difference of Kronecker products $\bar{A}=A \otimes \mathbf{1}^{n}-\mathbf{1}^{n} \otimes A$ where $\mathbf{1}^{n} \in \mathbb{R}^{n}$ is the all ones vector. Since $\bar{A}$ is not a Kronecker product itself, our earlier techniques for Kronecker product regression are not directly applicable. Therefore, we utilize new ideas, in addition to careful sketching techniques, to obtain an $\widetilde{O}(\operatorname{nnz}(A)+\operatorname{poly}(d / \epsilon))$ time algorithm for $p \in[1,2]$, which improves substantially on the $O\left(n^{2} d\right)$ time required to even compute $\bar{A}$, by a factor of at least $n$.

Our main technical contribution for both our $\ell_{p}$ regression algorithm and the rank regression problem is a novel and highly efficient $\ell_{p}$ sampling algorithm. Specifically, for the rank-regression problem we demonstrate, for a given $x \in \mathbb{R}^{d}$, how to independently sample $s$ entries of a vector $\bar{A} x=y \in \mathbb{R}^{n^{2}}$ from the $\ell_{p}$ distribution $\left(\left|y_{1}\right|^{p} /\|y\|_{p}^{p}, \ldots,\left|y_{n^{2}}\right|^{p} /\|y\|_{p}^{p}\right)$ in $\widetilde{O}(n d+\operatorname{poly}(d s))$ time. For the $\ell_{p}$ regression problem, we demonstrate the same result when
$y=\left(A_{1} \otimes \cdots \otimes A_{q}\right) x-b \in \mathbb{R}^{n_{1} \cdots n_{q}}$, and in time $\widetilde{O}\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+\mathrm{nnz}(b)+\operatorname{poly}(d s)\right)$. This result allows us to sample a small number of rows of the input to use in our sketch. Our algorithm draws from a large number of disparate sketching techniques, such as the dyadic trick for quickly finding heavy hitters [CM04, KNPW11, LNNT16, NS19], and the precision sampling framework from the streaming literature [AKO10].

For the Kronecker Product Low-Rank Approximation (LRA) problem, we give an input sparsity $O\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+\right.$ poly $\left.(d k / \epsilon)\right)$-time algorithm which computes a rank- $k$ matrix $B$ such that $\left\|B-\otimes_{i=1}^{q} A_{i}\right\|_{F}^{2} \leq(1+\epsilon) \min _{\text {rank }-k B^{\prime}}\left\|B^{\prime}-\otimes_{i=1}^{q} A_{i}\right\|_{F}^{2}$. Note again that the dominant term $\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)$ is substantially smaller than the $\mathrm{nnz}(A)=\prod_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)$ time required to write down the Kronecker Product $A$, which is also the running time of state-of-the-art general purpose LRA algorithms [CW13, MM13, NN13a]. Thus, our results demonstrate that substantially faster algorithms for approximate LRA are possible for inputs with a Kronecker product structure. Our technical contributions involve demonstrating that useful properties of known sketching matrices hold also for the Kronecker product of these matrices.

In addition, motivated by [VL00], we use our techniques to solve the low-trank approximation problem, where we are given an arbitrary matrix $A \in \mathbb{R}^{n^{q} \times n^{q}}$, and the goal is to output a trank- $k$ matrix $B \in \mathbb{R}^{n^{q} \times n^{q}}$ such that $\|B-A\|_{F}$ is minimized. Here, the trank of a matrix $B$ is the smallest integer $k$ such that $B$ can be written as a summation of $k$ matrices, where each matrix is the Kronecker product of $q$ matrices with dimensions $n \times n$. Compressing a matrix $A$ to a low-trank approximation yields many of the same benefits as LRA, such as compact representation, fast matrix-vector product, and fast matrix multiplication, and thus is applicable in many of the settings where LRA is used. Using similar
sketching ideas, we provide an $O\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+\operatorname{poly}\left(d_{1} \cdots d_{q} / \epsilon\right)\right)$ time algorithm for this problem under various loss functions. Our results for low-trank approximation can be found in Section 26.11.

### 26.2 Preliminaries

Notation. For a tensor $A \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$, we use $\|A\|_{p}$ to denote the entry-wise $\ell_{p}$ norm of $A$, i.e., $\|A\|_{p}=\left(\sum_{i_{1}} \sum_{i_{2}} \sum_{i_{3}}\left|A_{i_{1}, i_{2}, i_{3}}\right|^{p}\right)^{1 / p}$. For $n \in \mathbb{N}$, let $[n]=\{1,2, \ldots, n\}$. For a matrix $A$, let $A_{i, *}$ denote the $i$-th row of $A$, and $A_{*, j}$ the $j$-th column. For $a, b \in \mathbb{R}$ and $\epsilon \in(0,1)$, we write $a=(1 \pm \epsilon) b$ to denote $(1-\epsilon) b \leq a \leq(1+\epsilon) b$. We now define various sketching matrices used by our algorithms.
$p$-stable Transforms. We will utilize the well-known $p$-stable distribution, $\mathcal{D}_{p}$ (see [Nol07, Ind06] for further discussion), which exist for $p \in(0,2]$. For $p \in(0,2), X \sim \mathcal{D}_{p}$ is defined by its characteristic function $\mathbb{E}_{X}[\exp (\sqrt{-1} t X)]=\exp \left(-|t|^{p}\right)$, and can be efficiently generated to a fixed precision [Nol07, KNW10b]. For $p=2, \mathcal{D}_{2}$ is just the standard Gaussian distribution, and for $p=1, \mathcal{D}_{1}$ is the Cauchy distribution. The distribution $\mathcal{D}_{p}$ has the property that if $z_{1}, \ldots, z_{n} \sim D_{p}$ are i.i.d., and $a \in \mathbb{R}^{n}$, then $\sum_{i=1}^{n} z_{i} a_{i} \sim z\|a\|_{p}$ where $\|a\|_{p}=\left(\sum_{i=1}^{n}\left|a_{i}\right|^{p}\right)^{1 / p}$, and $z \sim \mathcal{D}_{p}$. This property will allow us to utilize sketches with entries independently drawn from $\mathcal{D}_{p}$ to preserve the $\ell_{p}$ norm.

Definition 26.2.1 (Dense $p$-stable Transform, [CDMI $\left.{ }^{+} 13\right]$ ). Let $p \in[1,2]$. Let $S=\sigma \cdot C \in$ $\mathbb{R}^{m \times n}$, where $\sigma$ is a scalar, and each entry of $C \in \mathbb{R}^{m \times n}$ is chosen independently from $\mathcal{D}_{p}$.

We will also need a sparse version of the above.
Definition 26.2.2 (Sparse $p$-Stable Transform, [MM13, CDMI $\left.{ }^{+} 13\right]$ ). Let $p \in[1,2]$. Let $\Pi=\sigma \cdot S C \in \mathbb{R}^{m \times n}$, where $\sigma$ is a scalar, $S \in \mathbb{R}^{m \times n}$ has each column chosen independently
and uniformly from the $m$ standard basis vectors of $\mathbb{R}^{m}$, and $C \in \mathbb{R}^{n \times n}$ is a diagonal matrix with diagonals chosen independently from the standard $p$-stable distribution. For any matrix $A \in \mathbb{R}^{n \times d}, \Pi A$ can be computed in $O(\mathrm{nnz}(A))$ time.

One nice property of $p$-stable transformations is that they provide low-distortion $\ell_{p}$ embeddings.

Lemma 26.2.1 (Theorem 1.4 of [WW19]; see also Theorem 2 and 4 of [MM13] for earlier work ${ }^{3}$ ). Fix $A \in \mathbb{R}^{n \times d}$, and let $S \in \mathbb{R}^{k \times n}$ be a sparse or dense $p$-stable transform for $p \in[1,2)$, with $k=\Theta\left(d^{2} / \delta\right)$. Then with probability $1-\delta$, for all $x \in \mathbb{R}^{d}$ :

$$
\|A x\|_{p} \leq\|S A x\|_{p} \leq O(d \log d)\|A x\|_{p}
$$

We simply call a matrix $S \in \mathbb{R}^{k \times n}$ a low distortion $\ell_{p}$ embedding for $A \in \mathbb{R}^{n \times d}$ if it satisfies the above inequality for all $x \in \mathbb{R}^{d}$.

Leverage Scores \& Well Condition Bases. We now introduce the notions of $\ell_{2}$ leverage scores and well-conditioned bases for a matrix $A \in \mathbb{R}^{n \times d}$.

Definition 26.2.3 ( $\ell_{2}$-Leverage Scores, [Woo14b, BSS12]). Given a matrix $A \in \mathbb{R}^{n \times d}$, let $A=Q \cdot R$ denote the QR factorization of matrix $A$. For each $i \in[n]$, we define $\sigma_{i}=\frac{\left\|\left(A R^{-1}\right)_{i}\right\|_{2}^{2}}{\left\|A R^{-1}\right\|_{F}^{2}}$, where $\left(A R^{-1}\right)_{i} \in \mathbb{R}^{d}$ is the $i$-th row of matrix $\left(A R^{-1}\right) \in \mathbb{R}^{n \times d}$. We say that $\sigma \in \mathbb{R}^{n}$ is the $\ell_{2}$ leverage score vector of $A$.

[^90]```
Algorithm 26.1 Our \(\ell_{2}\) Kronecker Product Regression Algorithm
    procedure \(\ell_{2} \operatorname{Kronecker} \operatorname{Regression}\left(\left(\left\{A_{i}, n_{i}, d_{i}\right\}_{i \in[q]}, b\right)\right) \quad \triangleright\) Theorem 26.3.1
        \(d \leftarrow \prod_{i=1}^{q} d_{i}, n \leftarrow \prod_{i=1}^{q} n_{i}, m \leftarrow \Theta\left(d /\left(\delta \epsilon^{2}\right)\right)\).
        Compute approximate leverage scores \(\widetilde{\sigma}_{i}\left(A_{j}\right)\) for all \(j \in[q], i \in\left[n_{j}\right]\). \(\triangleright\) Lemma 26.7.3
        Construct diagonal leverage score sampling matrix \(D \in \mathbb{R}^{n \times n}\), with \(m\) non-zero entries
    \(\triangleright\) Proposition 26.7.4
        Compute (via the psuedo-inverse)
                                    \(\widehat{x}=\arg \min _{x \in \mathbb{R}^{d}}\left\|D\left(A_{1} \otimes A_{2} \otimes \cdots \otimes A_{q}\right) x-D b\right\|_{2}\)
        return \(\widehat{x}\)
    end procedure
```

Definition 26.2.4 ( $\left(\ell_{p}, \alpha, \beta\right)$ Well-Conditioned Basis, [Cla05]). Given a matrix $A \in \mathbb{R}^{n \times d}$, we say $U \in \mathbb{R}^{n \times d}$ is an $\left(\ell_{p}, \alpha, \beta\right)$ well-conditioned basis for the column span of $A$ if the columns of $U$ span the columns of $A$, and if for any $x \in \mathbb{R}^{d}$, we have $\alpha\|x\|_{p} \leq\|U x\|_{p} \leq \beta\|x\|_{p}$, where $\alpha \leq 1 \leq \beta$. If $\beta / \alpha=d^{O(1)}$, then we simply say that $U$ is an $\ell_{p}$ well conditioned basis for $A$.

Fact 26.2.2 ([WW19, MM13]). Let $A \in \mathbb{R}^{n \times d}$, and let $S A \in \mathbb{R}^{k \times d}$ be a low distortion $\ell_{p}$ embedding for $A$ (see Lemma 26.2.1), where $k=O\left(d^{2} / \delta\right)$. Let $S A=Q R$ be the $Q R$ decomposition of $S A$. Then $A R^{-1}$ is an $\ell_{p}$ well-conditioned basis with probability $1-\delta$.

### 26.3 Kronecker Product Regression

We first introduce our algorithm for $p=2$. Our algorithm for $1 \leq p<2$ is given in Section 26.3.1. Our regression algorithm for $p=2$ is formally stated in Algorithm 26.1. Recall that our input design matrix is $A=\otimes_{i=1}^{q} A_{i}$, where $A_{i} \in \mathbb{R}^{n_{i} \times d_{i}}$, and we are also given $b \in \mathbb{R}^{n_{1} \cdots n_{q}}$. Let $n=\prod_{i=1}^{q} n_{i}$ and $d=\prod_{i=1}^{q} d_{i}$. The crucial insight of the algorithm is that one can approximately compute the leverage scores of $A$ given only good approximations to the leverage scores of each $A_{i}$. Applying this fact gives a efficent algorithm for sampling rows of $A$ with probability proportional to the leverage scores. Following standard arguments, we

```
Algorithm 26.2 Our \(\ell_{p}\) Kronecker Product Regression Algorithm, \(1 \leq p<2\)
    procedure \(O(1)\)-approximate \(\ell_{p}\) Regression \(\left(\left\{A_{i}, n_{i}, d_{i}\right\}_{i \in[q]}\right) \quad \triangleright\) Theorem 26.3.2
        \(d \leftarrow \prod_{i=1}^{q} d_{i}, n \leftarrow \prod_{i=1}^{q} n_{i}\).
        for \(i=1, \ldots, q\) do
            \(s_{i} \leftarrow O\left(q d_{i}^{2}\right)\)
            Generate sparse \(p\)-stable transform \(S_{i} \in \mathbb{R}^{s_{i} \times n}\) (def 26.2.2) \(\quad \triangleright\) Lemma 26.2.1
            Take the QR factorization of \(S_{i} A_{i}=Q_{i} R_{i}\) to obtain \(R_{i} \in \mathbb{R}^{d_{i} \times d_{i}} \quad \triangleright\) Fact 26.2.2
            Let \(Z \in \mathbb{R}^{d \times \tau}\) be a dense \(p\)-stable transform for \(\tau=\Theta(\log (n)) \triangleright\) Definition 26.2.1
            for \(j=1, \ldots, n_{i}\) do
                \(a_{i, j} \leftarrow \operatorname{median}_{\eta \in[\tau]}\left\{\left(\left|\left(A_{i} R_{i}^{-1} Z\right)_{j, \eta}\right| / \theta_{p}\right)^{p}\right\}\), where \(\theta_{p}\) is the median of \(\mathcal{D}_{p}\).
            end for
        end for
        Define a distribution \(\mathcal{D}=\left\{q_{1}^{\prime}, q_{1}^{\prime}, \ldots, q_{n}^{\prime}\right\}\) by \(q_{\sum_{i=1}^{q} j_{i} \prod_{l=1}^{j-1} n_{l}}^{\prime}=\prod_{i=1}^{q} a_{i, j_{i}}\).
        Let \(\Pi \in \mathbb{R}^{n \times n}\) denote a diagonal sampling matrix, where \(\Pi_{i, i}=1 / q_{i}^{1 / p}\) with probability
    \(q_{i}=\min \left\{1, r_{1} q_{i}^{\prime}\right\}\) and 0 otherwise, where \(r_{1}=\Theta\left(d^{3} / \epsilon^{2}\right) . \quad \triangleright\left[\mathrm{DDH}^{+} 09\right]\)
        Let \(x^{\prime} \in \mathbb{R}^{d}\) denote the solution of
                \(\min _{x \in \mathbb{R}^{d}}\left\|\Pi\left(A_{1} \otimes A_{2} \otimes \cdots \otimes A_{q}\right) x-\Pi b\right\|_{p}\)
        return \(x^{\prime} \quad \triangleright x^{\prime}\) is an \(O(1)\) approx: Lemma 26.8.4
    end procedure
    procedure \((1+\epsilon)\)-Approximate \(\ell_{p}\) REGRESSION \(\left(x^{\prime} \in \mathbb{R}^{d}\right)\)
        Implicitly define \(\rho=\left(A_{1} \otimes A_{2} \otimes \cdots \otimes A_{q}\right) x^{\prime}-b \in \mathbb{R}^{n}\)
        Via Lemma 26.8.7, compute a diagonal sampling matrix \(\Sigma \in \mathbb{R}^{n \times n}\) such that \(\Sigma_{i, i}=\)
    \(1 / \alpha_{i}^{1 / p}\) with probability \(\alpha_{i}=\min \left\{1, \max \left\{q_{i}, r_{2}\left|\rho_{i}\right|^{p} /\|\rho\|_{p}^{p}\right\}\right\}\) where \(r_{2}=\Theta\left(d^{3} / \epsilon^{3}\right)\).
        Compute \(\widehat{x}=\arg \min _{x \in \mathbb{R}^{d}}\left\|\Sigma\left(A_{1} \otimes A_{2} \otimes \cdots \otimes A_{q}\right)-\Sigma b\right\|_{p}\) (via convex optimization
    methods, e.g., [BCLL18])
        return \(\widehat{x}\)
    end procedure
```

will show that by restricting the regression problem to the sampled rows, we can obtain our desired ( $1 \pm \epsilon$ )-approximate solution efficiently.

Our main theorem for this section is stated below. A full proof of the theorem can be found in the supplementary, namely, Section 26.7.

Theorem 26.3.1 (Kronecker product $\ell_{2}$ regression). Let $D \in \mathbb{R}^{n \times n}$ be the diagonal row sampling matrix generated via Proposition 26.7.4, with $m=\Theta\left(1 /\left(\delta \epsilon^{2}\right)\right)$ non-zero entries, and let $A=\otimes_{i=1}^{q} A_{i}$, where $A_{i} \in \mathbb{R}^{n_{i} \times d_{i}}$, and $b \in \mathbb{R}^{n}$, where $n=\prod_{i=1}^{q} n_{i}$ and $d=\prod_{i=1}^{q} d_{i}$. Then let $\widehat{x}=\arg \min _{x \in \mathbb{R}^{d}}\|D A x-D b\|_{2}$, and let $x^{*}=\arg \min _{x^{\prime} \in \mathbb{R}^{d}}\|A x-b\|_{2}$. Then with probability $1-\delta$, we have

$$
\|A \widehat{x}-b\|_{2} \leq(1+\epsilon)\left\|A x^{*}-b\right\|_{2}
$$

Moreover, the total running time required to compute $\widehat{x}$ is $\widetilde{O}\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+(d q /(\delta \epsilon))^{O(1)}\right)$.

### 26.3.1 Kronecker Product $\ell_{p}$ Regression

We now consider $\ell_{p}$ regression for $1 \leq p<2$. Our algorithm is stated formally in Algorithm 26.2. Below is our main theorem, which demonstrates that Algorithm 26.2 approximately solves the Kronecker $\ell_{p}$ regression problem with input sparsity running time.

Theorem 26.3.2 (Main result, $\ell_{p}(1+\epsilon)$-approximate regression). Fix $1 \leq p<2$. Then for any constant $q=O(1)$, given matrices $A_{1}, A_{2}, \cdots, A_{q}$, where $A_{i} \in \mathbb{R}^{n_{i} \times d_{i}}$, let $n=\prod_{i=1}^{q} n_{i}$, $d=\prod_{i=1}^{q} d_{i}$. Let $\widehat{x} \in \mathbb{R}^{d}$ be the output of Algorithm 26.2. Then

$$
\left\|\left(A_{1} \otimes A_{2} \otimes \cdots \otimes A_{q}\right) \widehat{x}-b\right\|_{p} \leq(1+\epsilon) \min _{x \in \mathbb{R}^{n}}\left\|\left(A_{1} \otimes A_{2} \otimes \cdots \otimes A_{q}\right) x-b\right\|_{p}
$$

holds with probability at least $1-\delta$. In addition, our algorithm takes $\widetilde{O}\left(\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+\mathrm{nnz}(b)+(d / \epsilon)^{O(1)}\right)\right.$ time to output $\widehat{x} \in \mathbb{R}^{d}$.

We give a complete proof of the above theorem in the supplementary material, namely, in Section 26.8. Our high level approach follows that of [ $\left.\mathrm{DDH}^{+} 09\right]$. Namely, we first obtain
a vector $x^{\prime}$ which is an $O(1)$-approximate solution to the optimal solution. This is done by first constructing (implicitly) a matrix $U \in \mathbb{R}^{n \times d}$ that is a well-conditioned basis for the design matrix $A_{1} \otimes \cdots \otimes A_{q}$. We then efficiently sample rows of $U$ with probability proportional to their $\ell_{p}$ norm (which must be done without even explicitly computing most of $U)$. We then use the results of $\left[\mathrm{DDH}^{+} 09\right]$ to demonstrate that solving the regression problem constrained to these sampled rows gives a solution $x^{\prime} \in \mathbb{R}^{d}$ such that $\left\|\left(A_{1} \otimes \cdots \otimes A_{q}\right) x^{\prime}-b\right\|_{p} \leq$ $8 \min _{x \in \mathbb{R}^{d}}\left\|\left(A_{1} \otimes \cdots \otimes A_{q}\right) x^{\prime}-b\right\|_{p}$.

We define the residual error $\rho=\left(A_{1} \otimes \cdots \otimes A_{q}\right) x^{\prime}-b \in \mathbb{R}^{n}$ of $x^{\prime}$. Our goal is to sample additional rows $i \in[n]$ with probability proportional to their residual error $\left|\rho_{i}\right|^{p} /\|\rho\|_{p}^{p}$, and solve the regression problem restricted to the sampled rows. However, we cannot afford to compute even a small fraction of the entries in $\rho$ (even when $b$ is dense, and certainly not when $b$ is sparse). So to carry out this sampling efficiently, we design an involved, multi-part sketching and sampling routine (described in Section 26.8). This sampling technique is the main technical contribution of this section, and relies on a number of techniques, such as the Dyadic trick for quickly finding heavy hitters in the streaming literature, and a careful pre-processing step to avoid a poly $(d)$-blow up in the runtime. Given these samples, we can obtain the solution $\widehat{x}$ after solving the regression problem on the sampled rows, and the fact that this gives a $(1+\epsilon)$ approximate solution will follow from Theorem 6 of [ $\mathrm{DDH}^{+} 09$ ].

### 26.4 All-Pairs Regression

Given a matrix $A \in \mathbb{R}^{n \times d}$ and $b \in \mathbb{R}^{n}$, let $\bar{A} \in \mathbb{R}^{n^{2} \times d}$ be the matrix such that $\bar{A}_{i+(j-1) n, *}=A_{i, *}-A_{j, *}$, and let $\bar{b} \in \mathbb{R}^{n^{2}}$ be defined by $\bar{b}_{i+(j-1) n}=b_{i}-b_{j}$. Thus, $\bar{A}$ consists of all pairwise differences of rows of $A$, and $\bar{b}$ consists of all pairwise differences of rows of $b$,

The $\ell_{p}$ all pairs regression problem on the inputs $A, b$ is to solve $\min _{x \in \mathbb{R}^{d}}\|\bar{A} x-\bar{b}\|_{p}$.
First note that this problem has a close connection to Kronecker product regression. Namely, the matrix $\bar{A}$ can be written $\bar{A}=A \otimes \mathbf{1}^{n}-\mathbf{1}^{n} \otimes A$, where $\mathbf{1}^{n} \in \mathbb{R}^{n}$ is the all 1 's vector. Similarly, $\bar{b}=b \otimes \mathbf{1}^{n}-\mathbf{1}^{n} \otimes b$. For simplicity, we now drop the superscript and write $1=1^{n}$.

Our algorithm is given formally in Figure 26.3. We generate sparse $p$-stable sketches $S_{1}, S_{2} \in \mathbb{R}^{k \times n}$, where $k=(d /(\epsilon \delta))^{O(1)}$. We compute $M=\left(S_{1} \otimes S_{2}\right)(F \otimes \mathbf{1}-\mathbf{1} \otimes F)=$ $S_{1} F \otimes S_{2} \mathbf{1}-S_{1} \mathbf{1} \otimes S_{2} F$, where $F=[A, b]$. We then take the $Q R$ decomposition $M=Q R$. Finally, we sample rows of $(F \otimes \mathbf{1}-\mathbf{1} \otimes F) R^{-1}$ with probability proportional to their $\ell_{p}$ norms. This is done by an involved sampling procedure described in Lemma 26.4.2, which is similar to the sampling procedure used in the proof of Theorem 26.3.2. Finally, we solve the regression problem $\min _{x}\|\Pi(\bar{A} x-\bar{b})\|_{p}$, where $\Pi$ is the diagonal row-sampling matrix constructed by the sampling procedure. We summarize the guarantee of our algorithm in the following theorem.

Theorem 26.4.1. Given $A \in \mathbb{R}^{n \times d}$ and $b \in \mathbb{R}^{n}$, for $p \in[1,2]$, let $\bar{A}=A \otimes \mathbf{1}-\mathbf{1} \otimes A \in \mathbb{R}^{n^{2} \times d}$ and $\bar{b}=b \otimes \mathbf{1}-\mathbf{1} \otimes b \in \mathbb{R}^{n^{2}}$. Then there is an algorithm for that outputs $\widehat{x} \in \mathbb{R}^{d}$ such that with probability $1-\delta$ we have $\|\bar{A} \widehat{x}-\bar{b}\|_{p} \leq(1+\epsilon) \min _{x \in \mathbb{R}^{d}}\|\bar{A} x-\bar{b}\|_{p}$. The running time is $\widetilde{O}\left(\mathrm{nnz}(A)+(d /(\epsilon \delta))^{O(1)}\right)$.

The theorem crucially utilizes our fast $\ell_{p}$ sampling routine, which is described in Figure 26.5 in the supplementary. A full discussion and proof of the lemma can be found in the supplementary material 26.9.1.

```
Algorithm 26.3 Our All-Pairs Regression Algorithm
    procedure All-Pairs Regression \((A, b)\)
        \(F=[A, b] \in \mathbb{R}^{n \times d+1} . r \leftarrow \operatorname{poly}(d / \epsilon)\)
        Generate \(S_{1}, S_{2} \in \mathbb{R}^{k \times n}\) sparse \(p\)-stable transforms for \(k=\operatorname{poly}(d /(\epsilon \delta))\).
        Sketch \(\left(S_{1} \otimes S_{2}\right)(F \otimes \mathbf{1}-\mathbf{1} \otimes F)\).
        Compute \(Q R\) decomposition: \(\left(S_{1} \otimes S_{2}\right)(F \otimes \mathbf{1}-\mathbf{1} \otimes F)=Q R\).
        Let \(M=(F \otimes \mathbf{1}-\mathbf{1} \otimes F) R^{-1}\), and \(\sigma_{i}=\left\|M_{i, *}\right\|_{p}^{p} /\|M\|_{p}^{p}\).
        Obtain row sampling diagonal matrix \(\Pi \in \mathbb{R}^{n \times n}\) such that \(\Pi_{i, i}=1 / \widetilde{q}_{i}^{1 / p}\) indepen-
    dently with probability \(q_{i} \geq \min \left\{1, r \sigma_{i}\right\}\), where \(\widetilde{q}_{i}=\left(1 \pm \epsilon^{2}\right) q_{i}\). \(\quad\) Lemma
    26.4.2
        return \(\widehat{x}\), where \(\widehat{x}=\arg \min _{x \in \mathbb{R}^{d}}\|\Pi(\bar{A} x-\bar{b})\|_{p}\).
    end procedure
```

Lemma 26.4.2 (Fast $\ell_{p}$ sampling). Given $R \in \mathbb{R}^{d+1 \times d+1}$ and $F=[A, b] \in \mathbb{R}^{n \times d+1}$, there is an algorithm that, with probability $1-n^{-c}$ for any constant $c$, produces a diagonal matrix $\Pi \in \mathbb{R}^{n^{2} \times n^{2}}$ such that $\Pi_{i, i}=1 / \widetilde{q}_{i}^{1 / p}$ with probability $q_{i} \geq \min \left\{1, r\left\|M_{i, *}\right\|_{p}^{p} /\|M\|_{p}^{p}\right\}$ and $\Pi_{i, i}=0$ otherwise, where $r=\operatorname{poly}(d / \epsilon)$ and $M=(F \otimes \mathbf{1}-\mathbf{1} \otimes F) R^{-1}$, and $\widetilde{q}_{i}=\left(1 \pm \epsilon^{2}\right) q_{i}$ for all $i \in\left[n^{2}\right]$. The total time required is $\widetilde{O}(\mathrm{nnz} A+\operatorname{poly}(d / \epsilon))$.

### 26.5 Low Rank Approximation of Kronecker Product Matrices

We now consider low rank approximation of Kronecker product matrices. Given $q$ matrices $A_{1}, A_{2}, \ldots, A_{q}$, where $A_{i} \in \mathbb{R}^{n_{i} \times d_{i}}$, the goal is to output a rank- $k$ matrix $B \in$ $\mathbb{R}^{n \times d}$, where $n=\prod_{i=1}^{q} n_{i}$ and $d=\prod_{i=1}^{q} d_{i}$, such that $\|B-A\|_{F} \leq(1+\epsilon) \mathrm{OPT}_{k}$, where $\mathrm{OPT}_{k}=\min _{\text {rank }-k} A^{\prime}\left\|A^{\prime}-A\right\|_{F}$, and $A=\otimes_{i=1}^{q} A_{i}$. Our approach employs the Count-Sketch distribution of matrices [CW13, Woo14b]. A count-sketch matrix $S$ is generated as follows. Each column of $S$ contains exactly one non-zero entry. The non-zero entry is placed in a uniformly random row, and the value of the non-zero entry is either 1 or -1 chosen uniformly
at random.

Our algorithm is as follows. We sample $q$ independent Count-Sketch matrices $S_{1}, \ldots S_{q}$, with $S_{i} \in \mathbb{R}^{k_{i} \times n_{i}}$, where $k_{1}=\cdots=k_{q}=\Theta\left(q k^{2} / \epsilon^{2}\right)$. We then compute $M=\left(\otimes_{i=1}^{q} S_{i}\right) A$, and let $U \in \mathbb{R}^{k \times d}$ be the top $k$ right singular vectors of $M$. Finally, we output $B=A U^{\top} U$ in factored form (as $q+1$ separate matrices, $A_{1}, A_{2}, \ldots, A_{q}, U$ ), as the desired rank- $k$ approximation to $A$. The following theorem demosntrates the correctness of this algorithm. A full proof can be found in Section 26.10.

Theorem 26.5.1. For any constant $q \geq 2$, there is an algorithm which runs in time $O\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+d \operatorname{poly}(k / \epsilon)\right)$ and outputs a rank $k$-matrix $B$ in factored form such that $\|B-A\|_{F} \leq(1+\epsilon) \mathrm{OPT}_{k}$ with probability $9 / 10 .{ }^{4}$

### 26.6 Numerical Simulations

In our numerical simulations, we compare our algorithms to two baselines: (1) brute force, i.e., directly solving regression without sketching, and (2) the methods based sketching developed in [DSSW18]. All methods were implemented in Matlab on a Linux machine. We remark that in our implementation, we simplified some of the steps of our theoretical algorithm, such as the residual sampling algorithm (Alg. 26.4). We found that in practice, even with these simplifications, our algorithms already demonstrated substantial improvements over prior work.

Following the experimental setup in [DSSW18], we generate matrices $A_{1} \in \mathbb{R}^{300 \times 15}$,

[^91]$A_{2} \in \mathbb{R}^{300 \times 15}$, and $b \in \mathbb{R}^{300^{2}}$, such that all entries of $A_{1}, A_{2}, b$ are sampled i.i.d. from a normal distribution. Note that $A_{1} \otimes A_{2} \in \mathbb{R}^{90000 \times 225}$. We define $T_{\mathrm{bf}}$ to be the time of the brute force algorithm, $T_{\text {old }}$ to be the time of the algorithms from [DSSW18], and $T_{\text {ours }}$ to be the time of our algorithms. We are interested in the time ratio with respect to the brute force algorithm and the algorithms from [DSSW18], defined as, $r_{t}=T_{\text {ours }} / T_{\text {bf }}$, and $r_{t}^{\prime}=T_{\text {ours }} / T_{\text {old }}$. The goal is to show that our methods are significantly faster than both baselines, i.e., both $r_{t}$ and $r_{t}^{\prime}$ are significantly less than 1 .

We are also interested in the quality of the solutions computed from our algorithms, compared to the brute force method and the method from [DSSW18]. Denote the solution from our method as $x_{\text {our }}$, the solution from the brute force method as $x_{\mathrm{bf}}$, and the solution from the method in [DSSW18] as $x_{\text {old }}$. We define the relative residual percentage $r_{e}$ and $r_{e}^{\prime}$ to be:

$$
r_{e}=100 \frac{\mid\left\|\mathcal{A} x_{\mathrm{ours}}-b\right\|-\left\|\mathcal{A} x_{\mathrm{bf}}-b\right\| \|}{\left\|\mathcal{A} x_{\mathrm{bf}}-b\right\|}, \quad r_{e}^{\prime}=100 \frac{\left|\left\|\mathcal{A} x_{\mathrm{old}}-b\right\|-\left\|\mathcal{A} x_{\mathrm{bf}}-b\right\|\right|}{\left\|\mathcal{A} x_{\mathrm{bf}}-b\right\|}
$$

Where $\mathcal{A}=A_{1} \otimes A_{2}$. The goal is to show that $r_{e}$ is close zero, i.e., our approximate solution is comparable to the optimal solution in terms of minimizing the error $\|\mathcal{A} x-b\|$.

Throughout the simulations, we use a moderate input matrix size so that we can accommodate the brute force algorithm and to compare to the exact solution. We consider varying values of $m$, where $M$ denotes the size of the sketch (number of rows) used in either the algorithms of [DSSW18] or the algorithms in this paper. We also include a column $m / n$ in the table, which is the ratio between the size of the sketch and the original matrix $A_{1} \otimes A_{2}$. Note in this case that $n=90000$.

Table 26.1: Results for $\ell_{2}$ and $\ell_{1}$-regression with respect to different sketch sizes $m$.

|  | $m$ | $m / n$ | $r_{e}$ | $r_{e}^{\prime}$ | $r_{t}$ | $r_{t}^{\prime}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\ell_{2}$ | 8100 | .09 | $2.48 \%$ | $1.51 \%$ | 0.05 | 0.22 |
|  | 12100 | .13 | $1.55 \%$ | $0.98 \%$ | 0.06 | 0.24 |
|  | 16129 | .18 | $1.20 \%$ | $0.71 \%$ | 0.07 | 0.08 |
|  | 2000 | .02 | $7.72 \%$ | $9.10 \%$ | 0.02 | 0.59 |
|  | 4000 | .04 | $4.26 \%$ | $4.00 \%$ | 0.03 | 0.75 |
|  | 8000 | .09 | $1.85 \%$ | $1.6 \%$ | 0.07 | 0.83 |
|  | 12000 | .13 | $1.29 \%$ | $0.99 \%$ | 0.09 | 0.79 |
|  | 16000 | .18 | $1.01 \%$ | $0.70 \%$ | 0.14 | 0.90 |

Simulation Results for $\ell_{2}$ We first compare our algorithm, Alg. 26.1, to baselines under the $\ell_{2}$ norm. In our implementation, $\min _{x}\|A x-b\|_{2}$ is solved by Matlab backslash $A \backslash b$. Table 26.1 summarizes the comparison between our approach and the two baselines. The numbers are averaged over 5 random trials. First of all, we notice that our method in general provides slightly less accurate solutions than the method in [DSSW18], i.e., $r_{e}>r_{e}^{\prime}$ in this case. However, comparing to the brute force algorithm, our method still generates relatively accurate solutions, especially when $m$ is large, e.g., the relative residual percentage w.r.t. the optimal solution is around $1 \%$ when $m \approx 16000$. On the other hand, as suggested by our theoretical improvements for $\ell_{2}$, our method is significantly faster than the method from [DSSW18], consistently across all sketch sizes $m$. Note that when $m \approx 16000$, our method is around 10 times faster than the method in [DSSW18]. For small $m$, our approach is around 5 times faster than the method in [DSSW18].

Simulation Results for $\ell_{1}$ We compare our algorithm, Alg. 26.2, to two baselines under the $\ell_{1}$-norm. The first is a brute-force solution, and the second is the algorithm for [DSSW18].

For $\min _{x}\|A x-b\|_{1}$, the brute for solution is obtained via a Linear Programming solver in Gurobi [GO16]. Table 26.1 summarizes the comparison of our approach to the two baselines under the $\ell_{1}$-norm. The statistics are averaged over 5 random trials. Compared to the Brute Force algorithm, our method is consistently around 10 times faster, while in general we have relative residual percentage around $1 \%$. Compared to the method from [DSSW18], our approach is consistently faster (around 1.3 times faster). Note our method has slightly higher accuracy than the one from [DSSW18] when the sketch size is small, but slightly worse accuracy when the sketch size increases.

### 26.7 Missing Proofs from Section 26.3

In this section, we prove correctness of our $\ell_{2}$ Kronecker product regression algorithm. Specifically, we prove Theorem 26.3.1. To prove correctness, we need to establish several facts about the leverage scores of a Kronecker product.

Proposition 26.7.1. Let $U_{i} \in \mathbb{R}^{n_{i} \times d_{i}}$ be an orthonormal basis for $A_{i} \in \mathbb{R}^{n_{i} \times d_{i}}$. Then $U=\otimes_{i=1}^{q} U_{i}$ is an orthonormal basis for $A=\otimes_{i=1}^{q} A_{i}$.

Proof. Note that the column norm of each column of $U$ is the product of column norms of the $U_{i}$ 's, which are all 1 . Thus $U$ has unit norm columns. It suffices then to show that all the singular values of $U$ are 1 or -1 , but this follows from the fact that the singular values of $U$ are the product of singular values of the $U_{i}{ }^{\prime}$ 's, which completes the proof.

Corollary 26.7.2. Let $A=\otimes_{i=1}^{q} A_{i}$, where $A_{i} \in \mathbb{R}^{n_{i} \times d_{i}}$. Fix any $\vec{i}=\left(i_{1}, \ldots, i_{q}\right) \in\left[n_{1}\right] \times$ $\left[n_{2}\right] \times \cdots \times\left[n_{q}\right]$, and let $\vec{i}$ index into a row of $A$ in the natural way. Then the $\vec{i}$-th leverage score of $A$ is equal to $\prod_{j=1}^{q} \sigma_{i_{j}}\left(A_{j}\right)$, where $\sigma_{t}(B)$ is the $t$-th leverage score of a matrix $B$.

Proof. Note $U=\otimes_{i=1}^{q} U_{i}$ is an orthonormal basis for $A=\otimes_{i=1}^{q} A_{i}$ by the prior Proposition. Now if $U_{\vec{i}, *}$ is the $\vec{i}$-th row of $U$, then by fundamental properties of Kronecker products [VL00], we have $\left\|U_{\vec{i}, *}\right\|_{2}=\prod_{j=1}^{q}\left\|\left(U_{j}\right)_{i_{j}, *}\right\|_{2}$, which completes the proof. Note here that we used the fact that leverage scores are independent of the choice of orthonormal basis [Woo14b].

Proposition 26.7.3 (Theorem 29 of [CW13]). Given a matrix $A \in \mathbb{R}^{n \times d}$, let $\sigma \in \mathbb{R}^{n}$ be the $\ell_{2}$ leverage scores of $A$ (see definition 26.2.3). Then there is an algorithm which computes values $\widetilde{\sigma}_{1}, \widetilde{\sigma}_{2}, \ldots, \widetilde{\sigma}_{n}$ such that $\widetilde{\sigma}_{i}=(1 \pm \epsilon) \sigma_{i}$ simultaneously for all $i \in[n]$ with probability $1-1 / n^{c}$ for any constant $c \geq 1$. The runtime is $\widetilde{O}(\mathrm{nnz}(A)+\operatorname{poly}(d / \epsilon))$.

Proposition 26.7.4. Given $A=\otimes_{i=1}^{q} A_{i}$, where $A_{i} \in \mathbb{R}^{n_{i} \times d_{i}}$, there is an algorithm which, with probability $1-1 / n^{c}$ for any constant $c \geq 1$, outputs a diagonal matrix $D \in \mathbb{R}^{n \times n}$ with $m$ non-zeros entries, such that $D_{i, i}=1 /\left(m \widetilde{\sigma}_{i}\right)$ is non-zero with probability $\widetilde{\sigma}_{i} \in(1 \pm 1 / 10) \sigma_{i}(A)$. The time required is $\widetilde{O}\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+\operatorname{poly}(d q / \epsilon)+m q\right)$.

Proof. By Proposition 26.7.3, we can compute approximate leverage scores of each $A_{i}$ up to error $\Theta(1 / q)$ in time $\widetilde{O}\left(\operatorname{nnz}\left(A_{i}\right)+\operatorname{poly}(d / \epsilon)\right)$ with high probability. To sample a leverage score from $A$, it suffices to sample one leverage score from each of the $A_{i}$ 's by Corollary 26.7.2. The probability that a given row $\vec{i}=\left(i_{1}, \ldots, i_{q}\right) \in\left[n_{1}\right] \times\left[n_{2}\right] \times \cdots \times\left[n_{q}\right]$ of $A$ is chosen is $\prod_{j=1}^{q} \widetilde{\sigma}\left(A_{j}\right)_{i_{j}}=(1 \pm \Theta(1 / q))^{q} \sigma_{\vec{i}}(A)=(1 \pm 1 / 10) \sigma_{\vec{i}}(A)$ as needed. Obtaining a sample takes $\widetilde{O}(1)$ time per $A_{i}$ (since a random number needs to be generated to $O(\log (n))$ bits of precision in expectation and with high probability to obtain this sample), thus $O(q)$ time overall, so repeating the sampling $M$ times gives the desired additive $m q$ runtime.

The $q=1$ version of the following result can be found in [CW13, SWZ19b].

Proposition 26.7.5. Let $D \in \mathbb{R}^{n \times n}$ be the diagonal row sampling matrix generated via Proposition 26.7.4, with $m=\Theta\left(1 /\left(\delta \epsilon^{2}\right)\right)$ non-zero entries. Let $A=\otimes_{i=1}^{q} A_{i}$ as above, and let $U \in \mathbb{R}^{n \times r}$ be an orthonormal basis for the column span of $A$, where $r=\operatorname{rank}(A)$. Then for any matrix $B$ with $n$ rows, we have

$$
\operatorname{Pr}\left[\left\|U^{\top} D^{\top} D B-U^{\top} B\right\|_{F} \leq \epsilon\|U\|_{F}\|B\|_{F}\right] \geq 1-\delta
$$

Proof. By definition of leverage scores and Proposition 26.7.4, $D$ is a matrix which sample each row $U_{i, *}$ of $U$ with probability at least $(9 / 10)\left\|U_{i, *}\right\|_{2} /\|U\|_{F}$. Taking the average of $m$ such rows, we obtain the approximate matrix product result with error $O(1 / \sqrt{\delta m})$ with probability $1-\delta$ by Theorem 2.1 of [KV17].

We our now ready to prove the main theorem of this section, Theorem 26.3.1
Theorem 26.3.1(Kronecker product $\ell_{2}$ regression) Let $D \in \mathbb{R}^{n \times n}$ be the diagonal row sampling matrix generated via Proposition 26.7.4, with $m=\Theta\left(1 /\left(\delta \epsilon^{2}\right)\right)$ non-zero entries, and let $A=\otimes_{i=1}^{q} A_{i}$, where $A_{i} \in \mathbb{R}^{n_{i} \times d_{i}}$, and $b \in \mathbb{R}^{n}$, where $n=\prod_{i=1}^{q} n_{i}$ and $d=\prod_{i=1}^{q} d_{i}$. Then we have let $\widehat{x}=\arg \min _{x \in \mathbb{R}^{d}}\|D A x-D b\|_{2}$, and let $x^{*}=\arg \min _{x^{\prime} \in \mathbb{R}^{d}}\|A x-b\|_{2}$. Then with probability $1-\delta$, we have

$$
\|A \widehat{x}-b\|_{2} \leq(1+\epsilon)\left\|A x^{*}-b\right\|_{2}
$$

Moreover, the total runtime requires to compute $\widehat{x}$ is

$$
\widetilde{O}\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+(d q /(\delta \epsilon))^{O(1)}\right) .
$$

Proof. Let $U$ be an orthonormal basis for the column span of $A$. By Lemma 3.3 of [CW09], we have $\left\|A\left(\widehat{x}-x^{*}\right)\right\|_{2} \leq 2 \sqrt{\epsilon}\left\|A x^{*}-b\right\|_{2}$. Note that while Lemma 3.3 of [CW09] uses a different sketching matrix $D$ than us, the only property required for the proof of Lemma 3.3 is that $\mid U^{\top} D^{\top} D B-U^{\top} B\left\|_{F} \leq \sqrt{\epsilon / d}\right\| A\left\|_{F}\right\| B \|_{F}$ with probability at least $1-\delta$ for any fixed matrix $B$, which we obtain by Proposition 26.7 .5 by having $O\left(d /\left(\delta \epsilon^{2}\right)\right)$ non-zeros on the diagonal of $D)$. By the normal equations, we have $A^{\top}\left(A x^{*}-b\right)=0$, thus $\left\langle A\left(\widehat{x}-x^{*}\right),\left(A x^{*}-b\right)\right\rangle=0$, and so by the Pythagorean theorem we have

$$
\|A \widehat{x}-b\|_{2}^{2}=\left\|A x^{*}-b\right\|_{2}^{2}+\left\|A\left(\widehat{x}-x^{*}\right)\right\|_{2}^{2} \leq(1+4 \epsilon)\left\|A x^{*}-b\right\|_{2}^{2}
$$

Which completes the proof after rescaling of $\epsilon$. The runtime required to obtain the matrix $D$ is $\widetilde{O}\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+\operatorname{poly}(d q / \epsilon)\right)$ by Proposition 26.7.4, where we set $D$ to have $m=$ $\Theta\left(d /\left(\delta \epsilon^{2}\right)\right)$ non-zero entries on the diagonal. Once $D$ is obtained, one can compute $D(A+b)$ in time $O(m d)$, thus the required time is $O\left(\delta^{-1}(d / \epsilon)^{2}\right)$. Finally, computing $\widehat{x}$ once $D A, D b$ are computed requires a single pseudo-inverse computation, which can be carried out in $O\left(\delta^{-1} d^{3} / \epsilon^{2}\right)$ time (since $D A$ now has only $O\left(\delta^{-1}(d / \epsilon)^{2}\right)$ rows).

### 26.8 Missing Proofs from Section 26.3.1

We now give a complete proof of Theorem 26.3.2. Our high level approach follows that of [ $\left.\mathrm{DDH}^{+} 09\right]$. Namely, we first obtain a vector $x^{\prime}$ which is a $O(1)$ approximate solution to the optimal, and then use the residual error $\rho \in \mathbb{R}^{d}$ of $x^{\prime}$ to refine $x^{\prime}$ to a ( $1 \pm \epsilon$ ) approximation $\widehat{x}$. The fact that $x^{\prime}$ is a constant factor approximation follows from our Lemma 26.8.4. Given $x^{\prime}$, by Lemma 26.8.7 we can efficiently compute the matrix $\Sigma$ which samples from
the coordinates of the residual error $\rho=\left(A_{1} \otimes \cdots \otimes A_{q}\right) x^{\prime}-b$ in the desired runtime. The sampling lemma is the main technical lemma, and requires a careful multi-part sketching and sampling routine. Given this $\Sigma$, the fact that $\widehat{x}$ is a $(1+\epsilon)$ approximate solution follows directly from Theorem 6 of $\left[\mathrm{DDH}^{+} 09\right]$. Our main theorem and its proof is stated below. The proof will utilize the lemmas and sampling algorithm developed in the secitons which follow.

Theorem 26.3.2 (Main result, $\ell_{p}, 1+\epsilon$-approximation). Fix $1 \leq p<2$. Then for any constant $q=O(1)$, given matrices $A_{1}, A_{2}, \cdots, A_{q}$, where $A_{i} \in \mathbb{R}^{n_{i} \times d_{i}}$, let $n=\prod_{i=1}^{q} n_{i}$, $d=\prod_{i=1}^{q} d_{i}$. Let $\widehat{x} \in \mathbb{R}^{d}$ be the output of Algorithm 26.2. Then

$$
\left\|\left(A_{1} \otimes A_{2} \otimes \cdots \otimes A_{q}\right) \widehat{x}-b\right\|_{p} \leq(1+\epsilon) \min _{x \in \mathbb{R}^{n}}\left\|\left(A_{1} \otimes A_{2} \otimes \cdots \otimes A_{q}\right) x-b\right\|_{p}
$$

holds with probability at least $1-\delta$. In addition, our algorithm takes

$$
\widetilde{O}\left(\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+\mathrm{nnz}(b)+(d / \epsilon)^{O(1)}\right) \log (1 / \delta)\right)
$$

time to output $\widehat{x} \in \mathbb{R}^{d}$.

Proof. By Lemma 26.8.4, the output $x^{\prime}$ in line 16 of algorithm 26.3 .1 is an 8 approximation of the optimal solution, and $x^{\prime}$ is obtained in time $\widetilde{O}\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+(d q / \epsilon)^{O(1)}\right)$. We then obtain the residual error $\rho=\left(A_{1} \otimes \cdots \otimes A_{q}\right) x^{\prime}-b$ (implicitly). By Theorem 6 of [ $\left.\mathrm{DDH}^{+} 09\right]$, if we let $\Sigma \in \mathbb{R}^{n \times n}$ be a row sampling matrix where $\Sigma_{i, i}=1 / \alpha_{i}^{1 / p}$ with probability $\alpha_{i}=$ $\min \left\{1, \max \left\{q_{i}, r_{2} \frac{\left|\rho_{i}\right|^{p}}{\|\rho\|_{p}^{p}}\right\}\right.$, where $q_{i}$ is the row sampling probability used in the sketch $\Pi$ from which $x^{\prime}$ was obtained, and $r_{2}=O\left(d^{3} / \epsilon^{2} \log (1 / \epsilon)\right)$, then the solution to $\min _{x} \| \Sigma\left(A_{1} \otimes \cdots \otimes\right.$ $\left.A_{q}\right) x-\Sigma b \|_{p}$ will be a $(1+\epsilon)$ approximately optimal solution. By Lemma 26.8.7, we can obtain such a matrix $\Sigma$ in time $\widetilde{O}\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+q \mathrm{nnz}(b)+\left(d \log (n) /(\epsilon \delta)^{O\left(q^{2}\right)}\right)\right.$, which completes the proof of correctness. Finally, note that we can solve the sketched regression problem
$\min _{x}\left\|\Sigma\left(A_{1} \otimes \cdots \otimes A_{q}\right) x-\Sigma b\right\|_{p}$ which has $O\left((d \log (n) / \epsilon)^{O\left(q^{2}\right)}(1 / \delta)\right)$ constraints and $d$ variables in time $O\left((d \log (n) / \epsilon)^{O\left(q^{2}\right)}(1 / \delta)\right)$ using linear programming for $p=1$ (see [CLS19] for the state of the art linear program solver), or more generally interior point methods for convex programming for $p>1$ (see [BCLL18] for the state of the art $\ell_{p}$ solver).

Now to boost the failure probability from a $O(1 / \delta)$ to $\log (1 / \delta)$ dependency, we do the following. We run the above algorithm with $\delta=1 / 10$, so that our output $\widehat{x} \in \mathbb{R}^{d}$ is a $(1+\epsilon)$ approximation with probability $9 / 10$. Now note that we actually have an $(1+\epsilon)$ estimate of the cost $\left\|\left(A_{1} \otimes A_{2} \otimes \cdots \otimes A_{q}\right) \widehat{x}-b\right\|_{p}$ of the solution $\widehat{x}$, which is simply given by $\left\|\Sigma\left(A_{1} \otimes A_{2} \otimes \cdots \otimes A_{q}\right) \widehat{x}-\Sigma b\right\|_{p}$ where $\Sigma$ is the sampling matrix used to compute $\widehat{x}$. Thus we can simply repeat the above process $O(\log (1 / \delta))$ times, and take the solution with the minimal cost overall.

We start by defining a tensor operation which will be useful for our analysis.

Definition 26.8.1 $(((\cdot, \ldots, \cdot), \cdot)$ operator for tensors and matrices). Given tensor $A \in$ $\mathbb{R}^{d_{1} \times d_{2} \times \cdots \times d_{q}}$ and matrices $B_{i} \in \mathbb{R}^{n_{i} \times d_{i}}$ for $i \in[q]$, we define the tensor $\left(\left(B_{1}, B_{2}, \ldots, B_{q}\right), A\right) \in$ $\mathbb{R}^{n_{1} \times n_{2} \times \cdots \times n_{q}}:$

$$
\left(\left(B_{1}, B_{2}, \ldots, B_{q}\right), A\right)_{i_{1}, \ldots, i_{q}}=\sum_{i_{1}^{\prime}=1}^{d_{1}} \sum_{i_{2}^{\prime}=1}^{d_{2}} \cdots \sum_{i_{q}^{\prime}=1}^{d_{q}} A_{i_{1}^{\prime}, i_{2}^{\prime}, \ldots, i_{q}^{\prime}} \prod_{\ell=1}^{q}\left(B_{\ell}\right)_{i_{\ell}, i_{\ell}^{\prime}}
$$

Observe for the case of $q=2$, we just have $\left(\left(B_{1}, B_{2}\right), A\right)=B_{1} A B_{2}^{\top} \in \mathbb{R}^{n_{1} \times n_{2}}$.

Using the above notation, we first prove a result about reshaping tensors.
Lemma 26.8.1 (Reshaping). Given matrices $A_{1}, A_{2}, \cdots, A_{q} \in \mathbb{R}^{n_{i} \times d_{i}}$ and a tensor $B \in$ $\mathbb{R}^{n_{1} \times n_{2} \times \cdots \times n_{q}}$, let $n=\prod_{i=1}^{q} n_{i}$ and let $d=\prod_{i=1}^{d} d_{i}$. Let $b$ denote the vectorization of $B$.

For any tensor $X \in \mathbb{R}^{d_{1} \times d_{2} \times \cdots \times d_{q}}$, we have $\left\|\left(\left(A_{1}, A_{2}, \cdots, A_{q}\right), X\right)-B\right\|_{\xi}$ is equal to $\|\left(A_{1} \otimes\right.$ $\left.A_{2} \otimes \cdots \otimes A_{q}\right) x-b \|_{\xi}$ where $\xi$ is any entry-wise norm (such as an $\ell_{p}$-norm) and $x$ is the vectorization of $X$. See Definition 26.8.1 of the $((\cdot, \ldots, \cdot), \cdot)$ tensor operator.

Observe, for the case of $q=2$, this is equivalent to the statement that $\| A_{1} X A_{2}^{\top}-$ $B\left\|_{\xi}=\right\|\left(A_{1} \otimes A_{2}\right) x-b \|_{\xi}$.

Proof. For the pair $x \in \mathbb{R}^{d}, X \in \mathbb{R}^{d_{1} \times d_{2} \times \cdots \times d_{q}}$, the connection is the following: $\forall i_{1} \in$ $\left[d_{1}\right], \ldots, i_{q} \in\left[d_{q}\right]$,

$$
x_{i_{1}+\sum_{l=2}^{q}\left(i_{l}-1\right) \cdot \prod_{t=1}^{l-1} d_{t}}=X_{i_{1}, \cdots, i_{q}} .
$$

Similarly, for $b \in \mathbb{R}^{n}, B \in \mathbb{R}^{n_{1} \times n_{2} \times \cdots \times n_{q}}$, for any $j_{1}, \in\left[n_{1}\right], \ldots, j_{q} \in\left[n_{q}\right]$,

$$
b_{j_{1}+\sum_{l=2}^{q}\left(j_{l}-1\right) \cdot \prod_{t=1}^{l-1} n_{t}}=B_{j_{1}, j_{2}, \cdots, j_{q}} .
$$

For simplicity, for any $\left(i_{1}, \ldots, i_{q}\right) \in\left[d_{1}\right] \times \cdots \times\left[d_{q}\right]$ and $\left(j_{1}, \ldots, j_{q}\right) \in\left[n_{1}\right] \times \cdots \times\left[n_{q}\right]$ we define $\vec{i}=i_{1}+\sum_{l=2}^{q}\left(i_{l}-1\right) \cdot \prod_{t=1}^{l-1} d_{t}$ and similarly $\vec{j}=j_{1}+\sum_{l=2}^{q}\left(j_{l}-1\right) \cdot \prod_{t=1}^{l-1} n_{t}$. Then we can simplify the above relation and write $x_{\vec{i}}=X_{i_{1}, i_{2}, \cdots, i_{q}}$, and $b_{\vec{j}}=B_{j_{1}, j_{2}, \cdots, j_{q}}$.

For a matrix $Z$, let $Z_{i, *}$ denote the $i$-th row of $Z$. We consider the $\vec{j}$-th entry of $\left(A_{1} \otimes A_{2} \otimes \cdots \otimes A_{q}\right) x$,

$$
\begin{aligned}
\left(\left(A_{1} \otimes A_{2} \otimes \cdots \otimes A_{q}\right) x\right)_{\vec{j}} & =\left\langle\left(A_{1} \otimes A_{2} \otimes \cdots \otimes A_{q}\right)_{\vec{j}, *} \cdot x\right\rangle \\
& =\sum_{i_{1}=1}^{d_{1}} \sum_{i_{2}=1}^{d_{2}} \cdots \sum_{i_{q}=1}^{d_{q}}\left(\prod_{l=1}^{q}\left(A_{l}\right)_{j_{l}, i_{l}}\right) \cdot x_{\vec{i}} \\
& =\sum_{i_{1}=1}^{d_{1}} \sum_{i_{2}=1}^{d_{2}} \cdots \sum_{i_{q}=1}^{d_{q}}\left(\prod_{l=1}^{q}\left(A_{l}\right)_{j_{l}, i_{l}}\right) \cdot X_{i_{1}, i_{2}, \cdots, i_{q}} \\
& =\left(\left(A_{1}, A_{2}, \cdots, A_{q}\right), X\right)_{j_{1}, \ldots, j_{q}} .
\end{aligned}
$$

Where the last equality is by Definition (26.8.1). Since we also have $b_{\vec{j}}=B_{j_{1}, \ldots, j_{q}}$, this completes the proof of the Lemma.

### 26.8.1 Sampling from an $\ell_{p}$-Well-Conditioned Base

In this Section, we discuss the first half of Algorithm 26.2 which computes $x^{\prime} \in \mathbb{R}^{d}$, which we will show is a $O(1)$-approximate solution to the optimal. First note that by Lemma 26.2.1 together with fact 26.2.2, we know that $A_{i} R_{i}^{-1}$ is an $\ell_{p}$ well conditioned basis for $A_{i}$ (recall this means that $A_{i} R_{i}^{-1}$ is a $(\alpha, \beta, p)$ well conditioned basis for $A$, and $\beta / \alpha=d_{i}^{O(1)}$ ) with probability $1-O(1 / q)$, and we can then union bound over this occurring for all $i \in[q]$. Given this, we now prove that $\left(A_{1} R_{1}^{-1} \otimes A_{2} R_{2}^{-1} \otimes \cdots \otimes A_{q} R_{q}^{-1}\right)$ is a well conditioned basis for $\left(A_{1} \otimes A_{2} \otimes \cdots \otimes A_{q}\right)$.

Lemma 26.8.2. Let $A_{i} \in \mathbb{R}^{n_{i} \times d_{i}}$ and $R_{i} \in \mathbb{R}^{d_{i} \times d_{i}}$. Then if $A_{i} R_{i}^{-1}$ is a $\left(\alpha_{i}, \beta_{i}, p\right)$ wellconditioned basis for $A_{i}$ for $i=1,2, \ldots, q$, we have for all $x \in \mathbb{R}^{d_{1} \cdots d_{q}}$ :

$$
\prod_{i=1}^{q} \alpha_{i}\|x\|_{p} \leq\left\|\left(A_{1} R_{1}^{-1} \otimes A_{2} R_{2}^{-1} \otimes \cdots \otimes A_{q} R_{q}^{-1}\right) x\right\|_{p} \leq \prod_{i=1}^{q} \beta_{i}\|x\|_{p}
$$

Proof. We first consider the case of $q=2$. We would like to prove

$$
\alpha_{1} \alpha_{2}\|x\|_{p} \leq\left\|\left(A_{1} R_{1}^{-1} \otimes A_{2} R_{2}^{-1}\right) x\right\|_{p} \leq \beta_{1} \beta_{2}\|x\|_{p}
$$

First note, by the reshaping Lemma 26.8.1, this is equivalent to

$$
\alpha_{1} \alpha_{2}\|X\|_{p} \leq\left\|A_{1} R_{1}^{-1} X\left(R_{2}^{-1} A_{2}\right)^{\top}\right\|_{p} \leq \beta_{1} \beta_{2}\|X\|_{p} .
$$

Where $X \in \mathbb{R}^{d_{1} \times d_{2}}$ is the tensorization of $x$. We first prove one direction. Let $U_{1}=A_{1} R_{1}^{-1}$
and $U_{2}=A_{2} R_{2}^{-1}$. We have

$$
\begin{aligned}
\left\|U_{1} X U_{2}^{\top}\right\|_{p}^{p} & =\sum_{i_{2}=1}^{n_{2}}\left\|U_{1}\left(X U_{2}^{\top}\right)_{i_{2}}\right\|_{p}^{p} \\
& \leq \sum_{i_{2}=1}^{n_{2}} \beta_{1}^{p}\left\|\left(X U_{2}^{\top}\right)_{i_{2}}\right\|_{p}^{p} \\
& =\beta_{1}^{p}\left\|X U_{2}^{\top}\right\|_{p}^{p} \\
& \leq \beta_{1}^{p} \beta_{2}^{p}\|X\|_{p}^{p}
\end{aligned}
$$

where the first step follows from rearranging, the second step follows from the well-conditioned property of $U_{1}$, the third step follows from rearranging again, the last step follows from the well-conditioned property of $U_{2}$. Similarly, we have

$$
\begin{aligned}
\left\|U_{1} X U_{2}^{\top}\right\|_{p}^{p} & =\sum_{i_{2}=1}^{n_{2}}\left\|U_{1}\left(X U_{2}^{\top}\right)_{i_{2}}\right\|_{p}^{p} \\
& \geq \sum_{i_{2}=1}^{n_{2}} \alpha_{1}^{p}\left\|\left(X U_{2}^{\top}\right)_{i_{2}}\right\|_{p}^{p} \\
& =\alpha_{1}^{p}\left\|X U_{2}^{\top}\right\|_{p}^{p} \\
& \geq \alpha_{1}^{p} \alpha_{2}^{p}\|X\|_{p}^{p}
\end{aligned}
$$

where again the first step follows from rearranging, the second step follows from the wellconditioned property of $U_{1}$, the third step follows from rearranging again, the last step follows from the well-conditioned property of $U_{2}$.

In general, for arbitrary $q \geq 2$, similarly using our reshaping lemma, we have

$$
\begin{aligned}
& \left\|\left(\otimes_{i=1}^{q}\left(A_{i} R_{i}^{-1}\right)\right) x\right\|_{p} \geq \prod_{i=1}^{q} \alpha_{i}\|x\|_{p} \\
& \left\|\left(\otimes_{i=1}^{q}\left(A_{i} R_{i}^{-1}\right)\right) x\right\|_{p} \leq \prod_{i=1}^{q} \beta_{i}\|x\|_{p}
\end{aligned}
$$

Putting this together with fact 26.2.2, and noting $d=d_{1} \cdots d_{q}$, we have

Corollary 26.8.3. Let $A_{i} R_{i}^{-1}$ be as in algorithm 26.2. Then we have for all $x \in \mathbb{R}^{d_{1} \cdots d_{q}}$ :

$$
(1 / d)^{O(1)}\|x\|_{p} \leq\left\|\left(A_{1} R_{1}^{-1} \otimes \cdots \otimes A_{q} R_{q}^{-1}\right) x\right\|_{p} \leq d^{O(1)}\|x\|_{p}
$$

In other words, $\left(A_{1} R_{1}^{-1} \otimes \cdots \otimes A_{q} R_{q}^{-1}\right)$ is a well conditioned $\ell_{p}$ basis for $\left(A_{1} \otimes \cdots \otimes A_{q}\right)$

From this, we can obtain the following result.
Lemma 26.8.4. Let $x^{\prime} \in \mathbb{R}^{d}$ be the output of the $O(1)$-Approximate $\ell_{p}$ Regression Procedure in Algorithm 26.2. Then with probability 99/100 we have

$$
\left\|\left(A_{1} \otimes \cdots \otimes A_{q}\right) x^{\prime}-b\right\|_{p} \leq 8 \min _{x}\left\|\left(A_{1} \otimes \cdots \otimes A_{q}\right) x-b\right\|_{p}
$$

Moreover, the time required to compute $x^{\prime}$ is $\widetilde{O}\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+(d q / \epsilon)^{O(1)}\right)$.

Proof. By Theorem 6 of [ $\left.\mathrm{DDH}^{+} 09\right]$, if we let $\Pi$ be a diagonal row sampling matrix such that $\Pi_{i, i}=1 / q_{i}^{1 / p}$ with probability $q_{i} \geq \min \left\{1, r_{1} \frac{\left\|U_{i, *}\right\|_{p}^{p}}{\|U\|_{p}^{p}}\right\}$, where $U$ is a $\ell_{p}$ well-conditioned basis for $\left(A_{1} \otimes \cdots \otimes A_{q}\right)$ and $r_{1}=O\left(d^{3}\right)$, then the solution $x^{\prime}$ to

$$
\min _{x} \| \Pi\left(\left(A_{1} \otimes \cdots \otimes A_{q}\right) x-b \|\right.
$$

will be a 8-approximation. Note that we can solve the sketched regression problem $\min _{x} \| \Pi\left(\left(A_{1} \otimes\right.\right.$ $\left.\cdots \otimes A_{q}\right) x^{\prime}-b \|$ which has $O(\operatorname{poly}(d / \epsilon))$ constraints and $d$ variables in time poly $(d / \epsilon)$ using linear programming for $p=1$ (see [CLS19] for the state of the art linear program solver), or
more generally interior point methods for convex programming for $p>1$ (see [BCLL18] for the state of the art $\ell_{p}$ solver).

Then by Corollary 26.8.3, we know that setting $U=\left(A_{1} R_{1}^{-1} \otimes \cdots \otimes A_{q} R_{q}^{-1}\right)$ suffices, so now we must sample rows of $U$. To do this, we must approximately compute the norms of the rows of $U$. Here, we use the fact that $\|\cdot\|_{p}^{p}$ norm of a row of $\left(A_{1} R_{1}^{-1} \otimes \cdots \otimes A_{q} R_{q}^{-1}\right)$ is the product of the row norms of the $A_{i} R_{i}^{-1}$ that correspond to that row. Thus it suffices to sample a row $j_{i}$ from each of the $A_{i} R_{i}^{-1}$ 's with probability at least $\min \left\{1, r_{1}\left\|\left(A_{i} R_{i}^{-1}\right)_{j_{i}, *}\right\|_{p}^{p} /\left\|A_{i} R_{i}^{-1}\right\|_{p}^{p}\right\}$ for each $i \in[q]$.

To do this, we must estimate all the row norms $\left\|\left(A_{i} R_{i}^{-1}\right)_{j_{i}, *}\right\|_{p}^{p}$ to $(1 \pm 1 / 10)$ error. This is done in steps $7-10$ of Algorithm 26.2, which uses dense $p$-stable sketches $Z \in \mathbb{R}^{d \times \tau}$, and computes $\left(A_{i} R_{i}^{-1} Z\right)$, where $\tau=\Theta(\log (n))$. Note that computing $R_{i}^{-1} Z \in \mathbb{R}^{d \times \tau}$ requires $\widetilde{O}\left(d^{2}\right)$. Once computed, $A_{i}\left(R_{i}^{-1} Z\right)$ can be computed in $\widetilde{O}\left(\mathrm{nnz}\left(A_{i}\right)\right)$ time. We then take the median of the coordinates of $\left(A_{i} R_{i}^{-1} Z\right)$ (normalized by the median of the $p$-stable distribution $\mathcal{D}_{p}$, which can be efficiently approximated to $(1 \pm \epsilon)$ in $O(\operatorname{poly}(1 / \epsilon))$ time, see Appendix A. 2 of [KNW10b] for details) as our estimates for the row norms. This is simply the Indyk median estimator [Ind06], and gives a $(1 \pm 1 / 10)$ estimate $a_{i, j}$ of all the row norms $\left\|\left(A_{i} R_{i}^{-1}\right)_{j, *}\right\|_{p}^{p}$ with probability $1-1 / \operatorname{poly}(n)$. Then it follows by Theorem 6 of $\left[\mathrm{DDH}^{+} 09\right]$ that $x^{\prime}$ is a 8-approximation of the optimal solution with probability $99 / 100$ (note that we amplified the probability by increasing the sketch sizes $S_{i}$ by a constant factor), which completes the proof.

### 26.8.2 $\quad \ell_{p}$ Sampling From the residual of a $O(1)$-factor approximation

By Lemma 26.8.4 in the prior section, we know that the $x^{\prime}$ first returned by the in algorithm 26.2 is a 8 -approximation. We now demonstrate how we can use this $O(1)$ approximation to obtain a $(1+\epsilon)$ approximation. The approach is again to sample rows of $\left(A_{1} \otimes \cdots \otimes A_{q}\right)$. But instead of sampling rows with the well-conditioned leverage scores $q_{i}$, we now sample the $i$-th row with probability $\alpha_{i}=\min \left\{1, \max \left\{q_{i}, r_{2}\left|\rho_{i}\right|^{p} /\|\rho\|_{p}^{p}\right\}\right\}$, where $\rho=\left(A_{1} \otimes \cdots \otimes A_{q}\right) x^{\prime}-b \in \mathbb{R}^{n}$ is the residual error of the $O(1)$-approximation $x^{\prime}$. Thus we must now determine how to sample quickly from the residuals $\left|\rho_{i}\right|^{p} /\|\rho\|_{p}^{p}$. Our sampling algorithm will need a tool originally developed in the streaming literature.

Count-sketch for heavy hitters with the Dyadic Trick. We now introduce a sketch $S$ which finds the $\ell_{2}$ heavy hitters in a vector $x$ efficently. This sketch $S$ is known as countsketch for heavy hitters with the Dyadic Trick. To build $S$ we first stack $\Theta(\log (n))$ copies of the count sketch matrix $S^{i} \in \mathbb{R}^{k^{\prime} \times n}$ [CW13]. The matrix $S^{i}$ is constructed as follows. $S^{i}$ has exactly one non-zero entry per column, which is placed in a uniformly random row, and given the value 1 or -1 uniformly at random. For $S^{i}$, let $h_{i}:[n] \rightarrow\left[k^{\prime}\right]$ be such that $h_{i}(t)$ is the row with the non-zero entry in the $t$-th column of $S^{i}$, and let $g_{i}:[n] \rightarrow\{1,-1\}$ be such that the value of that non-zero entry is $g_{i}(t)$. Note that the $h_{i}, g_{i}$ can be implemented as 4 -wise independent hash functions. Fix any $x \in \mathbb{R}^{n}$. Then given $S^{1} x, S^{2} x, \cdots, S^{\Theta(\log (n))} x$, we can estimate the value of any coordinate $x_{j}$ by $\operatorname{median}_{i \in \Theta \log (n)}\left\{g_{i}(j)\left(S^{i} x\right)_{h_{i}(j)}\right\}$.

It is well-known that this gives an estimate of $x_{j}$ with additive error $\Theta\left(1 / \sqrt{k^{\prime}}\right)\|x\|_{2}$ with probability $1-1 / \operatorname{poly}(n)$ for all $j \in[n][C C F C 04]$. However, naively, to find the heaviest coordinates in $x$, that is all coordinates $x_{j}$ with $\left|x_{j}\right| \geq \Theta\left(1 / \sqrt{k^{\prime}}\right)\|x\|_{2}$, one would
need to query $O(n)$ estimates. This is where the Dyadic trick comes in [CM04]. We repeat the above process $\Theta(\log (n))$ times, with matrices $S^{(i, j)}$, for $i, j \in \Theta(\log (n))$. Importantly, however, in $S^{(i, j)}$, for all $t, t^{\prime} \in[n]$ such that the first $j$ most significant bits in their binary identity representation are the same, we set $h_{(i, j)}(t)=h_{(i, j)}\left(t^{\prime}\right)$, effectively collapsing these identities to one. To find a heavy item, we can then query the values of the *two* identities from $S^{(1,1)}, S^{(2,1)}, \cdots, S^{(\Theta(\log (n)), 1)}$, and recurse into all the portions which have size at least $\Theta\left(1 / \sqrt{k^{\prime}}\right)\|x\|_{2}$. It is easy to see that we recurse into at most $O\left(k^{\prime}\right)$ such pieces in each of the $\Theta(\log (n))$ levels, and it takes $O(\log (n))$ time to query a single estimate, from which the desired runtime of $O\left(k^{\prime} \log ^{2}(n)\right)$ is obtained. For a further improvement on size $k$ of the overall sketched required to quickly compute $Q$, see [LNNT16]. We summarize this construction below in definition 26.8.2.

Definition 26.8.2 (Count-sketch for heavy hitters with Dyadic Trick [CCFC04, LNNT16]). There is a randomized sketch $S \in \mathbb{R}^{k \times n}$ with $k=O\left(\log ^{2}(n) / \epsilon^{2}\right)$ such that, for a fixed vector $x \in \mathbb{R}^{n}$, given $S x \in \mathbb{R}^{k}$, one can compute a set $Q \subset[n]$ with $|Q|=O\left(1 / \epsilon^{2}\right)$ such that $\left\{i \in[n]\left|\left|x_{i}\right| \geq \epsilon\|x\|_{2}\right\} \subseteq Q\right.$ with probability $1-1 / \operatorname{poly}(n)$. Moreover, $S x$ can be computed in $O\left(\log ^{2}(n) \mathrm{nnz}(x)\right)$ time. Given $S x$, the set $Q$ can be computed in time $O(k)$.

We begin with some notation. For a vector $y \in \mathbb{R}^{n}$, where $n=n_{1} \cdots n_{q}$, one can index any entry of $y_{i}$ via $\vec{i}=\left(i_{1}, i_{2}, \cdots, i_{q}\right) \in\left[n_{1}\right] \times \cdots \times\left[n_{q}\right]$ via $i=i_{1}+\sum_{j=2}^{q}\left(i_{j}-1\right) \prod_{l=1}^{i_{j}-1} n_{l}$. It will useful to index into such a vector $y$ interchangably via a vector $y_{\vec{i}}$ and an index $y_{j}$ with $j \in[n]$. For any set of subsets $T_{i} \subset\left[n_{i}\right]$, we can define $y_{T_{1} \times \cdots T_{q}} \in \mathbb{R}^{n}$ as $y$ restricted to the $\vec{i} \in T_{1} \times \cdots \times T_{q}$. Here, by restricted, we mean the coordinates in $y$ that are not in this set are set equal to 0 . Similarly, for a $y \in \mathbb{R}^{n_{i}}$ and $S \subset\left[n_{i}\right]$, we can define $y_{S}$ as $y$ restricted
to the coordinates in $S$. Note that in Algorithm 26.4, $\mathbb{I}_{n}$ denotes the $n \times n$ identity matrix for any integer $n$. We first prove a proposition on the behavior of Kronecker products of $p$-stable vectors, which we will need in our analysis.

Proposition 26.8.5. Let $Z_{1}, Z_{2}, \cdots, Z_{q}$ be independent vectors with entries drawn i.i.d. from the $p$-stable distribution, with $Z_{i} \in \mathbb{R}^{n_{i}}$. Now fix any $i \in[q]$, and any $x \in \mathbb{R}^{n}$, where $n=n_{1} n_{2} \cdots n_{q}$. Let $e_{j} \in \mathbb{R}^{n_{i}}$ be the $j$-th standard basis column vector for any $j \in\left[n_{i}\right]$. Let $\Gamma(i, j)=\left[n_{1}\right] \times\left[n_{2}\right] \times \cdots \times\left[n_{i-1}\right] \times\{j\} \times\left[n_{i+1}\right] \times \cdots \times\left[n_{q}\right]$. Define the random variable

$$
X_{i, j}(x)=\left|\left(Z_{1} \otimes Z_{1} \otimes \cdots \otimes Z_{i-1} \otimes e_{j}^{\top} \otimes Z_{i+1} \otimes \cdots \otimes Z_{q}\right) x\right|^{p} .
$$

Then for any $\lambda>1$, with probability at least $1-O(q / \lambda)$ we have

$$
\left\|x_{\Gamma(i, j)}\right\|_{p}^{p} / \lambda^{q} \leq X_{i, j}(x) \leq(\lambda \log (n))^{q}\left\|x_{\Gamma(i, j)}\right\|_{p}^{p}
$$

Proof. First observe that we can reshape $y=x_{\Gamma} \in \mathbb{R}^{m}$ where $m=n / n_{i}$, and re-write this random variable as $X_{i, j}(x)=\left|\left(Z_{1} \otimes Z_{2} \otimes \cdots \otimes Z_{q-1}\right) y\right|^{p}$. By reshaping Lemma 26.8.1, we can write this as $\left|\left(Z_{1} \otimes Z_{2} \otimes \cdots \otimes Z_{q-2}\right) Y Z_{q-1}^{\top}\right|^{p}$, where $Y \in \mathbb{R}^{m / n_{q-1} \times n_{q-1}}$. We first prove a claim. In the following, for a matrix $A$, let $\|A\|_{p}^{p}=\sum_{i, j}\left|A_{i, j}\right|^{p}$.

Claim 26.8.6. Let $Z$ be any $p$-stable vector and $X$ a matrix. Then for any $\lambda>1$, with probability $1-O(1 / \lambda)$, we have

$$
\lambda^{-1}\|X\|_{p}^{p} \leq\|X Z\|_{p}^{p} \leq \log (n) \lambda\|X\|_{p}^{p} .
$$

Proof. By $p$-stability, each entry of $\left|(X Z)_{i}\right|^{p}$ is distributed as $\left|z_{i}\right|^{p}\left\|X_{i, *}\right\|_{p}^{p}$, where $z_{i}$ is again $p$-stable (but the $z_{i}^{\prime} s$ are not independent). Now $p$-stables have tails that decay at the rate $\Theta\left(1 / x^{p}\right)$ (see Chapter 1.5 of [Nol07]), thus $\operatorname{Pr}\left[\left|z_{i}\right|^{p}>x\right]=O(1 / x)$ for any $x>0$. We
can condition on the fact that $z_{i}<\lambda \cdot n^{10}$ for all $i$, which occurs with probability at least $1-n^{-9} / \lambda$ by a union bound. Conditioned on this, we have $\mathbb{E}\left[\left|z_{i}\right|^{p}\right]=O(\log (n))$ (this can be seen by integrating over the truncated tail $O(1 / x)$ ), and the upper bound then follows from a application of Markov's inequality.

For the lower bound Let $Y_{i}$ be an indicator random variable indicating the event that $\left|z_{i}\right|^{p}<2 / \lambda$. Now $p$-stables are anti-concentrated, namely, their pdf is upper bounded by a constant everywhere. It follows that $\operatorname{Pr}\left[Y_{i}\right]<c / \lambda$ for some constant $c$. By Markov's inequality $\operatorname{Pr}\left[\sum_{i} Y_{i}\left\|X_{i, *}\right\|_{p}^{p}>\|X\|_{p}^{p} / 2\right]<O(1 / \lambda)$. Conditioned on this, the remaining $\|X\|_{p}^{p} / 2$ of the $\ell_{p}$ mass shrinks by less than a $2 / \lambda$ factor, thus $\|X Z\|_{p}^{p}>\left(\|X\|_{p}^{p} / 2\right)(2 / \lambda)=\|X\|_{p}^{p} / \lambda$ as needed.

By the above claim, we have $\|Y\|_{p} / \lambda^{1 / p} \leq\left\|Y Z_{q-1}^{\top}\right\|_{p} \leq(\log (n) \lambda)^{1 / p}\|Y\|_{p}$ with probability $1-O(1 / \lambda)$. Given this, we have $X_{i, j}(x)=\left|\left(Z_{1} \otimes Z_{2} \otimes \cdots \otimes Z_{q-2}\right) y^{\prime}\right|^{p}$, where $\|Y\|_{p} / \lambda^{1 / p} \leq\left\|y^{\prime}\right\|_{p} \leq(\log (n) \lambda)^{1 / p}\|Y\|_{p}$. We can inductively apply the above argument, each time getting a blow up of $(\log (n) \lambda)^{1 / p}$ in the upper bound and $(1 / \lambda)^{p}$ in the lower bound, and a failure probability of $(1 / \lambda)$. Union bounding over all $q$ steps of the induction, the proposition follows.

Lemma 26.8.7. Fix any $r_{2} \geq 1$, and suppose that $x^{\prime}=\min _{x}\left\|\Pi\left(A_{1} \otimes \cdots \otimes A_{q}\right) x-\Pi b\right\|_{p}$ and $\Pi \in \mathbb{R}^{n \times n}$ is a row sampling matrix such that $\Pi_{i, i}=1 / q_{i}^{1 / p}$ with probability $q_{i}$. Define the residual error $\rho=\left(A_{1} \otimes \cdots \otimes A_{q}\right) x^{\prime}-b \in \mathbb{R}^{n}$. Then Algorithm 26.4, with probability $1-\delta$, succeeds in outputting a row sampling matrix $\Sigma \in \mathbb{R}^{n \times n}$ such that $\Sigma_{i, i}=1 / \alpha_{i}^{1 / p}$ with probability $\alpha_{i}=\min \left\{1, \max \left\{q_{i}, r_{3}\left|\rho_{i}\right|^{p} /\|\rho\|_{p}^{p}\right\}\right\}$ for some $r_{3} \geq r_{2}$, and otherwise $\Sigma_{i, i}=0$. The
algorithm runs in time

$$
\widetilde{O}\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+q \mathrm{nnz}(b)+\left(r_{2} \log (n) / \delta\right)^{O\left(q^{2}\right)}\right) .
$$

Proof. The algorithm is given formally in figure 26.4. We analyze the runtime and correctness here.

Proof of Correctness. The approach of the sampling algorithm is as follows. Recall that we can index into the coordinates of $\rho \in \mathbb{R}^{n}$ via $\vec{a}=\left(a_{1}, \ldots, a_{q}\right)$ where $a_{i} \in\left[n_{i}\right]$. We build the coordinates of $\vec{a}$ one by one. To sample a $\vec{a} \in \prod_{i=1}^{q}\left[n_{i}\right]$, we can first sample $a_{1} \in\left[n_{1}\right]$ from the distribution $\operatorname{Pr}\left[a_{1}=j\right]=\sum_{\vec{u}: u_{1}=j}\left|\rho_{\vec{u}}\right|^{p} /\left(\sum_{\vec{u}}\left|\rho_{\vec{u}}\right|^{p}\right)$. Once we fix $a_{1}$, we can sample $a_{2}$ from the conditional distribution distribution $\operatorname{Pr}\left[a_{2}=j\right]=\sum_{\vec{u}: u_{2}=j, u_{1}=a_{1}}\left|\rho_{\vec{u}}\right|^{p} /\left(\sum_{\vec{u}: u_{1}=a_{1}}\left|\rho_{\vec{u}}\right|^{p}\right)$, and so on. For notation, given a vector $\vec{a}=\left(a_{1}, \ldots, a_{i-1}\right)$, let $\Delta(\vec{a})=\left\{\vec{u} \in\left[n_{1}\right] \times \cdots \times\left[n_{q}\right] \mid a_{j}=\right.$ $y_{j}$ for all $\left.j=1,2, \ldots, i-1\right\}$. Then in general, when we have sampled $\vec{a}=\left(a_{1}, \ldots, a_{i-1}\right)$ for some $i \leq q$, we need to sample $a_{i} \leftarrow j \in\left[n_{k}\right]$ with probability

$$
\operatorname{Pr}\left[a_{i}=j\right]=\sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=j}\left|\rho_{\vec{u}}\right|^{p} /\left(\sum_{\vec{u} \in \Delta(\vec{a})}\left|\rho_{\vec{u}}\right|^{p}\right) .
$$

We repeat this process to obtain the desired samples. Note that to sample efficiently, we will have to compute these aforementioned sampling probabilities approximately. Because of the error in approximating, instead of returning $r_{2}$ samples, we over-sample and return $r_{3}=\Theta\left(r_{2} \log ^{q^{2}}(n)\right)$ samples.

The first step is of the algorithm is to generate the $p$-stable vectors $Z^{i, j} \in \mathbb{R}^{n_{i}}$ for $i \in[q]$ and $j=1,2, \ldots, \Theta(\log (n))$. We can pre-compute and store $Z^{i, j} A_{i}$ for $i \in[q]$, which takes $\widetilde{O}\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)\right)$ time. We set $w^{j} \leftarrow\left(\left(\mathbb{I}_{n_{1}}\right) \otimes\left(\bigotimes_{k=2}^{q} Z^{k, j}\right) \rho\right) \in \mathbb{R}^{n_{1}}$ and define $w \in \mathbb{R}^{n_{1}}$
by $w_{l}=\operatorname{median}_{j \in[\tau]}\left\{\left|w_{l}^{j}\right|\right\}$ for $l \in\left[n_{1}\right]$. Observe that $w_{l}^{j}$ is an estimate of $\sum_{\vec{u}: u_{1}=l}\left|\rho_{\vec{u}}\right|^{p}$. By Proposition 26.8.5, it is a $(c \log (n))^{q}$ approximation with probability at least $3 / 4$ for some constant $c$. Taking the median of $\Theta(\log (n))$ repetitions, we have that

$$
c^{-q} \cdot \sum_{\vec{u}: u_{1}=l}\left|\rho_{\vec{u}}\right|^{p} \leq\left|w_{l}\right|^{p} \leq(c \log (n))^{q} \cdot \sum_{\vec{u}: u_{1}=l}\left|\rho_{\vec{u}}\right|^{p}
$$

with probability $1-1 / \operatorname{poly}(n)$, and we can then union bound over all such estimates every conducted over the course of the algorithm. We call the above estimate $\left|w_{l}\right|^{p}$ a $O\left((c \log (n))^{q}\right)-$ error estimate of $\sum_{\vec{u}: u_{1}=l}\left|\rho_{\vec{u}}\right|^{p}$. Given this, we can correctly and independently sample the first coordinate of each of the $\Theta\left(r_{3}\right)$ samples. We now describe how to sample the $i$-th coordinate. So in general, suppose we have sampled $\left(a_{1}, \ldots, a_{i-1}\right)$ so far, and we need to now sample $a_{i} \in\left[n_{i}\right]$ conditioned on $\left(a_{1}, \ldots, a_{i-1}\right)$. We first consider

$$
W^{i, k}=\left(\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(\mathbb{I}_{n_{i}}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) \rho\right) \in \mathbb{R}^{n_{i}}
$$

Note that the $j$-th coordinate $W_{j}^{i, k}$ for $W^{i, k}$ is an estimate of $\sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=j}\left|\rho_{\vec{u}}\right|^{p}$. Again by By Proposition 26.8.5, with probability $1-1 / \operatorname{poly}(n)$, we will have $\left|W_{j}^{i, k}\right|^{p}$ is a $O\left((c \log (n))^{q}\right)$-error estimate of $\sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=j}\left|\rho_{\vec{u}}\right|^{p}$ or at least one $k \in[\tau]$. Our goal will now be to find all $j \in\left[n_{i}\right]$ such that $\sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=j}\left|\rho_{\vec{u}}\right|^{p} \geq \Theta\left((c \log (n))^{q} / r_{3}^{8}\right) \sum_{\vec{u} \in \Delta(\vec{a})}\left|\rho_{\vec{u}}\right|^{p}$. We call such a $j$ a heavy hitter.

Let $Q_{i} \subset\left[n_{i}\right]$ be the set of heavy hitters. To find all the heavy hitters, we use the count-sketch for heavy hitters with the Dyadic trick of definition 26.8.2. We construct this count-sketch of $\operatorname{def} 26.8 .2 S^{i} \in \mathbb{R}^{k^{\prime} \times n_{i}}$ where $k^{\prime}=O\left(\log ^{2}(n) r_{3}^{16}\right)$. We then compute $S^{i} W^{i, k}$, for $k=1,2, \ldots, \tau$, and obtain the set of heavy hitters $h \in H_{i, k} \subset\left[n_{i}\right]$ which satisfy $\left|W_{j}^{i, k}\right|^{p} \geq \Theta\left(1 / r_{3}^{8}\right)\left\|W^{i, k}\right\|_{p}^{p}$. By the above discussion, we know that for each $j \in Q_{i}$, we
will have $\left|W_{j}^{i, k}\right|^{p} \geq \Theta\left(1 / r_{3}^{16}\right)\left\|W^{i, k}\right\|_{p}^{p}$ for at least one $k \in[\tau]$ with high probability. Thus $H_{i}=\cup_{k=1}^{\tau} H_{i, k} \supseteq Q_{i}$.

We now will decide to either sample a heavy hitter $\xi \in H_{i}$, or a non-heavy hitter $\xi \in\left[n_{i}\right] \backslash H_{i}$. By Proposition 26.8.5, we can compute a $O\left((c \log (n))^{-q}\right)$-error estimate $\beta_{i}=\operatorname{median}_{j \in \tau}\left|\left(\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(\bigotimes_{k=i}^{q} Z^{k, j}\right) \rho\right)\right|^{p}$ of $\sum_{\vec{u} \in \Delta(\vec{a})}\left|\rho_{\vec{u}}\right|^{p}$, meaning:

$$
O\left(c^{-q}\right) \sum_{\vec{u} \in \Delta(\vec{a})}\left|\rho_{\vec{u}}\right|^{p} \leq \beta_{i} \leq O\left((c \log n)^{q}\right) \sum_{\vec{u} \in \Delta(\vec{a})}\left|\rho_{\vec{u}}\right|^{p} .
$$

Again, by Proposition 26.8.5, we can compute a $O\left((c \log (n))^{-q}\right)$-error estimate $\gamma_{i} \leftarrow \operatorname{median}_{j \in[\tau]}\left(\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right)\right.$, of $\sum_{h \in\left[n_{i}\right] \backslash H_{i}} \sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=j}\left|\rho_{\vec{u}}\right|^{p}$. It follows that

$$
O\left(c^{-2 q}\right) \frac{\sum_{h \in\left[n_{i}\right] \backslash H_{i}} \sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=j}\left|\rho_{\vec{u}}\right|^{p}}{\sum_{\vec{u} \in \Delta(\vec{a})}\left|\rho_{\vec{u}}\right|^{p}} \leq \frac{\gamma_{i}}{\beta_{i}} \leq O\left((c \log n)^{2 q}\right) \frac{\sum_{h \in\left[n_{i}\right] \backslash H_{i}} \sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=j}\left|\rho_{\vec{u}}\right|^{p}}{\sum_{\vec{u} \in \Delta(\vec{a})}\left|\rho_{\vec{u}}\right|^{p}}
$$

In other words, $\gamma_{i} / \beta_{i}$ is a $O\left((c \log (n))^{2 q}\right)$-error approximation of the true probability that we should sample a non-heavy item. Thus with probability $1-\gamma_{i} / \beta_{i}$, we choose to sample a heavy item.

To sample a heavy item, for each $\xi \in H_{i}$, by Proposition 26.8.5, we can compute an $O\left((c \log (n))^{-q}\right)$-error estimate $\operatorname{median}_{j \in \tau}\left|\left(\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(e_{\xi}^{\top}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) \rho\right)\right|^{p}$ of $\sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=\xi}\left|\rho_{\vec{u}}\right|^{p}$, meaning

$$
\begin{aligned}
\left(O\left(c^{-q}\right) \sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=\xi}\left|\rho_{\vec{u}}\right|^{p}\right) & \leq \operatorname{median}_{j \in \tau}\left|\left(\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(e_{\xi}^{\top}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) \rho\right)\right|^{p} \\
& \leq\left(O\left((c \log n)^{q}\right) \sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=\xi}\left|\rho_{\vec{u}}\right|^{p}\right)
\end{aligned}
$$

Thus we can choose to sample a heavy item $\xi \in H_{i}$ from the distribution given by

$$
\operatorname{Pr}\left[\text { sample } a_{i} \leftarrow \xi\right]=\frac{\operatorname{median}_{j \in \tau}\left|\left(\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(e_{\xi}^{\top}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) \rho\right)\right|^{p}}{\sum_{\xi^{\prime} \in H_{i}} \operatorname{median}_{j \in \tau}\left|\left(\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(e_{\xi^{\prime}}^{\top}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) \rho\right)\right|^{p}}
$$

Which gives a $O\left((c \log (n))^{2 q}\right)$-error approximation to the correct sampling probability for a heavy item.

In the second case, with probability $\gamma_{i} / \beta_{i}$, we choose to not sample a heavy item. In this case, we must now sample a item from $\left[n_{i}\right] \backslash H_{i}$. To do this, we partition $\left[n_{i}\right]$ randomly into $\Omega_{1}, \ldots, \Omega_{\eta}$ for $\eta=1 / r_{3}^{2}$. Now there are two cases. First suppose that we have

$$
\frac{\sum_{j \in\left[n_{i}\right] \backslash H_{i}} \sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=j}\left|\rho_{\vec{u}}\right|^{p}}{\sum_{\vec{u} \in \Delta(\vec{a})}\left|\rho_{\vec{u}}\right|^{p}} \leq \Theta\left(1 / r_{3}^{3}\right)
$$

Now recall that $\gamma_{i} / \beta_{i}$ was a $O\left((c \log (n))^{2 q}\right)$-error estimate of the ratio on the left hand side of the above equation, and $\gamma_{i} / \beta_{i}$ was the probability with which we choose to sample a non-heavy hitter. Since we only repeat the sampling process $r_{3}$ times, the probability that we ever sample a non-heavy item in this case is at most $\Theta\left(q(c \log (n))^{2 q} / r_{3}^{2}\right)<\Theta\left(q / r_{3}\right)$, taken over all possible repetitions of this sampling in the algorithm. Thus we can safely ignore this case, and condition on the fact that we never sample a non-heavy item in this
case. Otherwise, $\sum_{j \in\left[n_{i}\right] \backslash H_{i}} \sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=j}\left|\rho_{\vec{u}}\right|^{p}>\Theta\left(1 / r_{3}^{3}\right) \sum_{\vec{u} \in \Delta(\vec{a})}\left|\rho_{\vec{u}}\right|^{p}$, and it follows that $\sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=j^{\prime}}\left|\rho_{\vec{u}}\right|^{p} \leq \Theta\left(1 / r_{3}^{5} \sum_{j \in\left[n_{i}\right] \backslash H_{i}}\right) \sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=j}\left|\rho_{\vec{u}}\right|^{p}$ for all $j^{\prime} \in\left[n_{i}\right] \backslash H_{i}$, since we removed all $\Theta\left(1 / r_{3}^{8}\right)$ heavy hitters from $\left[n_{i}\right]$ originally. Thus by Chernoff bounds, with high probability we have that $\sum_{j \in \Omega_{i} \backslash H_{i}}\left(\sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=j}\left|\rho_{\vec{u}}\right|^{p}\right)=\Theta\left(1 / \eta \sum_{j \in\left[n_{i}\right] \backslash H_{i}} \sum_{\vec{u} \in \Delta(\vec{a}): u_{i}=j}\left|\rho_{\vec{u}}\right|^{p}\right)$, which we can union bound over all repetitions.

Given this, by choosing $t \sim[\eta]$ uniformly at random, and then choosing $j \in \Omega_{t} \backslash H_{i}$ with probability proportional to its mass in $\Omega_{t} \backslash H_{i}$, we get a $\Theta(1)$ approximation of the true sampling probability. Since we do not know its exact mass, we instead sample from the distribution $\left\{\frac{\theta_{j}}{\sum_{j^{\prime} \in \Omega_{t} \backslash H_{i}} \theta_{j^{\prime}}}\right\}_{j \in \Omega_{t} \backslash H_{i}}$, where

$$
\theta_{j}=\underset{l \in[\tau]}{\operatorname{median}}\left(\left|\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(e_{j}^{\top}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, l}\right) \rho\right|^{p}\right)
$$

Again by Proposition 26.8.5, this gives a $O\left((c \log (n))^{2 q}\right)$-error approximation to the correct sampling probability. Note that at each step of sampling a coorindate of $\vec{a}$ we obtained at most $O\left((c \log (n))^{2 q}\right)$-error in the sampling probability. Thus, by oversampling by a $O\left((c \log (n))^{2 q^{2}}\right)$ factor, we can obtain the desired sampling probabilities. This completes the proof of correctness. Note that to improve the failure probability to $1-\delta$, we can simply scale $r_{3}$ by a factor of $1 / \delta$.

Proof of Runtime. We now analyze the runtime. At every step $i=1,2, \ldots, q$ of the sampling, we compute $v_{i}^{j} \leftarrow S^{i}\left(\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(\mathbb{I}_{n_{i}}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) \rho\right) \in \mathbb{R}^{n_{i}}$ for $j=$ $1,2, \ldots \Theta(\log (n))$. This is equal to

$$
S^{i}\left(\left(\bigotimes_{k=1}^{i-1}\left(A_{k}\right)_{a_{k}, *}\right) \otimes\left(A_{i}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j} A_{k}\right) x^{\prime}-\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(\mathbb{I}_{n_{i}}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) b\right)
$$

We first consider the term inside of the parenthesis (excluding $S^{i}$ ). Note that the term $\left(\bigotimes_{k=i+1}^{q} Z^{k, j} A_{k}\right)$ was already pre-computed, and is a vector of length at most $d$, this this requires a total of $\widetilde{O}\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+d\right)$ time. Note that these same values are used for every sample. Given this pre-computation, we can rearrage the first term to write $\left(\bigotimes_{k=1}^{i-1}\left(A_{k}\right)_{a_{k}, *}\right) \otimes$ $\left(A_{i}\right) X^{\prime}\left(\bigotimes_{k=i+1}^{q} Z^{k, j} A_{k}\right)^{\top}$ where $X^{\prime}$ is a matrix formed from $x^{\prime}$ so that $x^{\prime}$ is the vectorization of $X^{\prime}$ (this is done via reshaping Lemma 26.8.1). The term $y=X^{\prime}\left(\bigotimes_{k=i+1}^{q} Z^{k, j} A_{k}\right)^{\top}$ can now be computed in $O(d)$ time, and then we reshape again to write this as $\left(\bigotimes_{k=1}^{i-1}\left(A_{k}\right)_{a_{k}, *}\right) Y A_{i}^{\top}$ where $Y$ again is a matrix formed from $y$. Observe that $\zeta=\operatorname{vec}\left(\bigotimes_{k=1}^{i-1}\left(A_{k}\right)_{a_{k}, *} Y\right) \in \mathbb{R}^{d_{i}}$ can be computed in time $O(q d)$, since each entry is a dot product of a column $Y_{*, j} \in$ $\mathbb{R}^{d_{1} \cdot d_{2} \cdots d_{i-1}}$ of $Y$ with the $d_{1} \cdot d_{2} \cdots d_{i-1}$ dimensional vector $\bigotimes_{k=1}^{i-1}\left(A_{k}\right)_{a_{k}, *}$, which can be formed in $O\left(d_{1} \cdot d_{2} \cdots d_{i-1} q\right)$ time, and there are a total of $d_{i}$ columns of $Y$.

Given this, The first entire term $S^{i}\left(\bigotimes_{k=1}^{i-1}\left(A_{k}\right)_{a_{k}, *}\right) \otimes\left(A_{i}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j} A_{k}\right) x^{\prime}$ can be rewritten as $S^{i} A_{i} \zeta$, where $\zeta=\zeta_{\vec{a}} \in \mathbb{R}^{d_{i}}$ can be computed in $O(d q)$ time for each sample $\vec{a}$. Thus if we recompute the value $S_{i} A_{i} \in \mathbb{R}^{k \times n}$, where $k=\widetilde{O}\left(r_{3}^{16}\right)$, which can be done in time $\widetilde{O}\left(\operatorname{nnz} A_{i}\right)$, then every time we are sampling the $i$-th coordinate of some $\vec{a}$, computing the value of $S^{i} A_{i} \zeta_{\vec{a}}$ can be done in time $O\left(k d_{i}^{2}\right)=r_{3}^{O(1)}$.

We now consider the second term. We do a similar trick, reshaping $b \in \mathbb{R}^{n}$ into $B \in \mathbb{R}^{\left(n_{1} \cdots n_{i}\right) \times\left(n_{i} \cdots n_{q}\right)}$ and writing this term as $\left(\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(\mathbb{I}_{n_{i}}\right)\right) B\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right)^{\top}$ and computing $b^{\prime}=B\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right)^{\top} \in \mathbb{R}^{\left(n_{1} \cdots n_{i}\right)}$ in $n n z(B)=n n z(b)$ time. Let $B^{\prime} \in \mathbb{R}^{\left(n_{1} \cdots n_{i-1}\right) \times n_{i}}$ be such that $\operatorname{vec}\left(B^{\prime}\right)=b^{\prime}$, and we reshape again to obtain $\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) B^{\prime}\left(\mathbb{I}_{n_{i}}\right)=\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) B^{\prime}$ Now note that so far, the value $B^{\prime}$ did not depend on the sample $\vec{a}$ at all. Thus for each $i=1,2, \ldots, q, B^{\prime}$ (which depends only on $i$ ) can be pre-computed in nnz $(b)$ time. Given this, the value $\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) B^{\prime}$ is just a row $B_{\left(a_{1}, \ldots, a_{k}\right), *}^{\prime}$ of $B^{\prime}$ (or a column of $\left.\left(B^{\prime}\right)^{\top}\right)$. We
first claim that $n n z\left(B^{\prime}\right) \leq n n z(b)=n n z(B)$. To see this, note that each entry of $B^{\prime}$ is a dot product $B_{j, *}\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right)^{\top}$ for some row $B_{j, *}$ of $B$, and moreover there is a bijection between these dot products and entries of $B^{\prime}$. Thus for every non-zero entry of $B^{\prime}$, there must be a unique non-zero row (and thus non-zero entry) of $B$. This gives a bijection from the support of $B^{\prime}$ to the support of $B$ (and thus $b$ ) which completes the claim. Since $S^{i}\left(B_{\left(a_{1}, \ldots, a_{k}\right), *}^{\prime}\right)^{\top}$ can be computed in $\widetilde{O}\left(\operatorname{nnz}\left(B_{\left(a_{1}, \ldots, a_{k}\right), *}^{\prime}\right)\right)$ time, it follows that $S^{i}\left(B_{\left(a_{1}, \ldots, a_{k}\right), *}^{\prime}\right)^{\top}$ can be computed for all rows $\left(B_{\left(a_{1}, \ldots, a_{k}\right), *}^{\prime}\right)$ of $B$ in $\widetilde{O}(\mathrm{nnz}(b))$ time. Given this precomputation, we note that $\left(\mathbb{I}_{n_{i}}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) b$ is just $S^{i}\left(B_{\left(a_{1}, \ldots, a_{k}\right), *}^{\prime}\right)^{\top}$ for some $\left(a_{1}, \ldots, a_{k}\right)$, which has already been pre-computed, and thus requires no addition time per sample. Thus, given a total of $\widetilde{O}\left(\sum_{i=1}^{q} \mathrm{nnz}\left(A_{i}\right)+q \mathrm{nnz}(b)+r_{3}^{O(1)}\right)$ pre-processing time, for each sample we can compute $v_{i}^{j}$ for all $i \in[q]$ and $j \in[\tau]$ in $\widetilde{O}\left(r_{3}^{O(1)}\right)$ time, and thus $\widetilde{O}\left(r_{3}^{O(1)}\right)$ time over all $r_{3}$ samples.

Given this, the procedure to compute the heavy hitters $H_{i, j}$ takes $\widetilde{O}\left(r_{3}^{16}\right)$ time by Definition 26.8.2 for each sample and $i \in[q], j \in[\tau]$. By a identical pre-computation and rearrangement argument as above, each $\beta_{i}^{j}$ (and thus $\beta_{i}$ ) can be computed in $\widetilde{O}\left(r_{3}^{O(1)}\right)$ time per sample after pre-computation. Now note that $\gamma_{i}$ is simply equal to median $j \in[\tau]\left(\beta_{i}^{j}-\right.$ $\left.\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(Z_{H_{i}}^{k, j}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) \rho\right)$. Since $\left(Z_{H_{i}}^{k, j}\right)$ is sparse, the above can similar be computed in $O\left(d\left|H_{i}\right|\right)=\widetilde{O}\left(r_{3}^{O(1)}\right)$ time per sample after pre-computation. To see this, note that the $b$ term of $\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(Z_{H_{i}}^{k, j}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) \rho$ can be written as $\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) B^{\prime \prime \prime}\left(Z_{H_{i}}^{k, j}\right)^{\top}$, where $B^{\prime \prime \prime} \in \mathbb{R}^{n_{1} \cdots n_{i-1} \times n_{i}}$ is a matrix that has already been pre-computed and does not depend on the given sample. Then this quantity is just the dot product of a row of $B^{\prime \prime \prime}$ with $\left(Z_{H_{i}}^{k, j}\right)^{\top}$, but since $\left(Z_{H_{i}}^{k, j}\right)$ is $\left|H_{i}\right|$-sparse, so the claim for the $b$ term follows. For the $\left(A_{1} \otimes \cdots \otimes A_{q}\right)$ term, just as we demonstrated in the discussion of computing $v_{i}^{j}$, note that
this can be written as $\left(\bigotimes_{k=1}^{i-1}\left(A_{k}\right)_{a_{k}, *}\right) Y\left(\left(A_{i}\right)_{H_{i}, *}\right)^{\top}$ for some matrix $Y \in \mathbb{R}^{d_{1} \cdots d_{i} \times d_{i-1}}$ that has already been precomputed. Since $\left(A_{i}\right)_{H_{i}, *}$ only has $O\left(\left|H_{i}\right|\right)$ non-zero rows, this whole product can be computed in time $O\left(d\left|H_{i}\right|\right)$ as needed.

Similarly, we can compute the sampling probabilities

$$
\operatorname{Pr}\left[\text { sample } a_{i} \leftarrow j\right]=\frac{\operatorname{median}_{j \in \tau}\left|\left(\left(\otimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(e_{\xi}^{\top}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) \rho\right)\right|^{p}}{\sum_{\xi^{\prime} \in H_{i}} \operatorname{median}_{j \in \tau}\left|\left(\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(e_{\xi^{\prime}}^{\top}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) \rho\right)\right|^{p}}
$$

for each every item $\zeta \in H_{i}$ in $\widetilde{O}\left(r_{3}^{O(1)}\right)$ time after pre-computation, and note $\left|H_{i}\right|=\widetilde{O}\left(r_{3}^{O(1)}\right)$ by definition 26.8.2. Thus the total time to sample a heavy hitter in a given coordinate $i \in[q]$ for each sample $\widetilde{O}\left(r_{3}^{O(1)}\right)$ per sample, for an overall time of $\widetilde{O}\left(q r_{3}^{O(1)}\right)$ over all samples and $i \in[q]$.

Finally, we consider the runtime for sampling a non-heavy item. Note that $\left|\Omega_{t}\right|=$ $O\left(n_{i} / \eta\right)$ with high probability for all $t \in[\eta]$ by chernoff bounds. Computing each

$$
\theta_{j}=\underset{l \in[\tau]}{\operatorname{median}}\left(\left|\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(e_{j}^{\top}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, l}\right) \rho\right|^{p}\right)
$$

takes $O(q d)$ time after pre-computation, and so we spend a total of $O\left(q d n_{i} / \eta\right)$ time sampling an item from $\Omega_{t} \backslash H_{i}$. Since we only ever sample a total of $r_{3}$ samples, and $\eta=\Theta\left(r_{3}^{2}\right)$, the total time for sampling non-heavy hitters over the course of the algorithm in coordinate $i$ is $o\left(n_{i}\right)=o\left(\mathrm{nnz}\left(A_{i}\right)\right)$ as needed, which completes the proof of the runtime.

Computing the Sampling Probabilities $\alpha_{i}$ The above arguments demonstrate how to sample efficiently from the desired distribution. We now must describe how the sampling probabilities $\alpha_{i}$ can be computed. First note, for each sample that is sampled in the above
way, at every step we compute exactly the probability with which we decide to sample a coordinate to that sample. Thus we know exactly the probability that we choose a sample, and moreover we can compute each $q_{i}$ in $O(d)$ time as in Lemma 26.8.4. Thus we can compute the maximum of $q_{i}$ and this probability exactly. For each item sampled as a result of the leverage score sampling probabilities $q_{i}$ as in Lemma 26.8.4, we can also compute the probability that this item was sampled in the above procedure, by using the same sketching vectors $Z^{i, k}$ and count-sketches $S^{i}$. This completes the proof of the Lemma.

### 26.9 Missing Proofs from Section 26.4

In this section, we prove the correctness of our all-pairs regression algorithm 26.3. Our main theorem, Theorem 26.4.1, relies crucially on the sample routine developed in Section 26.9.1. We first prove the theorem which utilizes this routine, and defer the description and proof of the routine to Section 26.9.1.

Recall first the high level description of our algorithm (given formally in Figure 26.3). We pick $S_{1}, S_{2} \in \mathbb{R}^{k \times n}$ and $S$ are sparse $p$-stable sketches. We then compute $M=\left(S_{1} \otimes\right.$ $\left.S_{2}\right)(F \otimes \mathbf{1}-\mathbf{1} \otimes F)=S_{1} F \otimes S_{2} \mathbf{1}-S_{1} \mathbf{1} \otimes S_{2} F$, where $F=[A, b]$. We then take the $Q R$ decomposition $M=Q R$. Finally, we sample rows of $(F \otimes \mathbf{1}-\mathbf{1} \otimes F) R^{-1}$ with probability proportional to their $\ell_{p}$ norms. This is done by the sampling procedure described in Section 26.9.1. Finally, we solve the regression problem $\min _{x}\|\Pi(\bar{A} x-\bar{b})\|_{p}$, where $\Pi$ is the diagonal row-sampling matrix constructed by the sampling procedure.

We begin by demonstrating that $S_{1} \otimes S_{2}$ is a poly $(d)$ distortion embedding for the
column span of $[\bar{A}, \bar{b}]$.

Lemma 26.9.1. Let $S_{1}, S_{2} \in \mathbb{R}^{k \times n}$ be sparse $p$-stable transforms, where $k=\operatorname{poly}(d /(\epsilon \delta))$.
Then for all $x \in \mathbb{R}^{d+1}$, with probability $1-\delta$ we have

$$
1 / O\left(d^{4} \log ^{4} d\right)\|[\bar{A}, \bar{b}] x\|_{p} \leq\left\|\left(S_{1} \otimes S_{2}\right)[\bar{A}, \bar{b}] x\right\|_{p} \leq O\left(d^{2} \log ^{2} d\right)\|[\bar{A}, \bar{b}] x\|_{p}
$$

Proof. Let $F=[A, b]$. Then a basis for the columns of $[\bar{A}, \bar{b}]$ is given by $F \otimes \mathbf{1}-\mathbf{1} \otimes F$. We first condition on both $S_{1}, S_{2}$ being a low-distortion embedding for the $d+2$ dimensional column-span of $[F, \mathbf{1}]$. Note that this holds with large constant probability by 26.2.1.

So for any $x \in \mathbb{R}^{d+1}$, we first show the upper bound

$$
\begin{aligned}
\left\|\left(S_{\mathbf{1}} \otimes S_{2}\right)(F \otimes \mathbf{1}-\mathbf{1} \otimes F) x\right\|_{p} & =\left\|\left(S_{1} F \otimes S_{2} \mathbf{1}\right) x-\left(S_{1} \mathbf{1} \otimes S_{2} F\right)\right\|_{p} \\
& =\left\|S_{1} F x \mathbf{1}^{\top} S_{2}^{\top}-S_{1} \mathbf{1} x^{\top} F^{\top} S_{2}^{\top}\right\|_{p} \\
& =\left\|S_{1}\left(F x \mathbf{1}^{\top}-\mathbf{1} x^{\top} F^{\top}\right) S_{2}^{\top}\right\|_{p} \\
& \leq O(d \log d)\left\|\left(F x \mathbf{1}^{\top}-\mathbf{1} x^{\top} F^{\top}\right) S_{2}^{\top}\right\|_{p} \\
& \leq O\left(d^{2} \log ^{2} d\right)\left\|F x \mathbf{1}^{\top}-\mathbf{1} x^{\top} F^{\top}\right\|_{p} \\
& =O\left(d^{2} \log ^{2} d\right)\|(F \otimes \mathbf{1}-\mathbf{1} \otimes F) x\|_{p}
\end{aligned}
$$

Where the first equality follows by properties of the Kronecker product [VL00], the second by reshaping Lemma 26.8.1. The first inequality follows from the fact that each column of $\left(F x \mathbf{1}^{\top}-\mathbf{1} x^{\top} F^{\top}\right) S_{2}^{\top}$ is a vector in the column span of $[F, \mathbf{1}]$, and then using that $S_{1}$ is a low distortion embedding. The second inequality follows from the fact that each row of $\left(F x \mathbf{1}^{\top}-\mathbf{1} x^{\top} F^{\top}\right)$ is a vector in the column span of $[F, \mathbf{1}]$, and similarly using that $S_{2}$ is a low
distortion embedding. The final inequality follows from reshaping. Using a similar sequence of inequalities, we get the matching lower bound as desired.

We now prove our main theorem.

Theorem 26.4.1 Given $A \in \mathbb{R}^{n \times d}$ and $b \in \mathbb{R}^{n}$, for $p \in[1,2]$ there is an algorithm for the All-Pairs Regression problem that outputs $\widehat{x} \in \mathbb{R}^{d}$ such that with probability $1-\delta$ we have

$$
\|\bar{A} \widehat{x}-\bar{b}\|_{p} \leq(1+\epsilon) \min _{x \in \mathbb{R}^{d}}\|\bar{A} x-\bar{b}\|_{p}
$$

Where $\bar{A}=A \otimes \mathbf{1}-\mathbf{1} \otimes A \in \mathbb{R}^{n^{2} \times d}$ and $\bar{b}=b \otimes \mathbf{1}-\mathbf{1} \otimes b \in \mathbb{R}^{n^{2}}$. For $p<2$, the running time is $\widetilde{O}\left(n d+(d /(\epsilon \delta))^{O(1)}\right)$, and for $p=2$ the running time is $O\left(\mathrm{nnz}(A)+(d /(\epsilon \delta))^{O(1)}\right)$.

Proof. We first consider the case of $p=2$. Here, we can use the fact that the TensorSketch random matirx $S \in \mathbb{R}^{k \times n}$ is a subspace embedding for the column span of $[\bar{A}, \bar{b}]$ when $k=\Theta\left(d / \epsilon^{2}\right)$ [DSSW18], meaning that $\|S[\bar{A}, \bar{b}]\|_{2}=(1 \pm \epsilon)\|[\bar{A}, \bar{b}] x\|_{2}$ for all $x \in \mathbb{R}^{d+1}$ with probability $9 / 10$. Moreover, $S \bar{A}$ and $S \bar{b}$ can be computed in $O(\mathrm{nnz}(A)+\mathrm{nnz}(b))=O(\mathrm{nnz}(A))$ by [DSSW18] since they are the difference of Kronecker products. As a result, we can simply solve the regression problem $\widehat{x}=\arg \min _{x}\|S \bar{A} x-S \bar{b}\|_{2}$ in poly $(k d)$ time to obtain the desired $\widehat{x}$.

For $p<2$, we use the algorithm in Figure 26.3, where the crucial leverage score sampling procedure to obtain $\Pi$ in step 7 of Figure 26.3 is described in Lemma 26.4.2. Our high level approach follows the general $\ell_{p}$ sub-space embedding approach of [ $\left.\mathrm{DDH}^{+} 09\right]$. Namely, we first compute a low-distortion embedding $\left(S_{1} \otimes S_{2}\right)(F \otimes \mathbf{1}-\mathbf{1} \otimes F)$. By Lemma 26.9.1, using sparse-p stable transformations $S_{1}, S_{2}$, we obtain the desired poly ( $d$ ) distortion
embedding into $\mathbb{R}^{k^{2}}$, where $k=\operatorname{poly}(d / \epsilon)$. Note that computing $\left(S_{1} \otimes S_{2}\right)(F \otimes \mathbf{1} \mathbf{- 1} \otimes F)$ can be done in $O(\mathrm{nnz}(A)+\mathrm{nnz}(b)+n)$ time using the fact that $\left(S_{1} \otimes S_{2}\right)(F \otimes \mathbf{1})=S_{1} F \otimes S_{2} \mathbf{1}$. As shown in $\left[\mathrm{DDH}^{+} 09\right]$, it follows that $M=(F \otimes \mathbf{1}-\mathbf{1} \otimes F) R^{-1}$ is an $\ell_{p}$ well-conditioned basis for the column span of $(F \otimes \mathbf{1}-\mathbf{1} \otimes F)$ (see definition 26.2.4). Then by Theorem 5 of [ $\left.\mathrm{DDH}^{+} 09\right]$, if we let $\widehat{\Pi}$ be the diagonal row sampling matrix such that $\widehat{\Pi}_{i, i}=1 / q_{i}^{1 / p}$ for each $i$ with probability $q_{i} \geq \min \left\{1, r\left\|M_{i, *}\right\|_{p}^{p} /\|M\|_{p}^{p}\right\}$ (and $\widehat{\Pi}_{i, i}=0$ otherwise) for $r=\operatorname{poly}(d \log (1 / \delta) / \epsilon)$, then with probability $1-\delta$ we have $\|\widehat{\Pi}(F \otimes \mathbf{1}-\mathbf{1} \otimes F) x\|_{p}=(1 \pm \epsilon)\|(F \otimes \mathbf{1}-\mathbf{1} \otimes F) x\|_{p}$ for all $x \in \mathbb{R}^{d+1}$. First assume that we had such a matrix.

Since $(\bar{A} x-\bar{b})$ is in the column span of $(F \otimes \mathbf{1}-\mathbf{1} \otimes F)$ for any $x \in \mathbb{R}^{d+1}$, it follows that $\|\widehat{\Pi}(\bar{A} x-\bar{b})\|_{p}=(1 \pm \epsilon)\|(\bar{A} x-\bar{b})\|_{p}$ for all $x \in \mathbb{R}^{d}$, which completes the proof of correctness. By Lemma 26.4.2, we can obtain a row sampling matrix $\Pi$ in time $\widetilde{O}(n d+\operatorname{poly}(d / \epsilon))$, except that the entries of $\Pi$ are instead equal to either 0 or $1 / \widetilde{q}_{i}^{1 / p}$ where $\widetilde{q}_{i}=\left(1 \pm \epsilon^{2}\right) q_{i}$. Now let $\widehat{\Pi}$ be the idealized row sampling matrices from above, with entries either 0 or $1 / q_{i}^{1 / p}$ as needed for Theorem 5 of $\left[\mathrm{DDH}^{+} 09\right]$. Note that for any matrix $Z$ each row of $\widehat{\Pi} Z x$ is equal to $\Pi Z x$ times some constant $1-\epsilon^{2}<c<1+\epsilon^{2}$. It follows that $\|\Pi(\bar{A} x-\bar{b})\|_{p}=\left(1 \pm \epsilon^{2}\right)\|\widehat{\Pi}(\bar{A} x-\bar{b})\|_{p}$ for all $x \in \mathbb{R}^{d}$, and thus the objective function is changed by at most a $\left(1 \pm \epsilon^{2}\right)$ term, which is simply handled by a constant factor rescaling of $\epsilon$.

Finally, we can solve the sketched regression problem $\|\Pi(\bar{A} x-\bar{b})\|_{p}$ which has poly $(d / \epsilon)$ constraints and $d$ variables in time poly $(d / \epsilon)$ using linear programming for $p=1$ (see [CLS19] for the state of the art linear program sovler), or more generally interior point methods for convex programming for $p>1$ (see [BCLL18] for the state of the art $\ell_{p}$ solver. Finally, the failure probability bound holds by union bounding over all the aforementioned results, and noting that the lowest probability event was the even that $S_{1} \otimes S_{2}$ was a low distortion
embedding via Lemma 26.9.1. This completes the proof of the theorem.

### 26.9.1 Proof of Fast Sampling Lemma 26.4.2

We now provide a full proof of the main technical lemma of Section 26.4. The sampling algorithm is given formally in Algorithm 26.5. The following proof of Lemma 26.4.2 analyzes each step in the process, demonstrating both correctness and the desired runtime bounds.

Lemma 26.4.2 Given $R \in \mathbb{R}^{(d+1) \times(d+1)}$ and $F=[A, b] \in \mathbb{R}^{n \times(d+1)}$, there is an algorithm that, with probability $1-\delta$ for any $\delta>n^{-c}$ for any constant $c$, produces a diagonal matrix $\Pi \in \mathbb{R}^{n^{2} \times n^{2}}$ such that $\Pi_{i, i}=1 / \widetilde{q}_{i}^{1 / p}$ with probability $q_{i} \geq \min \left\{1, r\left\|M_{i, *}\right\|_{p}^{p} /\|M\|_{p}^{p}\right\}$ and $\Pi_{i, i}=0$ otherwise, where $r=\operatorname{poly}(d / \epsilon)$ and $M=(F \otimes \mathbf{1}-\mathbf{1} \otimes F) R^{-1}$, and $\widetilde{q}_{i}=\left(1 \pm \epsilon^{2}\right) q_{i}$ for all $i \in\left[n^{2}\right]$. The total time required is $\widetilde{O}(\mathrm{nnz} A+\operatorname{poly}(d / \epsilon))$.

Proof. Our proof proceeds in several steps. We analyze the runtime concurrently with out analysis of correctness.

Reducing the number of Columns of $R^{-1}$ We begin by generating a matrix $G \in$ $\mathbb{R}^{(d+1) \times \xi}$ of i.i.d. $\mathcal{N}(0,1 / \sqrt{\xi})$ Gaussian random variables. We then compute $Y \leftarrow R^{-1} G$ in $\widetilde{O}\left(d^{2}\right)$ time. We first claim that it suffices to instead $\ell_{p}$ sample rows of $C=(F \otimes \mathbf{1}-\mathbf{1} \otimes F) Y=$ $M G$. Note that each entry $\left|C_{i, j}\right|^{p}$ is distributed as $g^{p}\left\|M_{i, *}\right\|_{2}^{p}$ where $G \mathcal{N}(0,1 / \sqrt{\xi})$ Gaussian, which holds by the 2-stability of Gaussian random variables. Note that $\mathbb{E}\left[|g|^{p}\right]=\Theta(1 / \xi)$, so $\mathbb{E}\left[\left\|C_{i, *}\right\|_{p}^{p}\right]=\left\|M_{i, *}\right\|_{2}^{p}$, and by sub-exponential concentration (see Chapter 2 of [Wai19]), we have that $\left\|C_{i, *}\right\|_{p}^{p}=(1 \pm 1 / 10)\left\|M_{i, *}\right\|_{2}^{p}$ with probability $1-1 / \operatorname{poly}(n)$, and we can union
bound over this holding for all $i \in\left[n^{2}\right]$. By relationships between the $p$ norms, we have $\left\|M_{i, *}\right\|_{p}^{p} / d<\left\|M_{i, *}\right\|_{2}^{p}<\left\|M_{i, *}\right\|_{p}^{p}$, thus this changes the overall sampling probabilities by a factor between $\Theta\left(1 / d^{2}\right)$ and $\Theta\left(d^{2}\right)$. Thus, we can safely oversample by this factor (absorbing it into the value of $r$ ) to compensate for this change in sampling probabilities.

Sampling a row from $C$. To sample a row from $C$, the approach will be to sample an entry $C_{i, j}$ of $C$ with probability proportional to $\left\|C_{i, j}\right\|_{p}^{p} /\|C\|_{p}^{p}$. For every $(i, j)$ sampled, we sample the entire $i$-th row of $j$, so that the $j$-th row is indeed sampled with probability proportional to its norm. Thus, it suffices to sample entries of $C$ such that each $C_{i, j}$ is chosen with probability at least $\min \left\{1, r\left\|C_{i, j}\right\|_{p}^{p} /\|C\|_{p}^{p}\right\}$. First note that the $i$-th column of $C=(F \otimes \mathbf{1}-\mathbf{1} \otimes F) Y$ can be rearranged into a $n \times n$ matrix via Lemma 26.8.1, given by $\left(F Y_{*, i} \mathbf{1}^{\top}-\mathbf{1} Y_{*, i}^{\top} F^{\top}\right)$. To $\ell_{p}$ sample a coordinate from $C$, it suffices to first $\ell_{p}$ sample a column of one of the above matrices, and then $\ell_{p}$ sample an entry from that column.

To do this, we first compute $F Y \in \mathbb{R}^{n \times \xi}$, which can be done in time $\widetilde{O}(\mathrm{nnz} A)$ because $Y$ only has $\xi=\Theta(\log (n))$ columns. We then compute $Z\left(F Y_{*, i} \mathbf{1}^{\top}-\mathbf{1} Y_{*, i}^{\top} F^{\top}\right) \in \mathbb{R}^{1 \times n}$ for all $i \in[d]$, where $Z \in \mathbb{R}^{1 \times n}$ is a fixed vector of i.i.d. $p$-stable random variables. Once $F Y$ has been computed, for each $i \in[\xi]$ it takes $O(n)$ time to compute this $n$-dimensional vector, thus the total time required to compute all $\xi$ vectors is $\widetilde{O}(n)$. We repeat this process $t=O(\log (n))$ times with different $p$-stable vectors $Z^{1}, \ldots, Z^{\top}$, and take the median of each coordinate of $Z^{j}\left(F Y_{*, i} \mathbf{1}^{\top}-\mathbf{1} Y_{*, i}^{\top} F^{\top}\right) \in \mathbb{R}^{n}, j \in[t]$, divided by the median of the $p$-stable distribution (which can be approximated to ( $1 \pm \epsilon$ ) error in poly $(1 / \epsilon)$ time, see Appendix A. 2 of [KNW10b] for details of this). This is done in Step 7 of Algorithm 26.5. It is standard this this gives a $(1 \pm 1 / 10)$ approximation the the norm $\left\|\left(F Y_{*, i} \mathbf{1}^{\top}-\mathbf{1} Y_{*, i}^{\top} F^{\top}\right)_{*, l}\right\|_{p}$ for each
$i \in[d], l \in[n]$ with probability $1-1 / \operatorname{poly}(n)$ (See the Indyk median estimator [Ind06]).
Now let $\sigma_{i, l}$ be our estimate of the norm $\left\|\left(F Y_{*, i} \mathbf{1}^{\top}-\mathbf{1} Y_{*, i}^{\top} F^{\top}\right)_{*, l}\right\|_{p}$, for all $i \in[\xi]$ and $l \in[n]$. We now sample a columns $(i, l) \in[\xi] \times[n]$, where each $(i, l)$ is chosen with probability $\sigma_{i, l} /\left(\sum_{i^{\prime}, l^{\prime}} \sigma_{i^{\prime}, l^{\prime}}\right)$. We repeat this process $\Theta(r)$ times, to obtain a multi-set $T \subset[\xi] \times[n]$ of sampled columns $(i, l)$. We stress that $T$ is a multi-set, because the same column $(i, l)$ may have been chosen for multiple samples, and each time it is chosen we must independently sample one of the entries of that column. For any $(i, l) \in T$, we define $W^{(i, l)}=\left(F Y_{*, i} \mathbf{1}^{\top}-\right.$ $\left.\mathbf{1} Y_{*, i}^{\top} F^{\top}\right)_{*, l}=\left(F Y_{*, i}-\mathbf{1}(F Y)_{l, i}\right)$.
$\ell_{p}$ Sampling an entry from $W^{(i, l)}$. Now fix any $(i, l) \in T$. We show how to $\ell_{p}$ sample an entry from the vector $W^{(i, l)} \in \mathbb{R}^{n}$. In other words, for a given $j \in[n]$, we want to sample $W_{j}^{(i, l)} \in[n]$ with probability at least $r\left|W_{j}^{(i, l)}\right|^{p} /\left\|W^{(i, l)}\right\|_{p}^{p}$. We do this in two steps. First, let $S_{0} \in \mathbb{R}^{k \times n}$ be the count-sketch for heavy hitters of definition 26.8.2, where $k=\operatorname{poly}(r)$. Note that we can compute $S_{0} F Y$ and $S_{0} \mathbf{1}$ in time $\widetilde{O}(n)$, since $F Y \in \mathbb{R}^{n \times \xi}$. Once this is done, for each $(i, l) \in T$ we can compute $S_{0} W^{(i, l)}$ in $O(k)$ time by computing $\left(S_{0} \mathbf{1}(F Y)_{l, i}\right)$ (note that $F Y$ and $S_{0} \mathbf{1}$ are already computed), and subtracting it off from the $i$-th column of $S_{0} F Y$, so the total time is $\widetilde{O}(n+\operatorname{poly}(d / \epsilon))$ to compute $S_{0} W^{(i, l)}$ for all $(i, l) \in|T|$. Now we can obtain the set $Q_{0}^{(i, l)} \subset[n]$ containing all the $\widetilde{\Omega}(1 / \sqrt{k})$ heavy hitters in $W^{(i, l)}$ with high probability. We can then explicitly compute the value of $W_{j}^{(i, l)}$ for all $j \in Q_{0}^{(i, l)}$, and exactly compute the set

$$
H^{(i, l)}=\left\{\left.j \in[n]| | W_{j}^{(i, l)}\right|^{p}>\beta / r^{16}\left\|W^{(i, l)}\right\|_{p}^{p}\right\}
$$

all in $\widetilde{O}(k)$ time via definition 26.8.2, where $\beta>0$ is a sufficiently small constant (here we use the fact that $|x|_{p} \geq|x|_{2}$ for $p \leq 2$ ). Note that we use the same sketch $S_{0}$ to compute all
sets $Q_{0}^{(i, l)}$, and union bound the event that we get the heavy hitters over all poly $(d / \epsilon)$ trails.
We are now ready to show how we sample an index from $W^{(i, l)}$. First, we estimate the total $\ell_{p}$ norm of the items in $\left[n_{i}\right] \backslash H^{(i, l)}$ (again with the Indyk median estimator), and call this $\alpha_{(i, l)}$ as in Algorithm 26.5, which can be computed in $O\left(\left|H^{(i, l)}\right|\right)$ additional time (by subtracting off the $\left|H^{(i, l)}\right|$ coordinates $Z W_{\zeta}^{(i, l)}$ for all heavy hitters $\zeta \in H^{(i, l)}$ from our estimate $\left.\sigma_{(i, l)}\right)$, and with probability $\alpha_{(i, l)} / \sigma_{(i, l)}$, we choose to sample one of the items of $H^{(i, l)}$, which we can then sample from the distribution $\left|W_{j}^{(i, l)}\right|^{p} /\left(\sum_{j \in H^{(i, l)}}\left|W_{j}^{(i, l)}\right|^{p}\right)$. Since all the $\sigma_{(i, l)}, \alpha_{(i, l)}$ 's were constant factor approximations, it follows that we sampled such an item with probability $\Omega\left(r\left|W_{j^{\prime}}^{(i, l)}\right|^{p} /\|C\|_{p}^{p}\right)$ as needed. Otherwise, we must sample an entry from $[n] \backslash H^{(i, l)}$. To do this, we first randomly partition $[n]$ into $\eta=\Theta\left(r^{4} / \epsilon^{4}\right)$ subsets $\Omega_{1}, \Omega_{2}, \ldots, \Omega_{\eta}$.

We now make the same argument made in the proof of Lemma 26.8.7, considering two cases. In the first case, the $\ell_{p}$ mass of $[n] \backslash H^{(i, l)}$ drops by a $1 / r^{2}$ factor after removing the heavy hitters. In this case, $\alpha_{(i, l)} / \sigma_{(i, l)}=O\left(1 / r^{2}\right)$, thus we will never not sample a heavy hitter with probability $1-O(1 / r)$, which we can safely ignore. Otherwise, the $\ell_{p}$ drops by less than a $1 / r^{2}$ factor, and it follows that all remaining items must be at most a $\beta / r^{14}$ heavy hitter over the remaining coordinates $[n] \backslash H^{(i, l)}$ (since if they were any larger, they would be $\beta / r^{16}$ heavy hitters in $[n]$, and would have been removed in $H^{(i, l)}$. Thus we can assume we are in the second case. So by Chernoff bounds, we have $\sum_{j \in \Omega_{t}}\left|W_{j}^{(i, l)}\right|_{p}=\Theta\left(\frac{1}{\eta} \sum_{j \in[n] \backslash H^{(i, l)}}\left|W_{j}^{(i, l)}\right|_{p}\right)$ with probability greater than $1-\exp (-\Omega(r))$. We can then union bound over this event occurring for all $t \in[\eta]$ and all $(i, l) \in T$. Given this, if we uniformly sample a $t \sim[\eta]$, and then $\ell_{p}$ sample a coordinate $j \in \Omega_{t}$, we will have sampled this coordinate with the correct probability up to a constant factor. We now sample such a $t$ uniformly from $\eta$.

To do this, we generate a diagonal matrix $D \in \mathbb{R}^{n \times n}$, where $D_{i, i}=1 / u_{i}^{1 / p}$, where $u_{1}, \ldots, u_{n}$ are i.i.d. exponential random variables. For any set $\Gamma \subset[n]$, let $D_{\Gamma}$ be $D$ with all diagonal entries $(j, j)$ such that $j \notin \Gamma$ set equal to 0 . Now let $S \in \mathbb{R}^{k^{\prime} \times n}$ be a second instance of count-sketch for heavy hitters of definition 26.8.2, where we set $k^{\prime}=\operatorname{poly}(k)$ from above. It is known that returning $j^{*}=\arg \max _{j \in \Omega_{t} \backslash H^{(i, l)}}\left|\left(D W^{(i, l)}\right)_{j}\right|$ is a perfect $\ell_{p}$ sample from $\Omega_{t} \backslash H^{(i, l)}[J W 18]$. Namely, $\operatorname{Pr}\left[j^{*}=j\right]=\left|W_{j}^{(i, l)}\right|^{p} /\left\|W_{\Omega_{t} \backslash H^{(i, \ell)}}\right\|_{p}^{p}$ for any $j \in \Omega_{t} \backslash H^{(i, \ell)}$ . Thus it will suffice to find this $j^{*}$. To find $j^{*}$, we compute $S(D W)_{\Omega_{t} \backslash H^{(i, \ell)}}$. Note that since $F Y$ has already been computed, to do this we need only compute $S D_{\Omega_{t} \backslash H^{(i, \ell)}} F Y_{*, i}$ and $S D_{\Omega_{t} \backslash H^{(i, \ell)}} \mathbf{1}(F Y)_{\ell, i}$, which takes total time $\widetilde{O}\left(\left|\Omega_{t} \backslash H^{(i, \ell)}\right|\right)=\widetilde{O}(n / \eta)$. We then obtain a set $Q^{(i, l)} \subset \Omega_{t} \backslash H^{(i, \ell)}$ which contains all $j$ with $\left|\left(D W^{(i, l)}\right)_{j}\right| \geq \widetilde{\Omega}\left(1 / \sqrt{k^{\prime}}\right)\left\|(D W)_{\Omega_{t} \backslash H^{(i, \ell)}}\right\|_{2}$.

As noted in [JW18], the value $\max _{j \in \Omega_{t} \backslash H^{(i, l)}}\left|\left(D W^{(i, l)}\right)_{j}\right|$ is distributed identically to $\left\|W_{\Omega_{t} \backslash H^{(i, \ell)}}\right\|_{p} / u^{1 / p}$ where $u$ is again an exponential random variable. Since exponential random variables have tails that decay like $e^{-\Omega(x)}$, it follows that with probability $1-\exp (-\Omega(r))$ that we have $\max _{j \in \Omega_{t} \backslash H^{(i, l)}}\left|\left(D W^{(i, l)}\right)_{j}\right|=\Omega\left(\left\|W_{\Omega_{t} \backslash H^{(i, \ell)}}\right\|_{p} / r\right)$, and we can then union bound over the event that this occurs for all $(i, l) \in T$ and $\Omega_{t}$. Given this it follows that $\left(D W^{(i, l)}\right)_{j^{*}}=\Omega\left(\left\|W_{\Omega_{t} \backslash H^{(i, \ell)}}\right\|_{p} / r\right)$. Next, for any constant $c \geq 2$, by Proposition 1 of [JW18], we have $\left\|\left((D W)_{\Omega_{t} \backslash H^{(i, \ell)}}\right)_{\operatorname{tail}(c \log (n))}\right\|_{2}=\widetilde{O}\left(\left\|W_{\Omega_{t} \backslash H^{(i, \ell)}}^{(i, l)}\right\|_{p}\right)$ with probability $1-n^{-c}$, where for a vector $x, x_{(\overline{[t]})}$ is $x$ but with the top $t$ largest (in absolute value) entries set equal to 0 . Since there are at most $c \log (n)$ coordinates in $(D W)_{\Omega_{t} \backslash H^{(i, \ell)}}$ not counted in $\left((D W)_{\Omega_{t} \backslash H^{(i, \ell)}}\right)_{\operatorname{tail}(\operatorname{clog}(n))}$, and since $(D W)_{j^{*}}$ is the largest coordinate in all of $(D W)_{\Omega_{t} \backslash H^{(i, \ell)}}$, by putting together all of the above it follows that $(D W)_{j^{*}}$ is a $\widetilde{\Omega}(1 / r)$-heavy hitter in $(D W)_{\Omega_{t} \backslash H^{(i, \ell)}}$. Namely, that $\left|(D W)_{j^{*}}\right| \geq \widetilde{\Omega}\left(\left\|(D W)_{\Omega_{t} \backslash H^{(i, \ell)}}\right\|_{2} / r\right)$. Thus, we conclude that $j^{*} \in Q^{(i, l)}$.

Given that $j^{*} \in Q^{(i, l)}$, we can then compute the value $\left(D W^{(i, l)}\right)_{j}=D_{j, j}\left(F Y_{j, i}-F Y_{l, i}\right)$ in $O(1)$ time to find the maximum coordinate $j^{*}$. Since $\left|Q^{(i, l)}\right|=O\left(k^{\prime}\right)=O(\operatorname{poly}(d / \epsilon))$, it follows that the total time required to do this is $\widetilde{O}(n / \eta+\operatorname{poly}(d / \epsilon))$. Since we repeat this process for each $(i, l) \in T$, and $|T|=\Theta(r)$ whereas $\eta=\Theta\left(r^{4}\right)$, it follows that the total runtime for this step is $\widetilde{O}\left(n / r^{3}+\operatorname{poly}(d / \epsilon)\right)$. By [JW18], the result is a perfect $\ell_{p}$ sample from $(D W)_{\Omega_{t} \backslash H^{(i, \ell)}}$, which is the desired result. To complete the proof, we note that the only complication that remains is that we utilize the same scaling matrix $D$ to compute the sampled used in each of the columns $W^{(i, l)}$ for each $(i, l) \in T$. However, note that for $t \neq t^{\prime}$, we have that $D_{\Omega_{t}}$ and $D_{\Omega_{t}}$ are independent random variables. Thus it suffices to condition on the fact that the $t \in[\eta]$ that is sampled for each of the $|T|$ repetitions of sampling a $\Omega_{t}$ are distinct. But this occurs with probability at least $1 / r$, since $|T|=\Theta(r)$ and $\eta=\Theta\left(r^{4}\right)$. Conditioned on this, all $|T|$ samples are independent, and each sample is an entry $C_{i, j}$ of $C$ such that the probability that a given $(i, j)$ is chosen is $\left|C_{i, j}\right|^{p} /\|C\|_{p}^{p}$. Repeating this sampling $\Theta(r)$ times, we get that each $C_{i, j}$ is sampled with probability at least $\min \left\{1, r\left|C_{i, j}\right|^{p} /\|C\|_{p}^{p}\right\}$, which completes the proof of correctness. Note that the dominant runtime of the entire procedure was $\widetilde{O}(\mathrm{nnz}(A)+\operatorname{poly}(d / \epsilon))$ as stated, and the probability of success was $1-\exp (-r)+1 / \operatorname{poly}(n)$, which we can be amplified to any $1-\delta$ for $\delta>1 / n^{c}$ for some constant $c$ by increasing the value of $r$ by $\log (1 / \delta)$ and the number of columns of the sketch $G$ to $\log (1 / \delta)$, which does not effect the $\widetilde{O}(\mathrm{nnz}(A)+\operatorname{poly}(d / \epsilon))$ runtime.

Computing approximations $\widetilde{q}_{i}$ for $q_{i}$. It remains now how to compute the approximate sampling probabilities $\widetilde{q}_{i}$ for $\Theta(r)$ rows of $C$ that were sampled. Note that to sample an entry, in $C$, we first sampled the $n \times 1$ submatrix $W^{(i, l)}$ of $C$ which contained it, where the
probability that we sample this submatrix is known to us. Next, if the entry of $C$ was a heavy hitter in $W^{(i, l)}$, we exactly compute the probability that we sample this entry, and sample it with this probability. If the entry $j$ of $W^{(i, l)}$ is not a heavy hitter, we first sample an $\Omega_{t}$ uniformly with probability exactly $1 / \eta$. The last step is sampling a coordinate from $W_{\Omega_{t} \backslash H^{(i, l)}}^{(i, l)}$ via exponential scaling. However, we do not know the exact probability of this sampling, since this will be equal to $\left|W_{j}^{(i, l)}\right|^{p} /\left\|W_{\Omega_{t} \backslash H^{(i, l)}}^{(i, l)}\right\|_{p}^{p}$, and we do not know $\left\|W_{\Omega_{t} \backslash H^{(i, l)}}^{(i, l)}\right\|_{p}^{p}$ exactly. Instead, we compute it approximately to error $\left(1 \pm \epsilon^{2}\right)$ as follows. For each $(i, l) \in T$ and $\alpha=1,2, \ldots, \Theta\left(\log (n) / \epsilon^{4}\right)$, we compute $Z^{(\alpha)} W_{\Omega_{t} \backslash H^{(i, l)}}^{(i, l)}$, where $Z \in \mathbb{R}^{1 \times\left|\Omega_{t} \backslash H^{(i, l)}\right|}$ is a vector of $p$-stable random variables. Again, we use the Indyk median estimator [Ind06], taking the median of these $\Theta\left(\log (n) / \epsilon^{4}\right)$ repetitions, to obtain an estimate of $\left\|W_{\Omega_{t} \backslash H^{(i, l)}}^{(i, l)}\right\|_{p}^{p}$ with high probability to $\left(1 \pm \epsilon^{2}\right)$ relative error. Each repetition requires $O\left(\left|\Omega_{t} \backslash H^{(i, l)}\right|\right)$ additional time, and since $\left|\Omega_{t} \backslash H^{(i, l)}\right||T|=o\left(\epsilon^{4} n / r^{3}\right)$, it follows that the total computational time is at most an additive $o(n)$, thus computing the $\widetilde{q}_{i}$ 's to error $\left(1 \pm \epsilon^{2}\right)$ does not effect the overall runtime.

### 26.10 Missing Proofs from Section 26.5

We first give the proof of the main theorem. The proof relies on Lemma 26.10.2, which the rest of this section will be devoted to proving.

Theorem 26.5.1 For any constant $q \geq 2$, there is an algorithm which runs in time $O\left(\sum_{i=1}^{q} \operatorname{nnz}\left(A_{i}\right)+d \operatorname{poly}(k / \epsilon)\right)$ and outputs a rank $k$-matrix $B$ in factored form such that $\|B-A\|_{F} \leq(1+\epsilon) \mathrm{OPT}_{k}$ with probability 9/10.

Proof. By Lemma 26.10.2, we have $(1-\epsilon)\|A-A P\|_{F}^{2} \leq\|M-M P\|_{F}^{2}+c \leq(1+\epsilon)\|A-A P\|_{F}^{2}$ for all rank $k$ projection matrices $P$. In particular, we have

$$
\min _{P}(1+\epsilon)\|A-A P\|_{F}^{2}+c=(1+\epsilon) \mathrm{OPT}_{k}^{2}
$$

where the minimum is taken over all rank $k$ projection matrices. The minimizer $P$ on the LHS is given by the projection onto the top $k$ singular space of $M$. Namely, $M P=M U^{\top} U$ where $U$ is the top $k$ singular row vectors of $M$. Thus $\left\|M-M U^{\top} U\right\|_{F}^{2}+c \leq(1+\epsilon) \mathrm{OPT}_{k}^{2}$. Moreover, we have $\left\|A-A U^{\top} U\right\|_{F}^{2} \leq(1+2 \epsilon)\left(\left\|M-M U^{\top} U\right\|_{F}^{2}+c\right) \leq(1+4 \epsilon) \mathrm{OPT}_{k}^{2}$. Thus $\left\|A-A U^{\top} U\right\|_{F} \leq(1+O(\epsilon)) \mathrm{OPT}_{k}$ as needed.

For runtime, note that we first must compute $M=\left(\otimes_{i=1}^{q} S_{i}\right)\left(A_{1} \otimes A_{2}\right)=S_{1} A_{1} \otimes \cdots \otimes$ $S_{q} A_{q}$. Now $S_{i} A_{i}$ can be computed in $O\left(\mathrm{nnz}\left(A_{i}\right)\right)$ time for each $i$ [CW13]. One all $S_{i} A_{i}$ are computed, their Kronecker product can be computed in time $O\left(q k_{1} k_{2} \cdots k_{q} d\right)=\operatorname{poly}(k d / \epsilon)$. Given $M \in \mathbb{R}^{k_{1} \cdots k_{q} \times d}$, the top $k$ singular vectors $U$ can be computed by computing the SVD of $M$, which is also done in time poly $(k d / \epsilon)$. Once $U$ is obtained, the algorithm can terminate, which yields the desired runtime.

To complete the proof of the main theorem, we will need to prove Lemma 26.10.2. To do this, we begin by introducing two definitions.

Definition 26.10.1. A random matrix $S$ is called a $\epsilon$-subspace embedding for a rank $k$ subspace $\mathcal{V}$ we have simultaneously for all $x \in \mathcal{V}$ that $\|S x\|_{2}=(1 \pm \epsilon)\|x\|_{2}$.

Definition 26.10.2. A random matrix $S$ satisfies the $\epsilon$-approximate matrix product property if, for any fixed matrices $A, B$, of the appropriate dimensions, we have $\operatorname{Pr}\left[\| A^{\top} S^{\top} S B-\right.$ $\left.A^{\top} B\left\|_{F} \leq \epsilon\right\| A\left\|_{F}\right\| B \|_{F}\right] \geq 9 / 10$.

We now show that $S$ is both a subspace embedding and satisfies approximate matrix product, where $S=\otimes_{i=1}^{q} S_{i}$ and $S_{i} \in \mathbb{R}^{k_{i} \times n_{i}}$ are count-sketch matrices.

Lemma 26.10.1. If $S=\left(\otimes_{i=1}^{q} S_{i}\right)$ with $S_{i} \in \mathbb{R}^{k_{i} \times n_{i}}, k_{1}=k_{2}=\cdots=k_{q}=\Theta\left(q k^{2} / \epsilon^{2}\right)$, then $S$ is an $\epsilon$-subspace embedding for any fixed $k$ dimensional subspace $\mathcal{V} \subset \mathbb{R}^{n}$ with probability $9 / 10$, and also satisfies the ( $\epsilon / k$ )-approximate matrix product property.

Proof. We first show that $S$ satisfies the $O(\epsilon / k, 1 / 10,2)$-JL moment property. Here, the $(\epsilon, \delta, \ell)$-JL moment property means that for any fixed $x \in \mathbb{R}^{n}$ with $\|x\|_{2}=1$, we have $\mathbb{E}\left[\left(\|S x\|_{2}^{2}-1\right)^{2}\right] \leq \epsilon^{\ell} \delta$, which will imply approximate matrix product by the results of [KN14].

We prove this by induction on $q$. Let $\bar{k}=k_{1}$. First suppose $S=(Q \otimes T)$, where $Q \in$ $\mathbb{R}^{k_{1} \times n_{1}}$ is a count-sketch, and $T \in \mathbb{R}^{k^{\prime} \times n^{\prime}}$ is any random matrix which satisfies $\mathbb{E}\left[\|T x\|_{2}^{2}\right]=$ $\|x\|_{2}^{2}\left(T \in \mathbb{R}^{k^{\prime} \times n^{\prime}}\right.$ is unbiased), and $\mathbb{E}\left[\left(\|T x\|_{2}-1\right)^{2}\right] \leq 1+c / \bar{k}$ for some value $c<\bar{k}$. Note that both of these properties are satisfied with $c=4$ if $T \in \mathbb{R}^{k_{2} \times n_{2}}$ is itself a count-sketch matrix [CW13]. Moreover, these are the only properties we will need about $T$, so we will. We now prove that $\mathbb{E}\left[\|(S \otimes T) x\|_{2}^{2}\right]=1$ and $\mathbb{E}\left[\|(S \otimes T) x\|_{2}^{4}\right] \leq 1+(c+4) / \bar{k}$ for any unit vector $x$.

Fix any unit $x \in \mathbb{R}^{n}$ now (here $n=n_{1} n^{\prime}$ ), and let $x^{j} \in \mathbb{R}^{n^{\prime}}$ be the vector obtained by restricted $x$ to the coordinates $j n_{1}+1$ to $(j+1) n_{1}$. For any $i \in\left[k_{1}\right], j \in\left[k^{\prime}\right]$, let $i_{j}=(i-1) k^{\prime}+j$. Let $h_{Q}(i) \in\left[k_{1}\right]$ denote the row where the non-zero entry in the $i$-th column is placed in $Q$. Let $\sigma_{Q}(i) \in\{1,-1\}$ denote the sign of the entry $Q_{h_{Q}(i), i}$. Let $\delta_{Q}(i, j)$
indicate the event that $h_{Q}(i)=j$. First note that

$$
\begin{aligned}
\mathbb{E}\left[\sum_{i, j}((Q \otimes T) x)_{i_{j}}^{2}\right] & =\mathbb{E}\left[\sum_{i=1}^{k_{1}} \sum_{j=1}^{k^{\prime}}\left(\sum_{\tau=1}^{n_{1}} \delta_{Q}(\tau, i) \sigma_{Q}(\tau)\left(T x^{\tau}\right)_{j}\right)^{2}\right] \\
& =\mathbb{E}\left[\sum_{i=1}^{k_{1}} \sum_{j=1}^{k^{\prime}} \sum_{\tau=1}^{n_{1}} \delta_{Q}(\tau, i)\left(T x^{\tau}\right)_{j}^{2}\right] \\
& =\mathbb{E}\left[\sum_{\tau=1}^{n_{1}} \sum_{i=1}^{k_{1}} \sum_{j=1}^{k^{\prime}} \delta_{Q}(\tau, i)\left(T x^{\tau}\right)_{j}^{2}\right] \\
& =\mathbb{E}\left[\sum_{\tau=1}^{n_{1}}\left\|T x^{\tau}\right\|_{2}^{2}\right] \\
& =\|x\|_{2}^{2}
\end{aligned}
$$

Where the last equality follows because count-sketch $T$ is unbiased for the base case, namely that $\mathbb{E}\left[\|T x\|_{2}^{2}\right]=\|x\|_{2}^{2}$ for any $x$ [Woo14b], or by induction. We now compute the second moment,

$$
\begin{aligned}
\mathbb{E}\left[\left(\sum_{i, j}((Q \otimes T) x)_{i_{j}}^{2}\right)^{2}\right] & =\mathbb{E}\left[\left(\sum_{i, j}\left(\sum_{\tau=1}^{n_{1}} \delta_{Q}(\tau, i) \sigma_{Q}(\tau)\left(T x^{\tau}\right)_{j}\right)^{2}\right)^{2}\right] \\
& =\mathbb{E}\left[\left(\sum_{i, j} \sum_{\tau_{1}, \tau_{2}}^{n_{1}} \delta_{Q}\left(\tau_{1}, i\right) \sigma_{Q}\left(\tau_{1}\right)\left(T x^{\tau_{1}}\right)_{j} \delta_{Q}\left(\tau_{2}, i\right) \sigma_{Q}\left(\tau_{2}\right)\left(T x^{\tau_{2}}\right)_{j}\right)^{2}\right] \\
& =\sum_{\tau_{1}, \tau_{2}, \tau_{3}, \tau_{4}}^{n_{1}} \mathbb{E}\left[\left(\sum_{i, j} \delta_{Q}\left(\tau_{1}, i\right) \sigma_{Q}\left(\tau_{1}\right)\left(T x^{\tau_{1}}\right)_{j} \delta_{Q}\left(\tau_{2}, i\right) \sigma_{Q}\left(\tau_{2}\right)\left(T x^{\tau_{2}}\right)_{j}\right)\right. \\
& \left.\cdot\left(\sum_{i, j} \delta_{Q}\left(\tau_{3}, i\right) \sigma_{Q}\left(\tau_{3}\right)\left(T x^{\tau_{3}}\right)_{j} \delta_{Q}\left(\tau_{4}, i\right) \sigma_{Q}\left(\tau_{4}\right)\left(T x^{\tau_{4}}\right)_{j}\right)\right] .
\end{aligned}
$$

We now analyze the above expectation. There are several cases for the expectation of each
term. First, we bound the sum of the expectations when $t_{1}=t_{2}=t_{3}=t_{4}$ by

$$
\begin{aligned}
& \sum_{\tau=1}^{n_{1}} \mathbb{E}\left[\left(\sum_{i, j} \delta_{Q}(\tau, i) \sigma_{Q}(\tau)\left(T x^{\tau}\right)_{j} \delta_{Q}(\tau, i) \sigma_{Q}(\tau)\left(T x^{\tau}\right)_{j}\right)\right. \\
& \left.\cdot\left(\sum_{i, j} \delta_{Q}(\tau, i) \sigma_{Q}(\tau)\left(T x^{\tau}\right)_{j} \delta_{Q}(\tau, i) \sigma_{Q}(\tau)\left(T x^{\tau}\right)_{j}\right)\right] \\
\leq & \sum_{\tau=1}^{n_{1}} \mathbb{E}\left[\left\|T x^{\tau}\right\|_{2}^{4}\right]=1+c / \bar{k}
\end{aligned}
$$

Where the last equation follows from the variance of count-sketch [CW13] for the base case, or by induction for $q \geq 3$. We now bound the sum of the expectations when $t_{1}=t_{2} \neq t_{3}=t_{4}$ by

$$
\begin{aligned}
& \sum_{\tau_{1} \neq \tau_{2}} \mathbb{E}\left[\left(\sum_{i, j} \delta_{Q}\left(\tau_{1}, i\right) \sigma_{Q}\left(\tau_{1}\right)\left(T x^{\tau_{1}}\right)_{j} \delta_{Q}\left(\tau_{1}, i\right) \sigma_{Q}\left(\tau_{1}\right)\left(T x^{\tau_{1}}\right)_{j}\right)\right. \\
& \left.\cdot\left(\sum_{i, j} \delta_{Q}\left(\tau_{2}, i\right) \sigma_{Q}\left(\tau_{2}\right)\left(T x^{\tau_{2}}\right)_{j} \delta_{Q}\left(\tau_{2}, i\right) \sigma_{Q}\left(\tau_{2}\right)\left(T x^{\tau_{2}}\right)_{j}\right)\right] \\
& \leq \sum_{\tau_{1} \neq \tau_{2}} \mathbb{E}\left[\left\|T x^{\tau_{1}}\right\|_{2}^{2}\left\|T x^{\tau_{2}}\right\|_{2}^{2} / k_{1}\right] \\
& \leq \mathbb{E}\left[\|T x\|_{2}^{4} / k_{1}\right] \leq(1+c / \bar{k}) / k_{1} .
\end{aligned}
$$

We can similarly bound the sum of the terms with $t_{1}=t_{3} \neq t_{2}=t_{4}$ and $t_{1}=t_{4} \neq t_{3}=t_{2}$ by $(1+c / \bar{k}) / k_{1}$, giving a total bound on the second moment of

$$
\left.\mathbb{E}\left[\|(Q \otimes T) x\|_{2}^{4}\right] \leq 1+c / \bar{k}+3(1+c / \bar{k}) / k_{1}\right) \leq 1+(4+c) / \bar{k}
$$

since any term with a $t_{i} \notin\left\{t_{1}, t_{2}, t_{3}, t_{4}\right\} \backslash\left\{t_{i}\right\}$ immediately has expectation 0 . By induction, it follows that $\mathbb{E}\left[\left(\otimes_{i=1}^{q} S_{i}\right) x \|_{2}^{2}\right]=1$ for any unit $x$, and $\mathbb{E}\left[\left(\otimes_{i=1}^{q} S_{i}\right) x \|_{2}^{4}\right] \leq 1+(4 q+c) / \bar{k}$, where $c$ is the constant from the original variance of count-sketch. Setting $\bar{k}=k_{1}=\cdots=$ $k_{q}=\Theta\left(q k^{2} / \epsilon^{2}\right)$ with a large enough constant, this completes the proof that $S=\left(\otimes_{i=1}^{q} S_{i}\right)$
has the $O(\epsilon / k, 1 / 10,2)$-JL moment property. Then by Theorem 21 of [KN14], we obtain the approximate matrix product property:

$$
\operatorname{Pr}\left[\left\|A^{\top} S^{\top} S B-A^{\top} B\right\|_{F} \leq O(\epsilon / k)\|A\|_{F}\|B\|_{F}\right] \geq 9 / 10
$$

for any two matrices $A, B$. Letting $A=B^{\top}=U$ where $U \in \mathbb{R}^{n \times k}$ is a orthogonal basis for any $k$-dimensional subspace $\mathcal{V} \subset \mathbb{R}^{n}$, it follows that

$$
\left\|U^{\top} S^{\top} S U-I_{k}\right\|_{F} \leq O(\epsilon / k)\|U\|_{F}^{2} \leq O(\epsilon)
$$

where the last step follows because $U$ is orthonormal, so $\|U\|_{F}^{2}=k$. Since the Frobenius norm upper bounds the spectral norm $\|\cdot\|_{2}$, we have $\left\|U^{\top} S^{\top} S U-I_{k}\right\|_{2} \leq O(\epsilon)$, from which it follows that all the eigenvalues of $U^{\top} S^{\top} S U$ are in $(1-O(\epsilon), 1+O(\epsilon)$ ), which implies $\|S U x\|_{2}=(1 \pm O(\epsilon))\|x\|_{2}$ for all $x \in \mathbb{R}^{n}$, so for any $y \in \mathcal{V}$, let $x_{y}$ be such that $y=U x_{y}$, and then $\|S y\|_{2}=\left\|S U x_{y}\right\|_{2}=(1 \pm O(\epsilon))\left\|x_{y}\right\|_{2}=(1 \pm O(\epsilon))\left\|U x_{y}\right\|_{2}=(1 \pm O(\epsilon))\|y\|_{2}$, which proves that $S$ is a subspace embedding for $\mathcal{V}$ (not the second to last inequality holds because $U$ is orthonormal).

Finally, we are ready to prove Lemma 26.10.2.

Lemma 26.10.2. Let $S=\left(\otimes_{i=1}^{q} S_{i}\right)$ with $S_{i} \in \mathbb{R}^{k_{i} \times n_{i}}, k_{1}=k_{2}=\cdots=k_{q}=\Theta\left(q k^{2} / \epsilon^{2}\right)$. Then with probability 9/10 SA is a Projection Cost Preserving Sketch (PCP) for A, namely for all rank $k$ orthogonal projection matrix $P \in \mathbb{R}^{d \times d}$,

$$
(1-\epsilon)\|A-A P\|_{F}^{2} \leq\|S A-S A P\|_{F}^{2}+c \leq(1+\epsilon)\|A-A P\|_{F}^{2}
$$

where $c \geq 0$ is some fixed constant independent of $P$ (but may depend on $A$ and $S A$ ).

Proof. To demonstrate that $S A$ is a PCP, we show that the conditions of Lemma 10 of $\left[\mathrm{CEM}^{+} 15\right]$ hold, which imply this result. Our result follows directly from Theorem 12 of $\left[\mathrm{CEM}^{+} 15\right]$. Note that all that is needed (as discussed below the theorem) for the proof is that $S$ is an $\epsilon$-subspace embedding for a fixed $k$-dimensional subspaces, and that $S$ satisfies the $(\epsilon / \sqrt{k})$ approximate matrix product property. By Lemma 26.10.1, we have both $\epsilon$ subspace embedding for $S$ as well as a stronger $(\epsilon / k)$ approximate matrix product property. Thus Theorem 12 holds for the random matrix $S$ when $k_{1}=k_{2}=\cdots=k_{q}=\Theta\left(q k^{2} / \epsilon^{2}\right)$, which completes the proof.

### 26.11 Entry-wise Norm Low T-rank Approximation

We now demonstrate our results for low trank approximation of arbitrary input matrices. Specifically, we study the following problem, defined in [VL00]: given $A \in \mathbb{R}^{n^{2} \times n^{2}}$, the goal is to output a trank- $k$ matrix $B \in \mathbb{R}^{n^{2} \times n^{2}}$ such that

$$
\begin{equation*}
\|B-A\|_{\xi} \leq \alpha \cdot \mathrm{OPT} \tag{26.1}
\end{equation*}
$$

for some $\alpha \geq 1$, where $\mathrm{OPT}=\min _{\text {trank }-k A^{\prime}}\left\|A^{\prime}-A\right\|_{\xi}$, where the trank of a matrix $B$ is defined as the smallest integer $k$ such that $B$ can be written as a summation of $k$ matrices, where each matrix is the Kronecker product of $q$ matrices with dimensions $n \times n$ : $B=$ $\sum_{i=1}^{k} U_{i} \otimes V_{i}$, where $U_{i}, V_{i} \in \mathbb{R}^{n \times n}$.

Using Lemma 26.8.1, we can rearrange the entries in $A \in \mathbb{R}^{n^{2} \times n^{2}}$ to obtain $\bar{A} \in \mathbb{R}^{n^{2} \times n^{2}}$, where the $(i+n(j-1))^{\prime}$ 'th row of $\bar{A}$ is equal to $\operatorname{vec}\left(\left(A_{1}\right)_{i, j} A_{2}\right)$, and also vectorize the matrix $U_{i} \in \mathbb{R}^{n \times n}$ and $V_{i} \in \mathbb{R}^{n \times n}$ to obtain vectors $u_{i} \in \mathbb{R}^{n^{2}}, v_{i} \in \mathbb{R}^{n^{2}}$. Therefore, for any entry-wise
norm $\xi$ we have

$$
\left\|\sum_{i=1}^{k} U_{i} \otimes V_{i}-A\right\|_{\xi}=\left\|\sum_{i=1}^{k} u_{i} v_{i}^{\top}-\bar{A}\right\|_{\xi}
$$

Lemma 26.11.1 (Reshaping for Low Rank Approximation). There is a one-to-one mapping $\pi:[n] \times[n] \times[n] \times[n] \rightarrow\left[n^{2}\right] \times\left[n^{2}\right]$ such that for any pairs $(U, u) \in \mathbb{R}^{n \times n} \times \mathbb{R}^{n^{2}}$ and $(V, v) \in \mathbb{R}^{n \times n} \times \mathbb{R}^{n^{2}}$, if $U_{i_{1}, j_{1}}=u_{i_{1}+n\left(j_{1}-1\right)}$ and $V_{i_{1}, j_{1}}=v_{i_{1}+n\left(j_{1}-1\right)}$, then we have for $i_{1}, i_{2}, j_{1}, j_{2}$

$$
(U \otimes V)_{i_{1}+n\left(i_{2}-1\right), j_{1}+n\left(j_{2}-1\right)}=\left(u \cdot v^{\top}\right)_{\pi\left(i_{1}, i_{2}, j_{1}, j_{2}\right)}
$$

where $U \otimes V \in \mathbb{R}^{n^{2} \times n^{2}}$ and $u v^{\top} \in \mathbb{R}^{n^{2} \times n^{2}}$.
Proof. We have

$$
\begin{aligned}
(U \otimes V)_{i_{1}+n\left(i_{2}-1\right), j_{1}+n\left(j_{2}-1\right)} & =U_{i_{1}, j_{1}} V_{i_{2}, j_{2}} \\
& =u_{i_{1}+n\left(j_{1}-1\right)} \cdot v_{i_{2}+n\left(j_{2}-1\right)} \\
& =\left(u v^{\top}\right)_{i_{1}+n\left(j_{1}-1\right), i_{2}+n\left(j_{2}-1\right)}
\end{aligned}
$$

where the first step follows from the definition of $\otimes$ product, the second step follows from the connection between $U, V$ and $u, v$, the last step follows from the outer product.

Therefore, instead of using trank to define low-rank approximation of the $\otimes$ product of two matrices, we can just use the standard notion of rank to define it since both $B$ and $A^{\prime}$ can be rearranged to have rank $k$.

Definition 26.11.1 (Based on Standard Notion of Rank). Given two matrices $A_{1}, A_{2} \times \mathbb{R}^{n \times n}$, let $\bar{A} \in \mathbb{R}^{n^{2} \times n^{2}}$ denote the re-shaping of $A_{1} \otimes A_{2}$. The goal is to output a rank- $k$ matrix $\bar{B}$ such that

$$
\|\bar{B}-\bar{A}\|_{\xi} \leq \alpha \mathrm{OPT}_{\xi, k}
$$

where $\mathrm{OPT}_{\xi, k}=\min _{\text {rank }-k} \bar{A}^{\prime}\left\|\bar{A}^{\prime}-\bar{A}\right\|_{\xi}$.

In other words, $\bar{B}$ can be written as $\bar{B}=\sum_{i=1}^{k} u_{i} v_{i}^{\top}$ where $u_{i}, v_{i}$ are length $n^{2}$ vectors.
Combining the low-rank reshaping Lemma 26.11.1 with the main input-sparsity lowrank approximation of [CW13], we obtain our Frobenius norm low rank approximation result.

Theorem 26.11.2 (Frobenius norm low rank approximation, $p=2$ ). For any $\epsilon \in(0,1 / 2)$, there is an algorithm that runs in $n^{2}$ poly $(k / \epsilon)$ and outputs a rank-k matrix $\bar{B}$ such that $\|\bar{B}-\bar{A}\|_{F} \leq(1+\epsilon) \mathrm{OPT}_{F, k}$ holds with probability at least $9 / 10$, where $\mathrm{OPT}_{p}$ is cost achieved by best rank-k solution under the $\ell_{p}$-norm.

Similarly, using the main $\ell_{p}$ low rank approximation algorithm of [SWZ17], we have

Theorem 26.11.3 (Entry-wise $\ell_{p}$-norm low rank approximation, $1 \leq p \leq 2$ ). There is an algorithm that runs in $n^{2} \operatorname{poly}(k)$ and outputs a rank- $k$ matrix $\bar{B}$ such that $\|\bar{B}-\bar{A}\|_{p} \leq$ $\operatorname{poly}(k \log n) \mathrm{OPT}_{p, k}$ holds with probability at least $9 / 10$, where $\mathrm{OPT}_{p}$ is cost achieved by best rank-k solution under the $\ell_{p}$-norm.

Applying the bi-criteria algorithm of [CGK $\left.{ }^{+} 17 \mathrm{c}\right]$ gives us:

Theorem 26.11.4 (General $p>1$, bicriteria algorithm). There is an algorithm that runs in $\operatorname{poly}(n, k)$ and outputs a rank-poly $(k \log n)$ matrix $\bar{B}$ such that $\|\bar{B}-\bar{A}\|_{p} \leq \operatorname{poly}(k \log n) \operatorname{OPT}_{p}$ holds with probability at least $9 / 10$, where $\mathrm{OPT}_{p, k}$ is cost achieved by best rank-k solution under the $\ell_{p}$-norm.

Finally using the low-rank approximation algorithm for general loss functions given in [SWZ18], we obtain a very general result. The parameters for the loss function described in the following theorem are discussed in Section 26.12.

Theorem 26.11.5 (General loss function $g$ ). For any function $g$ that satisfies Definition 26.12.1, 26.12.2, 26.12.3, there is an algorithm that runs in $O\left(n^{2} \cdot T_{\text {reg }, g, n^{2}, k, n^{2}}\right)$ time and outputs a rank- $O(k \log n)$ matrix $\bar{B} \in \mathbb{R}^{n^{2} \times n^{2}}$ such that

$$
\|\bar{B}-\bar{A}\|_{g} \leq \operatorname{ati}_{g, k} \cdot \operatorname{mon}_{g} \cdot \operatorname{reg}_{g, k} \cdot O(k \log k) \cdot \mathrm{OPT}_{g, k},
$$

holds with probability $1-1 / \operatorname{poly}(n)$.

Hence, overall, the strategy is to first reshape $A=A_{1} \otimes A_{2}$ into $\bar{A}$, then compute $\bar{B}=\sum_{i=1}^{k} u_{i} v_{i}^{\top}$ using any of the above three theorems depending on the desired norm, and finally reshape $u_{i}$ and $v_{i}$ back to $U_{i} \in \mathbb{R}^{n \times n}$ and $V_{i} \in \mathbb{R}^{n \times n}$. It is easy to verify that the guarantees from Theorems 26.11.5, 26.11.3, 26.11.2 are directly transferable to the guarantee of the trank - $k$ approximation shown in Eq 26.1.

### 26.12 Properties for General Loss Functions for Low trank Approximation

We re-state three general properties (defined in [SWZ18]), the first two of which are structural properties and are necessary and sufficient for obtaining a good approximation from a small subset of columns. The third property is needed for efficient running time.

Definition 26.12.1 (Approximate triangle inequality). For any positive integer $n$, we say a function $g(x): \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ satisfies the $\operatorname{ati}_{g, n}$-approximate triangle inequality if for any $x_{1}, x_{2}, \cdots, x_{n} \in \mathbb{R}$ we have

$$
g\left(\sum_{i=1}^{n} x_{i}\right) \leq \operatorname{ati}_{g, n} \cdot \sum_{i=1}^{n} g\left(x_{i}\right) .
$$

Definition 26.12.2 (Monotone property). For any parameter $\operatorname{mon}_{g} \geq 1$, we say function $g(x): \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ is mon $_{g}$-monotone if for any $x, y \in \mathbb{R}$ with $0 \leq|x| \leq|y|$, we have $g(x) \leq$ $\operatorname{mon}_{g} \cdot g(y)$.

Definition 26.12.3 (Regression property). We say function $g(x): \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ has the $\left(\operatorname{reg}_{g, d}, T_{\text {reg }, g, n, d, m}\right)$-regression property if the following holds: given two matrices $A \in \mathbb{R}^{n \times d}$ and $B \in \mathbb{R}^{n \times m}$, for each $i \in[m]$, let $\mathrm{OPT}_{i}$ denote $\min _{x \in \mathbb{R}^{d}}\left\|A x-B_{i}\right\|_{g}$. There is an algorithm that runs in $T_{\text {reg }, g, n, d, m}$ time and outputs a matrix $X^{\prime} \in \mathbb{R}^{d \times m}$ such that

$$
\left\|A X_{i}^{\prime}-B\right\|_{g} \leq \operatorname{reg}_{g, d} \cdot \mathrm{OPT}_{i}, \forall i \in[m]
$$

and outputs a vector $v \in \mathbb{R}^{d}$ such that

$$
\mathrm{OPT}_{i} \leq v_{i} \leq \operatorname{reg}_{g, d} \cdot \mathrm{OPT}_{i}, \forall i \in[m]
$$

The success probability is at least $1-1 / \operatorname{poly}(n m)$.

```
Algorithm 26.4 Algorithm to \(\ell_{p}\) sample \(\Theta\left(r_{2}\right)\) entires of \(\rho=\left(A_{1} \otimes \cdots \otimes A_{q}\right) x^{\prime}-b\)
    procedure Residual \(\ell_{p} \operatorname{SAMPLE}\left(\rho, r_{2}\right)\)
        \(r_{3} \leftarrow \Theta\left(r_{2} \log ^{q^{2}}(n) / \delta\right)\).
        Generate i.i.d. \(p\)-stable vectors \(Z^{1, j}, Z^{2, j}, \ldots, Z^{q, j} \in \mathbb{R}^{n}\) for \(j \in[\tau]\) for \(\tau=\Theta(\log (n))\)
        \(T \leftarrow \emptyset \quad \triangleright\) sample set to return
        Pre-compute and store \(Z^{i, j} A_{i} \in \mathbb{R}^{1 \times d_{i}}\) for all \(i \in[q]\) and \(j \in[\tau]\)
        Generate count-sketches for heavy hitters \(S^{i} \in \mathbb{R}^{k \times n_{i}}\) of Definition 26.8.2 for all
    \(i \in[q]\), where \(k=O\left(\log ^{2}(n) r_{3}^{O(1)}\right)\).
            for \(t=1,2, \ldots, r_{3}\) do
                \(s=\left(s_{1}, \ldots, s_{q}\right) \leftarrow(\emptyset, \ldots, \emptyset) \quad \triangleright\) next sample to return
            \(w^{j} \leftarrow\left(\left(\mathbb{I}_{n_{1}}\right) \otimes\left(\bigotimes_{k=2}^{q} Z^{k, j}\right) \rho\right) \in \mathbb{R}^{n_{1}} \quad \triangleright \mathbb{I}_{n} \in \mathbb{R}^{n \times n}\) is identity
            Define \(w \in \mathbb{R}^{n_{1}}\) by \(w_{l}=\operatorname{median}_{j \in[\tau]}\left\{\left|w_{l}^{j}\right|\right\}\) for \(l \in\left[n_{1}\right]\)
            Sample \(j^{*} \in\left[n_{1}\right]\) from the distribution \(\left(\frac{\left|w_{1}\right|^{p}}{\|w\|_{p}^{p}}, \frac{\left|w_{2}\right|^{p}}{\|w\|_{p}^{p}}, \ldots, \frac{\left|w_{n_{1}}\right|^{p}}{\|w\|_{p}^{p}}\right)\)
            \(s_{1} \leftarrow j^{*}\)
            for \(i=2, \ldots, q\) do
                for \(j \in[\tau]\) do
                    Write \(e_{a_{k}}^{\top} \in \mathbb{R}^{1 \times n_{k}}\) as the standard basis vector
                    \(v_{i}^{j} \leftarrow S^{i}\left(\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(\mathbb{I}_{n_{i}}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) \rho\right) \in \mathbb{R}^{k}\)
                    Compute heavy hitters \(H_{i, j} \subset\left[n_{i}\right]\) from \(v_{i}^{j} \triangleright\) Definition 26.8.2
                    \(\beta_{i}^{j} \leftarrow\left(\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(\bigotimes_{k=i}^{q} Z^{k, j}\right) \rho\right) \in \mathbb{R}\)
                end for
                Define \(\beta_{i} \in \mathbb{R}^{k^{\prime}}\) by \(\beta_{i}=\operatorname{median}_{j \in[\tau]}\left\{\left|\beta_{i}^{j}\right|^{p}\right\}\)
                \(H_{i}=\cup_{j=1}^{\tau} H_{i, j}\)
                \(\gamma_{i} \leftarrow \operatorname{median}_{j \in[\tau]}\left(\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes Z_{\left[n_{i} \backslash H_{i}\right.}^{i, j} \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) \rho\right) \in \mathbb{R}\)
                if with probability \(1-\gamma_{i} / \beta_{i}\) then
                    Draw \(\xi \in H_{i}\) with probability
                    \(\frac{\operatorname{median}_{j \in \tau}\left|\left(\left(\bigotimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(e_{\xi}^{\top}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) \rho\right)\right|^{p}}{\sum_{\xi^{\prime} \in H_{i}} \operatorname{median}_{j \in \tau}\left|\left(\left(\otimes_{k=1}^{i-1} e_{a_{k}}^{\top}\right) \otimes\left(e_{\xi^{\prime}}^{\top}\right) \otimes\left(\bigotimes_{k=i+1}^{q} Z^{k, j}\right) \rho\right)\right|^{p}}\)
                            \(s_{i} \leftarrow \xi\)
                else \(\triangleright s_{i}\) was not sampled as a heavy hitter
            Randomly partition \(\left[n_{i}\right]\) into \(\Omega_{1}^{i}, \Omega_{2}^{i}, \ldots, \Omega_{\eta}^{i}\) with \(\eta=\Theta\left(r_{3}^{2}\right)\)
            Sample \(t \sim[\eta]\) uniformly at random
            for \(j \in \Omega_{t} \backslash H_{i}\) do
```



```
                end for
            Sample \(s_{i} \leftarrow j^{*}\) from the distribution \(\left\{\frac{\theta_{j}}{\sum_{j^{\prime} \in \Omega_{t} \backslash H_{i}} \theta_{j^{\prime}}}\right\}_{j \in \Omega_{t} \backslash H_{i}}\)
                end if
            end for
            \(T \leftarrow S \cup s\) where \(s=\left(s_{1}, \ldots, s_{q}\right)\)
        end for
```

```
Algorithm 26.5 Algorithm to \(\ell_{p}\) sample \(\Theta(r)\) rows of \(M=(F \otimes \mathbf{1}-\mathbf{1} \otimes F) R^{-1}\)
    procedure \(\ell_{p} \operatorname{SAMPLE}\left(F=[A, b] \in \mathbb{R}^{n \times d}, R^{-1} \in \mathbb{R}^{d+1 \times d+1}, r\right)\)
        Generate a matrix \(G \in \mathbb{R}^{d+1 \times \xi}\) of i.i.d. \(\mathcal{N}(0,1 / \sqrt{\xi})\) Gaussian random variables, with
    \(\xi=\Theta(\log (n))\)
        \(Y \leftarrow R^{-1} G \in \mathbb{R}^{d+1 \times \xi}\)
        \(C \leftarrow(F \otimes \mathbf{1}-\mathbf{1} \otimes F) Y\)
        Reshape i-th column \(C_{*, i}\) into \(\left(F Y_{*, i} \mathbf{1}^{\top}-\mathbf{1}\left(Y_{*, i}\right)^{\top} F^{\top}\right) \in \mathbb{R}^{n \times n}\)
        Generate \(Z \in \mathbb{R}^{t \times n}\) i.i.d. \(p\)-stable for \(t=\Theta(\log (n)) \quad \triangleright\) Definition 26.2.1
        For all \((i, l) \in[\xi] \times[n]\), set
            \(\sigma_{i, l} \leftarrow \underset{\tau \in[t]}{\operatorname{median}}\left(\frac{\mid\left(\left.Z\left(F Y_{*, i} \mathbf{1}^{\top}-\mathbf{1}\left(Y_{*, i}\right)^{\top} F^{\top}\right)_{\tau, l}\right|^{p}\right.}{\left(\operatorname{median}\left(\mathcal{D}_{p}\right)\right)^{p}}\right)\)
                                    \(\triangleright\) Indyk Estimator [Ind06]
            Set \(W^{(i, l)} \leftarrow\left(F Y_{*, i} \mathbf{1}^{\top}-\mathbf{1} Y_{*, i}^{\top} F^{\top}\right)_{*, l}=F Y_{*, i}-\mathbf{1}(F Y)_{l, i} \in \mathbb{R}^{n}\)
            for \(j=1, \ldots, \Theta(r)\) do
            Sample \((i, l)\) from distribution \(\sigma_{i, l} /\left(\sum_{i^{\prime}, l^{\prime}} \sigma_{i^{\prime}, l^{\prime}}\right)\).
        end for
            \(T \leftarrow\) multi-set of samples \((i, l)\)
            Generate \(S_{0} \in \mathbb{R}^{k \times n} S \in \mathbb{R}^{k^{\prime} \times n}\) count-sketches for heavy hitters with \(k=r^{O(1)}, k^{\prime}=\)
    \(k^{O(1)}\).
                                    \(\triangleright\) Definition 26.8.2
            Generate \(u_{1}, \ldots, u_{n}\) i.i.d. exponential variables.
            \(D \leftarrow \operatorname{Diag}\left(1 / u_{1}^{1 / p}, \ldots, 1 / u_{n}^{1 / p}\right) \in \mathbb{R}^{n \times n}\).
            for each sample \((i, l) \in T\) do
            Compute \(S_{0} W^{(i, l)}\) and obtain set of heavy hitters \(Q_{0}^{(i, l)} \subset[n]\)
            Compute \(W_{j}^{(i, l)}\) exactly for all \(j \in Q_{0}^{(i, l)}\), to obtain true heavy hitters \(H^{(i, l)}\).
            Compute
\[
\alpha_{i, l} \leftarrow \underset{\tau \in[t]}{\operatorname{median}}\left(\frac{\left|Z_{\tau, *} W^{(i, l)}-\sum_{\zeta \in H^{(i, l)}} Z_{\tau, \zeta} W_{\zeta}^{(i, l)}\right|^{p}}{\left(\operatorname{median}\left(\mathcal{D}_{p}\right)\right)^{p}}\right)
\]
if With prob \(1-\alpha_{(i, l)} / \sigma_{(i, l)}\), sample a heavy item \(j^{*} \leftarrow j\) then
Sample a heavy item \(j^{*} \leftarrow j\) from the distribution \(\left|W_{j}^{(i, l)}\right|^{p} / \sum_{j \in H_{(i, l)}}\left|W_{j}^{(i, l)}\right|^{p}\).
return The row \(\left((l-1) n+j^{*}\right) \quad\) Note that \(C_{(l-1) n+j^{*}, *}\) contains \(W_{j^{*}}^{(i, l)}\) else
Randomly partition [ \(n\) ] into \(\Omega_{1}, \Omega_{2}, \ldots, \Omega_{\eta}\) with \(\eta=\Theta\left(r^{4} / \epsilon^{4}\right)\).
Sample \(t \sim[\eta]\) uniformly at random.
Compute \(S\left(D W^{(i, l)}\right)_{\Omega_{t} \backslash H^{(i, l)}}\), an\& set \(Q^{(i, l)} \subset \Omega_{t} \backslash H^{(i, l)}\) of heavy hitters.
\(j^{*} \leftarrow \arg \max _{j \in Q^{(i, l)}}\left(D W^{(i, l)}\right)_{j}\)
return The row \(\left((l-1) n+j^{*}\right) \quad\) Note that \(C_{(l-1) n+j^{*}, *}\) contains \(W_{j^{*}}^{(i, l)}\) end if
end for
end procedure
```


## Chapter 27

## Dynamic $k$-means Clustering

We consider the $k$-means clustering problem in the dynamic streaming setting, where points from a discrete Euclidean space $\{1,2, \ldots, \Delta\}^{d}$ can be dynamically inserted to or deleted from the dataset. For this problem, we provide a one-pass coreset construction algorithm using space $\widetilde{O}(k \cdot \operatorname{poly}(d, \log \Delta))$, where $k$ is the target number of centers. To our knowledge, this is the first dynamic geometric data stream algorithm for $k$-means using space polynomial in dimension and nearly optimal (linear) in $k$.

This part is based upon the following previous publication

- Wei Hu, Zhao Song, Lin F. Yang, Peilin Zhong

Nearly Optimal Dynamic $k$-Means Clustering for High-Dimensional Data.
Manuscript 2018 [HSYZ18]

### 27.1 Introduction

Clustering is one of the central problems in unsupervised learning. The idea is to partition data points into clusters in the hope that points in the same cluster are similar to each other and points in different clusters are dissimilar. One of the most important approaches to clustering is $k$-means, which has been extensively studied for more than 60 years and has a wide range of applications (see e.g. [Jai10] for a survey). Given a set of points $Q \subset \mathbb{R}^{d}$, the $k$-means problem asks for a set of $k$ centers $Z \subset \mathbb{R}^{d}$ such that the sum of squares of distances between data points to their closest centers is minimized, i.e., it tries to solve $\min _{Z \subset \mathbb{R}^{d},|Z|=k} \operatorname{cost}(Q, Z)$, where $\operatorname{cost}(Q, Z)$ is a cost function defined as:

$$
\operatorname{cost}(Q, Z):=\sum_{q \in Q} \min _{z \in Z} \operatorname{dist}^{2}(q, z)
$$

Here $\operatorname{dist}(\cdot, \cdot)$ stands for the Euclidean distance.
A major challenge in dealing with massive datasets is that the entire input data can be too large to be stored. A standard model of study in such settings is the streaming model, where data points arrive and are processed one at a time, and only a small amount of useful information (i.e., a sketch) about the data is maintained. See e.g. [Mut05b] for an introduction to the streaming model.

In this paper we study the $k$-means problem over dynamic data streams [Ind04], where data points from a discrete space $\{1,2, \ldots, \Delta\}^{d}$ can be either inserted to or deleted from the dataset. A standard approach to solving $k$-clustering problems like $k$-means and $k$-median in the streaming setting is to maintain an $\epsilon$-coreset, which is a small number of (weighted) points whose cost with respect to any $k$ centers is a $(1+\epsilon)$-approximation to the cost of the entire dataset on the same $k$ centers. As a consequence, at the end of the stream, we
only need to find an approximate $k$-means solution on the coreset, which is automatically an approximate solution on the entire dataset. Hence our goal is to design an efficient method to maintain an $\epsilon$-coreset over a dynamic data stream using as small space as possible.

### 27.1.1 Our Result

Theorem 27.1.1 (Main theorem, restatement of Theorem 27.6.22 in Section 27.6). Let $\epsilon \in(0,1 / 2), k, \Delta \in \mathbb{N}_{+}$, and $L=\log \Delta$. For dynamic data stream consisting of insertions and deletions of points in $[\Delta]^{d}$, there is an algorithm which uses a single pass over the stream and on termination outputs a weighted set $S$ with a positive weight for each point therein, such that with probability at least $0.9, S$ is an $\epsilon$-coreset for $k$-means of size $O\left(k \epsilon^{-2} d^{4} L^{2} \log (k d L)\right)$. The algorithm uses $\widetilde{O}(k) \cdot \operatorname{poly}\left(d, L, \epsilon^{-1}\right)$ bits in the worst case.

To our knowledge, this is the first algorithm for $k$-means in dynamic data streams that uses space polynomial in data dimension $d$ and nearly optimal (linear) ${ }^{1}$ in the number of clusters $k$. Previous algorithms for streaming $k$-means either require space exponential in $d$ or only work for insertion-only streams. ${ }^{2}$ See Section 27.1.3 and Section 27.7 for detailed discussions of previous results.

### 27.1.2 Our Techniques

At a high level our algorithm is based on a framework called sensitivity sampling, which was proposed by [FL11]. For a set $Q \subseteq[\Delta]^{d}$, the sensitivity of every point $q \in Q$ is

[^92]defined as
$$
s(q):=\max _{Z \subset \mathbb{R}^{d},|Z|=k} \frac{\operatorname{dist}^{2}(q, Z)}{\sum_{p \in Q} \operatorname{dist}^{2}(p, Z)}
$$

Namely, $s(q)$ represents how "sensitive" the cost can be to the removal of point $q$. A crucial result shown by [FL11, BFL16] is that once we know a good upper bound on each point's sensitivity, there is a sampling method to construct an $\epsilon$-coreset. Specifically, if we know an upper bound $s^{\prime}(q) \geq s(q)$ for each $q \in Q$, we can sample $q$ with probability $s^{\prime}(q) /\left(\sum_{p \in Q} s^{\prime}(p)\right)$. Let $R$ be a set of i.i.d. samples from this procedure with $|R| \geq \widetilde{\Omega}\left(\sum_{q \in Q} s^{\prime}(q) / \epsilon^{2}\right)$, and each sample $q$ is assigned a weight $\frac{\sum_{p \in Q} s^{\prime}(p)}{|R| s^{\prime}(q)}$. Then with high probability $R$ is an $\epsilon$-coreset for $Q$. Note that if $\sum_{q \in Q} s^{\prime}(q)=\widetilde{O}(k \cdot \operatorname{poly}(d))$, then an $\widetilde{O}(k \cdot \operatorname{poly}(d))$-size $\epsilon$-coreset can be constructed in this way. The formal description of this result is given in Theorem 27.2.1.

We give an efficient method to obtain sensitivity upper bounds $s^{\prime}(\cdot)$ such that: (i) $\sum_{q \in Q} s^{\prime}(q)$ is small, (ii) we can implement the sensitivity sampling procedure in the dynamic streaming setting. Then we are able to construct a coreset according to the previous paragraph.

The key intuition in our sensitivity estimation is the following. Imagine that there is a small region that is very dense, i.e., it contains a lot of points. Then the sensitivity of every point in that region must be low, because that point can be well represented by other points in the same small region. Therefore, the problem of finding sensitivity upper bound for a point boils down to figuring out the "right" region this point belongs to that can be considered "dense." Intuitively the sensitivity of this point depends on the size of this dense region - the smaller the size, the smaller the sensitivity.

We make this intuition formal by using a hierarchical grid structure similar to [FS05, $\left.B F L^{+} 17\right]$. This structure is illustrated in Figure 27.1. The top-level (level 0) grid consists of


Figure 27.1: The grid structure over the point set. From top to bottom, three levels of grids are shown. Each cell splits into $2^{d}$ cells in the next level.
cells that are $d$-dimensional cubes of side-length $\Delta$, and each cell in level $i-1$ splits into $2^{d}$ cells in level $i$. Each cell in level $i$ has side-length $\Delta / 2^{i}$. For a cell in level $i$, we say that it is heavy if it contains at least $T_{i}=\Theta\left(\frac{d^{2}}{k} \cdot \frac{\mathrm{OPT}}{\left(\Delta / 2^{2}\right)^{2}}\right)$ points in $Q$, where OPT is the optimal cost of the $k$-means problem. ${ }^{3}$ Since $T_{i}>T_{i-1}$, we know that if a cell in level $i$ is heavy, then its parent cell in level $i-1$ is heavy as well. Therefore the set of all heavy cells in all levels form a tree. Now for a point $p \in Q$, denote by $c_{i}(p)$ the cell in level $i$ that contains $p$, and then define $j$ to be the smallest level index such that $c_{j}(p)$ is not heavy; then we show an upper bound on the sensitivity $s(p)$ solely based on this index number $j$, namely $s(p) \leq s^{\prime}(p)=\Theta\left(d^{3} / T_{j}\right)$. Furthermore, we prove that the sum of our sensitivity upper bounds is small: $\sum_{p \in Q} s^{\prime}(p)=O\left(k d^{3} \log \Delta\right)$, which satisfies our requirement. To establish these bounds we need the total number of heavy cells to be small, for which we apply a random shift of grid at the beginning, as illustrated in Figure 27.2.

To implement the above sensitivity sampling method in the dynamic streaming setting when the dataset is updated by insertions and deletions of points, the key difficulties are: 1) we do not know the value of OPT, and it changes when the underlying dataset is updated; 2) we need to compute the sensitivity upper bounds and to sample points at the same time using limited space.

Let us first assume OPT is known and give an algorithm to implement our sensitivity sampling procedure in the dynamic streaming setting. Our algorithm makes crucial use of the $k$-set data structure in [Gan05] for counting distinct elements in a dynamic stream. The $k$-set data structure ensures that if the number of distinct elements is at most some

[^93]predetermined parameter, it will return all distinct elements and their frequencies; otherwise it will return FAIL. We summarize its guarantee in Lemma 27.6.1. Note that in order to implement sensitivity sampling, we need to know which cells are heavy. Our algorithm dynamically tracks all heavy cells, using the $k$-set structure as a building block. Then the sensitivity sampling method has two stages: first sample a level $i$ (with an appropriate probability for each level), and then uniformly sample a point from all points associated with level $i$, i.e., all points $p$ such that $c_{i}(p)$ is not heavy and $c_{i-1}(p)$ is heavy. (Note that for all points associated with level $i$, they have the same sensitivity upper bounds, which means uniformly sampling a point from them is enough.) In order to do uniform sampling, we also maintain for each level $i$ a uniformly random subset of points associated with $i$. Therefore it suffices to choose a point uniformly at random from this subset once $i$ is chosen.

For the issue of not knowing OPT, we run in parallel multiple copies of our sampling algorithm for different guesses of OPT: $1,2,4 \ldots, \Delta^{d} \cdot d \Delta$. Our sampling algorithm ensures that when the guessed value is less than OPT but not too far away, the required space is small. For other guesses, the required space might be a lot, but since we have a space budget, our algorithm can return FAIL when the space runs out. Since at least one guess is accurate, at least one copy of the algorithm will succeed and output a small $\epsilon$-coreset.

### 27.1.3 Related Work

It is well known that exactly solving $k$-means is NP-hard even for $k=2$ or $d=2$ [ADHP09, MNV09]. The most successful algorithm used in practice is Lloyd's algorithm, which is also known as "the $k$-means method" [Llo82]. Because of the NP-hardness, various attempts were made on approximation algorithms. $\left[\mathrm{KMN}^{+} 02\right]$ proved that a very simple


Figure 27.2: Random shift of grid brings down the number of heavy cells. In the left panel, we have a bad alignment of points and grids such that many cells contain lots of points. In the right panel, after the random shift, only two cells contain many points.
local search heuristic achieves $(9+\epsilon)$-approximation in polynomial time for any fixed $\epsilon>0$. When $d$ is a constant [FRS16, CAKM16] or $k$ is a constant [FMS07, KSS10, FL11], $(1+\epsilon)-$ approximation can be achieved in polynomial time.

There is a line of work studying $k$-means and $k$-median in insertion-only streams, e.g., [BS80a, GMMO00, COP03, BDMO03, AHPV04, HPM04, HPK05, Che09, FL11, FS12, $\mathrm{AMR}^{+} 12$, BFL16]. There also have been a lot of interests in dynamic streaming algorithms for other problems, e.g. $\left[\mathrm{BYJK}^{+} 02, \mathrm{FKM}^{+} 05\right.$, Bas08, KL11, AGM12a, AGM12b, GKP12, GKK12, AGM12b, BKS12a, CMS13, AGM13, Moi14a, BGS15, BHNT15, BS16, ACD ${ }^{+} 16$, $\mathrm{ADK}^{+} 16, \mathrm{BWZ16}, \mathrm{KLM}^{+} 17$, SWZ17, SWZ19b]. In addition, $k$-means and $k$-median were studied in various different settings, e.g., [CCGG98, IP11, BIRW16, BCMN14, SW18].

The most relevant papers are [FS05, FL11, BFL16, BFL+17]. [FS05] designed an algorithm to maintain an $\epsilon$-coreset of size $k \epsilon^{-O(d)}$ for $k$-means and $k$-median. [FL11] introduced the sensitivity sampling framework for coreset construction, and their approach was further improved by [BFL16], but both of them only work for insertion-only streams and do
not apply to dynamic streams. $\left[\mathrm{BFL}^{+} 17\right]$ focused on the $k$-median problem and constructed a coreset of size $O(k \cdot \operatorname{poly}(d, \log \Delta))$ in the dynamic streaming setting, but their technique heavily relies on $k$-median and cannot be extended to $k$-means. In Section 27.7 we explain in detail the limitations of previous approaches.

### 27.1.4 Roadmap

In Section 27.2, we introduce necessary notation and preliminaries. In Section 27.3, we give the offline version of our algorithm. Due to the page limit, missing proofs in Section 27.3 are given in Section 27.5, and we present our dynamic streaming algorithm and its analysis in Section 27.6.

### 27.2 Preliminaries

Notation. For $n \in \mathbb{N}_{+}$, let $[n]:=\{1,2, \ldots, n\}$. We define $\widetilde{O}(f)$ to be $O\left(f \cdot \log ^{O(1)}(f)\right)$. For any $x \in \mathbb{R}_{\geq 0}, \epsilon \in(0,1)$, we use $(1 \pm \epsilon) \cdot x$ to denote the interval $((1-\epsilon) \cdot x,(1+\epsilon) \cdot x)$. For any $x \in \mathbb{R}$ and $a \in \mathbb{R}_{>0}, x \pm a$ denotes the interval $(x-a, x+a)$.

We denote by $\operatorname{dist}(\cdot, \cdot)$ the Euclidean distance in $\mathbb{R}^{d}$, i.e., for $p, q \in \mathbb{R}^{d}$, $\operatorname{dist}(p, q):=$ $\|p-q\|_{2}$. For sets $P, Q \subseteq \mathbb{R}^{d}$ and a point $p \in \mathbb{R}^{d}$, we define $\operatorname{dist}(p, Q)=\operatorname{dist}(Q, p):=$ $\min _{q \in Q} \operatorname{dist}(p, q)$ and $\operatorname{dist}(P, Q):=\min _{p \in P, q \in Q} \operatorname{dist}(p, q)$. For any two sets $Q, Z \subseteq \mathbb{R}^{d}$, we define $\operatorname{cost}(Q, Z):=\sum_{q \in Q} \operatorname{dist}^{2}(q, Z)$. We define $\operatorname{diam}(Q)$ to be $Q$ 's diameter, i.e., $\operatorname{diam}(Q):=$ $\max _{p, q \in Q} \operatorname{dist}(p, q)$.

The dynamic streaming model. We consider the dynamic streaming model, defined below.

Definition 27.2.1 (Dynamic streaming model). Let $Q \subseteq[\Delta]^{d}$ initially be an empty set. In the dynamic streaming model, there is a stream of update operations such that the $t^{\text {th }}$ operation has the form $\left(p_{t}, \pm\right)$ which indicates that a point $p_{t} \in[\Delta]^{d}$ is inserted to or deleted from the set $Q$, where + denotes insertion and - denotes deletion. There is no invalid deletion during the stream. ${ }^{4}$ An algorithm is allowed a single pass over the stream. At the end of the stream, the algorithm stores some information regarding $Q$. The space complexity of an algorithm in this model is defined as the total number of bits used by the algorithm during the stream.

[^94]The goal of an algorithm in this model is to store some information which can be used for a certain computation task, while using as small space as possible. Although optimizing the running time is not required in this model, the algorithm in the current paper is actually efficient for each update.

In this paper we suppose that any two points in $Q$ have different locations ${ }^{5}$, i.e., $Q$ is not a multiset. Our algorithm can be easily extended to allow multiple copies of a point by blowing up the total space by an $O(\log M)$ factor, where $M$ is an upper bound on the number of copies.
$k$-means clustering. Now we introduce the $k$-means clustering problem and the notion of coreset.

Definition 27.2.2 ( $k$-means clustering). Given a point set $Q \subseteq[\Delta]^{d}$ and a parameter $k \in \mathbb{N}_{+}$ for the target number of centers, the goal of $k$-means clustering is to find a set of $k$ points $Z \subseteq \mathbb{R}^{d}$ such that the objective function, $\operatorname{cost}(Q, Z):=\sum_{q \in Q} \operatorname{dist}^{2}(q, Z)$, is minimized. Each point in $Z$ is called a center. OPT is defined to be the optimal cost of the $k$-means clustering problem.

However, solving the $k$-means problem exactly is NP-hard [ADHP09]. Oftentimes, we only need a good approximation. For the purpose of finding an approximate solution, an important concept is coreset, which is a small subset of (weighted) points whose $k$-means solution is a good approximate solution for the entire dataset. The formal definition is the following:

[^95]Definition 27.2.3 (Coreset for $k$-means). Given $Q \subseteq[\Delta]^{d}, k \in \mathbb{N}_{+}$and $\epsilon>0$, a set of point-weight pairs $S=\left\{\left(s_{1}, w_{1}\right),\left(s_{2}, w_{2}\right), \ldots,\left(s_{m}, w_{m}\right)\right\} \subset[\Delta]^{d} \times \mathbb{R}_{>0}$ is an $\epsilon$-coreset for $Q$, where $w_{i}$ is the weight of $s_{i}$, if $S$ satisfies

$$
\forall Z \subset \mathbb{R}^{d},|Z|=k:|\operatorname{cost}(S, Z)-\operatorname{cost}(Q, Z)| \leq \epsilon \cdot \operatorname{cost}(Q, Z)
$$

where $\operatorname{cost}(S, Z):=\sum_{i=1}^{m} w_{i} \operatorname{dist}^{2}\left(s_{i}, Z\right)$. The size of the coreset is $|S|$.

The main problem studied in this paper is how to construct a small coreset for $k$ means over a dynamic data stream. The formal description is the following.

Definition 27.2.4 (Coreset for $k$-means over a dynamic stream). Given a point set $Q \subseteq[\Delta]^{d}$ described by a dynamic stream of operations (Definition 27.2.1), a parameter $k \in \mathbb{N}_{+}$for the target number of centers, and an error parameter $\epsilon \in(0,0.5)$. The goal is to design an algorithm in the dynamic streaming model which can with probability at least 0.9 output a small size $k$-means $\epsilon$-coreset (Definition 27.2.3) for $Q$ using as small space as possible.

Sensitivity sampling based coreset construction. Let us briefly review the coreset construction framework proposed by [FL11, BFL16]. Given $Q \subseteq[\Delta]^{d}$ and $k \in \mathbb{N}_{+}$, the sensitivity of a point $p \in Q$ is defined as:

$$
s(p)=\max _{Z \in \mathbb{R}^{d},|Z|=k} \frac{\operatorname{dist}^{2}(p, Z)}{\sum_{q \in Q} \operatorname{dist}^{2}(q, Z)}
$$

The following theorem gives guarantee of a sensitivity sampling based coreset construction.
Theorem 27.2.1 ([FL11, BFL16]). Given a set of points $Q \subseteq[\Delta]^{d}$ and a parameter $k$, let $s(p)$ denote the sensitivity of each point $p \in Q$. For each $p \in Q$, let $s^{\prime}(p)$ be an upper
bound on the sensitivity of $p$, i.e., $s^{\prime}(p) \geq s(p)$, and let $t^{\prime}=\sum_{p \in Q} s^{\prime}(p)$. Consider a multiset $S$ of $m$ i.i.d. samples from $Q$, where each sample chooses $p \in Q$ with probability $s^{\prime}(p) / t^{\prime}$. For each sampled point $p$, a weight $w(p) \in(1 \pm \epsilon / 2) \cdot t^{\prime} /\left(m s^{\prime}(p)\right)$ is associated with $p$. If $m \geq \Omega\left(t^{\prime} \epsilon^{-2}\left(\log |Q| \log t^{\prime}+\log (1 / \delta)\right)\right)$, then with probability at least $1-\delta,\{(p, w(p)) \mid p \in S\}$ is an $\epsilon$-coreset (Definition 27.2.3) for $Q$.

According to the above theorem, if we can find a good sensitivity upper bound $s^{\prime}(p)$ for each point $p$, then we are able to construct a coreset with size nearly linear in $t^{\prime}=\sum_{p} s^{\prime}(p)$. In section 27.3, we give an offline algorithm which can estimate a good sensitivity upper bound for each point, which readily implies an efficient offline coreset construction algorithm. In Section 27.6, we show how to implement this sensitivity sampling procedure over a dynamic stream. Notice that [BFL16] gave a sensitivity sampling framework that works for clustering with general loss functions, and our method can be extended to those problems as well.

### 27.3 An Offline Sensitivity Sampling Procedure

In this section, we consider the offline setting in which all the data points are given. In this setting, we design a coreset construction algorithm based on sensitivity sampling. In Section 27.6, we will show how to implement this algorithm in the dynamic streaming setting.

### 27.3.1 Randomly Shifted Grids

We consider data points from $[\Delta]^{d}$ and assume without loss of generality that $\Delta=2^{L}$ for some positive integer $L$. The space $[\Delta]^{d}$ is partitioned by a hierarchical grid structure as follows. The first level (level 0) of the grid contains cells with side-length $\Delta$ such that all the data points are contained in a single cell. For each higher level, we refine the grid by splitting each cell into $2^{d}$ equal sized sub-cells. In the finest level, i.e., the $L$-th level, each cell contains a single point. We further randomly shift the boundary of the grids to achieve certain properties, which we will show later. Formally, our grid structure is defined as the following.

Definition 27.3.1 (Grids and cells). Let $g_{0}=\Delta$. Choose a vector $v$ uniformly at random from $[0, \Delta]^{d}$. Partition the space $\mathbb{R}^{d}$ into a regular Cartesian grid $G_{0}$ with side-length $g_{0}$ and translate $G_{0}$ such that a vertex of this grid falls on $v$. The grid $G_{0}$ can be regarded as an infinite set of disjoint cells, where each cell $C \in G_{0}$ can be expressed as

$$
\left[v_{1}+n_{1} g_{0}, v_{1}+\left(n_{1}+1\right) g_{0}\right) \times \cdots \times\left[v_{d}+n_{d} g_{0}, v_{d}+\left(n_{d}+1\right) g_{0}\right) \subset \mathbb{R}^{d}
$$

for some $\left(n_{1}, n_{2}, \ldots, n_{d}\right) \in \mathbb{Z}^{d}$. (Note that each cell is a Cartesian product of intervals.)

For $i \geq 1$, we define the regular grid $G_{i}$ as the grid with side-length $g_{i}=g_{0} / 2^{i}$ aligned such that each cell in $G_{i-1}$ contains $2^{d}$ cells in $G_{i}$. The finest grid is $G_{L}$ where $L=\log _{2} \Delta$. A cell of $G_{L}$ has side-length 1 and thus contains at most one data point.

For convenience, we also define $G_{-1}$ to be the regular grid with side-length $g_{-1}=2 \Delta$, and each cell in $G_{-1}$ is a union of $2^{d}$ cells in $G_{0}$. Since the data points are in $[\Delta]^{d}$, there must be a single cell in $G_{-1}$ which contains all the data points. Consider two cells $C \in G_{i}$ and $C^{\prime} \in G_{j}$ for some $i, j \in\{-1,0,1, \ldots, L\}$. If $C^{\prime} \subset C$, then $C$ is an ancestor of $C^{\prime}$. Furthermore, if $i=j-1$, then $C$ is the parent of $C^{\prime}$ and $C^{\prime}$ is a child of $C$. Thus every cell which is not from $G_{L}$ has exactly $2^{d}$ children cells. For a point $p$ (or a set $P$ of points), $c_{i}(p)$ (or $c_{i}(P)$ ) denotes the cell $C$ in grid $G_{i}$ which contains $p$ (or $P$ ). If $i$ is clear from the context, we will just use $c(p)$ (or $c(P)$ ) for short.

### 27.3.2 Sensitivity Estimation and Coreset Construction

In Algorithm 27.1 we describe how to assign a sensitivity upper bound for every point. It needs an estimate $o$ of the optimal $k$-means cost OPT. We will show how to enumerate the guesses $o$ later. According to Theorem 27.2.1, it directly gives an offline coreset construction algorithm.

We also give an alternative sampling procedure in Algorithm 27.2 which is useful for the dynamic streaming model.

Theorem 27.3.1. Suppose that for any $i \in\{0,1, \ldots, L\}$ and for any cell $C \in G_{i}$ with $C \cap Q \neq \emptyset$, the estimated value $z$ in line 27.1 of Algorithm 27.1 satisfies either $z \in \mid C \cap$ $Q \mid \pm 0.1 T_{i}(o)$ or $z \in(1 \pm 0.01) \cdot|C \cap Q|$, and for any $Q_{i}$, the estimated value $\widehat{q}_{i}$ in line 27.2

```
Algorithm 27.1 Sensitivity Estimation
    predetermined: a guess \(o \in\left[1, \Delta^{d} \cdot d \Delta\right]\) of the optimal \(k\)-means cost OPT
    input: a point set \(Q \subseteq[\Delta]^{d}\), a parameter \(k \in \mathbb{N}_{+}\)
    Impose randomly shifted grids \(G_{-1}, G_{0}, G_{1}, \ldots, G_{L}\) (Definition 27.3.1).
    Let \(C \in G_{-1}\) be the cell which contains \(Q\), i.e., \(C=c(Q)\). Mark \(C\) as heavy.
    for \(i:=0 \rightarrow L-1\) do
        Set the threshold value \(T_{i}(o)=\left(d / g_{i}\right)^{2} \cdot o / k \cdot 1 / 100\).
        for \(C \in G_{i}\) with \(C \cap Q \neq \emptyset\) do
            Let \(z\) be an estimated value of \(|C \cap Q|\) up to some precision.
            If \(z \geq T_{i}(o)\), mark \(C\) as heavy.
            Otherwise, if all the ancestors of \(C\) are marked as heavy, mark \(C\) as crucial.
            end for
    end for
    For \(C \in G_{L}\), if all the ancestors of \(C\) are marked as heavy, mark \(C\) as crucial.
    Initialize \(Q_{0}=Q_{1}=\cdots=Q_{L}=\emptyset\).
    For \(p \in Q\), if \(c_{i}(p)\) is marked as crucial, add \(p\) into set \(Q_{i}\) and set \(s^{\prime}(p)=10 d^{3} / T_{i}(o)\).
    output: \(Q_{0}, Q_{1}, \ldots, Q_{L}\) and \(s^{\prime}(\cdot)\)
```

of Algorithm 27.2 satisfies either $\widehat{q}_{i} \in\left|Q_{i}\right| \pm 0.1 \epsilon \gamma T_{i}(o)$ or $\widehat{q}_{i} \in(1 \pm 0.01 \epsilon) \cdot\left|Q_{i}\right|$. Suppose $o \in(0, \mathrm{OPT}]$. Then the set $S$ output by Algorithm 27.2 is an $\epsilon$-coreset (Definition 27.2.3) for $Q$. Furthermore, with probability at least $0.93,|S|$ is at most

$$
O\left(k \epsilon^{-2} d^{4} L^{2} \log (k d L) \cdot(\mathrm{OPT} / o+1)\right)
$$

### 27.3.3 Analysis

Now we give the proof of Theorem 27.3.1. All the missing proofs in this section are given in Section 27.5. Let us first state some simple facts.

Fact 27.3.2. The point sets $Q_{0}, Q_{1}, \ldots, Q_{L}$ obtained by Algorithm 27.1 form a partition of $Q$, i.e., for all $p \in Q$, there is exactly one $i \in\{0,1, \ldots, L\}$ such that $p \in Q_{i}$.

```
Algorithm 27.2 Sensitivity Sampling Based Coreset Construction
    predetermined: a guess \(o\) of the optimal \(k\)-means cost OPT, an error parameter \(\epsilon \in\)
    (0, 0.5)
    input: a point set \(Q \subset[\Delta]^{d}\), a parameter \(k \in \mathbb{N}_{+}\)
    Let \(Q_{0}, Q_{1}, \ldots, Q_{L}\) and \(s^{\prime}(\cdot)\) be the output of Algorithm 27.1.
    Let \(\widehat{q}_{0}, \widehat{q}_{1}, \ldots, \widehat{q}_{L}\) be the estimated values of \(\left|Q_{0}\right|,\left|Q_{1}\right|, \ldots,\left|Q_{L}\right|\) respectively.
    For \(i \in\{0,1, \ldots, L\}\), set \(T_{i}(o)=\left(d / g_{i}\right)^{2} \cdot o / k \cdot 1 / 100\) (same as in Algorithm 27.1).
    Set \(\gamma=\epsilon /\left(40^{2} L d^{3}\right)\).
    Let \(I=\left\{i \mid 0 \leq i \leq L, \widehat{q}_{i} \geq \gamma T_{i}(o)\right\}\).
    Let \(Q^{I}=\bigcup_{i \in I} Q_{i} . \quad \triangleright\) Only consider the levels with sufficient number of points.
    Set \(t^{\prime}=\sum_{i \in I} \widehat{q}_{i} \cdot 10 d^{3} / T_{i}(o)\). \(\triangleright\) Total estimated sensitivities.
    Set \(m=\Theta\left(t^{\prime} \epsilon^{-2} L d \log t^{\prime}\right)\) and initialize \(S=\emptyset . \quad \triangleright m\) is the total number of samples.
    for \(j=1 \rightarrow m\) do
        Choose a random level \(i \in I\) with probability \(\left(\widehat{q}_{i} \cdot 10 d^{3} / T_{i}(o)\right) / t^{\prime}\).
        Uniformly sample a point \(p\) from \(Q_{i}\).
        Add \(\left(p, t^{\prime} /\left(m s^{\prime}(p)\right)\right)\) to set \(S\).
    end for
    output: the set \(S\)
```

Fact 27.3.3. For $C \in G_{i}$, if $C$ is marked as heavy, then $|C \cap Q| \geq 0.9 T_{i}(o)$; otherwise $|C \cap Q| \leq 1.1 T_{i}(o)$. Similarly, if $i \in I$, then $\left|Q_{i}\right| \geq 0.9 \gamma T_{i}(o)$; otherwise $\left|Q_{i}\right| \leq 1.1 \gamma T_{i}(o)$.

Fact 27.3.4. For $Q_{i}$ output by Algorithm 27.1, every point $p \in Q_{i}$ is assigned the same sensitivity upper bound $s^{\prime}(p)=10 d^{3} / T_{i}(o)$.

In line 27.2 of Algorithm 27.2, we set $I$ to be the set of levels such that there are sufficient number of points in the crucial cells in those levels. The following lemma shows that the point set $Q^{I}$ (line 27.2 of Algorithm 27.2) is a good representative of the point set $Q$, i.e., for any set of $k$ centers $Z$, the $k$-means $\operatorname{cost} \operatorname{cost}(Q, Z)$ is close to the $\operatorname{cost}\left(Q^{I}, Z\right)$.

Lemma 27.3.5. Let $Q^{I}$ and $\epsilon$ be the same as in Algorithm 27.2. If $o \in(0, \mathrm{OPT}]$, then for
any $Z \subseteq \mathbb{R}^{d}$ with $|Z|=k$, we have:

$$
\operatorname{cost}\left(Q^{I}, Z\right) \leq \operatorname{cost}(Q, Z) \leq(1+\epsilon / 10) \operatorname{cost}\left(Q^{I}, Z\right)
$$

Next, instead of showing $s^{\prime}(p)$ (output by Algorithm 27.1) is a sensitivity upper bound with respect ot $Q$, we show that $s^{\prime}(p)$ is also an sensitivity upper bound with respect to $Q^{I}$. This is even stronger since $Q^{I}$ is a subset of $Q$ and we have:

$$
\forall Z \subseteq \mathbb{R}^{d}:|Z|=k, \frac{\operatorname{dist}^{2}(p, Z)}{\sum_{q \in Q} \operatorname{dist}^{2}(q, Z)} \leq \frac{\operatorname{dist}^{2}(p, Z)}{\sum_{q \in Q^{I}} \operatorname{dist}^{2}(q, Z)}
$$

Lemma 27.3.6. Let $Q^{I}$ be the same as in Algorithm 27.2. If $o \in(0, \mathrm{OPT}]$, then for all $i \in\{0,1,2, \ldots, L\}$ and $p \in Q_{i}$, we have:

$$
\max _{Z \subseteq \mathbb{R}^{d}:|Z|=k} \frac{\operatorname{dist}^{2}(p, Z)}{\sum_{q \in Q^{I}} \operatorname{dist}^{2}(q, Z)} \leq 10 \frac{d^{3}}{T_{i}(o)}=s^{\prime}(p)
$$

Now, we explain the reason of imposing randomly shifted grids. We fix an optimal set $Z^{*}=\left\{z_{1}^{*}, z_{2}^{*}, \ldots, z_{k}^{*}\right\} \subset \mathbb{R}^{d}$ of $k$ centers for the point set $Q$, i.e., $\operatorname{cost}\left(Q, Z^{*}\right)=$ OPT. We call a cell $C \in G_{i}$ a center cell if it is close to a center in $Z^{*}$, namely $\operatorname{dist}\left(C, Z^{*}\right) \leq g_{i} /(2 d)$. We claim that there will not be too many center cells since we randomly shift the grids. In other words, each center of $Z^{*}$ is far from the boundary of every gird.

Lemma 27.3.7. With probability at least 0.94, the total number of center cells is at most $100 k L$.

This lemma is similar to Lemma 2.2 in $\left[\mathrm{BFL}^{+} 17\right]$. For completeness, we also provide a proof in Section 27.5.

Consider the total estimated sensitivities, i.e., the sum of the sensitivity upper bounds over all the points. Due to Theorem 27.2.1, this sum determines the size of the coreset. We show that if the estimate $o$ of the optimal $k$-means cost is close to OPT, then the total estimated sensitivities can not be too large.

Lemma 27.3.8. Suppose the number of center cells is at most $100 k L$. Let $Q_{0}, Q_{1}, \ldots, Q_{L}, s^{\prime}(\cdot)$ be the output of Algorithm 27.1. Then the total estimated sensitivities satisfies

$$
\sum_{p \in Q} s^{\prime}(p) \leq 4000 k L d^{3} \cdot(\mathrm{OPT} / o+1)
$$

Since $Q^{I}$ is a subset of $Q$, according to the above lemma, we also have $\sum_{p \in Q^{I}} S^{\prime}(p) \leq$ $4000 d^{3} L k \cdot(\mathrm{OPT} / o+1)$. Now, we are ready to prove Theorem 27.3.1.

Proof of Theorem 27.3.1. By Lemma 27.3.7, with probability at least 0.94, the number of center cells is at most 100 kL . In the following, we condition on this event.

Algorithm 27.2 draws $m$ i.i.d. samples. For each sample, a point $p \in Q_{i} \subseteq Q^{I}$ is chosen with probability

$$
\frac{\widehat{q}_{i} / T_{i}(o)}{\sum_{j \in I} \widehat{q}_{j} / T_{j}(o)} \cdot \frac{1}{\left|Q_{i}\right|}=\frac{\frac{\widehat{q}_{i}}{\left|Q_{i}\right|} \cdot 20 \frac{d^{3}}{T_{i}(o)}}{\sum_{j \in I} \sum_{p^{\prime} \in Q_{j}} \frac{\widehat{q}_{j}}{\left|Q_{j}\right|} \cdot 20 \frac{d^{3}}{T_{j}(o)}}
$$

Each sample $p$ is given a weight

Let $s^{\prime \prime}(p)=\frac{\widehat{q_{i}}}{\left|Q_{i}\right|} \cdot 20 \frac{d^{3}}{T_{i}(o)}$. Since $\widehat{q_{i}} /\left|Q_{i}\right| \geq 1 / 2$, we know that $s^{\prime \prime}(p)$ is still a sensitivity upper bound of $p$ with respect to $Q^{I}$ by Lemma 27.3.6. According to Algorithm 27.2, we have $t^{\prime}=\frac{1}{2} \sum_{p \in Q^{I}} s^{\prime \prime}(p)$. Therefore, if we set $m$ to be a sufficiently large $\Omega\left(t^{\prime} \epsilon^{-2} L d \log t^{\prime}\right)=$ $\Omega\left(t^{\prime} \epsilon^{-2}\left(\log \left(\Delta^{d}\right) \log t^{\prime}+\log (1 / 0.01)\right)\right)$, then according to Theorem 27.2.1, $S$ output by Algorithm 27.2 is an $\epsilon / 2$-coreset for $Q^{I}$ with probability at least 0.99 . By Lemma 27.3.5, if $S$ is an $\epsilon / 2$-coreset for $Q^{I}$, then $S$ is also an $\epsilon$-coreset for $Q$. Thus, the correctness is proved, and the overall success probability is at least 0.93 obtained by a simple union bound. Now let us analyze the size of the coreset $S$. Since $\forall j \in I, \widehat{q}_{j} /\left|Q_{j}\right| \leq 2$, we have $\sum_{j \in I} \sum_{p^{\prime} \in Q_{j}} s^{\prime \prime}\left(p^{\prime}\right) \leq 2 t^{\prime} \leq 4 \cdot 4000 d^{3} L k \cdot(\mathrm{OPT} / o+1)$ by Lemma 27.3.8. Thus, the size of $S$ is $m=O\left(k \epsilon^{-2} d^{4} L^{2} \log (k d L) \cdot(\mathrm{OPT} / o+1)\right)$.

### 27.4 Conclusion

This paper gives the first $k$-means coreset construction in the dynamic streaming model using space polynomial in the dimension $d$ and nearly optimal (linear) in $k$. The algorithm is based on sensitivity sampling, which we believe is a powerful tool and can have broader applications.

### 27.5 Missing Details in Section 27.3

Proof of Fact 27.3.2. Consider an arbitrary point $p \in Q$. Let $C_{-1}, C_{0}, \ldots, C_{L}$ be the cells which contain $p$, where $\forall i \in\{-1,0,1, \ldots, L\}$, the cell $C_{i}$ is from the grid $G_{i}$. According to Algorithm 27.1, $C_{-1}$ is marked as heavy and $C_{L}$ cannot be marked as heavy. Let $l$ be the largest integer such that all the cells $C_{-1}, C_{0}, \ldots, C_{l-1}$ are marked as heavy. Then the cell $C_{l}$ must be marked as crucial, and all the cells $C_{l+1}, C_{l+2}, \ldots, C_{L}$ can not be crucial. Thus, we have $p \in Q_{l}$ and $\forall i \in\{0,1, \ldots, L\} \backslash\{l\}, p \notin Q_{i}$.

Facts 27.3.3 and 27.3.4 are obvious from the algorithms, so we omit the proofs.
The following claim is useful in the proofs.

Claim 27.5.1. Let $Q^{I}$ and $\gamma$ be the same as mentioned in Algorithm 27.2. For $i \in$ $\{0,1, \ldots, L\}$ and any heavy cell $C \in G_{i-1}$, if all the ancestors of $C$ are marked as heavy, then we have $\left|C \cap Q^{I}\right| \geq(1-5(L-i) \gamma) \cdot|C \cap Q|$.

Proof. The proof is by induction. When $i=L$, consider a heavy cell $C \in G_{L-1}$ whose ancestors are also heavy. If there is no such cell $C$, then the claim holds directly for $i=L$. Otherwise, according to the construction of $Q_{0}, Q_{1}, \ldots, Q_{L}, \forall j \in\{0,1, \ldots, L-1\}$, we have $Q_{j} \cap C=\emptyset$ and $(Q \cap C) \subseteq Q_{L}$. Since $C$ is marked as heavy, we know that $\left|Q_{L} \cap C\right|=|Q \cap C| \geq$ $0.9 T_{L-1}(o) \geq 0.9 T_{L}(o) / 4$. It implies that $\widehat{q}_{L} \geq \min \left(\left|Q_{L}\right|-0.1 \epsilon \gamma T_{L}(o),(1-0.01 \epsilon)\left|Q_{L}\right|\right) \geq$ $\min \left(|Q \cap C|-0.1 \epsilon \gamma T_{L}(o),(1-0.01 \epsilon)|Q \cap C|\right) \geq \gamma T_{L}(o)$. Thus, we have $Q_{L} \subseteq Q^{I}$ which implies that $\left|C \cap Q^{I}\right|=\left|C \cap Q_{L}\right|=|C \cap Q|$.

Now assume the claim is true for $i+1, i+2, \ldots, L$. Consider a heavy cell $C \in G_{i-1}$ whose ancestors are also marked as heavy. If there is no such cell $C$, the claim holds directly
for $i$. Now consider the case when $C$ exists. If level $i \in I$, i.e., $Q_{i} \subseteq Q^{I}$, we have

$$
\begin{aligned}
\left|C \cap Q^{I}\right| & =\sum_{C^{\prime} \in G_{i}: C^{\prime} \text { is a heavy child of } C}\left|C^{\prime} \cap Q^{I}\right|+\sum_{C^{\prime} \in G_{i}: C^{\prime} \text { is a crucial child of } C}\left|C^{\prime} \cap Q^{I}\right| \\
& =\sum_{C^{\prime} \in G_{i}: C^{\prime} \text { is a heavy child of } C}\left|C^{\prime} \cap Q^{I}\right|+\left|C \cap Q_{i}\right| \\
& \geq(1-5(L-i-1) \gamma) \sum_{C^{\prime} \in G_{i}: C^{\prime} \text { is a heavy child of } C}\left|C^{\prime} \cap Q\right|+\left|C \cap Q_{i}\right| \\
& \geq(1-5(L-i-1) \gamma)|C \cap Q| \\
& \geq(1-5(L-i) \gamma)|C \cap Q| .
\end{aligned}
$$

If level $i \notin I$, i.e., $Q_{i} \nsubseteq Q^{I}$, we have

$$
\begin{aligned}
\sum_{C^{\prime} \in G_{i}: C^{\prime} \text { is a heavy child of } C}\left|C^{\prime} \cap Q\right| & \geq|C \cap Q|-\left|Q_{i}\right| \\
& \geq|C \cap Q|-1.1 \gamma T_{i}(o) \\
& \geq(1-5 \gamma)|C \cap Q|
\end{aligned}
$$

Thus,

$$
\begin{aligned}
\left|C \cap Q^{I}\right| & \geq \sum_{C^{\prime} \in G_{i}: C^{\prime} \text { is a heavy child of } C}\left|C^{\prime} \cap Q^{I}\right| \\
& \geq(1-5(L-i-1) \gamma) \cdot \sum_{C^{\prime} \in G_{i}: C^{\prime} \text { is a heavy child of } C}\left|C^{\prime} \cap Q\right| \\
& \geq(1-5(L-i-1) \gamma)(1-5 \gamma)|C \cap Q| \\
& \geq(1-5(L-i) \gamma)|C \cap Q| .
\end{aligned}
$$

Proof of Lemma 27.3.5. Since $Q^{I}$ is a subset of $Q, \operatorname{cost}\left(Q^{I}, Z\right) \leq \operatorname{cost}(Q, Z)$ is trivial. In the following, we are trying to prove $\operatorname{cost}(Q, Z) \leq(1+\epsilon / 10) \operatorname{cost}\left(Q^{I}, Z\right)$.

We consider an level $i \notin I$, i.e., $Q_{i} \nsubseteq Q^{I}$. For a point $p \in Q_{i}$, all the ancestors of $c_{i}(p)$ must be heavy. By averaging argument, there must exist a point $q \in c_{i-1}(p) \cap Q^{I}$ such that

$$
\begin{align*}
\operatorname{dist}^{2}(p, Z) & \leq 2 \operatorname{dist}^{2}(p, q)+2 \operatorname{dist}^{2}(q, Z) \\
& \leq 2 d g_{i-1}^{2}+2 \operatorname{dist}^{2}(q, Z) \\
& \leq 2 d g_{i-1}^{2}+2 \frac{1}{\left|c_{i-1}(p) \cap Q^{I}\right|} \sum_{q \in c_{i-1}(p) \cap Q^{I}} \operatorname{dist}^{2}(q, Z) \tag{27.1}
\end{align*}
$$

where the first step follows from triangle inequality, the second step follows from definition of the grids, the last step follows from an averaging argument.

According to Claim 27.5.1, we have

$$
\left|c_{i-1}(p) \cap Q^{I}\right| \geq \frac{1}{2} T_{i-1}(o)
$$

Let $Q^{N}=Q \backslash Q^{I}$. We can lower bound $\operatorname{cost}(Q, Z)$ in the following sense,

$$
\begin{aligned}
& \operatorname{cost}(Q, Z)=\operatorname{cost}\left(Q^{I}, Z\right)+\operatorname{cost}\left(Q^{N}, Z\right) \\
&=\operatorname{cost}\left(Q^{I}, Z\right)+\sum_{i \notin I} \sum_{p \in Q_{i}} \operatorname{dist}^{2}(p, z) \\
& \leq \operatorname{cost}\left(Q^{I}, Z\right)+2 \sum_{i \notin I} \sum_{p \in Q_{i}}\left(d g_{i-1}^{2}+\frac{1}{\left|c_{i-1}(p) \cap Q^{I}\right|} \sum_{q \in c_{i-1}(p) \cap Q^{I}} \operatorname{dist}^{2}(q, Z)\right) \\
& \leq \operatorname{cost}\left(Q^{I}, Z\right)+2 \sum_{i \notin I} \sum_{p \in Q_{i}}\left(d g_{i-1}^{2}+\frac{2}{T_{i-1}(o)} \sum_{q \in c_{i-1}(p) \cap Q^{I}} \operatorname{dist}^{2}(q, Z)\right) \\
& \leq \operatorname{cost}\left(Q^{I}, Z\right)+2 \sum_{i \notin I} \sum_{p \in Q_{i}}\left(d g_{i-1}^{2}+\frac{2}{T_{i-1}(o)} \sum_{q \in Q^{I}} \operatorname{dist}^{2}(q, Z)\right) \\
& \leq \operatorname{cost}\left(Q^{I}, Z\right)+2 \sum_{i \notin I} 1.1 \gamma T_{i}(o) \cdot\left(d g_{i-1}^{2}+\frac{2}{T_{i-1}(o)} \sum_{q \in Q^{I}} \operatorname{dist}^{2}(q, Z)\right) \\
& \leq \operatorname{cost}\left(Q^{I}, Z\right)+5 L \gamma T_{i}(o) \cdot\left(d g_{i-1}^{2}+\frac{2}{T_{i-1}(o)} \sum_{q \in Q^{I}} \operatorname{dist}^{2}(q, Z)\right) \\
&=\operatorname{cost}\left(Q^{I}, Z\right)+5 L \gamma T_{i}(o) \cdot\left(d g_{i-1}^{2}+\frac{2}{T_{i-1}(o)} \operatorname{cost}\left(Q^{I}, Z\right)\right) \\
& \leq \operatorname{cost}\left(Q^{I}, Z\right)+5 L \gamma\left(d^{3} o /(25 k)+8 \operatorname{cost}\left(Q^{I}, Z\right)\right) \\
& \leq \operatorname{cost}\left(Q^{I}, Z\right)+5 L \gamma\left(d^{3} \operatorname{cost}(Q, Z) /(25 k)+8 \operatorname{cost}\left(Q^{I}, Z\right)\right) .
\end{aligned}
$$

where the second step follows from the definition of the cost, the third step follows from Eq. (27.1), the fourth step follows from $\left|c_{i-1}(p) \cap Q^{I}\right| \geq T_{i-1}(o) / 2$, the fifth step follows from $\left(c_{i-1}(p) \cap Q^{I}\right) \subset Q^{I}$, the sixth step follows from $\left|Q_{i}\right| \leq 1.1 \gamma T_{i}(o)$, the seventh step follows from $L+1-|I| \leq L+1 \leq 2 L$ and $2 \cdot 2 \cdot 1.1 \leq 5$, the ninth step follows from $T_{i}(o)=4 T_{i-1}(o)$ and $T_{i}(o)=\left(d / g_{i}\right)^{2} \cdot o / k \cdot 1 / 100$, and the last step follows from $o \leq \mathrm{OPT} \leq \operatorname{cost}(Q, Z)$.

It implies that

$$
\begin{aligned}
\frac{\operatorname{cost}(Q, Z)}{\operatorname{cost}\left(Q^{I}, Z\right)} & \leq \frac{1+40 L \gamma}{1-5 L \gamma d^{3} /(25 k)} \\
& \leq \frac{1+\epsilon / 40}{1-\epsilon / 40} \\
& \leq 1+\epsilon / 10
\end{aligned}
$$

where the second step follows from $\gamma \leq \epsilon /\left(40^{2} L d^{3}\right)$, and the last step follows from $\epsilon<1$.

Proof of Lemma 27.3.6. Let $Z \subseteq \mathbb{R}^{d}$ be an arbitrary set of $k$ centers. Fix a point $p \in Q$. Suppose $p$ is in $Q_{i}$, i.e., $p$ is in a crucial cell of $G_{i}$. Let $C=c_{i-1}(p)$, i.e., $C$ is the parent cell of the crucial cell that contains $p$. By Algorithm 27.1, $C$ and all of its ancestors must be heavy. By Claim 27.5.1, $C \cap Q^{I}$ cannot be empty. Thus, by an averaging argument, there is a point $p^{\prime} \in C$ such that

$$
\begin{equation*}
\operatorname{dist}^{2}\left(p^{\prime}, Z\right) \leq \frac{1}{\left|C \cap Q^{I}\right|} \sum_{q \in C \cap Q^{I}} \operatorname{dist}^{2}(q, Z) \tag{27.2}
\end{equation*}
$$

We have

$$
\begin{aligned}
\frac{\operatorname{dist}^{2}(p, Z)}{\sum_{q \in Q^{I}} \operatorname{dist}^{2}(q, Z)} & \leq 2 \frac{\operatorname{dist}^{2}\left(p^{\prime}, Z\right)}{\sum_{q \in Q^{I}} \operatorname{dist}^{2}(q, Z)}+2 \frac{\operatorname{dist}^{2}\left(p, p^{\prime}\right)}{\sum_{q \in Q^{I}} \operatorname{dist}^{2}(q, Z)} \\
& \leq 2 \frac{\sum_{q \in C \cap Q^{I}} \operatorname{dist}^{2}(q, Z)}{\left|C \cap Q^{I}\right| \sum_{q \in Q^{I}} \operatorname{dist}^{2}(q, Z)}+2 \frac{d g_{i-1}^{2}}{\sum_{q \in Q^{I}} \operatorname{dist}^{2}(q, Z)} \\
& \leq 2 \frac{\sum_{q \in Q^{I}} \operatorname{dist}^{2}(q, Z)}{\left|C \cap Q^{I}\right| \sum_{q \in Q^{I}} \operatorname{dist}^{2}(q, Z)}+2 \frac{d g_{i-1}^{2}}{\sum_{q \in Q^{I}} \operatorname{dist}^{2}(q, Z)} \\
& =2 \frac{1}{\left|C \cap Q^{I}\right|}+2 \frac{d g_{i-1}^{2}}{\sum_{q \in Q^{I}} \operatorname{dist}^{2}(q, Z)} \\
& \leq 2 \frac{1}{\left|C \cap Q^{I}\right|}+4 \frac{d g_{i-1}^{2}}{\mathrm{OPT}} \\
& \leq 2 \frac{1}{\left|C \cap Q^{I}\right|}+16 \frac{d g_{i}^{2}}{\mathrm{OPT}} \\
& \leq 2 \frac{1}{\left|C \cap Q^{I}\right|}+\frac{d^{3} o}{T_{i}(o) k \mathrm{OPT}} \\
& \leq 9 \frac{1}{T_{i}(o)}+\frac{d^{3} o}{T_{i}(o) k \mathrm{OPT}} \\
& \leq 9 \frac{1}{T_{i}(o)}+\frac{d^{3}}{T_{i}(o)} \\
& \leq 10 \frac{d^{3}}{T_{i}(o)}
\end{aligned}
$$

where the first step follows from triangle inequality, the second step follows from Eq. (27.2) and $p^{\prime} \in c_{i-1}(p)$, the fifth step follows from $\sum_{q \in Q^{I}} \operatorname{dist}^{2}(q, Z) \geq(1-\epsilon) \operatorname{cost}(Q, Z) \geq$ OPT $/ 2$ (Lemma 27.3.5), the sixth step follows from $g_{i-1}^{2} \leq 4 g_{i}^{2}$, the seventh step follows from $T_{i}(o)=$ $\frac{d^{2}}{g_{i}^{2}} \frac{o}{100 k}$, the eighth step follows from $T_{i}(o)=4 T_{i-1}(o) \leq 4.5\left|C \cap Q^{I}\right|$ (Claim 27.5.1 and $\left.|C \cap Q| \geq 0.9 T_{i-1}(o)\right)$, the ninth step follows from $1 / k \leq 1, o \leq \mathrm{OPT}$.

Proof of Lemma 27.3.7. Fix an $i \in\{0,1, \ldots, L\}$ and consider the grid $G_{i}$. For each optimal center $z_{j}^{*}$, we use $X_{j, \alpha}$ to denote the indicator random variable for the event that the distance
from $z_{j}^{*}$ to the boundary in dimension $\alpha$ of the grid $G_{i}$ is at most $g_{i} /(2 d)$. Since in each dimension, if the center is close to a boundary, it contributes a factor at most 2 to the total number of center cells. It follows that the number of cells that have distance at most $g_{i} /(2 d)$ to $z_{j}^{*}$ is at most

$$
N=2^{\sum_{\alpha=1}^{d} X_{j, \alpha}} .
$$

We denote $Y_{j, \alpha}$ to be $2^{X_{j, \alpha}}$, then

$$
\mathbb{E}[N]=\mathbb{E}\left[\prod_{\alpha=1}^{d} Y_{j, \alpha}\right]=\prod_{\alpha=1}^{d} \mathbb{E}\left[Y_{j, \alpha}\right] .
$$

By using $\operatorname{Pr}\left[X_{j, \alpha}=1\right] \leq\left(2 g_{i} /(2 d)\right) / g_{i}=1 / d$, we obtain

$$
\mathbb{E}\left[Y_{j, \alpha}\right] \leq \mathbb{E}\left[1+X_{j, \alpha}\right]=1+\mathbb{E}\left[X_{j, \alpha}\right] \leq 1+1 / d
$$

Thus $\mathbb{E}[N]=\prod_{\alpha=1}^{d} \mathbb{E}\left[Y_{j, \alpha}\right] \leq(1+1 / d)^{d} \leq e$. The expected number of center cells in a single grid is at most $(1+1 / d)^{d} k \leq e k \leq 3 k$. By linearity of expectation, the expected number of center cells in all grids is at most $e k(L+1) \leq 6 k L$. By Markov's inequality, the probability that we have more than $100 k L$ center cells in all grids is at most 0.06 .

Proof of Lemma 27.3.8. We have

$$
\begin{aligned}
\sum_{p \in Q} s^{\prime}(p) & =\sum_{i=0}^{L} \sum_{p \in Q_{i}} 10 \frac{d^{3}}{T_{i}(o)} \\
& =10 d^{3} \sum_{i=0}^{L} \frac{1}{T_{i}(o)}\left(\sum_{\text {center cell }} \mid C \in G_{i}\right. \\
& \leq 10 d^{3} \sum_{i=0}^{L} \frac{1}{T_{i}(o)}\left(\sum_{\text {center cell }} 1.1 T_{i}(o)+\frac{\mathrm{OPT}}{g_{i}^{2} /(2 d)^{2}}\right) \\
& \leq 11 d^{3} \cdot(\# \text { of center cells })+4000 d^{3} L k \cdot \frac{\mathrm{OPT}}{o} \\
& \leq 1100 d^{3} L k+4000 d^{3} L k \cdot \frac{\mathrm{OPT}}{o} \\
& \leq 4000 d^{3} L k \cdot\left(\frac{\mathrm{OPT}}{o}+1\right)
\end{aligned}
$$

where the first step follows from the definition of $s^{\prime}(p)$, the third step follows from

1. if $C \in G_{i}$ and $C \cap Q_{i} \neq \emptyset$, then $\left|C \cap Q_{i}\right|=|C \cap Q| \leq 1.1 T_{i}(o)$;
2. if $p$ is in a non-center cell $C \in G_{i}$, then $\operatorname{dist}^{2}\left(p, Z^{*}\right) \geq g_{i}^{2} /(2 d)^{2}$,
the forth step follows from $g_{i}^{2} / d^{2}=o /\left(100 k T_{i}(o)\right)$, the fifth step follows from that the total number of center cells is bounded by $100 k L$.

### 27.6 Coreset Construction over a Dynamic Stream

In this section, we show how to implement Algorithms 27.1 and 27.2 in the dynamic streaming setting.

First, we introduce a dynamic storage structure that allows us to insert and delete points or cells. We then use this data structure combined with hash functions to estimate the number of points falling into each cell. Lastly, we combine them with the sensitivity sampling procedure to obtain our final algorithm.

```
Algorithm 27.3 Point-cell storing procedure
    \(\operatorname{Storing}\left(G_{i}, \alpha, \beta, \delta\right)\) :
    Input: \(\left\{\left(\left(p_{1}, l_{1}\right), \pm\right),\left(\left(p_{2}, l_{2}\right), \pm\right), \ldots\right\} \triangleright l_{t} \in[\widehat{m}]\). Only Algorithm 27.5 uses the case for
    \(\widehat{m}>1\).
    Run Distinct \((\alpha, \delta / 4)\) on \(\left\{\left(c_{i}\left(p_{1}\right), \pm\right),\left(c_{i}\left(p_{2}\right), \pm\right), \ldots\right\}\) in parallel.
    Set \(r=\lceil\log (4 \alpha / \delta)\rceil\) and \(h_{1}, h_{2}, \ldots, h_{r}, \forall j \in[r], h_{j}: G_{i} \rightarrow[2 \alpha] . \quad \triangleright h_{j}\) is pairwise
    independent.
    Run \(r \cdot 2 \alpha\) copies of \(\operatorname{Distinct}(\beta, \delta /(2 \alpha))\) in parallel.
            \(\triangleright\) Each copy is indexed by a pair \((j, b) \in[r] \times[2 \alpha]\). The \((j, b)\)-th copy is run on the
    sub-stream \(\left\{\left(\left(p_{1}^{\prime}, l_{1}^{\prime}\right), \pm\right),\left(\left(p_{2}^{\prime}, l_{2}^{\prime}\right), \pm\right), \ldots\right\}\), where each \(\left(\left(p_{t}^{\prime}, l_{t}^{\prime}\right), \pm\right)\) satisfies \(h_{j}\left(c_{i}\left(p_{t}^{\prime}\right)\right)=b\).
    If line 27.3 returns FAIL, output FAIL; otherwise, let \(\mathcal{C}, f: \mathcal{C} \rightarrow \mathbb{N}_{+}\)be the output of
    line 27.3.
            \(\triangleright \mathcal{C} \subset G_{i}\) contains all the cells found, and \(f(C)\) denotes the number of points in \(C\).
    Initialize \(S \leftarrow \emptyset\).
    for \(C \in \mathcal{C}\) with \(f(C) \leq \beta\) do
            Find \(j \in[r]\) s.t. \(\forall C^{\prime} \in \mathcal{C}, h_{j}(C) \neq h_{j}\left(C^{\prime}\right)\) and the \(\left(j, h_{j}(C)\right)\)-th copy in line 27.3 does
    not FAIL.
            If such \(j\) does not exist, output FAIL; otherwise \(S \leftarrow S \cup S_{j, h_{j}(C)}\).
        \(\triangleright\) Here \(S_{j, h_{j}(C)}\) is the set of distinct points found by the \(\left(j, h_{j}(C)\right)\)-th copy in line 27.3.
    end for
    Output: \(\mathcal{C} \subset G_{i}, f: \mathcal{C} \rightarrow \mathbb{N}_{+}, S=\left\{\left(\widetilde{p}_{1}, \widetilde{l}_{1}\right),\left(\widetilde{p}_{2}, \widetilde{l}_{2}\right), \ldots\right\}\)
```


### 27.6.1 The Dynamic Point-Cell Storing Data Structure

We introduce an algorithm that maintains a set of points and cells in a dynamic data stream. Before that, let us recall [Gan05]'s result for finding distinct elements, which we use as a subroutine in our algorithm.

Lemma 27.6.1 (Distinct elements [Gan05]). Given parameters $M \geq 1, N \geq 1, s \geq 1, \delta \in$ $(0,1 / 2)$, there is an algorithm Distinct $(s, \delta)$ that requires $O(s(\log M+\log N) \log (s / \delta))$ bits to process a stream of insertion/deletion of data items. For each operation $(i, \pm)(i \in[N])$, the algorithm takes $O(\log (s / \delta))$ time. $M$ is an upper bound of the total frequency of all items during the stream. At the end of the stream, if the number of distinct elements is at most $s$, with probability at least $1-\delta$ it returns all the distinct elements and their frequencies. It returns FAIL otherwise.

We use Distinct as our sub-routine. We set the parameter $M$ and $N$ to be sufficiently large in our case, i.e., $M=N=\Delta^{2 d}$. In Algorithm 27.3, we describe a method which can with probability at least $1-\delta$ output all the non-empty cells in grid $G_{i}$ when the total number of non-empty cells is not too large (at most $\alpha$ ). Furthermore, if the number of points in a particular cell is not too large (at most $\beta$ ), the algorithm can output all the points in that cell. Notice that Algorithm 27.3 is only a subroutine of our final algorithm and will only work on some sub-stream of the entire data stream.

The following lemma shows the guarantee of Algorithm 27.3.
Lemma 27.6.2. Given parameters $i \in[L], \alpha, \beta \in \mathbb{N}_{+}, \delta \in(0,0.5), \operatorname{Storing}\left(G_{i}, \alpha, \beta, \delta\right)(A l-$ gorithm 27.3) uses $O\left(\alpha \beta d L \cdot \log ^{2}(\alpha \beta / \delta)\right)$ bits to process a stream $\left\{\left(\left(p_{1}, l_{1}\right), \pm\right),\left(\left(p_{2}, l_{2}\right), \pm\right), \ldots\right\}$ of insertion/deletion operations of data points. At the end of the stream, if the number of
non-empty cells in $G_{i}$ is at most $\alpha$, then with probability at least $1-\delta$ it returns the set $\mathcal{C}$ of all the non-empty cells, the number of points $f(C)$ in each cell $C \in \mathcal{C}$, and the set $S$ of points in all the non-empty cells that contain at most $\beta$ points. It returns FAIL otherwise.

Proof. If the number of non-empty cells of $G_{i}$ at the end of the stream is more than $\alpha$, then according to line 27.3 of Algorithm 27.3 and Lemma 27.6.1, Algorithm 27.3 must output FAIL.

Now consider the case when the total number of non-empty cells is at most $\alpha$. According to Lemma 27.6.1, with probability at least $1-\delta / 4$, line 27.3 of Algorithm 27.3 will return the set $\mathcal{C}$ of all the non-empty cells of $G_{i}$ and the number of points $f(C)$ for each cell $C \in \mathcal{C}$. For each $C \in \mathcal{C}$ with $|C| \leq \beta$, since $|\mathcal{C}| \leq \alpha$, the probability that $\exists j \in[r]$ such that $\forall C^{\prime} \in \mathcal{C}, h_{j}\left(C^{\prime}\right) \neq h_{j}(C)$ is at least $1-1 / 2^{r} \geq 1-\delta /(4 \alpha)$ and furthermore the probability that the $\left(j, h_{j}(C)\right)$-th copy of $\operatorname{Distinct}(\beta, \delta /(2 \alpha))$ in line 27.3 of Algorithm 27.3 will output all the points in $C$ is at least $1-\delta /(2 \alpha)$. By taking union bound, the overall probability that Algorithm 27.3 does not output FAIL is at most $1-\delta / 4-(\delta /(2 \alpha)+\delta /(4 \alpha)) \cdot \alpha=1-\delta$.

According to Lemma 27.6.1, the space needed by line 27.3 of Algorithm 27.3 is $O(\alpha d L$. $\log (\alpha / \delta))$ bits and the space needed of each copy in line 27.3 of Algorithm 27.3 is $O(\beta d L$. $\log (\alpha \beta / \delta))$ bits. Thus, the total space needed is at most $O\left(\alpha \beta d L \cdot \log ^{2}(\alpha \beta / \delta)\right)$.

### 27.6.2 Estimating the Number of Points in Each Cell

We use Algorithm 27.3 as a subroutine and design a dynamic streaming algorithm (Algorithm 27.4) that can estimate the number of points in each cell up to some precision. Furthermore, it also estimates the number of points $\left|Q_{i}\right|$ in crucial cells of each level $i$.

```
Algorithm 27.4 Estimating the number of points in each cell and in each level
    PointsEstimation \((o, \epsilon, \delta)\) :
    Input: a point set \(Q \subseteq[\Delta]^{d}\) described by a stream \(\left\{\left(p_{1}, \pm\right),\left(p_{2}, \pm\right), \ldots\right\}\)
    for \(i \in\{0,1, \ldots, L\}\), in parallel, do
        \(T_{i}(o) \leftarrow\left(d / g_{i}\right)^{2} \cdot o /(100 k), \alpha \leftarrow 10^{11} k L d \log (1 / \delta), \alpha^{\prime} \leftarrow 10^{16} \epsilon^{-3} k L^{2} d^{4}\).
        \(\triangleright T_{i}(o)\) : threshold for heaviness (Algortihm 27.1); \(\quad \alpha, \alpha^{\prime}\) : parameters for Storing.
        Let \(h_{i}:[\Delta]^{d} \rightarrow\{0,1\}\) be a \(\lambda\)-wise independent hash function.
                \(\triangleright \lambda=10\lceil(d L+\log (1 / \delta)+1)\rceil ; \quad \forall p \in[\Delta]^{d}, h_{i}(p)=1\) w.p. \(\min \left(4 \cdot 10^{4} \lambda / T_{i}(o), 1\right)\).
            Run \(\operatorname{Storing}\left(G_{i}, \alpha, 1,0.1 \delta / L\right)\) on a sub-stream \(\left\{\left(\left(p_{1}^{\prime}, 1\right), \pm\right),\left(\left(p_{2}^{\prime}, 1\right), \pm\right), \ldots\right\}\) in par-
    allel.
                            \(\triangleright\) Storing is defined in Algorithm 27.3. Here \(p_{j}^{\prime}\) satisfies \(h_{i}\left(p_{j}^{\prime}\right)=1\).
    If Storing returns FAIL, output FAIL.
            Let \(\gamma \leftarrow \epsilon /\left(40^{2} L d^{3}\right)\). \(\triangleright\) Threshold for discarding levels
            Let \(h_{i}^{\prime}:[\Delta]^{d} \rightarrow\{0,1\}\) be a \(\lambda\)-wise independent hash function.
                                    \(\triangleright \forall p \in[\Delta]^{d}, h_{i}^{\prime}(p)=1\) w.p. \(\min \left(4 \cdot 10^{4} \epsilon^{-2} \gamma^{-1} \lambda / T_{i}(o), 1\right)\).
            Run \(\operatorname{Storing}\left(G_{i}, \alpha^{\prime}, 1,0.1 \delta / L\right)\) on sub-stream \(\left\{\left(\left(p_{1}^{\prime \prime}, 1\right), \pm\right),\left(\left(p_{2}^{\prime \prime}, 1\right), \pm\right), \ldots\right\}\) in par-
    allel.
                        \(\triangleright\) Storing is defined in Algorithm 27.3. Here \(p_{j}^{\prime \prime}\) satisfies \(h_{i}^{\prime}\left(p_{j}^{\prime \prime}\right)=1\).
    If Storing returns FAIL, output FAIL.
            Let \(\mathcal{C}_{i}, f_{i}, S_{i} \leftarrow \operatorname{Storing}\left(G_{i}, \alpha, 1,0.1 \delta / L\right)\) in line 27.4. Let \(\widehat{f}(C)=f_{i}(C)\).
    \(\min \left\{T_{i}(o) /\left(4 \cdot 10^{4} \lambda\right), 1\right\}\).
    :
    \(|C \cap Q|\) and follow Algorithm 27.1 to determine whether \(C\) is marked as heavy, crucial
    or nothing. (And conceptually compute \(Q_{0}, \ldots, Q_{L}\) for analysis.)
            Let \(\mathcal{C}_{i}^{\prime}, f_{i}^{\prime}, S_{i}^{\prime} \leftarrow \operatorname{Storing}\left(G_{i}, \alpha^{\prime}, 1,0.1 \delta / L\right)\) in line 27.4.
            Let \(\widehat{q}_{i}=\min \left\{\epsilon^{2} \gamma T_{i}(o) /\left(4 \cdot 10^{4} \lambda\right) \cdot \sum_{C \in \mathcal{C}_{i}^{\prime}: C \text { is crucial }} f_{i}^{\prime}(C), 1\right\}\).
    end for
    Output: \(\widehat{q}_{0}, \widehat{q}_{1}, \ldots, \widehat{q}_{L}\) and \(\widehat{f}: \bigcup_{i=0}^{L} G_{i} \rightarrow \mathbb{R}_{+}\)
```

Now let us analyze Algorithm 27.4. We need the following high concentration bound in our analysis.

Theorem 27.6.3 ([BR94]). Let $\lambda$ be an even integer, and let $X$ be the sum of $n \lambda$-wise
independent random variables taking values in $[0,1]$. Let $\mu=\mathbb{E}[X]$ and $a>0$. Then we have

$$
\operatorname{Pr}[|X-\mu|>a] \leq 8 \cdot\left(\frac{\lambda \mu+\lambda^{2}}{a^{2}}\right)^{\lambda / 2}
$$

Lemma 27.6.4 (Samples from each cell). In PointsEstimation $(o, \epsilon, \delta)$ (Algorithm 27.4), with probability at least $1-\delta / 10, \forall i \in\{0,1, \ldots, L\}$ with $T_{i}(o) \geq 4 \cdot 10^{4} \lambda, \forall C \in G_{i}$, we have either

$$
\sum_{p \in C \cap Q} h_{i}(p) \in|C \cap Q| \cdot \frac{4 \cdot 10^{4} \lambda}{T_{i}(o)} \pm 4 \cdot 10^{3} \lambda
$$

or

$$
\sum_{p \in C \cap Q} h_{i}(p) \in|C \cap Q| \cdot \frac{4 \cdot 10^{4} \lambda}{T_{i}(o)} \cdot(1 \pm 0.01)
$$

Proof. Consider a cell $C \in G_{i}$. If $|C \cap Q| \leq T_{i}(o)$, then $\mu=\mathbb{E}\left[\sum_{p \in C \cap Q} h_{i}(p)\right]=4 \cdot 10^{4} \lambda \mid C \cap$ $Q \mid / T_{i}(o) \leq 4 \cdot 10^{4} \lambda$. According to Theorem 27.6.3, we have

$$
\operatorname{Pr}\left[\left|\sum_{p \in C \cap Q} h_{i}(p)-\mu\right|>4 \cdot 10^{3} \lambda\right] \leq 8 \cdot\left(\frac{4 \cdot 10^{4} \lambda^{2}+\lambda^{2}}{\left(4 \cdot 10^{3} \lambda\right)^{2}}\right)^{\lambda / 2} \leq 8 \cdot(1 / 2)^{\lambda / 2} \leq\left(\delta / \Delta^{d}\right)^{5}
$$

If $|C \cap Q|>T_{i}(o)$, then $\mu=\mathbb{E}\left[\sum_{p \in C \cap Q} h_{i}(p)\right]=4 \cdot 10^{4} \lambda|C \cap Q| / T_{i}(o)>4 \cdot 10^{4} \lambda$. According to Theorem 27.6.3, we have

$$
\operatorname{Pr}\left[\left|\sum_{p \in C \cap Q} h_{i}(p)-\mu\right|>0.01 \mu\right] \leq 8 \cdot\left(\frac{\lambda \mu+\lambda^{2}}{(0.01 \mu)^{2}}\right)^{\lambda / 2} \leq 8 \cdot(1 / 2)^{\lambda / 2} \leq\left(\delta / \Delta^{d}\right)^{5} .
$$

By taking union bound over all $i \in\{0,1, \ldots, L\}$ and all the cells in $G_{i}$, the claim is proved.

Lemma 27.6.5 (Estimating the number of points in each cell). If PointsEstimation $(o, \epsilon, \delta)$ (Algorithm 27.4) does not output FAIL, then with probability at least $1-\delta / 10$, $\forall i \in$ $\{0,1, \ldots, L\}, C \in G_{i}$, either $\widehat{f}(C) \in|C \cap Q| \pm 0.1 T_{i}(o)$ or $\widehat{f}(C) \in(1 \pm 0.01)|C \cap Q|$.

Proof. Suppose PointsEstimation $(o, \epsilon, \delta)$ does not output FAIL. According to Lemma 27.6.2, $\forall i \in\{0,1, \ldots, L\}$ and $\forall C \in G_{i}, f_{i}(C)=\sum_{p \in C \cap Q} h_{i}(p)$. If $T_{i}(o) \leq 4 \cdot 10^{4} \lambda$, then $\widehat{f}(C)=$ $f_{i}(C)=\sum_{p \in C \cap Q} h_{i}(p)=|C \cap Q|$.

Due to Lemma 27.6.4, with probability at least $1-\delta / 10, \forall i \in\{0,1, \ldots, L\}$ with $T_{i}(o)>4 \cdot 10^{4} \lambda$ and $\forall C \in G_{i}$, either $f_{i}(C) \in|C \cap Q| \cdot \frac{4 \cdot 10^{4} \lambda}{T_{i}(o)} \pm 4 \cdot 10^{3} \lambda$ or $f_{i}(C) \in|C \cap Q|$. $\frac{4 \cdot 10^{4} \lambda}{T_{i}(o)} \cdot(1 \pm 0.01)$. Thus, either $\widehat{f}(C) \in|C \cap Q| \pm 0.1 T_{i}(o)$ or $\widehat{f}(C) \in(1 \pm 0.01)|C \cap Q|$.

Lemma 27.6.6 (Samples from each $Q_{i}$ ). In PointsEstimation $(o, \epsilon, \delta)$ (Algorithm 27.4), with probability at least $1-\delta / 10, \forall i \in\{0,1, \ldots, L\}$ with $T_{i}(o) \geq 4 \cdot 10^{4} \epsilon^{-2} \gamma^{-1} \lambda, \forall C \in G_{i}$, either

$$
\sum_{p \in Q_{i}} h_{i}^{\prime}(p) \in\left|Q_{i}\right| \cdot \frac{4 \cdot 10^{4} \epsilon^{-2} \gamma^{-1} \lambda}{T_{i}(o)} \pm 4 \cdot 10^{3} \epsilon^{-1} \lambda
$$

or

$$
\sum_{p \in Q_{i}} h_{i}^{\prime}(p) \in\left|Q_{i}\right| \cdot \frac{4 \cdot 10^{4} \epsilon^{-2} \gamma^{-1} \lambda}{T_{i}(o)} \cdot(1 \pm 0.01 \epsilon)
$$

Proof. If $\left|Q_{i}\right| \leq \gamma T_{i}(o)$, then $\mu=\mathbb{E}\left[\sum_{p \in Q_{i}} h_{i}^{\prime}(p)\right]=4 \cdot 10^{4} \epsilon^{-2} \gamma^{-1} \lambda\left|Q_{i}\right| / T_{i}(o) \leq 4 \cdot 10^{4} \epsilon^{-2} \lambda$. According to Theorem 27.6.3, we have

$$
\operatorname{Pr}\left[\left|\sum_{p \in Q_{i}} h_{i}^{\prime}(p)-\mu\right|>4 \cdot 10^{3} \epsilon^{-1} \lambda\right] \leq 8 \cdot\left(\frac{4 \cdot 10^{4} \epsilon^{-2} \lambda^{2}+\lambda^{2}}{\left(4 \cdot 10^{3} \epsilon^{-1} \lambda\right)^{2}}\right)^{\lambda / 2} \leq 8 \cdot(1 / 2)^{\lambda / 2} \leq\left(\delta / \Delta^{d}\right)^{5}
$$

If $\left|Q_{i}\right|>\gamma T_{i}(o)$, then $\mu=\mathbb{E}\left[\sum_{p \in Q_{i}} h_{i}^{\prime}(p)\right]=4 \cdot 10^{4} \epsilon^{-2} \gamma^{-1} \lambda\left|Q_{i}\right| / T_{i}(o)>4 \cdot 10^{4} \epsilon^{-2} \lambda$. According to Theorem 27.6.3, we have

$$
\operatorname{Pr}\left[\left|\sum_{p \in Q_{i}} h_{i}^{\prime}(p)-\mu\right|>0.01 \epsilon \mu\right] \leq 8 \cdot\left(\frac{\lambda \mu+\lambda^{2}}{(0.01 \epsilon \mu)^{2}}\right)^{\lambda / 2} \leq 8 \cdot(1 / 2)^{\lambda / 2} \leq\left(\delta / \Delta^{d}\right)^{5} .
$$

By taking union bound over all $i \in\{0,1, \ldots, L\}$, the claim is proved.

Lemma 27.6.7 ( $\widehat{q}_{i}$ can estimate $\left|Q_{i}\right|$ well). If $\operatorname{PointsEsitmation~}(o, \epsilon, \delta)$ (Algorithm 27.4) does not output $\boldsymbol{F A I L}$, then with probability at least $1-\delta / 10, \forall i \in\{0,1, \ldots, L\}$, either $\widehat{q_{i}} \in\left|Q_{i}\right| \pm 0.1 \epsilon \gamma T_{i}(o)$ or $\widehat{q}_{i} \in(1 \pm 0.01 \epsilon)\left|Q_{i}\right|$.

Proof. Suppose PointsEstimation $(o, \epsilon, \delta)$ does not output FAIL. According to Lemma 27.6.2,
$\forall i \in\{0,1, \ldots, L\}, f_{i}^{\prime}(C)=\sum_{p \in Q_{i}} h_{i}^{\prime}(p)$. If $T_{i}(o) \leq 4 \cdot 10^{4} \epsilon^{-2} \gamma^{-1} \lambda$, then $\widehat{q}_{i}=\sum_{C \in G_{i}: C \text { is crucial }} f_{i}^{\prime}(C)=$ $\sum_{p \in Q_{i}} h_{i}^{\prime}(p)=\left|Q_{i}\right|$.

Due to Lemma 27.6.6, with probability at least $1-\delta / 10, \forall i \in\{0,1, \ldots, L\}$ with $T_{i}(o)>4 \cdot 10^{4} \epsilon^{-2} \gamma^{-1} \lambda$, either

$$
\sum_{C \in G_{i}: C \text { is crucial }} f_{i}^{\prime}(C) \in\left|Q_{i}\right| \cdot \frac{4 \cdot 10^{4} \epsilon^{-2} \gamma^{-1} \lambda}{T_{i}(o)} \pm 4 \cdot 10^{3} \epsilon^{-1} \lambda
$$

or

$$
\sum_{C \in G_{i}: C \text { is crucial }} f_{i}^{\prime}(C) \in\left|Q_{i}\right| \cdot \frac{4 \cdot 10^{4} \epsilon^{-2} \gamma^{-1} \lambda}{T_{i}(o)} \cdot(1 \pm 0.01 \epsilon) .
$$

Thus, either $\widehat{q}_{i} \in\left|Q_{i}\right| \pm 0.1 \epsilon \gamma T_{i}(o)$ or $\widehat{q}_{i} \in(1 \pm 0.01 \epsilon)\left|Q_{i}\right|$.

Lemma 27.6.8 (Number of points sampled from non-center cells). In PointsEsitmation $(o, \epsilon, \delta)$ (Algorithm 27.4), with probability at least $1-\delta / 10, \forall i \in\{0,1, \ldots, L\}$, we have

$$
\sum_{\text {non-center cell }} \sum_{C \in G_{i}} h_{p \in C \cap Q}(p) \leq 4 \cdot 10^{3} \lambda+1.01 \cdot \frac{4 \cdot 10^{4} \lambda}{T_{i}(o)} \sum_{\text {non-center cell } C \in G_{i}}|C \cap Q|
$$

and

$$
\sum_{\text {non-center cell }} \sum_{C \in G_{i}} h_{p \in C \cap Q}^{\prime}(p) \leq 4 \cdot 10^{3} \epsilon^{-1} \lambda+1.01 \cdot \frac{4 \cdot 10^{4} \epsilon^{-2} \gamma^{-1} \lambda}{T_{i}(o)} \sum_{\text {non-center cell } C \in G_{i}}|C \cap Q| .
$$

Proof. The proof is exactly the same as the proofs of Lemmas 27.6.4 and 27.6.6.

Lemma 27.6.9. $\forall i \in\{0,1, \ldots, L\}$,

$$
\sum_{\text {non-center cell } C \in G_{i}} \frac{|C \cap Q|}{T_{i}(o)} \leq 400 k \cdot(\mathrm{OPT} / o) .
$$

Proof.

$$
\begin{aligned}
& \sum_{\text {non-center cell } C \in G_{i}}|C \cap Q| \\
& \leq \frac{\mathrm{OPT}}{\left(g_{i} /(2 d)\right)^{2}}=400 k T_{i}(o) \cdot \frac{\mathrm{OPT}}{o} .
\end{aligned}
$$

Lemma 27.6.10 (The success probability). Condition on the number of center cells of all the grids is at most $100 k L$. If $o \in(\mathrm{OPT} / 16, \mathrm{OPT}]$, then $\operatorname{PointsEstimation~}(o, \epsilon, \delta)$ (Algorithm 27.4) does not output FAIL with probability at least $1-3 \delta / 10$.

Proof. Let $o \geq$ OPT /16, according to Lemma 27.6.9 and Lemma 27.6.8, with probability at least $1-\delta / 10, \forall i \in\{0,1, \ldots, L\}$, we have

$$
\sum_{\text {non-center cell }} \sum_{C \in G_{i}} h_{p \in C \cap Q}(p) \leq 4 \cdot 10^{3} \lambda+1.01 \cdot 4 \cdot 10^{4} \lambda \cdot 6400 k \leq 10^{11} k d L \log (1 / \delta)
$$

and

$$
\sum_{\text {non-center cell }} \sum_{C \in G_{i}} h_{p \in C \cap Q}^{\prime}(p) \leq 4 \cdot 10^{3} \epsilon^{-1} \lambda+1.01 \cdot 4 \cdot 10^{4} \epsilon^{-2} \gamma^{-1} \lambda \cdot 6400 k \leq 10^{15} \cdot \epsilon^{-3} k L^{2} d^{4}
$$

Since the number of center cells is at most $100 k L$, the total number of cells which contains some $p$ with $h_{i}(p)=1$ is at most $10^{12} k d L \log (1 / \delta) \leq \alpha$, and the total number of cells which contains some $p$ with $h_{i}^{\prime}(p)=1$ is at most $10^{16} \epsilon^{-3} k L^{2} d^{4} \leq \alpha^{\prime}$. According to Lemma 27.6.2, with probability at least $1-2 \delta / 10$, none the call of SToring will return FAIL. Thus the overall probability that Algorithm 27.4 does not output FAIL is at last $1-3 \delta / 10$.

Lemma 27.6.11 (Space of Algorithm 27.4). PointsEstimation $(o, \epsilon, \delta)$ (Algorithm 27.4) uses space at most $O\left(k \epsilon^{-3} L^{4} d^{5} \cdot \log ^{2}(k L d /(\epsilon \delta))\right)$ bits.

Proof. The total space used is dominated by the space needed to run $L+1$ copies of Stor$\operatorname{ING}\left(G_{i}, \alpha^{\prime}, 1,0.1 \delta / L\right)$ in line 27.4 of Algorithm 27.4. According to Lemma 27.6.2, the total space needed is $O\left(L \cdot \alpha^{\prime} d L \log ^{2}\left(L \alpha^{\prime} / \delta\right)\right)=O\left(k \epsilon^{-3} L^{4} d^{5} \cdot \log ^{2}(k L d /(\epsilon \delta))\right)$ bits.

### 27.6.3 Sensitivity Sampling over a Dynamic Stream

Since using Algorithm 27.4 we can estimate the number of points in each cell and the size of each $Q_{i}$, the only remaining thing for simulating Algorithm 27.2 is to draw samples based on their sensitivity upper bounds. In Algorithm 27.5, we show how to achieve this in a dynamic stream.

Now we analyze Algorithm 27.5.
Fact 27.6.12. If $\operatorname{SAmpling}(o, \epsilon, \delta)$ (Algorithm 27.5) does not output FAIL, then line 27.5 can be implemented, and $p$ is a uniform sample drawn from $Q_{i}$.

Proof. Although $Q_{i}$ cannot be stored explicitly, $\forall p \in Q$, we are able to determine whether $p \in Q_{i}$ since we can use $\widehat{f}$ to find all the crucial cells and check whether $p$ is in a crucial cell of $G_{i}$.

Suppose SAmpling $(o, \epsilon, \delta)$ does not output FAIL. According to Lemma 27.6.2 and the condition in line 27.5, $\forall j \in[\widehat{m}]$, we have $\left\{(p, j) \mid p \in Q_{i}, h_{i, j}(p)=1\right\} \subseteq S_{i}$. Since $\forall x, y \in Q_{i}, \operatorname{Pr}\left[h_{i, j}(x)=1\right]=\operatorname{Pr}\left[h_{i, j}(y)=1\right]$, then the sample $p$ in line 27.5 is drawn uniformly from $Q_{i}$.

```
Algorithm 27.5 Sensitivity sampling over a dynamic stream
    Sampling \((o, \epsilon, \delta)\) :
    Input: a point set \(Q \subseteq[\Delta]^{d}\) described by a stream \(\left\{\left(p_{1}, \pm\right),\left(p_{2}, \pm\right), \ldots\right\}\)
    Run PointsEstimation \((o, \epsilon, \delta / 2)\) in parallel on the input stream. If it returns FAIL, output
    FAIL.
    \(\forall i \in\{0,1, \ldots, L\}, T_{i}(o) \leftarrow\left(d / g_{i}\right)^{2} \cdot o / k \cdot 1 / 100 . \quad \triangleright\) Threshold for heaviness (Algorithm 27.1).
    \(\widehat{m} \leftarrow \Theta\left(k \epsilon^{-3} L^{4} d^{7} \log \left(\frac{d L k}{\delta}\right) \cdot \frac{1}{\delta}\right) . \quad \triangleright \widehat{m}\) hash functions needed for \(m\) independent samples.
    \(\lambda \leftarrow 10\lceil(d L+\log (1 / \delta)+1)\rceil\). \(\triangleright \lambda\) is the independence parameter.
    \(\forall i \in\{0,1, \ldots, L\}\), choose \(h_{i, 1}, h_{i, 2}, \ldots, h_{i, \widehat{m}}:[\Delta]^{d} \rightarrow\{0,1\}\).
        \(\triangleright \forall j \in[\widehat{m}], h_{i, j}\) is \(\lambda\)-wise independent, \(\forall p \in[\Delta]^{d}, h_{i, j}(p)=1\) w.p. \(\min \left\{1 /\left(10^{4} k L T_{i}(o)\right), 1\right\}\).
    \(\alpha \leftarrow \Theta\left(k \epsilon^{-3} L^{3} d^{7} \log (d L k / \delta) \delta^{-1}\right), \beta \leftarrow \Theta\left(\epsilon^{-3} L^{3} d^{7} \log (d L k / \delta) \delta^{-1}\right) . \quad \triangleright \alpha, \beta\) : for Storing.
    For \(i \in\{0,1, \ldots, L\}\), run \(\operatorname{Storing}\left(G_{i}, \alpha, \beta, 0.1 \delta / L\right)\) (Algorithm 27.3) in parallel.
    \(\triangleright\) Each instance is run on a new stream obtained by splitting each operation \(\left(p_{t}, \pm\right)\) from the
    original input stream into a set of new operations \(\left\{\left(\left(p_{t}, j\right), \pm\right) \mid j \in[\widehat{m}], h_{i, j}\left(p_{t}\right)=1\right\}\).
12: If any Storing returns FAIL, output FAIL.
13: \(\widehat{q}_{0}, \widehat{q}_{1}, \ldots, \widehat{q}_{L}, \widehat{f}: \bigcup_{i=0}^{L} G_{i} \rightarrow \mathbb{R}_{+} \quad \leftarrow \quad\) PointsEstimation \((o, \epsilon, \delta / 2)\) in line 27.5.
    \(\mathfrak{C}_{i}, f_{i}, S_{i} \leftarrow \operatorname{Storing}\left(G_{i}, \alpha, \beta, 0.1 \delta / L\right)\) in line 27.5.
    For \(i \in\{0,1, \ldots, L\}\), if \(\exists C \in \mathcal{C}_{i}\) marked as crucial in line 27.5, and \(f_{i}(C)>\beta\), output FAIL.
    \(\gamma \leftarrow \epsilon /\left(40^{2} L d^{3}\right)\).
    \(I \leftarrow\left\{i \mid 0 \leq i \leq L, \widehat{q}_{i} \geq \gamma T_{i}(o)\right\}\).
    \(t^{\prime} \leftarrow \sum_{i \in I} \widehat{q}_{i} \cdot 10 d^{3} / T_{i}(o) . \quad \triangleright\) Total estimated sensitivities.
    \(m \leftarrow \Theta\left(t^{\prime} \epsilon^{-2} L d \log \left(t^{\prime} / \delta\right)\right), S \leftarrow \emptyset . \quad \triangleright\) Total number of samples.
    For \(i \in\{0,1, \ldots, L\}\), set \(A_{i}=[\widehat{m}]\).
    for \(j=1 \rightarrow m\) do
        Choose a random level \(i \in I\) with probability \(\left(\widehat{q}_{i} \cdot 10 d^{3} / T_{i}(o)\right) / t^{\prime}\).
        Choose the minimum \(j \in A_{i}\) s.t. \(\exists p \in Q_{i},(p, j) \in S_{i}\). If no such \(j\), output FAIL.
        Uniformly choose a point \(p\) from the set \(\left\{q \in Q_{i} \mid(q, j) \in S_{i}\right\}\).
        Update \(A_{i} \leftarrow\{j+1, j+2, \ldots, \widehat{m}\}\).
        Add \(\left(p, t^{\prime} /\left(m s^{\prime}(p)\right)\right)\) to set \(S\). \(\triangleright s^{\prime}(p)=10 d^{3} / T_{i}(o)\).
    end for
    Output: the set \(S\)
```

Lemma 27.6.13 (Correctness of Algorithm 27.5). Suppose $o \in(0, \mathrm{OPT}]$. If $\operatorname{SAmPling}(o, \epsilon, \delta)$
(Algorithm 27.5) does not output FAIL, then with probability at least $1-\delta / 5$, the output $S$ by $\operatorname{Sampling}(o, \epsilon, \delta)$ is an $\epsilon$-coreset for $Q$. Furthermore, the size $|S|$ is at most $O\left(k \epsilon^{-2} L^{2} d^{4} \log (k L d) \cdot(\mathrm{OPT} / o+1)\right)$.

Proof. Due to Lemma 27.6.5 and Lemma 27.6.7, with probability at least $1-\delta / 10, \forall C \in$ $\bigcup_{i=0}^{L} G_{i}$, either $\widehat{f}(C) \in|C \cap Q| \pm 0.1 T_{i}(o)$ or $\widehat{f}(C) \in(1 \pm 0.01)|C \cap Q|$ and $\forall i \in\{0,1, \ldots, L\}$, either $\widehat{q}_{i} \in\left|Q_{i}\right| \pm 0.1 \epsilon \gamma T_{i}(o)$ or $\widehat{q}_{i} \in(1 \pm 0.01 \epsilon)\left|Q_{i}\right|$, where $Q_{0}, Q_{1}, \ldots, Q_{L}$ are defined by using the estimation $\widehat{f}(\cdot)$ (Algorithm 27.1 or Algorithm 27.4). According to Fact 27.6.12, the sampling procedure can be implemented. Then by Theorem ${ }^{6}$ 27.3.1, with probability at least $1-\delta / 10$, the output set $S$ is an $\epsilon$-coreset for $Q$. By taking union bound, with probability $1-\delta / 5$, the set $S$ is an $\epsilon$-coreset for $Q$.

Now let us consider the success probability of Algorithm 27.5. Since we know the success probability of PointsEstimation $(o, \epsilon, \delta / 2)$ in line 27.5 of Algortihm 27.5 and the success probability of $\operatorname{Storing}\left(G_{i}, \alpha, \beta, 0.1 \delta / L\right)$ in line 27.5 , we only need to analyze the success probability in line 27.5 of Algorithm 27.5. To make line 27.5 succeed, we need to find enough samples from $Q_{i}$, i.e., we hope that $\sum_{j=1}^{\widehat{m}} \mathbf{1}\left(\left|\left\{p \in Q_{i} \mid h_{i, j}(p)=1\right\}\right|>0\right)$ is large. In the following analysis, we will show that $\sum_{j=1}^{\widehat{m}} \mathbf{1}\left(\left|\left\{p \in Q_{i} \mid h_{i, j}(p)=1\right\}\right|>0\right)$ is large. First, we show that the number of samples drawn from level $i$ is bounded.

Lemma 27.6.14 (Number of samples from each level). Let I be the set computed in $\operatorname{Sampling}(o, \epsilon, \delta)$ (Algortihm 27.5). With probability at least $1-\delta / 10, \forall i \in I$, the number of times that $i$ is chosen in line 27.5 of Algorithm 27.5 is at most

$$
O\left(\epsilon^{-2} L d^{4} \log \left(t^{\prime} / \delta\right) \cdot \frac{L}{\delta} \cdot \frac{\widehat{q_{i}}}{T_{i}(o)}\right)
$$

Proof. For $i \in I$, the expected number of times that $i$ is chosen is $O\left(m \cdot \widehat{q}_{i} \cdot d^{3} / T_{i}(o) / t^{\prime}\right)$. By Markov's inequality, with probability at least $1-\delta /(20 L)$, the number of times that $i$

[^96]is chosen in line 27.5 of Algorithm 27.5 is at most $O\left(\epsilon^{-2} L d^{4} \log \left(t^{\prime} / \delta\right) \cdot \frac{L}{\delta} \cdot \frac{\widehat{q}_{i}}{T_{i}(o)}\right)$. By taking union bound over all $i \in I$, we complete the proof.

Lemma 27.6.15 (Bounding $t^{\prime}$ ). Consider $o \geq$ OPT /16. Condition on $\widehat{f}: \bigcup_{i=0}^{L} G_{i} \rightarrow \mathbb{R}_{+}$ (in Algorithm 27.5) is good (Lemma 27.6.4), $\widehat{q}_{0}, \widehat{q}_{1}, \ldots, \widehat{q}_{L}$ (in Algorithm 27.5) are good (Lemma 27.6.6), and the number of center cells is at most 100 kL , then we have $t^{\prime} \leq 10^{6} d^{3} L k$.

Proof. We have

$$
\begin{aligned}
t^{\prime} & =\sum_{i \in I} \widehat{q}_{i} \cdot 10 d^{3} / T_{i}(o) \\
& \leq \sum_{i \in I}\left|Q_{i}\right| \cdot 20 d^{3} / T_{i}(o) \\
& \leq \sum_{i=0}^{L}\left|Q_{i}\right| \cdot 20 d^{3} / T_{i}(o) \\
& \leq 2 \cdot 4000 d^{3} L k \cdot 20 \\
& \leq 10^{6} d^{3} L k
\end{aligned}
$$

where the first inequality follows by $\forall i \in I, \widehat{q}_{i} \geq \gamma T_{i}(o)$ and either $\widehat{q}_{i} \in\left|Q_{i}\right| \pm 0.1 \epsilon \gamma T_{i}(o)$ or $\widehat{q}_{i} \in(1 \pm 0.01 \epsilon)\left|Q_{i}\right|$, the second inequality follows by $I \subseteq\{0,1, \ldots, L\}$, the third inequality follows by Lemma 27.3.8 and $o \geq$ OPT $/ 16$.

Lemma 27.6.16 (Number of crucial points in each level). Consider $o \geq$ OPT /16. Conditioning on $\widehat{f}: \bigcup_{i=0}^{L} G_{i} \rightarrow \mathbb{R}_{+}$(in Algorithm 27.5) is good (Lemma 27.6.5) and the number of center cells is at most 100 kL , we have:

$$
\frac{\left|Q_{i}\right|}{T_{i}(o)} \leq 7000 k L
$$

Proof. We have

$$
\begin{aligned}
\frac{\left|Q_{i}\right|}{T_{i}(o)} & =\sum_{\text {center crucial cell } C \in G_{i}} \frac{|C \cap Q|}{T_{i}(o)}+\sum_{\text {non-center crucial cell } C \in G_{i}} \frac{|C \cap Q|}{T_{i}(o)} \\
& \leq 110 k L+400 k(\mathrm{OPT} / o) \\
& \leq 7000 k L
\end{aligned}
$$

where the first inequality follows by that the number of center cells is at most 100 kL , the number points in a crucial cell is at most $1.1 T_{i}(o)$ and Lemma 27.6.9.

Lemma 27.6.17 (The number of samples is large). Consider $o \geq$ OPT /16. Conditioning on $\widehat{f}: \bigcup_{i=0}^{L} G_{i} \rightarrow \mathbb{R}_{+}$(in Algorithm 27.5) is good (Lemma 27.6.4), $\widehat{q}_{0}, \widehat{q}_{1}, \ldots, \widehat{q}_{L}$ (in Algorithm 27.5) are good (Lemma 27.6.6), and the number of center cells is at most 100 kL , with probability at least $1-\delta / 10, \forall i \in I$, we have:

$$
\sum_{j=1}^{\widehat{m}} \mathbf{1}\left(\left|\left\{p \in Q_{i} \mid h_{i, j}(p)=1\right\}\right|>0\right) \geq \Omega\left(\epsilon^{-2} L d^{4} \log \left(\frac{d L k}{\delta}\right) \cdot \frac{L}{\delta} \cdot \frac{\widehat{q}_{i}}{T_{i}(o)}\right)
$$

Proof. Consider a fixed $i \in I . \forall j \in[\widehat{m}]$, by union bound, we have

$$
\underset{h_{i, j}}{\operatorname{Pr}}\left[\exists p \in Q_{i}, h_{i, j}(p)=1\right] \leq \frac{\left|Q_{i}\right|}{10^{4} k L T_{i}(o)}<1,
$$

where the inequality follows by Lemma 27.6.16. Thus $\left\lceil 10^{4} k L T_{i}(o) /\left|Q_{i}\right|\right\rceil \leq 2 \cdot 10^{4} k L T_{i}(o) /\left|Q_{i}\right|$. Let $b=10 \cdot\left\lceil 10^{4} k L T_{i}(o) /\left|Q_{i}\right|\right\rceil$. Let $r=\lfloor\widehat{m} / b\rfloor \geq \widehat{m} \cdot\left|Q_{i}\right| /\left(2 \cdot 10^{5} k L T_{i}(o)\right)-1$. Since $i \in I$, we have $\left|Q_{i}\right| / T_{i}(o) \geq \frac{1}{2} \gamma$. Since $\widehat{m} \geq 10^{9} k L / \gamma$, we have $r \geq \widehat{m} /\left(4 \cdot 10^{5} k L\right) \cdot\left|Q_{i}\right| / T_{i}(o)$. For $s \in[r]$, we can define a random variable $Y_{s}$,

$$
Y_{s}=\sum_{j=(s-1) \cdot b+1}^{s \cdot b} \sum_{p \in Q_{i}} h_{i, j}(p)
$$

We have $\mathbb{E}\left[Y_{s}\right]=b \cdot\left|Q_{i}\right| /\left(10^{4} k L T_{i}(o)\right)$. Thus, $E\left[Y_{s}\right] \in[10,20]$. Since $Y_{s}$ is a sum of several (at least) pairwise independent unit random variables, $\mathbb{V}\left[Y_{s}\right] \leq 20$. Thus, by Chebyshev's inequality, we have

$$
\operatorname{Pr}\left[\left|Y_{s}-\mathbb{E}\left[Y_{s}\right]\right| \geq 9\right] \leq 20 / 81 \leq 0.25
$$

Define $X_{s}$ be the random variable such that $X_{s}=\mathbf{1}\left(Y_{s} \geq 1\right)$. Then $\mathbb{E}\left[\sum_{s \in[r]} X_{s}\right] \geq 0.75 r$. By Chernoff bound, we know that

$$
\operatorname{Pr}\left[\sum_{s \in[r]} X_{s} \leq 0.5 r\right] \leq 2^{-r / 20} \leq 0.01 \delta / L
$$

where the last inequality follows by $r \geq \widehat{m} /\left(4 \cdot 10^{5} k L\right) \cdot \gamma / 2 \geq 20 \log (100 L / \delta)$.
Since $\widehat{m}$ is sufficiently large, i.e.,

$$
\widehat{m} \geq \Omega\left(k \epsilon^{-2} L^{3} d^{4} \log \left(\frac{d L k}{\delta}\right) \cdot \frac{1}{\delta \gamma}\right) \geq \Omega\left(k \epsilon^{-3} L^{4} d^{7} \log \left(\frac{d L k}{\delta}\right) \cdot \frac{1}{\delta}\right)
$$

we have

$$
r \geq \widehat{m} /\left(4 \cdot 10^{5} k L\right) \cdot\left|Q_{i}\right| / T_{i}(o) \geq \Omega\left(\epsilon^{-2} L d^{4} \log \left(\frac{d L k}{\delta}\right) \cdot \frac{L}{\delta} \cdot \frac{\widehat{q}_{i}}{T_{i}(o)}\right)
$$

where the last inequality follows by $\widehat{q}_{i}=\Theta\left(\left|Q_{i}\right|\right)$. Notice that

$$
\operatorname{Pr}\left[\sum_{j=1}^{\widehat{m}} \mathbf{1}\left(\left|\left\{p \in Q_{i} \mid h_{i, j}(p)=1\right\}\right|>0\right) \geq 0.5 r\right] \geq \operatorname{Pr}\left[\sum_{s \in[r]} X_{s} \geq 0.5 r\right] \geq 1-0.01 \delta / L
$$

Thus, with probability at least $1-0.01 \delta / L$,

$$
\sum_{j=1}^{\widehat{m}} \mathbf{1}\left(\left|\left\{p \in Q_{i} \mid h_{i, j}(p)=1\right\}\right|>0\right) \geq \Omega\left(\epsilon^{-2} L d^{4} \log \left(\frac{d L k}{\delta}\right) \cdot \frac{L}{\delta} \cdot \frac{\widehat{q}_{i}}{T_{i}(o)}\right)
$$

By taking the union bound over $i \in I$, we complete the proof.

Lemma 27.6.18 (Sampling stage succeeds). Consider $o \geq$ OPT /16. Conditioning on $\widehat{f}: \bigcup_{i=0}^{L} G_{i} \rightarrow \mathbb{R}_{+}$(in Algorithm 27.5) is good (Lemma 27.6.4), $\widehat{q}_{0}, \widehat{q}_{1}, \ldots, \widehat{q}_{L}$ (in Algorithm 27.5) are good (Lemma 27.6.6), and the number of center cells is at most 100 kL , if $\operatorname{Sampling}(o, \epsilon, \delta)$ does not output FAIL before line 27.5 of Algorithm 27.5, then with probability at least $1-\delta / 5$, it will not output $\boldsymbol{F A I L}$.

Proof. According to Lemma 27.6.14 and Lemma 27.6.15, with probability at least $1-\delta / 10$, $\forall i \in I$, the sampling procedure will not request too many samples from level $i$. According to Lemma 27.6.17, with probability at least $1-\delta / 10$, the number of samples needed for each level $i \in I$ is enough. Thus, with probability at least $1-\delta / 5$, the algorithm will not output FAIL.

Lemma 27.6.19 (Samples can fit into the space). Suppose $o \geq$ OPT /16. Conditioning on $\widehat{f}: \bigcup_{i=0}^{L} G_{i} \rightarrow \mathbb{R}_{+}$(in Algorithm 27.5) is good (Lemma 27.6.4), if the total number of center cells is at most $100 k L$, with probability at least $1-\delta / 5$, $\operatorname{SAmPLing}(o, \epsilon, \delta)$ (Algorithm 27.5) will not output FAIL in line 27.5 nor line 27.5.

Proof. Consider $i \in\{0,1, \ldots, L\}$ and a cell $C \in G_{i}$ which is marked as crucial by PointsEsTIMATION $(o, \epsilon, \delta / 2)$.

$$
\begin{aligned}
\mathbb{E}\left[\sum_{j=1}^{\widehat{m}} \sum_{p \in C \cap Q} h_{i, j}(p)\right] & =\widehat{m} \cdot|C \cap Q| /\left(10^{4} k L T_{i}(o)\right) \\
& \leq O\left(\epsilon^{-3} L^{3} d^{7} \log (d L k / \delta) \delta^{-1}\right)
\end{aligned}
$$

By Theorem 27.6.3,

$$
\operatorname{Pr}\left[\sum_{j=1}^{\widehat{m}} \sum_{p \in C \cap Q} h_{i, j}(p)>\mathbb{E}\left[\sum_{j=1}^{\widehat{m}} \sum_{p \in C \cap Q} h_{i, j}(p)\right]+\epsilon^{-3} L^{3} d^{7} \log (d L k / \delta) \delta^{-1}\right] \leq\left(\delta / \Delta^{d}\right)^{5} .
$$

Thus, by taking union bound, with probability at last $1-\delta / 10, \forall i \in\{0,1, \ldots, L\}$, any cell $C \in G_{i}$ which is marked as crucial, the total number of points sampled in $C$ is at most $\beta$. Thus, with probability at least $1-\delta / 20, \operatorname{Sampling}(o, \epsilon, \delta)$ (Algorithm 27.5) will not output FAIL in line 27.5.

Consider $i \in\{0,1 \ldots, L\}$. Let us analyze the number of cells in $G_{i}$ which contain at least 1 sample points. The number of points which are not in the center cell is at most

$$
\frac{\mathrm{OPT}}{\left(g_{i} /(2 d)\right)^{2}} \leq 400 k T_{i}(o) \cdot \mathrm{OPT} / o \leq 6400 k T_{i}(o)
$$

Thus, the expected number of sampled points in non-center cell is at most $O\left(k \epsilon^{-3} L^{3} d^{7} \log (d L k / \delta)\right.$. $1 / \delta)$. Since the number of center cells is at most $100 k L$, by Theorem 27.6 .3 , with probability at least $1-\delta /(100 L)$, the number of cells in $G_{i}$ which contain at least 1 sample points is $O\left(k \epsilon^{-3} L^{3} d^{7} \log (d L k / \delta) \cdot 1 / \delta\right)$. By taking union bound over all $i \in\{0,1, \ldots, L\}$, with probability at least $1-\delta / 20, \forall i \in\{0,1, \ldots, L\}$, the number of sampled cell in $G_{i}$ is at most $\alpha$.

Due to Lemma 27.6.2, with probability at least $1-\delta / 10$, none of the $\operatorname{Storing}\left(G_{i}, \alpha, \beta, 0.1 \delta / L\right)$ in line 27.5 of Algorithm 27.5 will output FAIL. By union bound over all the failure probabilities, we complete the proof.

Lemma 27.6.20 (The overall success probability). Suppose $o \geq$ OPT / 16 and the total number of center cells is at most $100 k L$. With probability at least $1-4 \delta / 5, \operatorname{Sampling}(o, \epsilon, \delta)$ (Algorithm 27.5) will not output FAIL.

Proof. Due to Lemma 27.6.10, PointsEstimation $(o, \epsilon, \delta / 2)$ in line 27.5 will not output FAIL with probability at least $1-3 \delta / 20$. By Lemma 27.6.5 and Lemma 27.6.6, with
probability at least $1-\delta / 5, \forall i \in\{0,1, \ldots, L\}$, either $\widehat{q}_{i} \in\left|Q_{i}\right| \pm 0.1 \epsilon \gamma T_{i}(o)$ or $\widehat{q}_{i} \in(1 \pm$ $0.01 \epsilon)\left|Q_{i}\right|$, and $\forall C \in G_{i}$, either $\widehat{f}(C) \in|C \cap Q| \pm 0.1 T_{i}(o)$ or $\widehat{f}(C) \in(1 \pm 0.01)|C \cap Q|$. Then by Lemma 27.6.19, with probability at least $1-\delta / 5, \operatorname{SAMPLing}(o, \epsilon, \delta)$ will not output FAIL in line 27.5 nor line 27.5. Finally, according to Lemma 27.6.18, with probability at least $1-\delta / 5$, the algorithm does not output FAIL. By taking union bound over all the bad events, with probability at least $1-3 \delta / 20-\delta / 5-\delta / 5-\delta / 5 \geq 1-4 \delta / 5, \operatorname{Sampling}(o, \epsilon, \delta)$ (Algorithm 27.5) will not output FAIL.

Lemma 27.6.21 (Total space of Algorithm 27.5). $\operatorname{Sampling}(o, \epsilon, \delta)$ uses space at most $O\left(k \epsilon^{-6} L^{8} d^{15} \delta^{-2} \cdot \log ^{4}(k L d /(\epsilon \delta))\right)$ bits.

Proof. According to Lemma 27.6.11, PointsEstimation $(o, \epsilon, \delta / 2)$ in line 27.5 of Algorithm 27.5 takes the space $O\left(k \epsilon^{-3} L^{4} d^{5} \cdot \log (k L d /(\epsilon \delta))\right)$ bits. According to Lemma 27.6.2, line 27.5 takes the total space:

$$
O\left(L \cdot \alpha \beta d L \cdot \log ^{2}(\alpha \beta / \delta)\right)=O\left(k \epsilon^{-6} L^{8} d^{15} \delta^{-2} \cdot \log ^{4}(k L d /(\epsilon \delta))\right)
$$

### 27.6.4 The Final Algorithm

Finally, we use exponential search to enumerate the guesses $o$. In Algorithm 27.6, we show the details of how to run Algorithm 27.5 with different guesses in parallel. Our main theorem is the following.

Theorem 27.6.22. Suppose a point set $Q \subseteq[\Delta]^{d}$ is given by a stream of insertion/deletion operations in the dynamic streaming model (Definition 27.2.1). Let $L=\log \Delta$. For given

```
Algorithm 27.6 Coreset construction over a dynamic stream
    DynamicCoreset \((\epsilon)\) :
    input: a point set \(Q \subseteq[\Delta]^{d}\) described by a stream \(\left\{\left(p_{1}, \pm\right),\left(p_{2}, \pm\right), \ldots\right\}\)
    Impose randomly shifted grids \(G_{-1}, G_{0}, G_{1}, \ldots, G_{L}\) (Definition 27.3.1).
    Run Distinct ( \(10000 k\), 0.001) (Lemma 27.6.1) over the input stream in parallel.
    If line 27.6 does not output FAIL, we output the entire point set \(Q\).
    For each \(u \in[2 d L]\), let \(o_{u}=2^{u} \cdot 50 k\) and run \(\operatorname{SAmpling}\left(o_{u}, \epsilon, 0.001 /(d L)\right)\) (Algorithm 27.5)
    over the input stream in parallel.
7: Set a threshold \(h=\Theta\left(k \epsilon^{-2} L^{2} d^{4} \log (k L d)\right)\).
8: Find the smallest \(u^{*}\) such that \(\operatorname{Sampling}\left(o_{u^{*}}, \epsilon, 0.001 /(d L)\right)\) in line 27.6 does not output FAIL,
    and the returned set \(S^{*}\) has size at most \(h\), i.e., \(\left|S^{*}\right| \leq h\).
9: If no such \(u^{*}\), output FAIL. Otherwise, output \(S^{*}\) returned by \(\operatorname{SAmPling}\left(o_{u^{*}}, \epsilon, 0.001 /(d L)\right)\).
```

$\epsilon \in(0,1 / 2)$, Algorithm 27.6 uses a single pass over the stream and on termination outputs a $k$-means $\epsilon$-coreset $S$ (Definition 27.2.3) for $Q$ with probability at least 0.9. Furthermore, the size of the coreset is at most $O\left(k \epsilon^{-2} d^{4} L^{2} \log (k d L)\right)$. The total space used by the algorithm is $\widetilde{O}(k) \cdot \operatorname{poly}(d L / \epsilon)$ bits.

We divide the proof of Theorem 27.6.22 into the following two lemmas.

Lemma 27.6.23 (Correctness and success probability). With probability at least 0.9, DynamicCoreset( $\epsilon$ (Algorithm 27.6) outputs an $\epsilon$-coreset for $Q$ and the size of the coreset is at most $O\left(k \epsilon^{-2} L^{2} d^{4} \log (k L d)\right)$.

Proof. If $|Q| \leq 10000 k$, then according to Lemma 27.6.1, with probability at least 0.999 , the entire data set $Q$ will be returned by line 27.6 in Algorithm 27.6.

Consider the case when $|Q| \geq 10000 k$. According to Lemma 27.3.7, with probability at least 0.94 , the total number of center cells is upper bounded by $100 k L$. Now, we condition on this happens. Since $|Q| \geq 10000 k$, we know that OPT $\geq 1000 k$. There exists $u \in[2 d L]$ such that $o_{u} \in[\mathrm{OPT} / 16, \mathrm{OPT}]$. According to Lemma 27.6.20, with probability
at least 0.999, SAmpling $\left(o_{u}, \epsilon, 0.001 /(d L)\right)$ will not output FAIL. By Lemma 27.6.13, with probability at least 0.999, the set $S$ returned by $\operatorname{SAMPLing}\left(o_{u}, \epsilon, 0.001 /(d L)\right)$ is an $\epsilon$-coreset and $|S| \leq O\left(k \epsilon^{-2} L^{2} d^{4}\right)$. Thus, with probability at least 0.998 , DynamicCoreset ( $\epsilon$ ) (Algorithm 27.6) will not output FAIL. Consider another $u^{\prime}<u$. If $\operatorname{Sampling}\left(o_{u^{\prime}}, \epsilon, 0.001 /(d L)\right)$ does not output FAIL, and the set $S^{\prime}$ returned has size at most $h$, then according to Lemma 27.6.13, with probability at least $1-0.001 /(d L), S^{\prime}$ is an $\epsilon$-coreset for $Q$. By taking union bound over all the such $u^{\prime}$, then with probability at least $0.999, S^{*}$ returned by $\operatorname{SAMPLING}\left(o_{u^{*}}, \epsilon, 0.001 /(d L)\right)$ is an $\epsilon$-coreset for $Q$. By taking union bound over all the bad events, we complete the proof.

Lemma 27.6.24 (Total space needed for Algortihm 27.6). DynamicCoreset( $\epsilon$ ) (Algorithm 27.6) uses space at most $O\left(k \epsilon^{-6} L^{11} d^{18} \log ^{4}(k L d / \epsilon)\right)$ bits.

Proof. DynamicCoreset $(\epsilon)$ (Algorithm 27.6) runs $\Theta(d L)$ copies of $\operatorname{SAmpling}\left(o_{u}, \epsilon, 0.001 /(d L)\right)$. By Lemma 27.6.21, the total space needed is

$$
O\left(d L \cdot k \epsilon^{-6} L^{8} d^{15} \cdot(d L)^{2} \cdot \log ^{4}(k L d / \epsilon)\right)=O\left(k \epsilon^{-6} L^{11} d^{18} \log ^{4}(k L d / \epsilon)\right)
$$

bits.

Proof of Theorem 27.6.22. Lemma 27.6.23 shows the correctness and the success probability of the algorithm. Lemma 27.6.24 shows the total space needed by the algorithm.

### 27.7 Why Do Previous Techniques Fail?

In this section, we describe some previous techniques in more detail and explain why they fail in our setting.

Uniform sampling method. [FIS05] is one of the early papers using sampling procedures to solve dynamic streaming geometric problems. They showed that it is possible to use point samples from a dynamic point set to solve several geometric problems, e.g., Euclidean Minimum Spanning Tree. However, they only showed how to implement uniform sampling by using counting distinct elements and subsampling procedure as subroutines. In our setting, we require different sampling probabilities for different points. Although the bottom-level uniform sampling scheme of ours is similar to theirs, our overall sampling method is more complicated.

Estimating the cost. [Ind04] used a critical observation to estimate the cost of $k$-median, that is: let $Z$ be a set of centers and $P$ be the point set, then $\operatorname{cost}(Z, P)=\int_{0}^{\infty} \mid P-$ $B(Z, r) \mid \mathrm{d} r$, where $B(Z, r)$ is the union of balls of radius $r$ centered at all points in $Z$. Then this integration is approximated by a summation with logarithmic levels, i.e., $\int_{0}^{\infty} \mid P-$ $B(Z, r)\left|\mathrm{d} r \approx \sum_{i=0}^{\infty}\right| P-B\left(Z, r^{(i+1)}\right) \mid \cdot\left(r^{(i+1)}-r^{(i)}\right)$, where $r^{(i)}=O\left(\epsilon(1+\epsilon)^{i}\right)$. The critical part is to estimate $\left|P-B\left(Z, r^{(i)}\right)\right|$. [Ind04] constructed a counting data structure based on grids with side length $O\left(r^{(i)}\right)$. Then every input point is snapped to a grid point. To obtain sufficiently accurate estimates for all $\left|P-B\left(Z, r^{(i)}\right)\right|$, the data structure needs to query $|Z| / \epsilon^{O(d)}$ many grid points per $Z$. Such a data structure is implemented using pairwise independent hash functions, and uses memory $|Z| / \epsilon^{O(d)}$. Notice that this method only
gives an estimate of the cost, and does not construct a coreset. In order to obtain a $k$ median solution, an exhaustive search is needed. Furthermore, this technique fails to extend to $k$-means which does not have an integration formula for the cost function.

Exponential size coreset. [FS05] constructed an $\epsilon$-coreset of size $k \epsilon^{-O(d)}$ for $k$-means and $k$-median. They also used the same grid structure as we use. A cell is marked as "heavy" if the cell contains enough points such that moving all points in the cell to the center of this cell incurs too much error in the optimal cost of $k$-means/median. Since the side-lengths of cells decrease as level increases, the number of points required to have this effect becomes larger. Eventually, all cells are non-heavy after some level. As such we also have a tree of heavy cells. The coreset is constructed by looking at each heavy cell and assigning each point in its non-heavy children cells to its center. It turns out that if we want an $\epsilon$-coreset, the threshold of non-heavy cells is exponential in $d$, i.e., each non-heavy cell in level $i$ cannot contain more than $\widetilde{O}\left(\epsilon^{O(d)} \cdot \mathrm{OPT} / 2^{i}\right)$ points. This small threshold gives rise to $\widetilde{O}\left(1 / \epsilon^{O(d)}\right)$ many heavy cells.

Insertion-only streams. Many of the previous insertion-only streaming coreset construction algorithms (e.g., [FS12]) heavily depend on a "merge-reduce" technique, i.e. reading some points in the stream, constructing a coreset, reading another part, constructing a new coreset, and then merging the two coresets. This procedure is repeated until the stream ends. This technique works well in the insertion-only streaming model, but it fails immediately when deletions are allowed. Although [BFL16] gave a new framework other than merge-reduce, their algorithm relies on a non-deleting structure of data streams as well.


Figure 27.3: Telescope sum $\left[\mathrm{BFL}^{+} 17\right]$ fails for $k$-means. In the $k$-median problem, for a fixed set of centers $Z$, the total cost can be written as a telescope sum $\sum_{p \in P}\left(\operatorname{dist}\left(c_{p}^{i}, Z\right)-\operatorname{dist}\left(c_{p}^{i-1}, Z\right)\right)$. For each piece, $\left|\operatorname{dist}\left(c_{p}^{i}, Z\right)-\operatorname{dist}\left(c_{p}^{i-1}, Z\right)\right|$ is always upper bounded by $\operatorname{dist}\left(c_{p}^{i-1}, c_{p}^{i}\right)$ which is independent from the choice of $Z$. However, in the $k$-means problem, the telescope sum of the total cost is $\sum_{p \in P}\left(\operatorname{dist}\left(c_{p}^{i}, Z\right)^{2}-\operatorname{dist}\left(c_{p}^{i-1}, Z\right)^{2}\right)$. For each piece, the upper bound of $\left|\operatorname{dist}\left(c_{p}^{i}, Z\right)^{2}-\operatorname{dist}\left(c_{p}^{i-1}, Z\right)^{2}\right|$ may depend on the location of $Z$, and it can be larger than $\Delta$ in the worst case.

Algorithm for $k$-median only. Although some $k$-median coreset construction techniques can be easily extended to $k$-means (see e.g. [FS12, BFL16]), those constructions can only be implemented in the insertion-only streaming model. [BFL $\left.{ }^{+} 17\right]$ gave a $k$-median coreset construction in the dynamic streaming model, but their construction cannot be extended to $k$-means directly, as we explain now. Their $k$-median algorithm heavily relies on writing the cost of each point as a telescope sum. For example, we consider the 1-median problem. let $z$ be a candidate center point and $p \in P$ be a point, $\operatorname{then} \operatorname{dist}(p, z)=\operatorname{dist}(p, z)-\operatorname{dist}\left(c_{p}^{L-1}, z\right)+$ $\operatorname{dist}\left(c_{p}^{L-1}, z\right)-\operatorname{dist}\left(c_{p}^{L-2}, z\right)+\cdots-\operatorname{dist}\left(c_{p}^{0}, z\right)$, where each $c_{p}^{i}$ is the center of the cell in the $i$-th level containing $p$. Therefore, the total 1-median cost $\sum_{p \in P} \operatorname{dist}(p, z)$ of point set $P$ on $z$ can be split into $L$ pieces, i.e., $\sum_{p \in P}\left(\operatorname{dist}\left(c_{p}^{i}, z\right)-\operatorname{dist}\left(c_{p}^{i-1}, z\right)\right)$ for each $i \in[L]$. [ $\left.\mathrm{BFL}^{+} 17\right]$ estimated each of the $L$ pieces by sampling points, i.e., let $S^{i}$ be the samples in the $i$-th level, then the estimator of $\sum_{p \in P}\left(\operatorname{dist}\left(c_{p}^{i}, z\right)-\operatorname{dist}\left(c_{p}^{i-1}, z\right)\right)$ is $\sum_{p \in S^{i}}\left(\operatorname{dist}\left(c_{p}^{i}, z\right)-\operatorname{dist}\left(c_{p}^{i-1}, z\right)\right) / \zeta_{p}^{i}$ where $\zeta_{p}^{i}$ is the probability that point $p$ is sampled. A crucial observation is that we have
$\left|\operatorname{dist}\left(c^{i}, z\right)-\operatorname{dist}\left(c^{i-1}, z\right)\right| \leq \Delta / 2^{i}$ - the cell size in level $i$ which is independent of the location of $z$. Using this nice upper bound on $\left|\operatorname{dist}\left(c^{i}, z\right)-\operatorname{dist}\left(c^{i-1}, z\right)\right|,\left[\mathrm{BFL}^{+} 17\right]$ applied Bernstein inequality to get high concentration of the estimator $\sum_{p \in S^{i}}\left(\operatorname{dist}\left(c_{p}^{i}, z\right)-\operatorname{dist}\left(c_{p}^{i-1}, z\right)\right) / \zeta_{p}^{i}$ with only $\widetilde{O}\left(1 / \epsilon^{2}\right)$ samples per level. However, this framework does not work for 1 -means even though one can still write the telescope sum structure $\sum_{p}\left(\operatorname{dist}^{2}\left(c_{p}^{i}, z\right)-\operatorname{dist}^{2}\left(c_{p}^{i-1}, z\right)\right)$ and can still setup an estimator $\sum_{p \in S^{i}}\left(\operatorname{dist}^{2}\left(c_{p}^{i}, z\right)-\operatorname{dist}^{2}\left(c_{p}^{i-1}, z\right)\right) / \zeta_{p}^{i}$. But $\left|\operatorname{dist}^{2}\left(c_{p}^{i}, z\right)-\operatorname{dist}^{2}\left(c_{p}^{i-1}, z\right)\right|$ is not upper bounded by the cell size anymore. Instead, it depends on the location of $z$. For example, it can be as large as $\left|\operatorname{dist}^{2}\left(c_{p}^{i}, z\right)-\operatorname{dist}^{2}\left(c_{p}^{i-1}, z\right)\right| \geq \Delta$. See Figure 27.3. If we apply Bernstein inequality here, we will need too many samples to save any space.

## Chapter 28

## Convergence Result of one-hidden Layer Neural Network

We improve the over-parametrization size over the two beautiful results [ Li and Liang' 2018] and [Du, Zhai, Poczos and Singh' 2019] in deep learning theory.

This part is based upon the following previous publication

- Zhao Song, Xin Yang

Quadratic Suffices for Over-parametrization via Matrix Chernoff Bound.
Manuscript 2019 [SY19]

### 28.1 Introduction

Over-parameterization theory for deep neural networks becomes extremely popular over the last few years. There is a long line (still growing very quickly) of work proving that (stochastic) gradient descent algorithm is able to find the global minimum if the network is wide enough [LL18, DZPS19, AZLS18, AZLS19, DLL ${ }^{+}$19, ZCZG18]. One fundamental question for over-parameterization is, how wide should the neural network be?

Formally speaking, the existing results show that as long as the width $m$ is at least polynomial of number of input data $n$, then (S)GD-type algorithm can work in the following sense: first randomly pick a weight matrix to be the initialization point, update the weight matrix according to gradient direction over each iteration, and eventually find the global minimum. It is conjectured [Lee18] that $m=\Omega(n \operatorname{poly}(\log (n / \delta)))$ is the right answer, where $\delta$ is the failure probability. The randomness is from the random initialization and also algorithm itself, but not from data. There are other work relied on input data to be random [BG17, Tia17, ZSJ ${ }^{+}$17, Son17, LY17, ZSD17, DLT ${ }^{+}$18, GLM18, BJW19], however overparameterization theory does not allow that assumption and it only needs to make very mild assumption on data, e.g. separable. The breakthrough result by Li and Liang [LL18] is the first one that is able to explain why the greedy algorithm works very well in practice for ReLU neural network from over-parameterization perspective. The state-of-the-art result for one-hidden-layer neural network with ReLU activation function is due to Du, Zhai, Poczos and Singh [DZPS19]. Their beautiful result proved that $m=\Omega\left(n^{6} \operatorname{poly}(\log n, 1 / \delta)\right)$ is sufficient. We improve the result [DZPS19] from two perspectives : one is the dependence on failure probability, and the other is the dependence on the number of input data. More precisely, we show that $m=\Omega\left(n^{4} \operatorname{poly}(\log (n / \delta))\right)$ is sufficient via a careful concentration analysis.

More interestingly, when the input data have certain property, we can improve the bound to $m=\Omega\left(n^{2} \operatorname{poly}(\log (n / \delta))\right)$ via a more careful concentration analysis for random variables.

The study of concentration of sums of random variables dates back to Central Limit Theorems, the first modern concentration bounds were probably proposed by Bernstein [Ber24]. Chernoff bound is an extremely popular variant, which was introduced by Rubin and published by Chernoff [Che52]. Chernoff bound is a fundamental tool in Theoretical Computer Science and has been used in almost every randomized algorithm paper without even stating it. One common statement is the following: given a list of independent random variables $x_{1}, \cdots, x_{m} \in[0,1]$ with mean $\mu$, then

$$
\operatorname{Pr}\left[\left|\frac{1}{m} \sum_{i=1}^{m} x_{i}-\mu\right|>\epsilon\right] \leq 2 \exp \left(-\Omega\left(m \epsilon^{2}\right)\right)
$$

In many applications, we are not just dealing with scalar random variables. A natural generalization of the Chernoff bound appeared in the works of Rudelson [Rud99], AhlswedeWinter [AW02], and Tropp [Tro12]. They proved that a similar concentration phenomenon is true even for matrix random variables. Given a list of independent complex Hermitian random matrices $X_{1}, \cdots, X_{m} \in \mathbb{C}^{n \times n}$ with mean $\mu$ and $\left\|X_{i}\right\| \leq 1, \forall i \in[m]$, then

$$
\operatorname{Pr}\left[\left\|\frac{1}{m} \sum_{i=1}^{m} X_{i}-\mu\right\|>\epsilon\right] \leq 2 n \exp \left(-\Omega\left(m \epsilon^{2}\right)\right)
$$

For a more detailed survey and recent progress on the topic Matrix Chernoff bound. We refer the readers to [Tro15, GLSS18, KS18].

In this work, we draw an interesting connection between deep learning theory and Matrix Chernoff bound : we can view the width of neural network as the number of independent random matrices.

### 28.1.1 Our Result

We start with the definition of Gram matrix, which can be found in [DZPS19].

Definition 28.1.1 (Data-dependent function $H$ ). Given a collection of data $\left\{x_{1}, \cdots, x_{n}\right\} \subset$ $\mathbb{R}^{d}$. For any vector $w \in \mathbb{R}^{d}$, we define symmetric matrix $H(w) \in \mathbb{R}^{n \times n}$ as follows

$$
H(w)_{i, j}=x_{i}^{\top} x_{j} \mathbf{1}_{w^{\top} x_{i} \geq 0, w^{\top} x_{j} \geq 0}, \forall(i, j) \in[n] \times[n] .
$$

Then we define continuous Gram matrix $H^{\text {cts }} \in \mathbb{R}^{n \times n}$ in the following sense

$$
H^{\mathrm{cts}}=\underset{w \sim N(0, I)}{\mathbb{E}}[H(w)]
$$

Similarly, we define discrete Gram matrix $H^{\text {dis }} \in \mathbb{R}^{n \times n}$ in the following sense

$$
H^{\mathrm{dis}}=\frac{1}{m} \sum_{r=1}^{m} H\left(w_{r}\right)
$$

We use $N(0, I)$ to denote Gaussian distribution. We use $\mathbb{E}_{w}$ to denote $\mathbb{E}_{w \sim N(0, I)}$ and $\operatorname{Pr}_{w}$ to denote $\operatorname{Pr}_{w \sim N(0, I)}$. We introduce some mild data-dependent assumption. Without loss of generality, we can assume that $\left\|x_{i}\right\|_{2} \leq 1, \forall i \in[n]$.

Assumption 28.1.1 (Data-dependent assumption). We made the following data-dependent assumption:

1. Let $\lambda=\lambda_{\min }\left(\mathbb{E}_{w}[H(w)]\right)$ and $\lambda \in(0,1]$.
2. Let $\alpha \in[1, n]$ and $\gamma \in[0,1)$ be the parameter such that ${ }^{1}$

$$
\underset{w}{\operatorname{Pr}}[\|H(w)-\underset{w}{\mathbb{E}}[H]\| \leq \alpha] \geq 1-\gamma .
$$

[^97]3. Let $\beta \in\left[1, n^{2}\right]$ be the parameter such that
$$
\left\|\underset{w}{\mathbb{E}}\left[(H(w)-\underset{w}{\mathbb{E}}[H(w)])(H(w)-\underset{w}{\mathbb{E}}[H(w)])^{\top}\right]\right\| \leq \beta
$$
4. Let $\theta \in[0, \sqrt{n}]$ be parameter such that
$$
\left|x_{i}^{\top} x_{j}\right| \leq \theta / \sqrt{n}, \forall i \neq j .
$$

The first assumption is from [DZPS19]. For more detailed discussion about that assumption, we refer the readers to [DZPS19]. The last assumption is similar to assumption in [AZLS18, AZLS19], where they assumed that for $i \neq j,\left\|x_{i}-x_{j}\right\|_{2} \geq \theta^{\prime}$. If we think of $\left\|x_{i}\right\|_{2}=1, \forall i \in[n]$, then we know that $\left(\theta^{\prime}\right)^{2} \leq 2-2 x_{i}^{\top} x_{j}$. It indicates $\left(\theta^{\prime}\right)^{2}+2 \theta / \sqrt{n} \leq 2$. The second and the third assumption are motivated by Matrix Chernoff bound. The reason for introducing these Matrix Chernoff-type assumption is, the goal is to bound the spectral norm of the sums of random matrices in several parts of the proof. One way is to relax the spectral norm to the Frobenious norm, and bound each entry of the matrix, and finally union bound over all entries in the matrix. This could potentially lose a $\sqrt{n}$ factor compared to applying Matrix Chernoff bound. We feel these assumptions can indicate how the input data affect the over-parameterization size $m$ in a more clear way.

We state our result for the concentration of sums of independent random matrices:

Proposition 28.1.2 (Informal of Theorem 28.5.1). Assume Part 1,2 and 3 of Assumption 28.1.1. If $m=\Omega\left(\left(\lambda^{-2} \beta+\lambda^{-1} \alpha\right) \log (n / \delta)\right)$, then

$$
\operatorname{Pr}_{w_{1}, \cdots, w_{m} \in N(0, I)}\left[\left\|H^{\mathrm{dis}}-H^{\mathrm{cts}}\right\|_{2} \leq \lambda / 4\right] \geq 1-\delta
$$

Table 28.1: Summary of Convergence Result

| Reference | $m$ | $\lambda$ | $\alpha$ | $\theta$ |
| :--- | :--- | :--- | :--- | :--- |
| $[$ DZPS19] | $\lambda^{-4} n^{6} \operatorname{poly}(\log n, 1 / \delta)$ | Yes | No | No |
| Theorem 28.1.3 | $\lambda^{-4} n^{4} \log ^{3}(n / \delta)$ | Yes | No | No |
| Theorem 28.1.4 | $\lambda^{-4} n^{3} \log ^{3}(n / \delta) \cdot \alpha$ | Yes | Yes | No |
| Theorem 28.1.5 | $\lambda^{-4} n^{2} \log ^{3}(n / \delta) \cdot \alpha\left(\alpha+\theta^{2}\right)$ | Yes | Yes | Yes |

Proposition 28.1.2 is a direct improvement compared to Lemma 3.1 in [DZPS19], which requires $m=\Omega\left(\lambda^{-2} n^{2} \log (n / \delta)\right)$. Proposition 28.1.2 is better when input data points have some good properties, e.g., $\beta, \alpha=o\left(n^{2}\right)$. However the result in [DZPS19] always needs to pay $n^{2}$ factor, no matter what the input data points are.

We state our convergence result as follows:

Theorem 28.1.3 (Informal of Theorem 28.4.5). Assume Part 1 of Assumption 28.1.1. Let $m$ denote the width of neural network, let $n$ denote the number of input data points. If $m=\Omega\left(\lambda^{-4} n^{4} \operatorname{poly}(\log (n / \delta))\right)$, then gradient descent is able to find the global minimum from a random initialization point with probability $1-\delta$.

Theorem 28.1.3 is a direct improvement compared to Theorem 4.1 in [DZPS19], which requires $m=\Omega\left(\lambda^{-4} n^{6} \operatorname{poly}(\log n, 1 / \delta)\right)$.

If we also allow Part 2 of Assumption 28.1.1, we can slightly improve Theorem 28.1.3 from $n^{4}$ to $n^{3}$,

Theorem 28.1.4 (Informal of Theorem 28.5.5). Assume Part 1 and 2 of Assumption 28.1.1. If $m=\Omega\left(\lambda^{-4} n^{3} \alpha \operatorname{poly}(\log (n / \delta))\right)$, then gradient descent is able to find the global minimum from a random initialization point with probability $1-\delta$.

Except for $m$, Theorem 4.1 in [DZPS19] requires step size $\eta$ to be $\Theta\left(\lambda / n^{2}\right)$. Theorem 28.1.4 only needs step size $\eta$ to be $\Theta(\lambda /(\alpha n))$.

Further, if we also allow Part 4 of Assumption 28.1.1, we can slightly improve Theorem 28.1.3 from $n^{4}$ to $n^{2}$,

Theorem 28.1.5 (Informal of Theorem 28.6.4). Assume Part 1, 2 and 4 of Assumption 28.1.1. If $m=\Omega\left(\lambda^{-4} n^{2} \alpha\left(\theta^{2}+\alpha\right) \operatorname{poly}(\log (n / \delta))\right)$, the gradient descent is able to find the global minimum from a random initialization point with probability $1-\delta$.

### 28.1.2 Technical Overview

We follow the exact same optimization framework as Du, Zhai, Poczos and Singh [DZPS19]. We improve the bound on $m$ by doing a careful concentration analysis for random variables without changing the high-level optimization framework.

We briefly summarize the optimization framework here: the minimal eigenvalue $\lambda$ of $H^{\text {cts }}$, as introduced in [DZPS19], turns out to be closely related with the convergence rate. As time evolves, the weights $w$ in the network may vary; however if $w$ stay in a ball of radius $R$ that only depends on the number of data $n$ and $\lambda$, and particularly does not depend on the number of neurons $m$, then we are still able to lower bound the minimal eigenvalue of $H(w)$. On the other hand, we want to upper bound $D$, the actual move of $w$, with high probability. It turns out $D$ is proportional to $\frac{1}{\sqrt{m}}$. We require $D<R$ in order to control the convergence rate. In this way we derive a lower bound of $m$.

In order to bound $\|H\|$, [DZPS19] relax it to Frobenius norm and then relax it to entry-wise L1 norm,

$$
\|H\| \leq\|H\|_{F} \leq\|H\|_{1} .
$$

Then they can bound each term of $H_{i, j}$ individually via Markov inequality.
One key observation is that $\|H\|_{1}$ is a quite loose bound for $\|H\|_{F}$, in the sense that $\|H\|_{1}=\|H\|_{F}$ holds only if $H$ contains at most 1 non-zero entry. This means we can work on the Frobenius norm directly, and we shall be able to obtain a tighter estimation. By definition of $H$, it can be written as a summation of $m$ independent matrices $A_{1}, A_{2}, \cdots, A_{m} \in \mathbb{R}^{n \times n}$,

$$
H=\frac{1}{m} \sum_{r=1}^{m} A_{r}
$$

In order to bound $\|H\|_{F}$, for each $i, j$, we regard each $H_{i j}$ as summation of $m$ independent random variables, then apply Bernstein bound to obtain experiential tail bound on the concentration of $H_{i j}$. Finally, by taking a union bound over all the $n^{2}$ pairs we obtain a tighter bound for $\|H\|_{F}$.

We shall mention that $\|H\|_{F}$ is also a loose upper bound of $\|H\|$, i.e., $\|H\|_{F}=\|H\|$ only if $H$ is a rank- 1 matrix. Hence, if the condition number of $H$ is small, which may happen as a property of the data, then we may benefit from bounding $\|H\|$ directly. We achieve this by apply matrix Chernoff bound, which states the spectral norm of summation of $m$ independent matrices concentrates under certain conditions.

We shall stress that mutually independence plays a very important role in our argument. Throughout the whole paper we are dealing with summations of the form $\sum_{r=1}^{m} y_{r}$ where $\left\{y_{m}\right\}_{r=1}^{m}$ are independent random variables. Previous argument mainly applies Markov inequality, which pays a factor of $1 / \delta$ around the mean for error probability $\delta$. But we can obtain much tighter concentration bound by taking advantage of independence as in Bernstein inequality and Hoeffding inequality. This allows us to improve the dependency on $\delta$ from $1 / \delta$ to $\log 1 / \delta$.

We also make use of matrix spectral norm to deal with summation of the form $\left\|\sum_{i=1}^{n} a_{i} x_{i}\right\|_{2}$ where $\left\{a_{i}\right\}_{i=1}^{n}$ are scalars and $\left\{x_{i}\right\}_{i=1}^{n}$ are vectors. Naively applying triangle inequality leads to an upper bound proportional to $\|a\|_{1}$, which can be as large as $\sqrt{n}\|a\|_{2}$. Instead, we observe that the matrix formed by $\left(\begin{array}{lll}x_{1} & \cdots & x_{n}\end{array}\right):=X$ has good singular value property, which allows us to obtain the bound $\|X\|_{2} \cdot\|a\|_{2}$. Therefore, this bound does not rely on number of inputs explicitly.

Roadmap We provide some basic definitions and probability tools in Section 28.2. We define the optimization problem in Section 28.3. We present our quartic result in Section 28.4. We improve it to cubic and quadratic in Section 28.5 and Section 28.6.

### 28.2 Preliminaries

### 28.2.1 Notation

We use $[n]$ to denote $\{1,2, \cdots, n\}$. We use $\phi$ to denote ReLU activation function, i.e., $\phi(x)=\max \{x, 0\}$. For an event $f(x)$, we define $\mathbf{1}_{f(x)}$ such that $\mathbf{1}_{f(x)}=1$ if $f(x)$ holds and $\mathbf{1}_{f(x)}=0$ otherwise. For a matrix $A$, we use $\|A\|$ to denote the spectral norm of $A$. We define $\|A\|_{F}=\left(\sum_{i} \sum_{j} A_{i, j}^{2}\right)^{1 / 2}$ and $\|A\|_{1}=\sum_{i} \sum_{j}\left|A_{i, j}\right|$.

### 28.3 Problem Formulation

Our problem formulation is the same as [DZPS19]. We consider a two-layer ReLU activated neural network with $m$ neurons in the hidden layer:

$$
f(W, x, a)=\frac{1}{\sqrt{m}} \sum_{r=1}^{m} a_{r} \phi\left(w_{r}^{\top} x\right)
$$

where $x \in \mathbb{R}^{d}$ is the input, $w_{1}, \cdots, w_{m} \in \mathbb{R}^{d}$ are weight vectors in the first layer, $a_{1}, \cdots, a_{m} \in$ $\mathbb{R}$ are weights in the second layer. For simplicity, we only optimize $W$ but not optimize $a$ and $W$ at the same time.

Recall that the ReLU function $\phi(x)=\max \{x, 0\}$. Therefore for $r \in[m]$, we have

$$
\begin{equation*}
\frac{f(W, x, a)}{\partial w_{r}}=\frac{1}{\sqrt{m}} a_{r} x \mathbf{1}_{w_{r}^{\top} x \geq 0} . \tag{28.1}
\end{equation*}
$$

We apply the gradient descent to optimize the weight matrix $W$ in the following standard way,

$$
\begin{equation*}
W(k+1)=W(k)-\eta \frac{\partial L(W(k))}{\partial W(k)} . \tag{28.2}
\end{equation*}
$$

We define objective function $L$ as follows

$$
L(W)=\frac{1}{2} \sum_{i=1}^{n}\left(y_{i}-f\left(W, x_{i}, a\right)\right)^{2}
$$

We can compute the gradient of $L$ in terms of $w_{r}$

$$
\begin{equation*}
\frac{\partial L(W)}{\partial w_{r}}=\frac{1}{\sqrt{m}} \sum_{i=1}^{n}\left(f\left(W, x_{i}, a_{r}\right)-y_{i}\right) a_{r} x_{i} \mathbf{1}_{w_{r}^{\top} x_{i} \geq 0} \tag{28.3}
\end{equation*}
$$

We consider the ordinary differential equation defined by

$$
\begin{equation*}
\frac{\mathrm{d} w_{r}(t)}{\mathrm{d} t}=-\frac{\partial L(W)}{\partial w_{r}} \tag{28.4}
\end{equation*}
$$

At time $t$, let $u(t)=\left(u_{1}(t), \cdots, u_{n}(t)\right) \in \mathbb{R}^{n}$ be the prediction vector where each $u_{i}(t)$ is defined as

$$
\begin{equation*}
u_{i}(t)=f\left(W(t), a, x_{i}\right) \tag{28.5}
\end{equation*}
$$

### 28.4 Quartic Suffices

### 28.4.1 Bounding the difference between continuous and discrete

Lemma 28.4.1 (Lemma 3.1 in [DZPS19]). We define $H^{\text {cts }}, H^{\text {dis }} \in \mathbb{R}^{n \times n}$ as follows

$$
\begin{aligned}
H_{i, j}^{\mathrm{cts}} & =\underset{w \sim \mathcal{N}(0, I)}{\mathbb{E}}\left[x_{i}^{\top} x_{j} \mathbf{1}_{w^{\top} x_{i} \geq 0, w^{\top} x_{j} \geq 0}\right], \\
H_{i, j}^{\mathrm{dis}} & =\frac{1}{m} \sum_{r=1}^{m}\left[x_{i}^{\top} x_{j} \mathbf{1}_{w_{r}^{\top} x_{i} \geq 0, w_{r}^{\top} x_{j} \geq 0}\right] .
\end{aligned}
$$

Let $\lambda=\lambda_{\min }\left(H^{\text {cts }}\right)$. If $m=\Omega\left(\lambda^{-2} n^{2} \log (n / \delta)\right)$, we have

$$
\left\|H^{\mathrm{dis}}-H^{\mathrm{cts}}\right\|_{F} \leq \frac{\lambda}{4}, \text { and } \lambda_{\min }\left(H^{\mathrm{dis}}\right) \geq \frac{3}{4} \lambda
$$

hold with probability at least $1-\delta$.

For the completeness, we still provide a proof here.

Proof. For every fixed pair $(i, j), H_{i, j}^{\text {dis }}$ is an average of independent random variables, i.e.

$$
H_{i, j}^{\mathrm{dis}}=\frac{1}{m} \sum_{r=1}^{m} x_{i}^{\top} x_{j} \mathbf{1}_{w_{r}^{\top} x_{i} \geq 0, w_{r}^{\top} x_{j} \geq 0} .
$$

Then the expectation of $H_{i, j}^{\text {dis }}$ is

$$
\begin{aligned}
\mathbb{E}\left[H_{i, j}^{\mathrm{dis}}\right] & =\frac{1}{m} \sum_{r=1}^{m} \underset{w_{r} \sim N\left(0, I_{d}\right)}{\mathbb{E}}\left[x_{i}^{\top} x_{j} \mathbf{1}_{w_{r}^{\top} x_{i} \geq 0, w_{r}^{\top} x_{j} \geq 0}\right] \\
& =\underset{w \sim N\left(0, I_{d}\right)}{\mathbb{E}}\left[x_{i}^{\top} x_{j} \mathbf{1}_{w^{\top} x_{i} \geq 0, w^{\top} x_{j} \geq 0}\right] \\
& =H_{i, j}^{\mathrm{cts}}
\end{aligned}
$$

 $\left\{z_{r}\right\}_{r \in[m]}$ are mutually independent. Moreover, $-\frac{1}{m} \leq z_{r} \leq \frac{1}{m}$. So by Hoeffding inequality(Lemma A.1.5) we have for all $t>0$,

$$
\operatorname{Pr}\left[\left|H_{i, j}^{\mathrm{dis}}-H_{i, j}^{\mathrm{cts}}\right| \geq t\right] \leq 2 \exp \left(-\frac{2 t^{2}}{4 / m}\right)=2 \exp \left(-m t^{2} / 2\right)
$$

Setting $t=\left(\frac{1}{m} 2 \log \left(2 n^{2} / \delta\right)\right)^{1 / 2}$, we can apply union bound on all pairs $(i, j)$ to get with probability at least $1-\delta$, for all $i, j \in[n]$,

$$
\left|H_{i, j}^{\mathrm{dis}}-H_{i, j}^{\mathrm{cts}}\right| \leq\left(\frac{2}{m} \log \left(2 n^{2} / \delta\right)\right)^{1 / 2} \leq 4\left(\frac{\log (n / \delta)}{m}\right)^{1 / 2} .
$$

Thus we have

$$
\left\|H^{\mathrm{dis}}-H^{\mathrm{cts}}\right\|^{2} \leq\left\|H^{\mathrm{dis}}-H^{\mathrm{cts}}\right\|_{F}^{2}=\sum_{i=1}^{n} \sum_{j=1}^{n}\left|H_{i, j}^{\mathrm{dis}}-H_{i, j}^{\mathrm{cts}}\right|^{2} \leq \frac{1}{m} 16 n^{2} \log (n / \delta)
$$

Hence if $m=\Omega\left(\lambda^{-2} n^{2} \log (n / \delta)\right)$ we have the desired result.

We define the event

$$
A_{i, r}=\left\{\exists u:\left\|u-\widetilde{w}_{r}\right\|_{2} \leq R, \mathbf{1}_{x_{i}^{\top} \tau} \tilde{w}_{r} \geq 0 \neq \mathbf{1}_{x_{i}^{\top} u \geq 0}\right\} .
$$

Note this event happens if and only if $\left|\widetilde{w}_{r}^{\top} x_{i}\right|<R$. Recall that $\widetilde{w}_{r} \sim \mathcal{N}(0, I)$. By anticoncentration inequality of Gaussian (Lemma A.1.11), we have

$$
\begin{equation*}
\operatorname{Pr}\left[A_{i, r}\right]=\underset{z \sim \mathcal{N}(0,1)}{\operatorname{Pr}}[|z|<R] \leq \frac{2 R}{\sqrt{2 \pi}} \tag{28.6}
\end{equation*}
$$

### 28.4.2 Bounding changes of $H$ when $w$ is in a small ball

We improve the Lemma 3.2 in [DZPS19] from the two perspective : one is the probability, and the other is upper bound on spectral norm.

Lemma 28.4.2 (perturbed $w)$. Let $R \in(0,1)$. If $\widetilde{w}_{1}, \cdots, \widetilde{w}_{m}$ are i.i.d. generated $\mathcal{N}(0, I)$. For any set of weight vectors $w_{1}, \cdots, w_{m} \in \mathbb{R}^{d}$ that satisfy for any $r \in[m],\left\|\widetilde{w}_{r}-w_{r}\right\|_{2} \leq R$, then the $H: \mathbb{R}^{m \times d} \rightarrow \mathbb{R}^{n \times n}$ defined

$$
H(w)_{i, j}=\frac{1}{m} x_{i}^{\top} x_{j} \sum_{r=1}^{m} \mathbf{1}_{w_{r}^{\top} x_{i} \geq 0, w_{r}^{\top} x_{j} \geq 0} .
$$

Then we have

$$
\|H(w)-H(\widetilde{w})\|_{F}<2 n R,
$$

holds with probability at least $1-n^{2} \cdot \exp (-m R / 10)$.

Proof. The random variable we care is

$$
\begin{aligned}
\sum_{i=1}^{n} \sum_{j=1}^{n}\left|H(\widetilde{w})_{i, j}-H(w)_{i, j}\right|^{2} & =\frac{1}{m^{2}} \sum_{i=1}^{n} \sum_{j=1}^{n}\left|x_{i}^{\top} x_{j} \sum_{r=1}^{m}\left(\mathbf{1}_{\widetilde{w}_{r}^{\top} x_{i} \geq 0, \widetilde{w}_{r}^{\top} x_{j} \geq 0}-\mathbf{1}_{w_{r}^{\top} x_{i} \geq 0, w_{r}^{\top} x_{j} \geq 0}\right)\right|^{2} \\
& \leq \frac{1}{m^{2}} \sum_{i=1}^{n} \sum_{j=1}^{n}\left(\sum_{r=1}^{m} \mathbf{1}_{\widetilde{w}_{r}^{\top} x_{i} \geq 0, \widetilde{w}_{r}^{\top} x_{j} \geq 0}-\mathbf{1}_{w_{r}^{\top} x_{i} \geq 0, w_{r}^{\top} x_{j} \geq 0}\right)^{2} \\
& =\frac{1}{m^{2}} \sum_{i=1}^{n} \sum_{j=1}^{n}\left(\sum_{r=1}^{m} s_{r, i, j}\right)^{2}
\end{aligned}
$$

where the last step follows from for each $r, i, j$, we define

$$
s_{r, i, j}:=\mathbf{1}_{\tilde{w}_{r}^{\top} x_{i} \geq 0, \widetilde{w}_{r}^{\top} x_{j} \geq 0}-\mathbf{1}_{w_{r}^{\top} x_{i} \geq 0, w_{r}^{\top} x_{j} \geq 0} .
$$

We consider $i, j$ are fixed. We simplify $s_{r, i, j}$ to $s_{r}$.

Then $s_{r}$ is a random variable that only depends on $\widetilde{w}_{r}$. Since $\left\{\widetilde{w}_{r}\right\}_{r=1}^{m}$ are independent, $\left\{s_{r}\right\}_{r=1}^{m}$ are also mutually independent.

If $\neg A_{i, r}$ and $\neg A_{j, r}$ happen, then

$$
\left|\mathbf{1}_{\tilde{w}_{r}^{\top} x_{i} \geq 0, \widetilde{w}_{r}^{\top} x_{j} \geq 0}-\mathbf{1}_{w_{r}^{\top} x_{i} \geq 0, w_{r}^{\top} x_{j} \geq 0}\right|=0 .
$$

If $A_{i, r}$ or $A_{j, r}$ happen, then

$$
\left|\mathbf{1}_{\widetilde{w}_{r}^{\top} x_{i} \geq 0, \widetilde{w}_{r}^{\top} x_{j} \geq 0}-\mathbf{1}_{w_{r}^{\top} x_{i} \geq 0, w_{r}^{\top} x_{j} \geq 0}\right| \leq 1 .
$$

So we have

$$
\underset{\widetilde{w}_{r}}{\mathbb{E}}\left[s_{r}\right] \leq \underset{\widetilde{w}_{r}}{\mathbb{E}}\left[\mathbf{1}_{A_{i, r} \vee A_{j, r}}\right] \leq \operatorname{Pr}\left[A_{i, r}\right]+\operatorname{Pr}\left[A_{j, r}\right] \leq \frac{4 R}{\sqrt{2 \pi}} \leq 2 R,
$$

and

$$
\underset{\widetilde{w}_{r}}{\mathbb{E}}\left[\left(s_{r}-\underset{\widetilde{w}_{r}}{\mathbb{E}}\left[s_{r}\right]\right)^{2}\right]=\underset{\widetilde{w}_{r}}{\mathbb{E}}\left[s_{r}^{2}\right]-\underset{\widetilde{w}_{r}}{\mathbb{E}}\left[s_{r}\right]^{2} \leq \underset{\widetilde{w}_{r}}{\mathbb{E}}\left[s_{r}^{2}\right] \leq \underset{\widetilde{w}_{r}}{\mathbb{E}}\left[\left(\mathbf{1}_{A_{i, r} \vee A_{j, r}}\right)^{2}\right] \leq \frac{4 R}{\sqrt{2 \pi}} \leq 2 R .
$$

We also have $\left|s_{r}\right| \leq 1$. So we can apply Bernstein inequality (Lemma A.1.2) to get for all $t>0$,

$$
\begin{aligned}
\operatorname{Pr}\left[\sum_{r=1}^{m} s_{r} \geq 2 m R+m t\right] & \leq \operatorname{Pr}\left[\sum_{r=1}^{m}\left(s_{r}-\mathbb{E}\left[s_{r}\right]\right) \geq m t\right] \\
& \leq \exp \left(-\frac{m^{2} t^{2} / 2}{2 m R+m t / 3}\right)
\end{aligned}
$$

Choosing $t=R$, we get

$$
\operatorname{Pr}\left[\sum_{r=1}^{m} s_{r} \geq 3 m R\right] \leq \exp \left(-\frac{m^{2} R^{2} / 2}{2 m R+m R / 3}\right) \leq \exp (-m R / 10)
$$

Thus, we can have

$$
\operatorname{Pr}\left[\frac{1}{m} \sum_{r=1}^{m} s_{r} \geq 2 R\right] \leq \exp (-m R / 10)
$$

Therefore, we complete the proof.

Table 28.2: Table of Parameters for the $m=\widetilde{\Omega}\left(n^{4}\right)$ result in Section 28.4. Nt. stands for notations.

| Nt. | Choice | Place | Comment |
| :--- | :--- | :--- | :--- |
| $\lambda$ | $:=\lambda_{\min }\left(H^{\text {cts }}\right)$ | Ass. 28.1.1 | Data-dependent |
| $R$ | $\lambda / n$ | Eq. $(28.11)$ | Maximal allowed movement of weight |
| $D_{\text {cts }}$ | $\frac{\sqrt{n}\\|y-u(0)\\|_{2}}{\sqrt{m} \lambda}$ | Lemma 28.4.3 | Actual moving distance of weight, continuous case |
| $D$ | $\frac{4 \sqrt{n}\\|y(0)\\|_{2}}{\sqrt{m} \lambda}$ | Lemma 28.4.6 | Actual moving distance of weight, discrete case |
| $\eta$ | $\lambda / n^{2}$ | Eq. $(28.11)$ | Step size of gradient descent |
| $m$ | $\lambda^{-2} n^{2} \log (n / \delta)$ | Lemma 28.4.1 | Bounding discrete and continuous |
| $m$ | $\lambda^{-4} n^{4} \log ^{3}(n / \delta)$ | Lemma 28.4.4 <br> Claim 28.4.7 | $D<R$ and $\\|y-u(0)\\|_{2}^{2}=\widetilde{O}(n)$ |

### 28.4.3 Loss is decreasing while weights are not changing much

For simplicity of notation, we provide the following definition.

Definition 28.4.1. For any $s \in[0, t]$, we define matrix $H(s) \in \mathbb{R}^{n \times n}$ as follows

$$
H(s)_{i, j}=\frac{1}{m} \sum_{r=1}^{m} x_{i}^{\top} x_{j} \mathbf{1}_{w_{r}(s)^{\top} x_{i} \geq 0, w_{r}(s)^{\top} x_{j} \geq 0} .
$$

With $H$ defined, it becomes more convenient to write the dynamics of predictions.

For each $i \in[n]$, we have

$$
\begin{aligned}
\frac{\mathrm{d}}{\mathrm{~d} t} u_{i}(t) & =\sum_{r=1}^{m}\left\langle\frac{\partial f\left(W(t), a, x_{i}\right)}{\partial w_{r}(t)}, \frac{\mathrm{d} w_{r}(t)}{\mathrm{d} t}\right\rangle \\
& =\sum_{r=1}^{m}\left\langle\frac{\partial f\left(W(t), a, x_{i}\right)}{\partial w_{r}(t)},-\frac{\partial L(w(t), a)}{\partial w_{r}(t)}\right\rangle \\
& =\sum_{r=1}^{m}\left\langle\frac{\partial f\left(W(t), a, x_{i}\right)}{\partial w_{r}(t)},-\frac{1}{\sqrt{m}} \sum_{i=1}^{n}\left(f\left(W, x_{i}, a_{r}\right)-y_{i}\right) a_{r} x_{i} \mathbf{1}_{w_{r}^{\top} x_{i} \geq 0}\right\rangle \\
& =\sum_{j=1}^{n}\left(y_{j}-u_{j}(t)\right)\left\langle\frac{\partial f\left(W(t), a, x_{i}\right)}{\partial w_{r}(t)}, \frac{\partial f\left(W(t), a, x_{j}\right)}{\partial w_{r}(t)}\right\rangle \\
& =\sum_{j=1}^{n}\left(y_{j}-u_{j}(t)\right) H(t)_{i, j}
\end{aligned}
$$

where the first step follows from (28.5) and the chain rule of derivatives, the second step uses (28.3), the third step uses (28.4), the fourth step uses (28.1) and (28.5), and the last step uses the definition of the matrix $H$.

Hence, we have the following compact expression for $\frac{\mathrm{d}}{\mathrm{d} t} u(t)$ as a whole vector:

$$
\frac{\mathrm{d}}{\mathrm{~d} t} u(t)=H(t) \cdot(y-u(t)) .
$$

Lemma 28.4.3 (Lemma 3.3 in [DZPS19]). Suppose for $0 \leq s \leq t, \lambda_{\min }(H(w(s))) \geq \lambda / 2$. Let $D_{\text {cts }}$ be defined as

$$
D_{\mathrm{cts}}:=\frac{\sqrt{n}\|y-u(0)\|_{2}}{\sqrt{m} \lambda} .
$$

Then we have

$$
\|y-u(t)\|_{2}^{2} \leq \exp (-\lambda t) \cdot\|y-u(0)\|_{2}^{2}
$$

and

$$
\left\|w_{r}(t)-w_{r}(0)\right\|_{2} \leq D_{\text {cts }} .
$$

For the completeness, we still provide a proof.

Proof. Recall we can write the dynamics of predictions as

$$
\frac{\mathrm{d}}{\mathrm{~d} t} u(t)=H(t) \cdot(y-u(t))
$$

We can calculate the loss function dynamics

$$
\begin{aligned}
\frac{\mathrm{d}}{\mathrm{~d} t}\|y-u(t)\|_{2}^{2} & =-2(y-u(t))^{\top} \cdot H(t) \cdot(y-u(t)) \\
& \leq-\lambda\|y-u(t)\|_{2}^{2}
\end{aligned}
$$

Thus we have $\frac{\mathrm{d}}{\mathrm{d} t}\left(\exp (\lambda t)\|y-u(t)\|_{2}^{2}\right) \leq 0$ and $\exp (\lambda t)\|y-u(t)\|_{2}^{2}$ is a decreasing function with respect to $t$.

Using this fact we can bound the loss

$$
\begin{equation*}
\|y-u(t)\|_{2}^{2} \leq \exp (-\lambda t)\|y-u(0)\|_{2}^{2} \tag{28.7}
\end{equation*}
$$

Now, we can bound the gradient norm. Recall for $0 \leq s \leq t$,

$$
\begin{align*}
\left\|\frac{\mathrm{d}}{\mathrm{~d} s} w_{r}(s)\right\|_{2} & =\left\|\sum_{i=1}^{n}\left(y_{i}-u_{i}\right) \frac{1}{\sqrt{m}} a_{r} x_{i} \cdot \mathbf{1}_{w_{r}(s)^{\top} x_{i} \geq 0}\right\|_{2} \\
& \leq \frac{1}{\sqrt{m}} \sum_{i=1}^{n}\left|y_{i}-u_{i}(s)\right| \\
& \leq \frac{\sqrt{n}}{\sqrt{m}}\|y-u(s)\|_{2}  \tag{28.8}\\
& \leq \frac{\sqrt{n}}{\sqrt{m}} \exp (-\lambda s)\|y-u(0)\|_{2} .
\end{align*}
$$

where the first step follows from (28.3), the second step follows from triangle inequality and $a_{r}= \pm 1$ for $r \in[m]$ and $\left\|x_{i}\right\|_{2}=1$ for $i \in[n]$, the third step follows from Cauchy-Schwartz inequality, and the last step follows from (28.7).

Integrating the gradient, we can bound the distance from the initialization

$$
\begin{aligned}
\left\|w_{r}(t)-w_{r}(0)\right\|_{2} & \leq \int_{0}^{t}\left\|\frac{\mathrm{~d}}{\mathrm{~d} s} w_{r}(s)\right\|_{2} \mathrm{~d} s \\
& \leq \frac{\sqrt{n}\|y-u(0)\|_{2}}{\sqrt{m} \lambda}
\end{aligned}
$$

Lemma 28.4.4 (Lemma 3.4 in [DZPS19]). If $D_{\text {cts }}<R$. then for all $t \geq 0, \lambda_{\min }(H(t)) \geq \frac{1}{2} \lambda$. Moreover, for all $r \in[m]$,

$$
\left\|w_{r}(t)-w_{r}(0)\right\| \leq D_{\mathrm{cts}}
$$

and

$$
\|y-u(t)\|_{2}^{2} \leq \exp (-\lambda t) \cdot\|y-u(0)\|_{2}^{2}
$$

For the completeness, we still provide a proof.

Proof. Assume the conclusion does not hold at time $t$. We argue there must be some $s \leq t$ so that $\lambda_{\min }(H(s))<\frac{1}{2} \lambda$.

If $\lambda_{\min }(H(t))<\frac{1}{2} \lambda$, then we can simply take $s=t$.
Otherwise since the conclusion does not hold, there exists $r$ so that

$$
\left\|w_{r}(t)-w_{r}(0)\right\| \geq D_{\mathrm{cts}} \quad \text { or } \quad\|y-u(t)\|_{2}^{2}>\exp (-\lambda t)\|y-u(0)\|_{2}^{2}
$$

Then by Lemma 28.4.3, there exists $s \leq t$ such that

$$
\lambda_{\min }(H(s))<\frac{1}{2} \lambda
$$

By Lemma 28.4.2, there exists $t_{0}>0$ defined as

$$
t_{0}=\inf \left\{t>0: \max _{r \in[m]}\left\|w_{r}(t)-w_{r}(0)\right\|_{2}^{2} \geq R\right\}
$$

Thus at time $t_{0}$, there exists $r \in[m]$ satisfying $\left\|w_{r}\left(t_{0}\right)-w_{r}(0)\right\|_{2}^{2}=R$.
By Lemma 28.4.2,

$$
\lambda_{\min }\left(H\left(t^{\prime}\right)\right) \geq \frac{1}{2} \lambda, \forall t^{\prime} \leq t_{0}
$$

However, by Lemma 28.4.3, this implies

$$
\left\|w_{r}\left(t_{0}\right)-w_{r}(0)\right\|_{2} \leq D_{\mathrm{cts}}<R
$$

which is a contradiction.

### 28.4.4 Convergence

Theorem 28.4.5. Recall that $\lambda=\lambda_{\min }\left(H^{\text {cts }}\right)>0$. Let $m=\Omega\left(\lambda^{-4} n^{4} \log (n / \delta)\right)$, we i.i.d. initialize $w_{r} \in \mathcal{N}(0, I)$, ar sampled from $\{-1,+1\}$ uniformly at random for $r \in[m]$, and we set the step size $\eta=O\left(\lambda / n^{2}\right)$ then with probability at least $1-\delta$ over the random initialization we have for $k=0,1,2, \cdots$

$$
\begin{equation*}
\|u(k)-y\|_{2}^{2} \leq(1-\eta \lambda / 2)^{k} \cdot\|u(0)-y\|_{2}^{2} \tag{28.9}
\end{equation*}
$$

Correctness We prove Theorem 28.4.5 by induction. The base case is $i=0$ and it is trivially true. Assume for $i=0, \cdots, k$ we have proved (28.9) to be true. We want to show (28.9) holds for $i=k+1$.

From the induction hypothesis, we have the following Lemma stating that the weights should not change too much.

Lemma 28.4.6 (Corollary 4.1 in [DZPS19]). If (28.9) holds for $i=0, \cdots, k$, then we have for all $r \in[m]$

$$
\left\|w_{r}(k+1)-w_{r}(0)\right\|_{2} \leq \frac{4 \sqrt{n}\|y-u(0)\|_{2}}{\sqrt{m} \lambda}:=D
$$

For the completeness, we still provide the proof

Proof. We use the norm of gradient to bound this distance,

$$
\begin{aligned}
\left\|w_{r}(k+1)-w_{r}(0)\right\|_{2} & \leq \eta \sum_{i=0}^{k}\left\|\frac{\partial L(W(i))}{\partial w_{r}(i)}\right\|_{2} \\
& \leq \eta \sum_{i=0}^{k} \frac{\sqrt{n}\|y-u(i)\|_{2}}{\sqrt{m}} \\
& \leq \eta \sum_{i=0}^{k} \frac{\sqrt{n}(1-\eta \lambda / 2)^{i / 2}}{\sqrt{m}}\|y-u(0)\|_{2} \\
& \leq \eta \sum_{i=0}^{\infty} \frac{\sqrt{n}(1-\eta \lambda / 2)^{i / 2}}{\sqrt{m}}\|y-u(0)\|_{2} \\
& =\frac{4 \sqrt{n}\|y-u(0)\|_{2}}{\sqrt{m} \lambda}
\end{aligned}
$$

where the first step follows from (28.2), the second step follows from (28.8), the third step follows from the induction hypothesis, the fourth step relaxes the summation to an infinite summation, and the last step follows from $\sum_{i=0}^{\infty}(1-\eta \lambda / 2)^{i / 2}=\frac{2}{\eta \lambda}$.

Thus, we complete the proof.

Next, we calculate the different of predictions between two consecutive iterations, analogue to $\frac{\mathrm{d} u_{i}(t)}{\mathrm{d} t}$ term in Lemma 28.4.3. For each $i \in[n]$, we have

$$
\begin{aligned}
u_{i}(k+1)-u_{i}(k) & =\frac{1}{\sqrt{m}} \sum_{r=1}^{m} a_{r} \cdot\left(\phi\left(w_{r}(k+1)^{\top} x_{i}\right)-\phi\left(w_{r}(k)^{\top} x_{i}\right)\right) \\
& =\frac{1}{\sqrt{m}} \sum_{r=1} a_{r} \cdot\left(\phi\left(\left(w_{r}(k)-\eta \frac{\partial L(W(k))}{\partial w_{r}(k)}\right)^{\top} x_{i}\right)-\phi\left(w_{r}(k)^{\top} x_{i}\right)\right) .
\end{aligned}
$$

Here we divide the right hand side into two parts. $v_{1, i}$ represents the terms that the pattern does not change and $v_{2, i}$ represents the term that pattern may changes. For each
$i \in[n]$, we define $v_{1, i}$ and $v_{2, i}$ as follows

$$
\begin{aligned}
& v_{1, i}:=\frac{1}{\sqrt{m}} \sum_{r \in S_{i}} a_{r} \cdot\left(\phi\left(\left(w_{r}(k)-\eta \frac{\partial L(W(k))}{\partial w_{r}(k)}\right)^{\top} x_{i}\right)-\phi\left(w_{r}(k)^{\top} x_{i}\right)\right) \\
& v_{2, i}:=\frac{1}{\sqrt{m}} \sum_{r \in \bar{S}_{i}} a_{r} \cdot\left(\phi\left(\left(w_{r}(k)-\eta \frac{\partial L(W(k))}{\partial w_{r}(k)}\right)^{\top} x_{i}\right)-\phi\left(w_{r}(k)^{\top} x_{i}\right)\right)
\end{aligned}
$$

Thus, we can rewrite $u(k+1)-u(k) \in \mathbb{R}^{n}$ in the following sense

$$
u(k+1)-u(k)=v_{1}+v_{2}
$$

In order to analyze $v_{1} \in \mathbb{R}^{n}$, we provide definition of $H$ and $H^{\perp} \in \mathbb{R}^{n \times n}$ first,

$$
\begin{aligned}
& H(k)_{i, j}=\frac{1}{m} \sum_{r=1}^{m} x_{i}^{\top} x_{j} \mathbf{1}_{w_{r}(k)^{\top} x_{i} \geq 0, w_{r}(k)^{\top} x_{j} \geq 0} \\
& H(k)_{i, j}^{\perp}=\frac{1}{m} \sum_{r \in \bar{S}_{i}} x_{i}^{\top} x_{j} \mathbf{1}_{w_{r}(k)^{\top} x_{i} \geq 0, w_{r}(k)^{\top} x_{j} \geq 0}
\end{aligned}
$$

Then, we can rewrite $v_{1, i} \in \mathbb{R}$

$$
\begin{aligned}
v_{1, i} & =-\frac{\eta}{m} \sum_{j=1}^{n} x_{i}^{\top} x_{j}\left(u_{j}-y_{j}\right) \sum_{r \in S_{i}} \mathbf{1}_{w_{r}(k)^{\top} x_{i} \geq 0, w_{r}(k)^{\top} x_{j} \geq 0} \\
& =-\eta \sum_{j=1}^{n}\left(u_{j}-y_{j}\right)\left(H_{i, j}(k)-H_{i, j}^{\perp}(k)\right),
\end{aligned}
$$

which means vector $v_{1} \in \mathbb{R}^{n}$ can be written as

$$
\begin{equation*}
v_{1}=\eta(y-u(k))^{\top}\left(H(k)-H^{\perp}(k)\right) . \tag{28.10}
\end{equation*}
$$

We are ready to prove the induction hypothesis. We can rewrite $\|y-u(k+1)\|_{2}^{2}$ as follows:

$$
\begin{aligned}
\|y-u(k+1)\|_{2}^{2} & =\|y-u(k)-(u(k+1)-u(k))\|_{2}^{2} \\
& =\|y-u(k)\|_{2}^{2}-2(y-u(k))^{\top}(u(k+1)-u(k))+\|u(k+1)-u(k)\|_{2}^{2}
\end{aligned}
$$

We can rewrite the second term in the above Equation in the following sense,

$$
\begin{aligned}
& (y-u(k))^{\top}(u(k+1)-u(k)) \\
= & (y-u(k))^{\top}\left(v_{1}+v_{2}\right) \\
= & (y-u(k))^{\top} v_{1}+(y-u(k))^{\top} v_{2} \\
= & \eta(y-u(k))^{\top} H(k)(y-u(k))-\eta(y-u(k))^{\top} H(k)^{\perp}(y-u(k))+(y-u(k))^{\top} v_{2},
\end{aligned}
$$

where the third step follows from Eq. (28.10).
We define

$$
\begin{aligned}
& C_{1}=-2 \eta(y-u(k))^{\top} H(k)(y-u(k)) \\
& C_{2}=2 \eta(y-u(k))^{\top} H(k)^{\perp}(y-u(k)) \\
& C_{3}=-2(y-u(k))^{\top} v_{2} \\
& C_{4}=\|u(k+1)-u(k)\|_{2}^{2}
\end{aligned}
$$

Thus, we have

$$
\begin{aligned}
\|y-u(k+1)\|_{2}^{2} & =\|y-u(k)\|_{2}^{2}+C_{1}+C_{2}+C_{3}+C_{4} \\
& \leq\|y-u(k)\|_{2}^{2}\left(1-\eta \lambda+8 \eta n R+8 \eta n R+\eta^{2} n^{2}\right)
\end{aligned}
$$

where the last step follows from Claim 28.4.8, 28.4.9, 28.4.10 and 28.4.11, which we will prove given later.

Choice of $\eta$ and $R$. Next, we want to choose $\eta$ and $R$ such that

$$
\begin{equation*}
\left(1-\eta \lambda+8 \eta n R+8 \eta n R+\eta^{2} n^{2}\right) \leq(1-\eta \lambda / 2) \tag{28.11}
\end{equation*}
$$

If we set $\eta=\frac{\lambda}{4 n^{2}}$ and $R=\frac{\lambda}{64 n}$, we have

$$
8 \eta n R+8 \eta n R=16 \eta n R \leq \eta \lambda / 4, \quad \text { and } \quad \eta^{2} n^{2} \leq \eta \lambda / 4
$$

This implies

$$
\|y-u(k+1)\|_{2}^{2} \leq\|y-u(k)\|_{2}^{2} \cdot(1-\eta \lambda / 2)
$$

holds with probability at least $1-2 n \exp (-m R)$.

Over-parameterization size, lower bound on $m$. We require

$$
D=\frac{4 \sqrt{n}\|y-u(0)\|_{2}}{\sqrt{m} \lambda}<R=\frac{\lambda}{64 n}
$$

and

$$
2 n \exp (-m R) \leq \delta
$$

By Claim 28.4.7, it is sufficient to choose $m=\Omega\left(\lambda^{-4} n^{4} \log (m / \delta) \log ^{2}(n / \delta)\right)$.

### 28.4.5 Technical claims

Claim 28.4.7. For $0<\delta<1$, with probability at least $1-\delta$,

$$
\|y-u(0)\|_{2}^{2}=O\left(n \log (m / \delta) \log ^{2}(n / \delta)\right)
$$

Proof.

$$
\begin{aligned}
\|y-u(0)\|_{2}^{2} & =\sum_{i=1}^{n}\left(y_{i}-f\left(W(0), a, x_{i}\right)\right)^{2} \\
& =\sum_{i=1}^{n}\left(y_{i}-\frac{1}{\sqrt{m}} \sum_{r=1}^{m} a_{r} \phi\left(w_{r}^{\top} x_{i}\right)\right)^{2} \\
& =\sum_{i=1}^{n} y_{i}^{2}-2 \sum_{i=1}^{n} \frac{y_{i}}{\sqrt{m}} \sum_{r=1}^{m} a_{r} \phi\left(w_{r}^{\top} x_{i}\right)+\sum_{i=1}^{n} \frac{1}{m}\left(\sum_{r=1}^{m} a_{r} \phi\left(w_{r}^{\top} x_{i}\right)\right)^{2} .
\end{aligned}
$$

Fix $r \in[m]$ and $i \in[n]$. Since $w_{r} \sim N(0, I)$ and $\left\|x_{i}\right\|_{2}=1, w_{r}^{\top} x_{i}$ follows distribution $N(0,1)$. From concentration of Gaussian distribution, we have

$$
\underset{w_{r}}{\operatorname{Pr}}\left[w_{r}^{\top} x_{i} \geq \sqrt{2 \log (2 m n / \delta)}\right] \leq \frac{\delta}{2 m n} .
$$

Let $E_{1}$ be the even that for all $r \in[m]$ and $i \in[n]$ we have

$$
\phi\left(w_{r}^{\top} x_{i}\right) \leq \sqrt{2 \log (2 m n / \delta)}
$$

Then by union bound, $\operatorname{Pr}\left[E_{1}\right] \geq 1-\frac{\delta}{2}$,
Fix $i \in[n]$. For every $r \in[m]$, we define random variable $z_{i, r}$ as

$$
z_{i, r}:=\frac{1}{\sqrt{m}} \cdot a_{r} \cdot \phi\left(w_{r}^{\top} x_{i}\right) \cdot \mathbf{1}_{w_{r}^{\top} x_{i} \leq \sqrt{2 \log (2 m n / \delta)}}
$$

Then $z_{i, r}$ only depends on $a_{r} \in\{-1,1\}$ and $w_{r} \sim N(0, I)$. Notice that $\mathbb{E}_{a_{r}, w_{r}}\left[z_{i, r}\right]=0$, and $\left|z_{i, r}\right| \leq \sqrt{2 \log (2 m n / \delta)}$. Moreover,

$$
\begin{aligned}
\underset{a_{r}, w_{r}}{\mathbb{E}}\left[z_{i, r}^{2}\right] & =\underset{a_{r}, w_{r}}{\mathbb{E}}\left[\frac{1}{m} a_{r}^{2} \phi^{2}\left(w_{r}^{\top} x_{i}\right) \mathbf{1}_{w_{r}^{\top} x_{i} \leq \sqrt{2 \log (2 m n / \delta)}}^{2}\right] \\
& =\frac{1}{m} \underset{a_{r}}{\mathbb{E}}\left[a_{r}^{2}\right] \cdot \underset{w_{r}}{\mathbb{E}}\left[\phi^{2}\left(w_{r}^{\top} x_{i}\right) \mathbf{1}_{w_{r}^{\top} x_{i} \leq \sqrt{2 \log (2 m n / \delta)}}^{2}\right] \\
& \leq \frac{1}{m} \cdot 1 \cdot \underset{w_{r}}{\mathbb{E}}\left[\left(w_{r}^{\top} x_{i}\right)^{2}\right] \\
& =\frac{1}{m}
\end{aligned}
$$

where the second step uses independence between $a_{r}$ and $w_{r}$, the third step uses $a_{r} \in\{-1,1\}$ and $\phi(t)=\max \{t, 0\}$, and the last step follows from $w_{r}^{\top} x_{i} \sim N(0,1)$.

Now we are ready to apply Bernstein inequality (Lemma A.1.2) to get for all $t>0$,

$$
\operatorname{Pr}\left[\sum_{r=1}^{m} z_{i, r}>t\right] \leq \exp \left(-\frac{t^{2} / 2}{m \cdot \frac{1}{m}+\sqrt{2 \log (2 m n / \delta)} \cdot t / 3}\right) .
$$

Setting $t=\sqrt{2 \log (2 m n / \delta)} \cdot \log (4 n / \delta)$, we have with probability at least $1-\frac{\delta}{4 n}$,

$$
\sum_{r=1}^{m} z_{i, r} \leq \sqrt{2 \log (2 m n / \delta)} \cdot \log (4 n / \delta)
$$

Notice that we can also apply Bernstein inequality (Lemma A.1.2) on $-z_{i, r}$ to get

$$
\operatorname{Pr}\left[\sum_{r=1}^{m} z_{i, r}<-t\right] \leq \exp \left(-\frac{t^{2} / 2}{m \cdot \frac{1}{m}+\sqrt{2 \log (2 m n / \delta)} \cdot t / 3}\right) .
$$

Let $E_{2}$ be the event that for all $i \in[n]$,

$$
\left|\sum_{r=1}^{m} z_{i, r}\right| \leq \sqrt{2 \log (2 m n / \delta)} \cdot \log (4 n / \delta)
$$

By applying union bound on all $i \in[n]$, we have $\operatorname{Pr}\left[E_{2}\right] \geq 1-\frac{\delta}{2}$.

If both $E_{1}$ and $E_{2}$ happen, we have

$$
\begin{aligned}
\|y-u(0)\|_{2}^{2} & =\sum_{i=1}^{n} y_{i}^{2}-2 \sum_{i=1}^{n} \frac{y_{i}}{\sqrt{m}} \sum_{r=1}^{m} a_{r} \phi\left(w_{r}^{\top} x_{i}\right)+\sum_{i=1}^{n} \frac{1}{m}\left(\sum_{r=1}^{m} a_{r} \phi\left(w_{r}^{\top} x_{i}\right)\right)^{2} \\
& =\sum_{i=1}^{n} y_{i}^{2}-2 \sum_{i=1}^{n} y_{i} \sum_{r=1}^{m} z_{i, r}+\sum_{i=1}^{n}\left(\sum_{r=1}^{m} z_{i, r}\right)^{2} \\
& \leq \sum_{i=1}^{n} y_{i}^{2}+2 \sum_{i=1}^{n}\left|y_{i}\right| \sqrt{2 \log (2 m n / \delta)} \cdot \log (4 n / \delta)+\sum_{i=1}^{n}(\sqrt{2 \log (2 m n / \delta)} \cdot \log (4 n / \delta))^{2} \\
& =O\left(n \log (m / \delta) \log ^{2}(n / \delta)\right),
\end{aligned}
$$

where the second step uses $E_{1}$, the third step uses $E_{2}$, and the last step follows from $\left|y_{i}\right|=$ $O(1), \forall i \in[n]$.

By union bound, this will happen with probability at least $1-\delta$.

Claim 28.4.8. Let $C_{1}=-2 \eta(y-u(k))^{\top} H(k)(y-u(k))$. We have

$$
C_{1} \leq-\|y-u(k)\|_{2}^{2} \cdot \eta \lambda
$$

Proof. By Lemma 28.4.2 and our choice of $R<\frac{\lambda}{8 n}$, We have $\|H(0)-H(k)\|_{F} \leq 2 n \cdot \frac{\lambda}{8 n}=\frac{\lambda}{4}$. Recall that $\lambda=\lambda_{\min }(H(0))$. Therefore

$$
\lambda_{\min }(H(k)) \geq \lambda_{\min }(H(0))-\|H(0)-H(k)\| \geq \lambda / 2 .
$$

Then we have

$$
(y-u(k))^{\top} H(k)(y-u(k)) \geq\|y-u(k)\|_{2}^{2} \cdot \lambda / 2
$$

Thus, we complete the proof.

Claim 28.4.9. Let $C_{2}=2 \eta(y-u(k))^{\top} H(k)^{\perp}(y-u(k))$. We have

$$
C_{2} \leq\|y-u(k)\|_{2}^{2} \cdot 8 \eta n R
$$

holds with probability $1-n \exp (-m R)$.

Proof. Note that

$$
C_{2} \leq 2 \eta\|y-u(k)\|_{2}^{2}\left\|H(k)^{\perp}\right\| .
$$

It suffices to upper bound $\left\|H(k)^{\perp}\right\|$. Since $\|\cdot\| \leq\|\cdot\|_{F}$, then it suffices to upper bound $\|\cdot\|_{F}$.

For each $i \in[n]$, we define $y_{i}$ as follows

$$
y_{i}=\sum_{r=1}^{m} \mathbf{1}_{r \in \bar{S}_{i}} .
$$

Then we have

$$
\begin{aligned}
\left\|H(k)^{\perp}\right\|_{F}^{2} & =\sum_{i=1}^{n} \sum_{j=1}^{n}\left(H(k)_{i, j}^{\perp}\right)^{2} \\
& =\sum_{i=1}^{n} \sum_{j=1}^{n}\left(\frac{1}{m} \sum_{r \in \bar{S}_{i}} x_{i}^{\top} x_{j} \mathbf{1}_{w_{r}(k)^{\top} x_{i} \geq 0, w_{r}(k)^{\top} x_{j} \geq 0}\right)^{2} \\
& =\sum_{i=1}^{n} \sum_{j=1}^{n}\left(\frac{1}{m} \sum_{r=1}^{m} x_{i}^{\top} x_{j} \mathbf{1}_{w_{r}(k)^{\top} x_{i} \geq 0, w_{r}(k)^{\top} x_{j} \geq 0} \cdot \mathbf{1}_{r \in \bar{S}_{i}}\right)^{2} \\
& =\sum_{i=1}^{n} \sum_{j=1}^{n}\left(\frac{x_{i}^{\top} x_{j}}{m}\right)^{2}\left(\sum_{r=1}^{m} \mathbf{1}_{w_{r}(k)^{\top} x_{i} \geq 0, w_{r}(k)^{\top} x_{j} \geq 0} \cdot \mathbf{1}_{r \in \bar{S}_{i}}\right)^{2} \\
& \leq \frac{1}{m^{2}} \sum_{i=1}^{n} \sum_{j=1}^{n}\left(\sum_{r=1}^{m} \mathbf{1}_{w_{r}(k)^{\top} x_{i} \geq 0, w_{r}(k)^{\top} x_{j} \geq 0} \cdot \mathbf{1}_{r \in \bar{S}_{i}}\right)^{2} \\
& =\frac{n}{m^{2}} \sum_{i=1}^{n}\left(\sum_{r=1}^{m} \mathbf{1}_{r \in \bar{S}_{i}}\right)^{2} \\
& =\frac{n}{m^{2}} \sum_{i=1}^{n} y_{i}^{2} .
\end{aligned}
$$

Fix $i \in[n]$. The plan is to use Bernstein inequality to upper bound $y_{i}$ with high probability.

First by Eq. (28.6) we have

$$
\mathbb{E}\left[\mathbf{1}_{r \in \bar{S}_{i}}\right] \leq R .
$$

We also have

$$
\begin{aligned}
\mathbb{E}\left[\left(\mathbf{1}_{r \in \bar{S}_{i}}-\mathbb{E}\left[\mathbf{1}_{r \in \bar{S}_{i}}\right]\right)^{2}\right] & =\mathbb{E}\left[\mathbf{1}_{r \in \bar{S}_{i}}^{2}\right]-\mathbb{E}\left[\mathbf{1}_{r \in \bar{S}_{i}}\right]^{2} \\
& \leq \mathbb{E}\left[\mathbf{1}_{r \in \bar{S}_{i}}^{2}\right] \\
& \leq R .
\end{aligned}
$$

Finally we have $\left|\mathbf{1}_{r \in \bar{S}_{i}}-\mathbb{E}\left[\mathbf{1}_{r \in \bar{S}_{i}}\right]\right| \leq 1$.
Notice that $\left\{\mathbf{1}_{r \in \bar{S}_{i}}\right\}_{r=1}^{m}$ are mutually independent, since $\mathbf{1}_{r \in \bar{S}_{i}}$ only depends on $w_{r}(0)$. Hence from Bernstein inequality (Lemma A.1.2) we have for all $t>0$,

$$
\operatorname{Pr}\left[y_{i}>m \cdot R+t\right] \leq \exp \left(-\frac{t^{2} / 2}{m \cdot R+t / 3}\right) .
$$

By setting $t=3 m R$, we have

$$
\operatorname{Pr}\left[y_{i}>4 m R\right] \leq \exp (-m R)
$$

Hence by union bound, with probability at least $1-n \exp (-m R)$,

$$
\left\|H(k)^{\perp}\right\|_{F}^{2} \leq \frac{n}{m^{2}} \cdot n \cdot(4 m R)^{2}=16 n^{2} R^{2}
$$

Putting all together we have

$$
\left\|H(k)^{\perp}\right\| \leq\left\|H(k)^{\perp}\right\|_{F} \leq 4 n R
$$

with probability at least $1-n \exp (-m R)$.

Claim 28.4.10. Let $C_{3}=-2(y-u(k))^{\top} v_{2}$. Then we have

$$
C_{3} \leq\|y-u(k)\|_{2}^{2} \cdot 8 \eta n R
$$

with probability at least $1-n \exp (-m R)$.

Proof. We have

$$
\text { LHS } \leq 2\|y-u(k)\|_{2} \cdot\left\|v_{2}\right\|_{2} .
$$

We can upper bound $\left\|v_{2}\right\|_{2}$ in the following sense

$$
\begin{aligned}
\left\|v_{2}\right\|_{2}^{2} & \leq \sum_{i=1}^{n}\left(\frac{\eta}{\sqrt{m}} \sum_{r \in \bar{S}_{i}}\left|\left(\frac{\partial L(W(k))}{\partial w_{r}(k)}\right)^{\top} x_{i}\right|\right)^{2} \\
& =\frac{\eta^{2}}{m} \sum_{i=1}^{n}\left(\sum_{r=1}^{m} \mathbf{1}_{r \in \bar{S}_{i}}\left|\left(\frac{\partial L(W(k))}{\partial w_{r}(k)}\right)^{\top} x_{i}\right|\right)^{2} \\
& \leq \frac{\eta^{2}}{m} \cdot \max _{r \in[m]}\left|\frac{\partial L(W(k))}{\partial w_{r}(k)}\right|^{2} \cdot \sum_{i=1}^{n}\left(\sum_{r=1}^{m} \mathbf{1}_{r \in \bar{S}_{i}}\right)^{2} \\
& \leq \frac{\eta^{2}}{m} \cdot\left(\frac{\sqrt{n}}{\sqrt{m}}\|u(k)-y\|_{2}\right)^{2} \cdot \sum_{i=1}^{n}\left(\sum_{r=1}^{m} \mathbf{1}_{r \in \bar{S}_{i}}\right)^{2} \\
& \leq \frac{\eta^{2}}{m} \cdot\left(\frac{\sqrt{n}}{\sqrt{m}}\|u(k)-y\|_{2}\right)^{2} \cdot \sum_{i=1}^{n}(4 m R)^{2} \\
& =16 n^{2} R^{2} \eta^{2}\|u(k)-y\|_{2}^{2},
\end{aligned}
$$

where the first step follows from definition of $v_{2}$, the fourth step follows from $\max _{r \in[m]}\left|\frac{\partial L(W(k))}{\partial w_{r}(k)}\right| \leq$ $\frac{\sqrt{n}}{\sqrt{m}} \cdot\|u(k)-y\|_{2}$, the fifth step follows from $\sum_{r=1}^{m} \mathbf{1}_{r \in \bar{S}_{i}} \leq 4 m R$ with probability at least $1-\exp (-m R)$.

Claim 28.4.11. Let $C_{4}=\|u(k+1)-u(k)\|_{2}^{2}$. Then we have

$$
C_{4} \leq \eta^{2} n^{2}\|y-u(k)\|_{2}^{2} .
$$

Proof. We have

$$
\begin{aligned}
\mathrm{LHS} & \leq \eta^{2} \sum_{i=1}^{n} \frac{1}{m}\left(\sum_{r=1}^{m}\left\|\frac{\partial L(W(k))}{\partial w_{r}(k)}\right\|_{2}\right)^{2} \\
& \leq \eta^{2} n^{2}\|y-u(k)\|_{2}^{2}
\end{aligned}
$$

### 28.5 Cubic Suffices

We prove a more general version of Lemma 28.4.1 in this section.
Theorem 28.5.1 (Data-dependent version, bounding the difference between discrete and continuous). Let $H^{\text {cts }}$ and $H^{\text {dis }}$ be defined as Definition 28.1.1. Let $\lambda, \alpha, \beta$ be satisfied Assumption 28.1.1. If

$$
m=\Omega\left(\left(\lambda^{-2} \beta+\lambda^{-1} \alpha\right) \log (n / \delta)\right)
$$

we have

$$
\left\|H^{\mathrm{dis}}-H^{\mathrm{cts}}\right\|_{2} \leq \lambda / 4, \text { and } \lambda_{\min }\left(H^{\mathrm{dis}}\right) \geq \frac{3}{4} \lambda
$$

holds with probability at least $1-\exp (-\Omega(\log (n / \delta)))$.

Proof. Recall the definition, we know

$$
H^{\mathrm{cts}}=\underset{w}{\mathbb{E}}[H(w)], \quad \text { and } \quad H^{\mathrm{dis}}=\frac{1}{m} \sum_{r=1}^{m} H\left(w_{r}\right)
$$

We define matrix $Y_{r}=H\left(w_{r}\right)-\mathbb{E}_{w}[H(w)]$. We know that, $Y_{r}$ are all independent,

$$
\mathbb{E}\left[Y_{r}\right]=0, \quad\left\|Y_{r}\right\| \leq \alpha, \quad\left\|\sum_{r=1}^{m} \mathbb{E}\left[Y_{r} Y_{r}^{\top}\right]\right\| \leq m \beta .
$$

Let $Y=\sum_{r=1}^{m} Y_{r}$. We apply Matrix Bernstein inequality (Lemma A.2.2) with $t=\sqrt{m \beta \log (n / \delta)}+$ $\alpha \log (n / \delta)$,

$$
\begin{aligned}
\operatorname{Pr}[\|Y\| \geq t] & \leq 2 n \exp \left(-\frac{t^{2} / 2}{m \beta+\alpha t / 3}\right) \\
& \leq 2 n \exp (-\log (n / \delta)) \\
& \leq \exp (-\Omega(\log (n / \delta)))
\end{aligned}
$$

Thus, we have

$$
\operatorname{Pr}\left[\left\|\frac{1}{m} \sum_{r=1}^{m} Y_{r}\right\| \geq \frac{1}{m}(\sqrt{m \beta \log (n / \delta)}+\alpha \log (n / \delta))\right] \leq \exp (-\Omega(\log (n / \delta)))
$$

In order to guarantee that $\frac{1}{m}(\sqrt{m \beta \log (n / \delta)}+\alpha \log (n / \delta)) \leq \lambda$, we need

$$
\sqrt{m} \geq \lambda^{-1} \sqrt{\beta \log (n / \delta)}
$$

when the first term is the dominated one; we need

$$
m \geq \lambda^{-1} \alpha \log (n / \delta)
$$

Overall, we need

$$
m \geq \Omega\left(\left(\lambda^{-2} \beta+\lambda^{-1} \alpha\right) \log (n / \delta)\right)
$$

Thus, we complete the proof.

Lemma 28.5.2 (Stronger version of Lemma 3.3 in [DZPS19]). Let Part 4 in Assumption 28.1.1 hold. Let $D_{\mathrm{cts}}=\frac{\sqrt{\alpha}\|y-u(0)\|_{2}}{\sqrt{m} \lambda}$. Suppose for $0 \leq s \leq t$, $\lambda_{\min }(H(s)) \geq \lambda / 2$. Then we have

$$
\|y-u(t)\|_{2}^{2} \leq \exp (-\lambda t) \cdot\|y-u(0)\|_{2}^{2},
$$

and

$$
\left\|w_{r}(t)-w_{r}(0)\right\|_{2} \leq D_{\text {cts }} .
$$

Proof. Recall we can write the dynamics of predictions as

$$
\frac{\mathrm{d}}{\mathrm{~d} t} u(t)=H(t) \cdot(y-u(t))
$$

Table 28.3: Table of Parameters for the $m=\widetilde{\Omega}\left(n^{3}\right)$ result in Section 28.5. Nt. stands for notations.

| Nt. | Choice | Place | Comment |
| :--- | :--- | :--- | :--- |
| $\lambda$ | $:=\lambda_{\min }\left(H^{\text {cts }}\right)$ | Part 1 of Ass. 28.1.1 | Data-dependent |
| $\alpha$ | Absolute | Part 2 of Ass. 28.1.1 | Data-dependent |
| $\beta$ | Variance | Part 3 of Ass. 28.1.1 | Data-dependent |
| $R$ | $\lambda / n$ | Eq. (28.11) | Maximal allowed movement of weight |
| $D_{\text {cts }}$ | $\frac{\sqrt{\alpha}\\|y-u(0)\\|_{2}}{\sqrt{m \lambda}}$ | Lemma 28.5.2 | Actual moving distance, continuous case |
| $D$ | $\frac{\sqrt{\alpha}\\|y-u(0)\\|_{2}}{\sqrt{m} \lambda}$ | Theorem 28.5.5 | Actual moving distance, discrete case |
| $\eta$ | $\lambda /(\alpha n)$ | Eq. (28.11) | Step size of gradient descent |
| $m$ | $\left(\lambda^{-2} \beta+\lambda^{-1} \alpha\right) \log (n / \delta)$ | Theorem 28.5.1 | Bounding discrete and continuous |
| $m$ | $\lambda^{-4} \alpha n^{3} \log ^{3}(n / \delta)$ | Lemma 28.4.4 <br> Claim 28.4.7 | $D<R$ and $\\|y-u(0)\\|_{2}^{2}=\widetilde{O}(n)$ |

We can calculate the loss function dynamics

$$
\begin{aligned}
\frac{\mathrm{d}}{\mathrm{~d} t}\|y-u(t)\|_{2}^{2} & =-2(y-u(t))^{\top} \cdot H(t) \cdot(y-u(t)) \\
& \leq-\lambda\|y-u(t)\|_{2}^{2}
\end{aligned}
$$

Thus we have $\frac{\mathrm{d}}{\mathrm{d} t}\left(\exp (\lambda t)\|y-u(t)\|_{2}^{2}\right) \leq 0$ and $\exp (\lambda t)\|y-u(t)\|_{2}^{2}$ is a decreasing function with respect to $t$.

Using this fact we can bound the loss

$$
\|y-u(t)\|_{2}^{2} \leq \exp (-\lambda t)\|y-u(0)\|_{2}^{2}
$$

Therefore, $u(t) \rightarrow y$ exponentially fast.
Now, we can bound the gradient norm. Recall for $0 \leq s \leq t$,

$$
\left\|\frac{\mathrm{d}}{\mathrm{~d} s} w_{r}(s)\right\|_{2}=\left\|\sum_{i=1}^{n}\left(y_{i}-u_{i}\right) \frac{1}{\sqrt{m}} a_{r} x_{i} \cdot \mathbf{1}_{w_{r}(s)^{\top} x_{i} \geq 0}\right\|_{2}
$$

Define matrix $X_{r} \in \mathbb{R}^{d \times n}$ by setting the $i$-th column to be $\mathbf{1}_{w_{r}(s)^{\top} x_{i} \geq 0} \cdot x_{i}$, then $X_{r}^{\top} X_{r}=$ $H\left(w_{r}(s)\right)$, where $H(\cdot)$ is the matrix defined in Definition 28.1.1. Then we have $\left\|X_{r}^{\top} X_{r}\right\|_{2} \leq \alpha$ by Part 4 in Assumption 28.1.1, which leads to $\left\|X_{r}\right\|_{2} \leq \sqrt{\alpha}$. So we have

$$
\begin{align*}
\left\|\frac{\mathrm{d}}{\mathrm{~d} s} w_{r}(s)\right\|_{2} & =\frac{1}{\sqrt{m}}\left\|X_{r}(y-u(s))\right\|_{2} \\
& \leq \frac{1}{\sqrt{m}}\left\|X_{r}\right\|_{2}\|(y-u(s))\|_{2} \\
& \leq \frac{\sqrt{\alpha}}{\sqrt{m}}\|y-u(s)\|_{2}  \tag{28.12}\\
& \leq \frac{\sqrt{\alpha}}{\sqrt{m}} \exp (-\lambda s)\|y-u(0)\|_{2} .
\end{align*}
$$

Integrating the gradient, we can bound the distance from the initialization

$$
\begin{aligned}
\left\|w_{r}(t)-w_{r}(0)\right\|_{2} & \leq \int_{0}^{t}\left\|\frac{\mathrm{~d}}{\mathrm{~d} s} w_{r}(s)\right\|_{2} \mathrm{~d} s \\
& \leq \frac{\sqrt{\alpha}\|y-u(0)\|_{2}}{\sqrt{m} \lambda}
\end{aligned}
$$

### 28.5.1 Technical claims

Claim 28.5.3. Let $C_{3}=-2(y-u(k))^{\top} v_{2}$. Then we have

$$
C_{3} \leq\|y-u(k)\|_{2}^{2} \cdot 8 \eta(\alpha n)^{1 / 2} R
$$

with probability at least $1-n \exp (-m R)$.

Proof. We have

$$
\text { LHS } \leq 2\|y-u(k)\|_{2} \cdot\left\|v_{2}\right\|_{2}
$$

We can upper bound $\left\|v_{2}\right\|_{2}$ in the following sense

$$
\begin{aligned}
\left\|v_{2}\right\|_{2}^{2} & \leq \sum_{i=1}^{n}\left(\frac{\eta}{\sqrt{m}} \sum_{r \in \bar{S}_{i}}\left|\left(\frac{\partial L(W(k))}{\partial w_{r}(k)}\right)^{\top} x_{i}\right|\right)^{2} \\
& =\frac{\eta^{2}}{m} \sum_{i=1}^{n}\left(\sum_{r=1}^{m} \mathbf{1}_{r \in \bar{S}_{i}}\left|\left(\frac{\partial L(W(k))}{\partial w_{r}(k)}\right)^{\top} x_{i}\right|\right)^{2} \\
& \leq \frac{\eta^{2}}{m} \cdot \max _{r \in[m]}\left|\frac{\partial L(W(k))}{\partial w_{r}(k)}\right|^{2} \cdot \sum_{i=1}^{n}\left(\sum_{r=1}^{m} \mathbf{1}_{r \in \bar{S}_{i}}\right)^{2} \\
& \leq \frac{\eta^{2}}{m} \cdot\left(\frac{\sqrt{\alpha}}{\sqrt{m}}\|u(k)-y\|_{2}\right)^{2} \cdot \sum_{i=1}^{n}\left(\sum_{r=1}^{m} \mathbf{1}_{r \in \bar{S}_{i}}\right)^{2} \\
& \leq \frac{\eta^{2}}{m} \cdot\left(\frac{\sqrt{\alpha}}{\sqrt{m}}\|u(k)-y\|_{2}\right)^{2} \cdot \sum_{i=1}^{n}(4 m R)^{2} \\
& =16 \alpha n R^{2} \eta^{2}\|u(k)-y\|_{2}^{2},
\end{aligned}
$$

where the first step follows from definition of $v_{2}$, the fourth step follows from (28.12) and

$$
\begin{aligned}
\max _{r \in[m]}\left|\frac{\partial L(W(k))}{\partial w_{r}(k)}\right| & =\max _{r \in[m]}\left|\frac{\mathrm{d} w_{r}(k)}{\mathrm{d} k}\right| \\
& \leq \frac{\sqrt{\alpha}}{\sqrt{m}}\|y-u(k)\|_{2},
\end{aligned}
$$

the fifth step follows from $\sum_{r=1}^{m} \mathbf{1}_{r \in \bar{S}_{i}} \leq 4 m R$ with probability at least $1-\exp (-m R)$.
Claim 28.5.4. Let $C_{4}=\|u(k+1)-u(k)\|_{2}^{2}$. Then we have

$$
C_{4} \leq \eta^{2} \alpha n\|y-u(k)\|_{2}^{2} .
$$

Proof. We have

$$
\begin{aligned}
\text { LHS } & \leq \eta^{2} \sum_{i=1}^{n} \frac{1}{m}\left(\sum_{r=1}^{m}\left\|\frac{\partial L(W(k))}{\partial w_{r}(k)}\right\|_{2}\right)^{2} \\
& \leq \eta^{2} \sum_{i=1}^{n} \frac{1}{m}\left(\sum_{r=1}^{m} \frac{\sqrt{\alpha}}{\sqrt{m}}\|u(k)-y\|_{2}\right)^{2} \\
& \leq \eta^{2} \alpha n\|y-u(k)\|_{2}^{2} .
\end{aligned}
$$

### 28.5.2 Main result

Theorem 28.5.5. Recall that $\lambda=\lambda_{\text {min }}\left(H^{\text {cts }}\right)>0$. Let $m=\Omega\left(\lambda^{-4} n^{3} \alpha \log ^{3}(n / \delta)\right)$, we i.i.d. initialize $w_{r} \in \mathcal{N}(0, I)$, a sampled from $\{-1,+1\}$ uniformly at random for $r \in[m]$, and we set the step size $\eta=O(\lambda /(\alpha n))$ then with probability at least $1-\delta$ over the random initialization we have for $k=0,1,2, \ldots$

$$
\begin{equation*}
\|u(k)-y\|_{2}^{2} \leq(1-\eta \lambda / 2)^{k} \cdot\|u(0)-y\|_{2}^{2} \tag{28.13}
\end{equation*}
$$

Proof. This proof, similar to the proof of Theorem 28.4.5, is again by induction. trivially holds when $k=0$, which is the base case.

If (28.13) holds for $k^{\prime}=0, \cdots, k$, then we claim that for all $r \in[m]$

$$
\begin{equation*}
\left\|w_{r}(k+1)-w_{r}(0)\right\|_{2} \leq \frac{4 \sqrt{\alpha}\|y-u(0)\|_{2}}{\sqrt{m} \lambda}:=D \tag{28.14}
\end{equation*}
$$

To see this, we use the norm of gradient to bound this distance,

$$
\begin{aligned}
\left\|w_{r}(k+1)-w_{r}(0)\right\|_{2} & \leq \eta \sum_{k^{\prime}=0}^{k}\left\|\frac{\partial L\left(W\left(k^{\prime}\right)\right)}{\partial w_{r}\left(k^{\prime}\right)}\right\|_{2} \\
& \leq \eta \sum_{k^{\prime}=0}^{k} \frac{\sqrt{\alpha}\left\|y-u\left(k^{\prime}\right)\right\|_{2}}{\sqrt{m}} \\
& \leq \eta \sum_{k^{\prime}=0}^{k} \frac{\sqrt{\alpha}(1-\eta \lambda / 2)^{k^{\prime} / 2}}{\sqrt{m}}\|y-u(0)\|_{2} \\
& \leq \eta \sum_{k^{\prime}=0}^{\infty} \frac{\sqrt{\alpha}(1-\eta \lambda / 2)^{k^{\prime} / 2}}{\sqrt{m}}\|y-u(0)\|_{2} \\
& =\frac{4 \sqrt{\alpha}\|y-u(0)\|_{2}}{\sqrt{m} \lambda},
\end{aligned}
$$

where the first step follows from (28.2), the second step follows from (28.12), the third step follows from the induction hypothesis, the fourth step relaxes the summation to an infinite summation, and the last step follows from $\sum_{k^{\prime}=0}^{\infty}(1-\eta \lambda / 2)^{k^{\prime} / 2}=\frac{2}{\eta \lambda}$.

Then from Claim 28.5.4, it is sufficient to choose $\eta=\frac{\lambda}{4 \alpha n}$ so that (28.13) holds for $k^{\prime}=k+1$. This completes the induction step.

## Over-parameterization size, lower bound on $m$.

We require

$$
D=\frac{4 \sqrt{\alpha}\|y-u(0)\|_{2}}{\sqrt{m} \lambda}<R=\frac{\lambda}{64 n},
$$

and

$$
2 n \exp (-m R) \leq \delta
$$

This implies that

$$
\begin{aligned}
m & =\Omega\left(\lambda^{-4} n^{2} \alpha\|y-u(0)\|_{2}^{2}\right) \\
& =\Omega\left(\lambda^{-4} n^{3} \alpha \log (m / \delta) \log ^{2}(n / \delta)\right)
\end{aligned}
$$

where the last step follows from Claim 28.4.7.

### 28.6 Quadratic Suffices

Lemma 28.6.1 (perturbed $w)$. Let $R \in(0,1)$. Let Assumption 4 in 28.1.1 hold, i.e. for all $i \neq j,\left|x_{i}^{\top} x_{j}\right| \leq \theta / \sqrt{n}$. If $\widetilde{w}_{1}, \cdots, \widetilde{w}_{m}$ are i.i.d. generated $\mathcal{N}(0, I)$. For any set of weight vectors $w_{1}, \cdots, w_{m} \in \mathbb{R}^{d}$ that satisfy for any $r \in[m],\left\|\widetilde{w}_{r}-w_{r}\right\|_{2} \leq R$, then the $H: \mathbb{R}^{m \times d} \rightarrow \mathbb{R}^{n \times n}$ defined

$$
H(w)_{i, j}=\frac{1}{m} x_{i}^{\top} x_{j} \sum_{r=1}^{m} \mathbf{1}_{w_{r}^{\top} x_{i} \geq 0, w_{r}^{\top} x_{j} \geq 0}
$$

Then we have

$$
\|H(w)-H(\widetilde{w})\|_{F}<2\left(n\left(1+\theta^{2}\right)\right)^{1 / 2} R
$$

holds with probability at least $1-n^{2} \cdot \exp (-m R / 10)$.

Proof. The random variable we care is

$$
\begin{aligned}
\sum_{i=1}^{n} \sum_{j=1}^{n}\left|H(\widetilde{w})_{i, j}-H(w)_{i, j}\right|^{2} & =\frac{1}{m^{2}} \sum_{i=1}^{n} \sum_{j=1}^{n}\left|x_{i}^{\top} x_{j} \sum_{r=1}^{m}\left(\mathbf{1}_{\widetilde{w}_{r}^{\top} x_{i} \geq 0, \widetilde{w}_{r}^{\top} x_{j} \geq 0}-\mathbf{1}_{w_{r}^{\top} x_{i} \geq 0, w_{r}^{\top} x_{j} \geq 0}\right)\right|^{2} \\
& =B_{1}+B_{2}
\end{aligned}
$$

where $B_{1}, B_{2}$ are defined as

$$
\begin{aligned}
& B_{1}=\frac{1}{m^{2}} \sum_{i=1}^{n}\left|\sum_{r=1}^{m}\left(\mathbf{1}_{\widetilde{w}_{r}^{\top} x_{i} \geq 0}-\mathbf{1}_{w_{r}^{\top} x_{i} \geq 0}\right)\right|^{2}, \\
& B_{2}=\frac{1}{m^{2}} \sum_{i=1}^{n} \sum_{j \in[n \backslash \backslash i\}}\left|x_{i}^{\top} x_{j} \sum_{r=1}^{m}\left(\mathbf{1}_{\widetilde{w}_{r}^{\top} x_{i} \geq 0, \widetilde{w}_{r}^{\top} x_{j} \geq 0}-\mathbf{1}_{w_{r}^{\top} x_{i} \geq 0, w_{r}^{\top} x_{j} \geq 0}\right)\right|^{2} .
\end{aligned}
$$

We can further bound $B_{2}$ as

$$
\begin{aligned}
B_{2} & \leq \frac{1}{m^{2}} \sum_{i=1}^{n} \sum_{j \in[n] \backslash\{i\}} \frac{\theta^{2}}{n}\left|\sum_{r=1}^{m}\left(\mathbf{1}_{\widetilde{w}_{r}^{\top} x_{i} \geq 0, \tilde{w}_{r}^{\top} x_{j} \geq 0}-\mathbf{1}_{w_{r}^{\top} x_{i} \geq 0, w_{r}^{\top} x_{j} \geq 0}\right)\right|^{2} \\
& =\frac{\theta^{2}}{n m^{2}} \sum_{i=1}^{n} \sum_{j \in[n \backslash \backslash\{i\}}\left|\sum_{r=1}^{m}\left(\mathbf{1}_{\widetilde{w}_{r}^{\top} x_{i} \geq 0, \tilde{w}_{r}^{\top} x_{j} \geq 0}-\mathbf{1}_{w_{r}^{\top} x_{i} \geq 0, w_{r}^{\top} x_{j} \geq 0}\right)\right|^{2} .
\end{aligned}
$$

For each $r, i, j$, we define

$$
s_{r, i, j}:=\mathbf{1}_{\tilde{w}_{r}^{\top} x_{i} \geq 0, \widetilde{w}_{r}^{\top} x_{j} \geq 0}-\mathbf{1}_{w_{r}^{\top} x_{i} \geq 0, w_{r}^{\top} x_{j} \geq 0} .
$$

Then we can rewrite $B_{1}$ and $B_{2}$ as

$$
\begin{aligned}
& B_{1}=\frac{1}{m^{2}} \sum_{i=1}^{n}\left(\sum_{r=1}^{m} s_{r, i, i}\right)^{2}, \\
& B_{2}=\frac{\theta^{2}}{n m^{2}} \sum_{i=1}^{n} \sum_{j \in[n] \backslash\{i\}}\left(\sum_{r=1}^{m} s_{r, i, j}\right)^{2} .
\end{aligned}
$$

Therefore it is sufficient to bound $\sum_{r=1}^{m} s_{r, i, j}$ simutaneously for all pair $i, j$. Using same technique in the proof of Theorem 28.4.2, we have

$$
\operatorname{Pr}\left[\frac{1}{m} \sum_{r=1}^{m} s_{r, i, j} \geq 2 R\right] \leq \exp (-m R / 10)
$$

By applying union bound on all $i, j$ pairs, we get with probability at least $1-\exp (-m R / 10)$,

$$
\|H(w)-H(\widetilde{w})\|_{F}^{2} \leq B_{1}+B_{2} \leq 4 n R^{2}(1+\theta)^{2}
$$

which is precisely what we need.

Table 28.4: Table of Parameters for the $m=\widetilde{\Omega}\left(n^{2}\right)$ result in Section 28.6. Nt. stands for notations.

| Nt. | Choice | Place | Comment |
| :--- | :--- | :--- | :--- |
| $\lambda$ | $:=\lambda_{\min }\left(H^{\text {cts }}\right)$ | Part 1 of Ass. 28.1.1 | Data-dependent |
| $\alpha$ | Absolute | Part 2 of Ass. 28.1.1 | Data-dependent |
| $\beta$ | Variance | Part 3 of Ass. 28.1.1 | Data-dependent |
| $\theta$ | Inner product | Part 4 of Ass. 28.1.1 | Data-dependent |
| $R$ | $\frac{\lambda}{\sqrt{n}} \cdot \min \left\{\frac{1}{\sqrt{\alpha}}, \frac{1}{\sqrt{1+\theta^{2}}}\right\}$ | Eq. (28.15) | Maximal allowed movement of weight |
| $D$ | $\frac{4 \sqrt{\alpha}\\|y-u(0)\\|_{2}}{\sqrt{m} \lambda}$ | Theorem 28.6.4 | Actual moving distance, discrete case |
| $\eta$ | $\lambda /(\alpha n)$ | Eq. (28.11) | Step size of gradient descent |
| $m$ | $\left(\lambda^{-2} \beta+\lambda^{-1} \alpha\right) \log ^{2}(n / \delta)$ | Theorem 28.5.1 | Bounding discrete and continuous |
| $m$ | $\lambda^{-4} \alpha\left(\alpha+\theta^{2}\right) n^{2} \log ^{3}(n / \delta)$ | Lemma 28.4.4 <br> Claim 28.4.7 | $D<R$ and $\\|y-u(0)\\|_{2}^{2}=\widetilde{O}(n)$ |

Claim 28.6.2. Assume $R \leq \frac{\lambda}{64 \sqrt{n}} \cdot \frac{1}{\sqrt{1+\theta^{2}}}$. Let $C_{1}=-2 \eta(y-u(k))^{\top} H(k)(y-u(k))$. We have

$$
C_{1} \leq-\|y-u(k)\|_{2}^{2} \cdot \eta \lambda
$$

Proof. By Lemma 28.6.1 and our choice of $R \leq \frac{\lambda}{64 \sqrt{n}} \cdot \frac{1}{\sqrt{1+\theta^{2}}}$, We have

$$
\|H(0)-H(k)\|_{F} \leq 2\left(n\left(1+\theta^{2}\right)\right)^{1 / 2} \cdot \frac{\lambda}{64 \sqrt{n}} \cdot \frac{1}{\sqrt{1+\theta^{2}}} \leq \frac{\lambda}{4}
$$

Recall that $\lambda=\lambda_{\min }(H(0))$. Therefore

$$
\lambda_{\min }(H(k)) \geq \lambda_{\min }(H(0))-\|H(0)-H(k)\| \geq \lambda / 2
$$

Then we have

$$
(y-u(k))^{\top} H(k)(y-u(k)) \geq\|y-u(k)\|_{2}^{2} \cdot \lambda / 2
$$

Thus, we complete the proof.

Claim 28.6.3. Let $C_{2}=2 \eta(y-u(k))^{\top} H(k)^{\perp}(y-u(k))$. We have

$$
C_{2} \leq\|y-u(k)\|_{2}^{2} \cdot 8 \eta R\left(n\left(1+\theta^{2}\right)\right)^{1 / 2} .
$$

holds with probability $1-n \exp (-m R)$.

Proof. Note that

$$
C_{2} \leq 2 \eta \cdot\|y-u(k)\|_{2}^{2} \cdot\left\|H(k)^{\perp}\right\| .
$$

It suffices to upper bound $\left\|H(k)^{\perp}\right\|$. Since $\|\cdot\| \leq\|\cdot\|_{F}$, then it suffices to upper bound $\|\cdot\|_{F}$.

For each $i \in[n]$, we define $y_{i}$ as follows

$$
y_{i}=\sum_{r=1}^{m} \mathbf{1}_{r \in \bar{S}_{i}} .
$$

Then we have

$$
\begin{aligned}
\left\|H(k)^{\perp}\right\|_{F}^{2} & =\sum_{i=1}^{n} \sum_{j=1}^{n}\left(H(k)_{i, j}^{\perp}\right)^{2} \\
& =\sum_{i=1}^{n} \sum_{j=1}^{n}\left(\frac{1}{m} \sum_{r \in \bar{S}_{i}} x_{i}^{\top} x_{j} \mathbf{1}_{w_{r}(k)^{\top} x_{i} \geq 0, w_{r}(k)^{\top} x_{j} \geq 0}\right)^{2} \\
& =\sum_{i=1}^{n} \sum_{j=1}^{n}\left(\frac{1}{m} \sum_{r=1}^{m} x_{i}^{\top} x_{j} \mathbf{1}_{w_{r}(k)^{\top} x_{i} \geq 0, w_{r}(k)^{\top} x_{j} \geq 0} \cdot \mathbf{1}_{r \in \bar{S}_{i}}\right)^{2} \\
& =\sum_{i=1}^{n} \sum_{j=1}^{n} \frac{\left|x_{i}^{\top} x_{j}\right|^{2}}{m^{2}}\left(\sum_{r=1}^{m} \mathbf{1}_{w_{r}(k)^{\top} x_{i} \geq 0, w_{r}(k)^{\top} x_{j} \geq 0} \cdot \mathbf{1}_{r \in \bar{S}_{i}}\right)^{2} \\
& =B_{1}+B_{2},
\end{aligned}
$$

where $B_{1}$ and $B_{2}$ are defined as:

$$
\begin{aligned}
B_{1} & :=\sum_{i=1}^{n} \frac{1}{m^{2}}\left(\sum_{r=1}^{m} \mathbf{1}_{w_{r}(k)^{\top} x_{i} \geq 0} \cdot \mathbf{1}_{r \in \bar{S}_{i}}\right)^{2}, \\
B_{2} & :=\sum_{i=1}^{n} \sum_{j \in[n] \backslash\{i\}} \frac{\left|x_{i}^{\top} x_{j}\right|^{2}}{m^{2}}\left(\sum_{r=1}^{m} \mathbf{1}_{w_{r}(k)^{\top} x_{i} \geq 0, w_{r}(k)^{\top} x_{j} \geq 0} \cdot \mathbf{1}_{r \in \bar{S}_{i}}\right)^{2} .
\end{aligned}
$$

We bound $B_{1}$ and $B_{2}$ separately.
We first bound $B_{1}$.

$$
\begin{aligned}
B_{1} & =\sum_{i=1}^{n} \frac{1}{m^{2}}\left(\sum_{r=1}^{m} \mathbf{1}_{w_{r}(k)^{\top} x_{i} \geq 0} \cdot \mathbf{1}_{r \in \bar{S}_{i}}\right)^{2} \\
& \leq \frac{1}{m^{2}} \sum_{i=1}^{n}\left(\sum_{r=1}^{m} \mathbf{1}_{r \in \bar{S}_{i}}\right)^{2} \\
& =\frac{1}{m^{2}} \sum_{i=1}^{n} y_{i}^{2} .
\end{aligned}
$$

Fix $i \in[n]$. The plan is to use Bernstein inequality to upper bound $y_{i}$ with high probability.

First by Eq. (28.6) we have

$$
\mathbb{E}\left[\mathbf{1}_{r \in \bar{S}_{i}}\right] \leq R
$$

We also have

$$
\begin{aligned}
\mathbb{E}\left[\left(\mathbf{1}_{r \in \bar{S}_{i}}-\mathbb{E}\left[\mathbf{1}_{r \in \bar{S}_{i}}\right]\right)^{2}\right] & =\mathbb{E}\left[\mathbf{1}_{r \in \bar{S}_{i}}^{2}\right]-\mathbb{E}\left[\mathbf{1}_{r \in \bar{S}_{i}}\right]^{2} \\
& \leq \mathbb{E}\left[\mathbf{1}_{r \in \bar{S}_{i}}^{2}\right] \\
& \leq R .
\end{aligned}
$$

Finally we have $\left|\mathbf{1}_{r \in \bar{S}_{i}}-\mathbb{E}\left[\mathbf{1}_{r \in \bar{S}_{i}}\right]\right| \leq 1$.

Notice that $\left\{\mathbf{1}_{r \in \bar{S}_{i}}\right\}_{r=1}^{m}$ are mutually independent, since $\mathbf{1}_{r \in \bar{S}_{i}}$ only depends on $w_{r}(0)$. Hence from Bernstein inequality (Lemma A.1.2) we have for all $t>0$,

$$
\operatorname{Pr}\left[y_{i}>m \cdot R+t\right] \leq \exp \left(-\frac{t^{2} / 2}{m \cdot R+t / 3}\right)
$$

By setting $t=3 m R$, we have

$$
\operatorname{Pr}\left[y_{i}>4 m R\right] \leq \exp (-m R)
$$

Hence by union bound, with probability at least $1-n \exp (-m R)$, for all $i \in[n]$,

$$
y_{i} \leq 4 m R
$$

If this happens, we have

$$
B_{1} \leq 16 n R^{2}
$$

Next we bound $B_{2}$. We have

$$
\begin{aligned}
B_{2} & =\sum_{i=1}^{n} \sum_{j \in[n] \backslash\{i\}} \frac{\left|x_{i}^{\top} x_{j}\right|^{2}}{m^{2}}\left(\sum_{r=1}^{m} \mathbf{1}_{w_{r}(k)^{\top} x_{i} \geq 0, w_{r}(k)^{\top} x_{j} \geq 0} \cdot \mathbf{1}_{r \in \bar{S}_{i}}\right)^{2} \\
& \leq \frac{1}{m^{2}} \sum_{i=1}^{n}\left(\sum_{j \in[n] \backslash\{i\}}\left(x_{i}^{\top} x_{j}\right)^{4}\right)^{1 / 2} \cdot\left(\sum_{j \in[n] \backslash\{i\}}\left(\sum_{r=1}^{m} \mathbf{1}_{w_{r}(k)^{\top} x_{i} \geq 0, w_{r}(k)^{\top} x_{j} \geq 0} \cdot \mathbf{1}_{r \in \bar{S}_{i}}\right)^{4}\right)^{1 / 2} \\
& \leq \frac{1}{m^{2}} \sum_{i=1}^{n}\left(\sum_{j \in[n] \backslash\{i\}}\left(x_{i}^{\top} x_{j}\right)^{4}\right)^{1 / 2}\left(\sum_{j \in[n] \backslash\{i\}} y_{i}^{4}\right)^{1 / 2} \\
& =\frac{\sqrt{n-1}}{m^{2}} \sum_{i=1}^{n}\left(\sum_{j \in[n \backslash \backslash\{i\}}\left(x_{i}^{\top} x_{j}\right)^{4}\right)^{1 / 2} y_{i}^{2} \\
& \leq 16 R^{2} \sqrt{n} \sum_{i=1}^{n}\left(\sum_{j \in[n] \backslash\{i\}}\left(x_{i}^{\top} x_{j}\right)^{4}\right)^{1 / 2}
\end{aligned}
$$

where the last step happens when $y_{i} \leq 4 m R$ for all $i \in[n]$.
Now, using the assumption $x_{i}^{\top} x_{j} \leq \frac{\theta}{\sqrt{n}}$ (Part 4 of Assumption 28.1.1), we have

$$
B_{2} \leq 16 n R^{2} \theta^{2} .
$$

Putting things together, we have with probability at least $1-n \exp (-m R)$,

$$
\begin{aligned}
\left\|H(k)^{\perp}\right\|_{F}^{2} & \leq B_{1}+B_{2} \\
& \leq 16 n R^{2}\left(1+\theta^{2}\right) .
\end{aligned}
$$

This gives us $\left\|H(k)^{\perp}\right\| \leq 4 R\left(n\left(1+\theta^{2}\right)\right)^{1 / 2}$, which is precisely what we need.

### 28.6.1 Main result

Theorem 28.6.4. Let $\lambda, \alpha, \beta, \theta$ be defined as Assumption 28.1.1. Let

$$
m=\Omega\left(\lambda^{-4} n^{2} \alpha \max \left\{1+\theta^{2}, \alpha\right\} \log ^{3}(n / \delta)\right)
$$

We i.i.d. initialize $w_{r} \in \mathcal{N}(0, I)$, $a_{r}$ sampled from $\{-1,+1\}$ uniformly at random for $r \in[m]$, and we set the step size $\eta=O(\lambda /(\alpha n))$ then with probability at least $1-\delta$ over the random initialization we have for $k=0,1,2, \cdots$

$$
\|u(k)-y\|_{2}^{2} \leq(1-\eta \lambda / 2)^{k} \cdot\|u(0)-y\|_{2}^{2} .
$$

Proof. Choice of $\eta$ and $R$. We want to choose $\eta$ and $R$ such that

$$
\begin{equation*}
\left(1-\eta \lambda+8 \eta R\left(n\left(1+\theta^{2}\right)\right)^{1 / 2}+8 \eta(\alpha n)^{1 / 2} R+\eta^{2} \alpha n\right) \leq(1-\eta \lambda / 2) . \tag{28.15}
\end{equation*}
$$

Now, if we set $\eta=\frac{\lambda}{4 \alpha n}$ and $R=\frac{\lambda}{64 \sqrt{n}} \cdot \min \left\{\frac{1}{\sqrt{1+\theta^{2}}}, \frac{1}{\sqrt{\alpha}}\right\}$, we have

$$
8 \eta R\left(n\left(1+\theta^{2}\right)\right)^{1 / 2}+8 \eta(\alpha n)^{1 / 2} R \leq \frac{1}{4} \eta \lambda
$$

and $\eta^{2} n^{2} \leq \frac{1}{4} \eta \lambda$. This gives us

$$
\|y-u(k+1)\|_{2}^{2} \leq\|y-u(k)\|_{2}^{2}(1-\eta \lambda / 2)
$$

with probability at least $1-2 n \exp (-m R)$.
Over-parameterization size, lower bound on $m$. By same analysis as in the proof of Theorem 28.5.5, we still have

$$
\left\|w_{r}(k+1)-w_{r}(0)\right\|_{2} \leq \frac{4 \sqrt{\alpha}\|y-u(0)\|_{2}}{\sqrt{m} \lambda}:=D
$$

We require

$$
D=\frac{4 \sqrt{\alpha}\|y-u(0)\|_{2}}{\sqrt{m} \lambda}<R=\frac{\lambda}{64 \sqrt{n}} \cdot \min \left\{\frac{1}{\sqrt{1+\theta^{2}}}, \frac{1}{\sqrt{\alpha}}\right\}
$$

and

$$
2 n \exp (-m R) \leq \delta
$$

This implies that

$$
\begin{aligned}
m & =\Omega\left(\lambda^{-4} n \alpha\|y-u(0)\|_{2}^{2} \max \left\{1+\theta^{2}, \alpha\right\}\right) \\
& =\Omega\left(\lambda^{-4} n^{2} \alpha \cdot \max \left\{1+\theta^{2}, \alpha\right\} \cdot \log (m / \delta) \log ^{2}(n / \delta)\right),
\end{aligned}
$$

where the last step follows from Claim 28.4.7.

## Chapter 29

## Longest Common Subsequence I

Given a pair of $n$-character strings, the problems of computing their Longest Common Subsequence and Edit Distance have been extensively studied for decades. For exact algorithms, LCS and Edit Distance (with character insertions and deletions) are equivalent; the state of the art running time is (almost) quadratic in $n$, and this is tight under plausible fine-grained complexity assumptions. But for approximation algorithms the picture is different: there is a long line of works with improved approximation factors for Edit Distance, but for LCS (with binary strings) only a trivial 1/2-approximation was known. In this work we give a reduction from approximate LCS to approximate Edit Distance, yielding the first efficient $(1 / 2+\epsilon)$-approximation algorithm for LCS for some constant $\epsilon>0$.

This part is based on previous following publication

- Aviad Rubinstein, Zhao Song

Reducing approximate Longest Common Subsequence to approximate Edit Distance. Manuscript 2019 [RS19b]

### 29.1 Introduction

In this work we consider two of the most ubiquitous measures of similarity between a pair of strings: the longest common subsequence (LCS) and the edit distance. The LCS of two $n$-character strings $A$ and $B$ is simply their longest (not necessarily contiguous) common substring. Edit distance is the minimum number of character insertions, deletions, and substitutions required to transform $A$ to $B$. In fact, under a slightly more restricted definition that does not allow substitutions ${ }^{1}$, the two measures are complements and the problems of computing them exactly are equivalent.

There is a textbook dynamic programming algorithm for computing LCS (or edit distance) than runs in $O\left(n^{2}\right)$ time, and a slightly faster $O\left(n^{2} / \log ^{2}(n)\right)$-time algorithm due to Masek and Paterson [MP80]. Finding faster algorithms is a central and long standing open problem both in theory and in practice (e.g. Problem 35 of [Knu72]). Under plausible fine-grained complexity assumptions such as SETH, neither problem can be computed much faster [AWW14, ABW15, BI15, BK15, AHWW16].

For (multiplicative) approximation, the two problems are no longer equivalent. For edit distance, there is a long sequence of approximation algorithms with improving factors [BYJKK04, BES06, AO12, AKO10, $\mathrm{BEG}^{+}$18]; in particular, [CDG ${ }^{+}$18] gives a constant factor approximation in truly sub-quadratic time. For LCS with alphabet size $|\Sigma|$, in contrast, there is a trivial $1 /|\Sigma|$-approximation, and no better algorithms are known (for large alphabet ${ }^{2}$ there are some hardness of approximation results [AB17, AR18, CGL ${ }^{+}$19] and also

[^98]approximation algorithms with non-trivial polynomial factors [HSSS19, RSSS19]).
In this work we focus on binary strings, where the trivial algorithm gives a $(1 /|\Sigma|=$ $1 / 2$ )-approximation. Breaking this $1 / 2$ barrier is a well-known open problem in this area. Our main result is a fine-grained reduction that implies obtaining a $1 / 2+\epsilon$-approximation for binary LCS (for some constant $\epsilon>0$ ) is no harder than approximating edit distance to within some constant factor.

Theorem 29.1.1 (Reduction: approximate ED implies approximate LCS).
Suppose that there exists a constant cand an approximate edit distance algorithm that runs in time $T(n)$ and, given two binary strings $A, B$ of length $n$, returns an estimate $\widetilde{\mathrm{ED}}(A, B) \in[\mathrm{ED}(A, B), c \cdot \mathrm{ED}(A, B)+o(n)]$. Then there exists a fixed constant $\epsilon=\epsilon(c) \in(0,1 / 2)$ and a deterministic approximation algorithm for longest common subsequence that runs in deterministic $T(n)+O(n)$ and approximates $\operatorname{LCS}(A, B)$ to within a $(1 / 2+\epsilon)$-approximation factor.

Remark 29.1.1. We state the above theorem in terms of estimating the edit distance or length of the LCS. If the edit distance algorithm can efficiently compute the transformation (this assumption is almost wlog by [CGKK18]), then our algorithm can also efficiently compute the common string.

As mentioned above, the recent breakthrough of $\left[\mathrm{CDG}^{+} 18\right]$ gives a constant factor approximation of edit distance in truly-subquadratic $\left(\widetilde{O}\left(n^{2-2 / 7}\right)\right)$ time. By plugging their algorithm into our reduction, we would obtain $(1 / 2+\epsilon)$-approximation algorithm for binary
our framework can give a $1 /|\Sigma|+\epsilon$-approximation in this case as well, but the proof may become more complicated.

LCS with the same running time. By applying our reduction to the even more recent approximation algorithms for edit distance ${ }^{3}$ that run in near-linear time [KS19, BR19], we obtain the following stronger corollary:

Corollary 29.1.2 (Approximate LCS). For every constant $\delta>0$ there exists a constant $\epsilon>0$ such that, given two binary strings $A, B \in\{0,1\}^{n}$, there is an algorithm that runs in $O\left(n^{1+\delta}\right)$ time and $\operatorname{LCS}(A, B)$ to within a $(1 / 2+\epsilon)$-factor.

## Technical preview

The crux of our algorithm is analyzing first order statistics (counts of 0s and 1s) of the input strings $(A, B)$ and their substrings. We begin with a few simple observations. Below, we normalize ED and LCS so that they're always between 0 and 1 (as opposed to 0 and $n$ ).

- If the strings are balanced, namely have the same number of 0 s and 1 s , we know that $\operatorname{LCS}(A, B) \in[1 / 2,1]$. If the strings are very close, say $\operatorname{LCS}(A, B) \geq(1-\delta)$ for sufficiently small $\delta>0$, we can use the assumed edit distance algorithm as a black box and find a common substring of length $\geq(1-O(\delta))$. On the converse if the substring returned by the algorithm is shorter than $(1-O(\delta))$, we know that $\operatorname{LCS}(A, B)<1-\delta$, and thus returning an all- 1 string of length $1 / 2$ is a $(1 / 2+2 \delta)$-approximation.
- If $A$ is balanced and $B$ has e.g. $10 \% 0$ s and $90 \% 1 \mathrm{~s}$, we know that $\operatorname{LCS}(A, B) \in[0.5,0.6]$, so simply returning the all- 1 string of length $1 / 2$ is a $5 / 6$-approximation. The same holds for most ways in which one or both strings are unbalanced.

[^99]- However there is one difficult case when the string are perfectly unbalanced, e.g. $A$ has $99 \%$ 0s and $B$ has $99 \%$ 1s. Now the first order statistics over the entire strings only tell us that $\operatorname{LCS}(A, B) \in[0.01,0.02]$, so the trivial approximation doesn't beat $1 / 2$. On the other hand, the edit distance is at least 0.98 , so even a 1.1-approximation algorithm for edit distance wouldn't give us a non-trivial guarantee for this case.

Our main technical contribution is a careful analysis of this last case (and its many sub-cases).

### 29.2 Preliminaries

For strings $x, y \in\{0,1\}^{m}$ for $m \leq n$, we use $1(x)$ to denote the number of 1 in $x, 0(x)$ to denote the number of 0 in $x$, and $\operatorname{LCS}(x, y)$ to denote the length of their longest common subsequence. All of these function are normalized w.r.t. the length of the original input to our main algorithm, $n$; in particular we always have $0(x), 1(x), \operatorname{LCS}(x, y) \in[0, m / n]$.

## Fact 29.2.1.

$$
\operatorname{LCS}(A, B) \leq \min \{0(A), 0(B)\}+\min \{1(A), 1(B)\}
$$

Parameters $\alpha, \beta, \gamma, \delta$
In the proof we consider the following parameters:
$\alpha$ We define $\alpha:=\min \{1(A), 1(B), 0(A), 0(B)\}$. Notice that $\alpha$ may be very small, and even approaching 0 as a function of $n$. We assume wlog that this minimum is attained by $1(A)=\alpha$.
$\beta$ The parameter $\beta$ will represent a robustness parameter for some of our bounds. We take $\beta=\Theta(\alpha)$, but it may be smaller by an arbitrary constant factor.
$\gamma$ The parameter $\gamma \in(0,1)$ is a constant that depends on the approximation factor $c$ of the approximation algorithm for edit distance that we assume. We choose $\beta$ sufficiently small such that $\gamma \alpha \gg \beta$.
$\delta$ The parameter $\delta$ represents the deviation from "perfectly unbalanced" case (see Lemma 29.3.1). It is an arbitrary small constant. In particular, $\delta \alpha \ll \beta$. It is sufficiently small
that for succinctness of representation we'll simply omit it (as if it were zero) after Lemma 29.3.1.

## Subroutines

Our reduction will assume the availability of an algorithm APPROXED which takes as input two strings $A, B$ of length $n$ and outputs $1-\widetilde{\mathrm{ED}}(A, B)$ where $\widetilde{\mathrm{ED}}(A, B) \in[\mathrm{ED}(A, B), c$. $\mathrm{ED}(A, B)+o(1)]$.

In addition, we also define three trivial algorithms; they all run in time linear in length of input string.

Definition 29.2.1 (Match). Given input string $A$ and $B$, and a symbol $\sigma \in \Sigma$. The algorithm $\operatorname{Match}(A, B, \sigma)$ will output a string $C$ where every character is $\sigma$ and the length of $C$ is $\min \{\sigma(A), \sigma(B)\}$. This algorithm takes $O(|A|+|B|)$ time.

Definition 29.2.2 (BestMatch). Given input string $A$ and $B$. The algorithm $\operatorname{BestMatch}(A, B)$ will take the longest one of $\operatorname{Match}(A, B, 0)$ and $\operatorname{Match}(A, B, 1)$. This algorithm also takes $O(|A|+|B|)$ time.

Definition 29.2.3 (Greedy). Given input string $A_{1}, A_{2}$ and $B$. The algorithm $\operatorname{Greedy}\left(A_{1}, A_{2}, B\right)$ will find the optimal contiguous partition $B=B_{1} \cup B_{2}$ so as to maximize $\operatorname{BestMatch}\left(A_{1}, B_{1}\right)+\operatorname{BestMatch}\left(A_{2}, B_{2}\right)$. This algorithm also takes $O(|A|+|B|)$ time.

Below, we slightly abuse notation and refer to the above algorithms (ApproxED, Match, BestMatch, Greedy) both when we want their output to be the actual common string, and the length. Which output we need will be clear from context.

### 29.3 Reducing to perfectly unbalanced case

In this section we formalize the intuition from the introduction that BESTMATCH gives a better-than- $1 / 2$-approximation unless $1(A) \approx 0(B)$.

Lemma 29.3.1 (Reduction to perfectly unbalanced case).

$$
\begin{aligned}
& \text { If }|1(A)-0(B)|>\delta \min \{0(A), 0(B), 1(A), 1(B)\} \text {, then } \\
& \\
& \quad \operatorname{BestMATCH}(A, B) \geq(1 / 2+\delta / 6) \operatorname{LCS}(A, B) .
\end{aligned}
$$

Proof. Assume wlog ${ }^{4}$ that $1(A)=\min \{0(A), 0(B), 1(A), 1(B)\}$. Then we have,

$$
\begin{align*}
& \operatorname{BestMatch}(A, B) \\
& =\operatorname{MATCH}(A, B, 0) \\
& =\min \{0(B), 0(A)\} \\
& =0(B) \\
& >(1+\delta) 1(A) \\
& \text { (By assumption } 1(A) \leq 1(B) \text { ) } \\
& \text { (By premise of lemma) } \\
& =(1+\delta) \min \{1(A), 1(B)\} \\
& \geq(1+\delta)(\operatorname{LCS}(A, B)-\min \{0(A), 0(B)\})  \tag{Fact29.2.1}\\
& =(1+\delta)(\operatorname{LCS}(A, B)-\operatorname{BestMatch}(A, B)) \text {. }
\end{align*}
$$

[^100]Therefore,

$$
\begin{aligned}
\operatorname{BestMatch}(A, B) & \geq \frac{1+\delta}{2+\delta} \operatorname{LCS}(A, B) \\
& \geq\left(1 / 2+\frac{\delta}{4+2 \delta}\right) \operatorname{LCS}(A, B) \\
& \geq(1 / 2+\delta / 6) \operatorname{LCS}(A, B)
\end{aligned}
$$

We henceforth assume wlog that

$$
\begin{equation*}
|0(A)-1(B)| \leq \delta \min \{0(A), 0(B), 1(A), 1(B)\} \tag{29.1}
\end{equation*}
$$

For ease of presentation, we henceforth omit $\delta$ from our calculations, i.e. we'll assume that $\delta=0$. It will be evident that modifying any of our inequalities by factors in $[ \pm \delta \alpha]$ will not affect the proofs.

Eq. (29.1) is very important in our analysis, but it does not rule out the perfectly balanced case, namely $0(A) \approx 0(B) \approx 1(A) \approx 1(B) \approx 1 / 2$.

Lemma 29.3.2 (Ruling out the perfectly balanced case).

$$
\begin{align*}
& \text { Let } \beta^{\prime}, \gamma>0 \text { be sufficiently small constants }{ }^{5} \text {. If } 0(A) \in\left[1 / 2 \pm \beta^{\prime}\right] \text {, then } \\
& \max \{\operatorname{BestMatch}(A, B), \operatorname{ApproxED}(A, B)\} \geq(1 / 2+\gamma) \operatorname{LCS}(A, B) \tag{29.2}
\end{align*}
$$

Proof. By Eq. (29.1), the premise implies that $1(B) \in\left[1 / 2 \pm \beta^{\prime}\right]$, and by symmetry also $1(A), 0(B) \in\left[1 / 2 \pm \beta^{\prime}\right]$. Therefore, $\operatorname{BestMatch}(A, B) \geq 1 / 2-\beta^{\prime}$.

[^101]Suppose that $\operatorname{BestMatch}(A, B)$ isn't big enough to satisfy Eq. (29.2) (otherwise we're done). Then,

$$
\operatorname{LCS}(A, B)>2(\operatorname{BestMAtch}(A, B)-\gamma) \geq 2\left(1 / 2-\left(\beta^{\prime}+\gamma\right)\right)=1-2\left(\beta^{\prime}+\gamma\right)
$$

Thus also

$$
\operatorname{ED}(A, B)=1-\operatorname{LCS}(A, B) \leq 2\left(\beta^{\prime}+\gamma\right)
$$

Therefore, by its approximation guarantee, we have that

$$
\begin{equation*}
\operatorname{ApproxED}(A, B) \geq 1-c \cdot \operatorname{ED}(A, B)-o(1) \geq 1-2 c\left(\beta^{\prime}+\gamma\right)-o(1) \geq 1 / 2+\gamma \tag{29.3}
\end{equation*}
$$

(The latter inequality follows by choosing $\beta^{\prime}$ and $\gamma$ sufficiently small.)

Setting $\beta^{\prime}=10 \beta$, we henceforth assume wlog that

$$
\begin{equation*}
0(A), 1(A), 0(B), 1(B) \notin[1 / 2 \pm 10 \beta] . \tag{29.4}
\end{equation*}
$$

### 29.4 Perfectly unbalanced strings

In this section we build on our assumptions from Eq. (29.1) and Eq. (29.4) from the previous section to complete the proof of our reduction.

Recall that we define $\alpha:=1(A)<1 / 2$, and by Eq. (29.1), we also have $0(B)=\alpha$. We partition each string into three contiguous substrings, where the extreme left and right substring are each of length $\alpha$ :

$$
\begin{aligned}
& A=L_{A} \cup M_{A} \cup R_{A} \\
& B=L_{B} \cup M_{B} \cup R_{B}
\end{aligned}
$$

$$
\left|L_{A}\right|=\left|R_{A}\right|=\left|L_{B}\right|=\left|R_{B}\right|=\alpha
$$

$$
\left|M_{A}\right|=\left|M_{B}\right|=1-2 \alpha
$$



Figure 29.1

We consider six cases for the proportions of 1's and 0's in $R_{A}, L_{A}, R_{B}, L_{B}$ as in Eq. (29.5). By Table 29.1, we know that those six cases cover all the possibilities.

Our six cases can be summarized in the following equation,

$$
\begin{cases}1\left(R_{B}\right) \leq \alpha / 2+2 \beta, 0\left(R_{A}\right) \leq \alpha / 2+2 \beta & \text { Case } 1  \tag{29.5}\\ 1\left(L_{B}\right) \leq \alpha / 2+2 \beta, 0\left(L_{A}\right) \leq \alpha / 2+2 \beta & \text { Case } 2 \\ 1\left(R_{B}\right) \leq \alpha / 2+\beta, 1\left(L_{B}\right) \leq \alpha / 2+\beta, 0\left(L_{A}\right)>\alpha / 2+2 \beta, 0\left(R_{A}\right)>\alpha / 2+2 \beta & \text { Case } 3 \\ 1\left(R_{B}\right)>\alpha / 2+2 \beta, 1\left(L_{B}\right)>\alpha / 2+2 \beta, 0\left(L_{A}\right) \leq \alpha / 2+\beta, 0\left(R_{A}\right) \leq \alpha / 2+\beta & \text { Case } 4 \\ 1\left(R_{B}\right)>\alpha / 2+\beta, 0\left(L_{A}\right)>\alpha / 2+\beta & \text { Case } 5 \\ 1\left(L_{B}\right)>\alpha / 2+\beta, 0\left(R_{A}\right)>\alpha / 2+\beta & \text { Case } 6\end{cases}
$$

Table 29.1: Fill all the six cases in Eq. (29.5) into the whole space. Note that $1+2+3$ means the combination of 1,2 and 3 covers it. 5,6 means any one of them covers it.

|  | $\begin{aligned} & 0\left(R_{A}\right) \quad \leq \\ & \alpha / 2+\beta, \\ & 0\left(L_{A}\right) \leq \alpha / 2+\beta \end{aligned}$ | $\begin{aligned} & 0\left(R_{A}\right) \\ & \alpha / 2+\quad \leq \\ & 0\left(L_{A}\right)>\alpha / 2+\beta \\ & \hline \end{aligned}$ | $\begin{aligned} & 0\left(R_{A}\right) \quad> \\ & \alpha / 2+\beta, \\ & 0\left(L_{A}\right) \leq \alpha / 2+\beta \\ & \hline \end{aligned}$ | $\begin{aligned} & 0\left(R_{A}\right) \\ & \alpha / 2+\beta \\ & 0\left(L_{A}\right)>\alpha / 2+\beta \\ & \hline \end{aligned}$ |
| :---: | :---: | :---: | :---: | :---: |
| $\begin{aligned} & 1\left(R_{B}\right) \leq \alpha / 2+ \\ & \beta, \quad 1\left(L_{B}\right) \leq \\ & \alpha / 2+\beta \end{aligned}$ | 1,2 | 1 | 2 | $1+2+3$ |
| $\begin{aligned} & 1\left(R_{B}\right) \quad \leq \\ & \alpha / 2+\quad+\quad \\ & 1\left(L_{B}\right)>\alpha / 2+\beta \end{aligned}$ | 1 | 1 | 6 | 6 |
| $\begin{array}{ll} 1\left(R_{B}\right)>\alpha / 2+ \\ \beta, \quad 1\left(L_{B}\right) & \leq \\ \alpha / 2+\beta \end{array}$ | 2 | 5 | 2 | 5 |
| $\begin{aligned} & 1\left(R_{B}\right) \\ & \alpha / 2+\quad+\quad \beta, \\ & 1\left(L_{B}\right)>\alpha / 2+\beta \end{aligned}$ | $1+2+4$ | 5 | 6 | 5,6 |

Case 1: $1\left(R_{B}\right) \leq \alpha / 2+2 \beta, 0\left(R_{A}\right) \leq \alpha / 2+2 \beta$
We split this case into three sub-cases, as follows:

$$
\begin{cases}1\left(R_{B}\right) \in[\alpha / 2 \pm 4 \beta], 0\left(R_{A}\right) \in[\alpha / 2 \pm 4 \beta] & \text { Case } 1(\mathrm{a})  \tag{29.6}\\ 1\left(R_{B}\right)<\alpha / 2-4 \beta, 0\left(R_{A}\right) \leq \alpha / 2+2 \beta & \text { Case 1(b) } \\ 1\left(R_{B}\right) \leq \alpha / 2+2 \beta, 0\left(R_{A}\right)<\alpha / 2-4 \beta & \text { Case 1(c) }\end{cases}
$$

Case 1(a): $1\left(R_{B}\right) \in[\alpha / 2 \pm 4 \beta], 0\left(R_{A}\right)=[\alpha / 2 \pm 4 \beta]$
At a high level, we want to split the original problem into two subproblems:
left-middle $\left(L_{A} \cup M_{A}, L_{B} \cup M_{B}\right) ;$
right $\left(R_{A}, R_{B}\right)$.

Running BestMatch on the left-middle subproblem gives a ( $1 / 2$ )-approx; the right subproblem is (approximately) balanced so Lemma 29.3 .2 (i.e. taking the better of BestMatch and ApproxED) gives better-than- $1 / 2$. The visualization of this case is presented in Figure 29.2.

We first want to upper bound $\operatorname{LCS}(A, B)$ as roughly the sum of LCSs of the two subproblems, but in general this may not be the case. Fix an optimal matching $\mu$ corresponding to a longest common substring between $A$ and $B$. Assume wlog (by symmetry w.r.t. switching $A$ and $B$ ) that $\mu\left(R_{A}\right) \subseteq R_{B}$, i.e. the LCS does not match any $R_{A}$ characters with characters from $L_{B} \cup M_{B}$. $\mu$ induces a new partition of $B$ into two ${ }^{6}$ contiguous

[^102]substrings $\widehat{L_{B}} \cup \widehat{R_{B}}$ such that $\mu\left(R_{A}\right) \subseteq \widehat{R_{B}} \subseteq R_{B}$. By optimality of $\mu$, we have
\[

$$
\begin{equation*}
\operatorname{LCS}(A, B)=\operatorname{LCS}\left(L_{A} \cup M_{A}, \widehat{L_{B}}\right)+\operatorname{LCS}\left(R_{A}, \widehat{R_{B}}\right) \tag{29.7}
\end{equation*}
$$

\]

Applying Fact 29.2.1 to both terms on the RHS, we have

$$
\begin{align*}
\operatorname{LCS}(A, B) \leq & \underbrace{\min \left\{1\left(L_{A} \cup M_{A}\right), 1\left(\widehat{L_{B}}\right)\right\}+\min \left\{0\left(L_{A} \cup M_{A}\right), 0\left(\widehat{L_{B}}\right)\right\}}_{=X} \\
& +\underbrace{\min \left\{1\left(R_{A}\right), 1\left(\widehat{R_{B}}\right)\right\}+\min \left\{0\left(R_{A}\right), 0\left(\widehat{R_{B}}\right)\right\}}_{=Y} . \tag{29.8}
\end{align*}
$$

We henceforth denote the left and right contributions to the bound on the LCS by $X$ and $Y$ respectively. (So $\operatorname{LCS}(A, B) \leq X+Y$.) We also define:

$$
Z:=\max \left\{\min \left\{1\left(L_{A} \cup M_{A}\right), 1\left(\widehat{L_{B}}\right)\right\}, \min \left\{0\left(L_{A} \cup M_{A}\right), 0\left(\widehat{L_{B}}\right)\right\}\right\} .
$$

(Observe that $Z \geq X / 2$.)
We now prove a lower bound on the LCS that our algorithm can find.

$$
\begin{align*}
\operatorname{GreEDY}\left(L_{A} \cup M_{A}, R_{A}, B\right) \geq & \max \left\{\min \left\{1\left(L_{A} \cup M_{A}\right), 1\left(\widehat{L_{B}}\right)\right\}, \min \left\{0\left(L_{A} \cup M_{A}\right), 0\left(\widehat{L_{B}}\right)\right\}\right\} \\
& +\max \left\{\min \left\{1\left(R_{A}\right), 1\left(\widehat{R_{B}}\right)\right\}, \min \left\{0\left(R_{A}\right), 0\left(\widehat{R_{B}}\right)\right\}\right\} \\
\geq & Z+Y / 2 . \tag{29.9}
\end{align*}
$$

We break into sub-cases, depending on the value of $Z$.

Case 1(a-i): $Z>\alpha / 2+10 \beta$ In this case, observe that

$$
\begin{aligned}
X-Z & =\min \left\{1\left(L_{A} \cup M_{A}\right), 1\left(\widehat{L_{B}}\right), 0\left(L_{A} \cup M_{A}\right), 0\left(\widehat{L_{B}}\right)\right\} & & \\
& \leq 1\left(L_{A} \cup M_{A}\right) & & \\
& =\alpha-1\left(R_{A}\right) & & (\alpha=1(A)) \\
& \leq \alpha / 2+4 \beta & & \text { (Case 1(a) assumption) } \\
& <Z-6 \beta & & \text { (Case 1(a-i) assumption) }
\end{aligned}
$$

Therefore, $Z>X / 2+3 \beta$. Combining with Eq. (29.8) and (29.9), we have that

$$
\operatorname{Greedy}\left(L_{A} \cup M_{A}, R_{A}, B\right) \geq Z+Y / 2>X / 2+Y / 2+3 \beta \geq \operatorname{LCS}(A, B) / 2+3 \beta
$$

Case 1(a-ii): $Z \leq \alpha / 2+10 \beta \quad$ By Eq. (29.7), we have

$$
\begin{array}{rlr}
\operatorname{LCS}(A, B) & \leq X+\operatorname{LCS}\left(R_{A}, \widehat{R_{B}}\right) & \\
& \leq 2 Z+\operatorname{LCS}\left(R_{A}, \widehat{R_{B}}\right) & (Z \geq X / 2) \\
& \leq \alpha+20 \beta+\operatorname{LCS}\left(R_{A}, \widehat{R_{B}}\right) & (\text { Case 1(a-ii) assumption) } \\
& \leq \alpha+20 \beta+\operatorname{LCS}\left(R_{A}, R_{B}\right) . & \left(\widehat{R_{B}} \subseteq R_{B}\right)
\end{array}
$$

For our purposes, this is effectively as good as bounding $\operatorname{LCS}(A, B)$ by the sum of LCSs of the left-middle and right subproblems.

We run BestMatch on the left-middle subproblem. We have that

$$
\begin{aligned}
& \operatorname{BESTMATCH}\left(L_{A} \cup M_{A}, L_{B} \cup M_{B}\right) \\
\geq & \min \left\{0\left(L_{A} \cup M_{A}\right), 0\left(L_{B} \cup M_{B}\right)\right\} \\
\geq & \alpha / 2-4 \beta
\end{aligned}
$$

We run BestMatch and ApproxED on $R_{A}, R_{B}$ and take the better of the two outcomes. We apply Lemma 29.3.2 to strings $R_{A}, R_{B}$ with $\beta^{\prime}=4 \beta / \alpha$. (Notice that by Case 1 (a) assumption, they are guaranteed to be approximately balanced to within $\pm 4 \beta$, or a relative $\pm 4 \beta / \alpha$.) We therefore have that

$$
\begin{array}{rlr} 
& \max \left\{\operatorname{BestMatch}\left(R_{A}, R_{B}\right), \operatorname{ApproxED}\left(R_{A}, R_{B}\right)\right\} & \\
\geq & (1 / 2+\gamma) \operatorname{LCS}\left(R_{A}, R_{B}\right) & \quad \text { (Lemma 29.3.2) }  \tag{Lemma29.3.2}\\
\geq & \operatorname{LCS}\left(R_{A}, R_{B}\right) / 2+\gamma \alpha / 2-O(\beta \gamma) \\
\geq & \operatorname{LCS}\left(R_{A}, R_{B}\right) / 2+\gamma \alpha / 2-O(\beta) . & \left(\operatorname{LCS}\left(R_{A}, R_{B}\right) \geq \alpha / 2-O(\beta)\right) \\
\quad(\gamma \leq 1)
\end{array}
$$

So in total, our algorithm finds a common substring of length at least

$$
\begin{array}{rlr} 
& \left(\alpha+\operatorname{LCS}\left(R_{A}, R_{B}\right)\right) / 2+\gamma \alpha-O(\beta) & \\
\geq & \operatorname{LCS}(A, B) / 2+\gamma \alpha / 2-O(\beta) & (\text { Eq. }(29.10))  \tag{29.10}\\
\geq & \operatorname{LCS}(A, B) / 2+\frac{2}{6} \gamma \alpha & (\gamma \alpha \gg \beta) \\
\geq & (1 / 2+\gamma / 6) \operatorname{LCS}(A, B) . & (\operatorname{LCS}(A, B) \leq 2 \alpha)
\end{array}
$$

Case 1(b): $1\left(R_{B}\right)<\alpha / 2-4 \beta, 0\left(R_{A}\right) \leq \alpha / 2+2 \beta$
Fix an optimal matching $\mu$. We further split this case into two sub-cases, depending on whether $\mu\left(R_{A}\right) \subseteq R_{B}$ or $R_{B} \subseteq \mu\left(R_{A}\right)$. (In Case 1(a) we could assume the former wlog by symmetry. Also notice that in general both may occur simultaneously.)

(b) $\operatorname{BestMAtch}\left(L_{A} \cup M_{A}, L_{B} \cup M_{B}\right)$ and max $\left(\operatorname{BestMatch}\left(R_{A}, R_{B}\right), \operatorname{ApproxED}\left(R_{A}, R_{B}\right)\right)$

Figure 29.2: Visualization of Case 1 (a) which is $0\left(R_{A}\right) \in[\alpha / 2 \pm 4 \beta]$ and $1\left(R_{B}\right) \in[\alpha / 2 \pm 4 \beta]$. If $0\left(\widehat{L_{B}}\right)>\alpha / 2+10 \beta$, we use GREEDY result. If $0\left(\widehat{L_{B}}\right) \leq \alpha / 2+10 \beta$, we use the result BestMatch + max(BestMatch, ApproxED).

If $\mu\left(R_{A}\right) \subseteq R_{B}$, define the partition $\widehat{L_{B}}, \widehat{R_{B}}$ as in Case 1 (a). We have

$$
\begin{align*}
\operatorname{LCS}(A, B) & =\operatorname{LCS}\left(L_{A} \cup M_{A}, \widehat{L_{B}}\right)+\operatorname{LCS}\left(R_{A}, \widehat{R_{B}}\right) \\
& \leq \underbrace{1\left(L_{A} \cup M_{A}\right)}_{\leq \alpha / 2+2 \beta}+\underbrace{0\left(\widehat{L_{B}}\right)+0\left(\widehat{R_{B}}\right)}_{\leq \alpha}+\underbrace{1\left(\widehat{R_{B}}\right)}_{\leq \alpha / 2-4 \beta}  \tag{Fact29.2.1}\\
& \leq 2 \alpha-2 \beta . \tag{29.11}
\end{align*}
$$

Similarly, if $\mu\left(R_{A}\right) \supseteq R_{B}$, we can define an analogous partition of $A$ into $\widehat{L_{A}}, \widehat{R_{A}}$ :

$$
\begin{align*}
\operatorname{LCS}(A, B) & =\operatorname{LCS}\left(\widehat{L_{A}}, L_{B} \cup M_{B}\right)+\operatorname{LCS}\left(\widehat{R_{A}}, R_{B}\right) \\
& \leq \underbrace{0\left(L_{B} \cup M_{B}\right)}_{\leq \alpha / 2-4 \beta}+\underbrace{1\left(\widehat{L_{A}}\right)+1\left(\widehat{R_{A}}\right)}_{\leq \alpha}+\underbrace{1\left(R_{B}\right)}_{\leq \alpha / 2-4 \beta}  \tag{Fact29.2.1}\\
& \leq 2 \alpha-8 \beta . \tag{29.12}
\end{align*}
$$

Either way, we have that $\operatorname{LCS}(A, B) \leq 2 \alpha-2 \beta$; therefore $\operatorname{Match}(A, B, 0)=\alpha$ guarantees a better-than-1/2-approximation.

Case $1(\mathrm{c}): 1\left(R_{B}\right) \leq \alpha / 2+2 \beta, 0\left(R_{A}\right)<\alpha / 2-4 \beta$
Follows analogously to Case 1(b).

Case 2: $1\left(L_{B}\right) \leq \alpha / 2+\beta, 0\left(L_{A}\right) \leq \alpha / 2+\beta$
We reverse the order of string $A$ and $B$, then the proof is the same as Case 1 .

Case 3: $1\left(R_{B}\right) \leq \alpha / 2+\beta, 1\left(L_{B}\right) \leq \alpha / 2+\beta, 0\left(L_{A}\right)>\alpha / 2+2 \beta$ and $0\left(R_{A}\right)>\alpha / 2+2 \beta$
We visualize this case in Figure 29.3a.
We show that simple applications of MATCH to the left, middle, and right substrings can guarantee a common string of at least $\alpha+2 \beta \geq \operatorname{LCS}(A, B) / 2+2 \beta$.

For the middle substrings, observe that $1\left(M_{A}\right)=1(A)-1\left(R_{A}\right)-1\left(L_{A}\right)>4 \beta$. We
can also lower bound $1\left(M_{B}\right)$ by:

$$
\begin{array}{rlrl}
1\left(M_{B}\right) & =1(B)-1\left(L_{B}\right)-1\left(R_{B}\right) & \\
& >\alpha+20 \beta-1\left(L_{B}\right)-1\left(R_{B}\right) & & \text { (Eq. (29.4)) }  \tag{29.4}\\
& \geq 18 \beta & & \text { (Premise of Case 3). }
\end{array}
$$

Therefore,

$$
\begin{equation*}
\operatorname{MATCH}\left(M_{A}, M_{B}, 1\right)=\min \left\{1\left(M_{A}\right), 1\left(M_{B}\right)\right\} \geq 4 \beta \tag{29.13}
\end{equation*}
$$

For the left substrings, observe that $0\left(L_{B}\right)=\left|L_{B}\right|-1\left(L_{B}\right)>\alpha / 2-\beta$. Therefore,

$$
\begin{equation*}
\operatorname{Match}\left(L_{A}, L_{B}, 0\right)=\min \left\{0\left(L_{A}\right), 0\left(L_{B}\right)\right\} \geq \alpha / 2-\beta \tag{29.14}
\end{equation*}
$$

Similarly,

$$
\begin{equation*}
\operatorname{MATCH}\left(R_{A}, R_{B}, 0\right)=\min \left\{0\left(R_{A}\right), 0\left(R_{B}\right)\right\} \geq \alpha / 2-\beta \tag{29.15}
\end{equation*}
$$

Summing up Eq. (29.13),(29.14),(29.15), our algorithm obtains a common string of length at least $\alpha+2 \beta$.

Case 4: $1\left(R_{B}\right)>\alpha / 2+2 \beta, 1\left(L_{B}\right)>\alpha / 2+2 \beta, 0\left(L_{A}\right) \leq \alpha / 2+\beta$ and $0\left(R_{A}\right) \leq \alpha / 2+\beta$
We visualize this case in Figure 29.3b.

If we switch $A$ and $B$, then the proof is the same as Case 3 .

Case 5: $1\left(R_{B}\right)>\alpha / 2+\beta$, and $0\left(L_{A}\right)>\alpha / 2+\beta$
We visualize this case in Figure 29.3c.

We apply Match to two subproblems to obtain a common substring of length greater than $\alpha+2 \beta \geq \operatorname{LCS}(A, B) / 2+2 \beta$.

Observe that $0\left(L_{B} \cup M_{B}\right)=0(B)-0\left(R_{B}\right)>\alpha / 2+\beta$.

$$
\operatorname{MATCH}\left(L_{A}, L_{B} \cup M_{B}, 0\right)=\min \left\{0\left(L_{A}\right), 0\left(L_{B} \cup M_{B}\right)\right\}>\alpha / 2+\beta
$$

By an analogous argument,

$$
\operatorname{MATCh}\left(M_{A} \cup R_{A}, R_{B}, 1\right)=\min \left\{1\left(M_{A} \cup R_{A}\right), 1\left(R_{B}\right)\right\}>\alpha / 2+\beta
$$

Case 6: $1\left(L_{B}\right)>\alpha / 2+\beta$, and $0\left(R_{A}\right)>\alpha / 2+\beta$

We visualize this case in Figure 29.3d.
We reverse the oder of string $A$ and $B$, then the proof is the same as Case 5 .

|  | $\alpha$ | $1-2 \alpha$ | $\alpha$ |
| :---: | :---: | :---: | :---: |
| A | $L_{A}$ | $M_{A}$ | $R_{A}$ |
| Case 3 | $\begin{gathered} \text { MATCH } 0 \\ 1\left(L_{B}\right)>\alpha / 2+2 \beta \\ \end{gathered}$ | Match 1 | $\begin{gathered} \text { MATCH } 0 \\ 1\left(R_{B}\right) \leq \alpha / 2+\beta=\alpha \\ \end{gathered}$ |
| $B$ | $L_{B}$ | $M_{B}$ | $R_{B}$ |

(a)

|  | $\alpha$ | $1-2 \alpha$ | ${ }^{\alpha}$ |
| :---: | :---: | :---: | :---: |
| A | $L_{A}$ | $M_{A}$ | $R_{A}$ |
| Case 4 | $\begin{aligned} & 0\left(L_{A}\right) \leq \alpha / 2+\beta \\ & \quad \mathrm{MATCH} 1 \\ & 1\left(L_{B}\right)>\alpha / 2+2 \beta^{2} \end{aligned}$ | Match 0 | $\begin{gathered} 0\left(R_{A}\right) \leq \alpha / 2+\beta \\ \text { МАТСН } 1 \\ 1\left(R_{B}\right)>\alpha / 2+2 \beta^{\prime} \end{gathered}$ |
| $B$ | $L_{B}$ | $M_{B}$ | $R_{B}$ |

(b)

(c)

(d)

Figure 29.3: Case 3-6.

```
Algorithm 29.1 Approximate LCS algorithm
    procedure ApproxLCS \((A, B, \alpha)\)
        Split \(A\) into three parts, \(L_{A}, M_{A}\) and \(R_{A}\) such that \(\left|L_{A}\right|=\left|R_{A}\right|=\alpha\), similarly for \(B\)
        Choose \(\beta\) to be sufficiently small constant
        if \(1\left(R_{B}\right) \leq \alpha / 2+2 \beta 0\left(R_{A}\right) \leq \alpha / 2+2 \beta\) then \(\quad \triangleright\) Case 1
            if \(1\left(R_{B}\right) \in[\alpha / 2 \pm 4 \beta], 0\left(R_{A}\right) \in[\alpha / 2 \pm 4 \beta]\) then \(\quad \triangleright\) Case 1(a)
                \(C, \widehat{L_{B}}, \widehat{R_{B}} \leftarrow \operatorname{Greedy}\left(L_{A} \cup M_{A}, R_{A}, B\right)\)
                \(Z \leftarrow \max \left\{\min \left\{1\left(L_{A} \cup M_{A}\right), 1\left(\widehat{L_{B}}\right)\right\}, \min \left\{0\left(L_{A} \cup M_{A}\right), 0\left(\widehat{L_{B}}\right)\right\}\right\}\)
                if \(Z \leq \alpha / 2+10 \beta\) then
                    \(C \leftarrow \operatorname{BestMatch}\left(L_{A} \cup M_{A}, L_{B} \cup M_{B}\right)\)
                        \(+\max \left\{\operatorname{BestMatch}\left(R_{A}, R_{B}\right), \operatorname{ApproxED}\left(R_{A}, R_{B}\right)\right\}\)
                end if
            else if \(1\left(R_{B}\right)<\alpha / 2-4 \beta, 0\left(R_{A}\right) \leq \alpha / 2+2 \beta\) then \(\quad \triangleright\) Case \(1(\mathrm{~b})\)
                \(C \leftarrow \operatorname{Match}(A, B, 0)\)
            else if \(1\left(R_{B}\right) \leq \alpha / 2+2 \beta, 0\left(R_{A}\right)<\alpha / 2-4 \beta\) then \(\quad \triangleright\) Case \(1(\mathrm{c})\)
                    Similar to Case 1(b)
            end if
        else if \(1\left(L_{B}\right) \leq \alpha / 2+\beta 0\left(L_{A}\right) \leq \alpha / 2+\beta\) then \(\quad \triangleright\) Case 2
            Similar to Case 1
        else if \(1\left(R_{B}\right), 1\left(L_{B}\right) \leq \alpha / 2+\beta, 0\left(L_{A}\right), 0\left(R_{A}\right)>\alpha / 2+2 \beta\) then \(\quad\) Case 3
            \(C \leftarrow \operatorname{Match}\left(L_{A}, L_{B}, 0\right)+\operatorname{Match}\left(M_{A}, M_{B}, 1\right)+\operatorname{Match}\left(R_{A}, R_{B}, 0\right)\)
        else if \(1\left(R_{B}\right), 1\left(L_{B}\right)>\alpha / 2+2 \beta, 0\left(L_{A}\right), 0\left(R_{A}\right) \leq \alpha / 2+\beta\) then \(\quad \triangleright\) Case 4
            Similar to Case 3
        else if \(1\left(R_{B}\right)>\alpha / 2+\beta, 0\left(L_{A}\right)>\alpha / 2+\beta\) then \(\triangleright\) Case 5
            \(C \leftarrow \operatorname{Match}\left(L_{A}, L_{B} \cup M_{B}, 0\right)+\operatorname{Match}\left(M_{A} \cup R_{A}, R_{B}, 1\right)\)
        else if \(1\left(L_{B}\right)>\alpha / 2+\beta, 0\left(R_{A}\right)>\alpha / 2+\beta\) then \(\quad \triangleright\) Case 6
            \(C \leftarrow \operatorname{Match}\left(L_{A} \cup M_{A}, L_{B}, 1\right)+\operatorname{Match}\left(R_{A}, M_{B} \cup R_{B}, 0\right)\)
        end if
        return \(C\)
    end procedure
```


## Chapter 30

## Longest Common Subsequence II

Longest common subsequence (LCS) is a classic and central problem in combinatorial optimization. While LCS admits a quadratic time solution, recent evidence suggests that solving the problem may be impossible in truly subquadratic time. A special case of LCS wherein each character appears at most once in every string is equivalent to the longest increasing subsequence problem (LIS) which can be solved in quasilinear time. In this work, we present novel algorithms for approximating LCS in truly subquadratic time and LIS in truly sublinear time. Our approximation factors depend on the ratio of the optimal solution size over the input size. We denote this ratio by $\lambda$ and obtain the following results for LCS and LIS without any prior knowledge of $\lambda$.

- A truly subquadratic time algorithm for LCS with approximation factor $O\left(\lambda^{3}\right)$.
- A truly sublinear time algorithm for LIS with approximation factor $O\left(\lambda^{3}\right)$.

Triangle inequality was recently used by Boroujeni et al. $\left[\mathrm{BEG}^{+} 18\right]$ and Chakraborty et al. $\left[\mathrm{CDG}^{+} 18\right]$ to present new approximation algorithms for edit distance. Our techniques for LCS extend the notion of triangle inequality to non-metric settings.

This part is based on previous following publication

- Aviad Rubinstein, Saeed Seddighin, Zhao Song, Xiaorui Sun

Approximation Algorithms for LCS and LIS with Truly Improved Running Times. FOCS 2019 [RSSS19]

### 30.1 Introduction

In this paper, we consider three of the most classic problems in combinatorial optimization: the longest common subsequence (LCS), the edit distance (ED), and the longest increasing subsequence (LIS). The LCS of two strings $A$ and $B$ is simply their longest (not necessarily contiguous) common subsequence. The edit distance is defined as the minimum number of character deletions, insertions, and substitutions required to transform $A$ into $B$. For the purpose of our discussion, we consider a more restricted definition of edit distance where substitutions are not allowed ${ }^{1}$. Longest increasing subsequence is equivalent to a special case of LCS where the input strings are permutations. All three problems are very fundamental and have been subject to a plethora of studies in the past few decades and specially in recent years [LMS98, BYJKK04, BES06, AO09, AKO10, SS10, BI15, ABW15, AHWW16, AB17, AR18, $\left.\mathrm{CGL}^{+} 19, \mathrm{BEG}^{+} 18, \mathrm{CDG}^{+} 18, \operatorname{HSSS} 19\right]$.

If the strings have length $n$, both LCS and ED can be solved in quadratic time $\left(O\left(n^{2}\right)\right)$ with dynamic programming. These running times are slightly improved to $O\left(n^{2} / \log ^{2}(n)\right)$ by Masek and Paterson [MP80], however, efforts to improve the running time to $O\left(n^{2-\Omega(1)}\right)$ for either edit distance or LCS were all futile.

In recent years, our understanding of the source of complexity for these problems tremendously improved thanks to a sequence of fine-grained complexity developments [ABW15, AHWW16]. We now know that assuming the strong exponential time hypothesis (SETH) [ABW15], or even weaker assumptions such as the orthogonal vectors conjecture (OVC) [ABW15] or branching-program-SETH [AHWW16], there are no truly sub-quadratic ${ }^{2}$ time algorithms for

[^103]LCS. Similar results also hold for edit distance [BI15].
Indeed, the classic approach to break the quadratic barrier for these problems is approximation algorithms. Note that for (multiplicative) approximations, LCS and edit distance are no longer equivalent (much like we have a 2-approximation algorithm for Vertex Cover, but Independent Set is NP-hard to approximate within near-linear factors).

For edit distance, an $\widetilde{O}\left(n+\Delta^{2}\right)$-time algorithm of [LMS98] (where $\Delta$ is the true edit distance between the strings) implies a linear-time $\sqrt{n}$-approximation algorithm. The approximation factor has been significantly improved in a series of works to $O\left(n^{3 / 7}\right)$ [BYJKK04], to $O\left(n^{0.34}\right)$ [BES06], to $O\left(2^{\tilde{O}(\sqrt{\log n})}\right)$ [AO09] ${ }^{3}$, and finally to polylogarithmic [AKO10]. A recent work of Boroujeni et al. $\left[\mathrm{BEG}^{+} 18\right]$ obtains a constant factor approximation quantum algorithm for edit distance that runs in truly subquadratic time. Finally, the breakthrough of Chakraborty et al. $\left[\mathrm{CDG}^{+} 18\right]$ gave a classic (randomized) constant factor approximation for edit distance in truly subquadratic time. A key component in both of the latest constant factor approximation algorithms is the application of triangle inequality (for edit distance between certain substrings of the input). A particular challenge in extending these ideas to LCS is that LCS is not a metric and in particular does not satisfy the triangle inequality.

Our understanding of the complexity of approximate solutions for LCS is embarrassingly limited. For general strings, there are several linear-time $1 / \sqrt{n}$-approximation algorithms based on sampling techniques. For alphabet size $|\Sigma|$, there is a trivial $1 /|\Sigma|-$ approximation algorithm that runs in linear time. Whether or not these approximation factors can be improved by keeping the running time subquadratic is one of the central

[^104]problems in fine-grained complexity. Very recently, both the general $1 / \sqrt{n}$-approximation factor, and, for binary strings, the $1 / 2$-approximation factor, have been slightly improved ([HSSS19] and [RS19b], respectively). There are a few fine-grained complexity results for approximate LCS, but they only hold against deterministic algorithms, and rely on very strong assumptions [AB17, AR18, CGL+19].

### 30.1.1 Our Results

For simplicity, we use $\operatorname{Ics}(A, B)$ to denote the size (not the whole sequence) of the longest common subsequence for two strings $A$ and $B$. Similarly, we use $\operatorname{ED}(A, B)$ to denote the edit distance and $\operatorname{lis}(A)$ for the size of the longest common subsequence. We sometimes normalize the solution by the length of the strings so that the size of the solution remains in the interval $[0,1]$. We refer to the normalized solutions by $\|\operatorname{lcs}(A, B)\|=\operatorname{Ics}(A, B) / n$ and $\|\mathrm{ED}(A, B)\|=\mathrm{ED}(A, B) / 2 n$ (here both strings have equal length $n$ ), and $\|\operatorname{lis}(A)\|=$ $\operatorname{lis}(A) / n$. In this way, $\|\operatorname{ED}(A, B)\|+\|\operatorname{lcs}(A, B)\|=1$ (assuming both strings have equal length).

As mentioned earlier, recent developments for edit distance are based on a simple but rather useful observation. Edit distance satisfies triangle inequality, or in other words, given three strings $A_{1}, A_{2}, A_{3}$ of length $n$ such that $\left\|\operatorname{ED}\left(A_{1}, A_{2}\right)\right\| \leq \delta$ and $\left\|\operatorname{ED}\left(A_{2}, A_{3}\right)\right\| \leq \delta$ hold, we can easily imply that $\left\|\operatorname{ED}\left(A_{1}, A_{3}\right)\right\| \leq 2 \delta$. While Ics does not satisfy the triangle inequality in any meaningful way, it does, on average, satisfy the following birthday-paradoxlike property that we call birthday triangle inequality.

Property 30.1.1 (birthday triangle inequality). Consider three equal-length strings $A_{1}, A_{2}$, and $A_{3}$ such that $\left\|\operatorname{lcs}\left(A_{1}, A_{2}\right)\right\| \geq \lambda$ and $\left\|\operatorname{lcs}\left(A_{2}, A_{3}\right)\right\| \geq \lambda$. If the common subsequences
correspond to random indices of each string, we expect that $\left\|\operatorname{lcs}\left(A_{1}, A_{3}\right)\right\| \geq \lambda^{2}$.

Of course, this is not necessarily the case in general. More precisely, it is easy to construct examples ${ }^{4}$ in which $\left\|\operatorname{lcs}\left(A_{1}, A_{2}\right)\right\|=1 / 2$ and $\left\|\operatorname{lcs}\left(A_{2}, A_{3}\right)\right\|=1 / 2$, but $\left\|\operatorname{lcs}\left(A_{1}, A_{3}\right)\right\|=$ 0 . Our first main result shows that while it only holds on average, we can algorithmically replace the triangle inequality for edit distance with the birthday triangle inequality on worst case inputs.

Theorem 30.1.2 (Main Theorem, formally stated as Theorem 30.2.1). Given strings $A, B$ both of length $n$ such that $\|\operatorname{LCS}(A, B)\|=\lambda$, we can approximate the length of the LCS between the two strings within an $O\left(\lambda^{3}\right)$ factor in subquadratic time. The approximation factor improves to $(1-\epsilon) \lambda^{2}$ when $1 / \lambda$ is constant.

We remark that our algorithm is actually able to output the whole sequence of the solution, but we only focus on estimating the size of the solution for simplicity. We begin by comparing our main theorem to previous work on edit distance. In this case, $1 / \lambda$ is constant w.l.o.g. ${ }^{5}$ and therefore the approximation factor of our algorithm is $(1-\epsilon) \lambda^{2}$. If $\delta=\|\mathrm{ED}(A, B)\|$, then our LCS algorithm outputs a transformation from $A$ to $B$ using at most $2 n\left(1-(1-\epsilon)(1-\delta)^{3}\right)$ operations. Observe that when the strings are not overly close and $\delta=\Omega(1)$ by scaling $\epsilon$, we already recover a $\left(3+\epsilon^{\prime}\right)$-approximation for edit distance in truly subquadratic time. For mildly far strings, say $\delta=0.1$, a more careful look at the expansion of $(1-\delta)^{3}$ reveals that we save an additive $\Theta\left(\delta^{2}\right)$ in the approximation factor. For example, with $\delta=0.1$ our approximation factor for edit distance is 2.71 instead of 3 .

[^105]An interesting implication of our main result is for LCS over a large alphabet $\Sigma$, where the optimum $\|\operatorname{LCS}(A, B)\|$ may be much smaller than 1 . This is believed to be the hardest regime for approximation algorithms (and indeed the only one for which we have any conditional hardness of approximation results [AB17, AR18, CGL $\left.{ }^{+} 19\right]$ ). Here, we consider instances that satisfy a mild balance assumption: we assume that there is a character that appears with frequency at least $1 /|\Sigma|$ in both strings ${ }^{6}$. Then, our main theorem implies an $O\left(1 /|\Sigma|^{3 / 4}\right)$-approximate solution in truly subquadratic time (the first improvement over the trivial $1 /|\Sigma|$ approximation in this regime).

Corollary 30.1.3 (LCS, formally stated as Corollary 30.2.3). Given a pair of strings $(A, B)$ of length $n$ over alphabet $\Sigma$ that satisfy the balance condition, we can approximate their LCS within an $O\left(|\Sigma|^{3 / 4}\right)$ factor in truly subquadratic time.

Next, we show that a similar result can be obtained for LIS. Perhaps coincidentally, the approximation factor of our algorithm is also $O\left(\lambda^{3}\right)$ which is same to LCS, but the technique is completely different. Although LIS can be solved exactly in time $O(n \log n)$, there have been several attempts to approximate the size of LIS and related problems in sublinear time [Sch61, Fre75, DGL+99, EKK ${ }^{+} 00$, Fis04, ACCL07, SS10]. The best known solution is due to the work of Saks and Seshadhri [SS10] that obtains a (1+ $)$-approximate algorithm for LIS in polylogarithmic time, when the solution size is at least a constant fraction of the input size ${ }^{7}$. In other words, if $\|\operatorname{lis}(A)\|=\lambda$ and $1 / \lambda$ is constant, their algorithm approximates

[^106]$\operatorname{lis}(A)$ in polylogarithmic time. However, this only works if $1 / \lambda$ is constant and even if $1 / \lambda$ is logarithmically large, their method fails to run in sublinear time ${ }^{8}$. We complement the work of Saks and Seshadhri [SS10] by presenting a result for LIS similar to our result for LCS. More precisely, we show that when $\|\operatorname{lis}(A)\|=\lambda$, an $O\left(\lambda^{3}\right)$ approximation of LIS can be obtained in truly sublinear time. Although our approximation factor is worse than that of [SS10], our result works for any (not necessarily constant) $\lambda$.

Theorem 30.1.4 (LIS, formally stated as Theorem 30.4.8). Given an array $A$ of $n$ integer numbers such that $\|\operatorname{lis}(A)\|=\lambda$. We can approximate the length of the LIS for $A$ in sublinear time within a factor $O\left(\lambda^{3}\right)$.

If one favors the running time over the approximation factor, it is possible to improve the exponent of $n$ in the running time down to any constant $\kappa>0$ at the expense of incurring a larger multiplicative factor to the approximation.

### 30.1.2 Preliminaries

In LCS or edit distance, we are given two strings $A$ and $B$ as input. We assume for simplicity that the two strings have equal length and refer to that by $n$. In LCS, the goal is to find the largest subsequence of the characters which is shared between the two strings. In edit distance, the goal is to remove as few characters as possible from the two strings such that the remainders for the two strings are the same. We use $\operatorname{lcs}(A, B)$ and $\operatorname{ED}(A, B)$ to denote the size of the longest common subsequence and the edit distance of two strings $A$ and $B$.

[^107]In LIS, the input contains an array $A$ of $n$ integer numbers and the goal is to find a sequence of elements of $A$ whose values (strictly) increase as their indices increase. For LIS, we denote the solution size for an array $A$ by lis $(A)$. We also use lis ${ }^{[\alpha, \beta]}(A)$ to denote the size of the longest increasing subsequence subject to elements whose values lie in range $[\alpha, \beta]$. Longest increasing subsequence is equivalent to LCS when the inputs are two permutations of $n$ distinct characters.

Finally, we define a notation to denote the normalized solution sizes. For LCS, we denote the normalized solution size by $\|\operatorname{lcs}(A, B)\|=\operatorname{Ics}(A, B) / \sqrt{|A||B|}$ for $A$ and $B$ and we use $\|\operatorname{ED}(A, B)\|=\operatorname{ED}(A, B) /(2 \sqrt{|A||B|})$ for edit distance. Note that, when the two strings have equal length we have $\|\operatorname{ED}(A, B)\|+\|\operatorname{lcs}(A, B)\|=1$. Similarly, for longest increasing subsequence, we denote by $\|\operatorname{lis}(A)\|=\operatorname{lis}(A) /|A|$ the normalized solution size. We usually refer to the size of the input array by $n$.

Throughout this paper, we call an algorithm $f(\lambda)$-approximation for LCS if it is able to distinguish the following two cases: i) $\|\operatorname{lcs}(A, B)\| \geq \lambda$ or ii) $\|\operatorname{lcs}(A, B)\|<\lambda f(\lambda)$. A similar definition carries over to LIS. Once an $f(\lambda)$-approximation algorithm is provided for either LCS or LIS, one can turn it into an algorithm that outputs a solution with size $f(\lambda)(1-\epsilon) \lambda n$ provided that the optimal solution has a size $\lambda n$. The algorithm is not aware of the value of $\lambda$ but will start with $\lambda_{0}=1$ and iteratively multiply $\lambda_{0}$ by $1-\epsilon$ until a solution is found.

### 30.1.3 Techniques Overview

Our algorithm for LCS is closely related to the recent developments for edit distance $\left[\mathrm{BEG}^{+} 18, \mathrm{CDG}^{+} 18\right]$. We begin by briefly explaining the previous ideas for approxi-
mating edit distance and then we show how we use these techniques to obtain a solution for LCS. Finally, in Section 30.1.3.2 we outline our algorithm for LIS.

### 30.1.3.1 Summary of Previous ED Techniques

Indeed, edit distance is hard only if the two strings are far $(\|\operatorname{ED}(A, B)\|=\delta$ and $\left.\delta=n^{-o(1)}\right)$ otherwise the $O\left(n+(n \delta)^{2}\right)$ algorithm of Landau et al. [LMS98] computes the solution in truly subquadratic time. The algorithm of Chakraborty et al. $\left[\mathrm{CDG}^{+} 18\right]$ for edit distance has three main steps that we briefly discuss in the following.

Step 0 (window-compatible solutions): In the first step, they construct a set of windows $W_{A}$ for string $A$ and a set of windows $W_{B}$ for string $B$. Each window is essentially a substring of the given string. Let $k$ denote the total number of windows of $W_{A} \cup W_{B}$. For simplicity, let all the windows have the same size $d$ and $n \simeq O(k d)^{9}$. The construction features two key properties: 1) provided that the edit distances of the windows between $W_{A}$ and $W_{B}$ are available, one can recover a $1+\epsilon$ approximation of edit distance in time $\widetilde{O}\left(n+k^{2}\right)$ via dynamic programming. 2) $k^{2} \times d^{2} \simeq O\left(n^{2}\right)$. That is, if we naively compute the edit distance of every pair of windows, the overall running time would still asymptotically be the same as that of the classic algorithm.

In order to obtain a solution for edit distance, it suffices to know the distances between the windows. However, Chakraborty et al. [CDG $\left.{ }^{+} 18\right]$ show that knowing the distances between most of the window pairs is enough to obtain an approximately optimal solution for edit distance. More precisely, if the distances are not correctly approximated for at

[^108]most $O\left(k^{2-\Omega(1)}\right)$ pairs, we can still obtain an approximate solution for edit distance. Step 1 provides estimates for the distances of the windows which is approximately correct except for $O\left(k^{2-\Omega(1)}\right)$ many pairs and Step 2 shows how this can be used to obtain a solution for edit distance. Discretization simplifies the problem substantially. For a fixed $0 \leq \delta \leq 1$, they introduce a graph $\mathrm{G}_{\delta}$ where the nodes correspond to the windows and an edge between window $w_{i} \in W_{A}$ and window $w_{j} \in W_{B}$ means that $\left\|\operatorname{ED}\left(w_{i}, w_{j}\right)\right\| \leq \delta$. If we are able to construct $\mathrm{G}_{\delta}$ for logarithmically different choices of $\delta$, we can as well estimate the distances within a $1+\epsilon$ factor for the windows. Therefore the problem boils down to constructing $\mathrm{G}_{\delta}$ for a fixed given $\delta$ without computing the edit distance between all pairs of windows.

Step 1 (sparsification via triangle inequality): This step is the heart of the algorithm. Suppose we choose a high-degree vertex $v$ from $\mathrm{G}_{\delta}$ and discover all its incident edges by computing its edit distance to the rest of the windows. Triangle inequality implies that every pair of windows in $N(v)$ has a distance bounded by $2 \delta$. Therefore by losing a factor 2 in the approximation, one can put all these edges in $\mathrm{G}_{\delta}$ and not compute the edit distances explicitly. Although this does save some running time, in order to make sure the running time is truly subquadratic, we need to make a similar argument for paths of length 3 and thereby lose a factor 3 in the approximation. This method sparsifies the graph and what remains is to discover the edges of a sparse graph.

Step 2 (discovering the edges of the sparse graph): Step 1 uses triangle inequality and discovers many edges between the vertices of $\mathrm{G}_{\delta}$. However, it may not discover all the edges completely. When in the remainder graph, the degrees are small (and hence the graph
is sparse) triangle inequality does not offer an improvement and thus a different approach is required. Roughly speaking, Chakraborty et al. $\left[\mathrm{CDG}^{+} 18\right]$ subsample the windows of $W_{A}$ into a smaller set $S$ and discover all pairs of windows $w_{i} \in S$ and $w_{j} \in W_{B}$ such that edge $(i, j)$ is not discovered in Step 1. Next, they compute the edit distance of each pair of windows $\left(w_{i}, w_{j}\right), w_{i} \in W_{A}, w_{j} \in W_{B}$ such that there exist two nearby windows ( $w_{a}, w_{b}$ ) satisfying $w_{a} \in S, w_{b} \in W_{B}$ and the edge between $w_{a}$ and $w_{b}$ was missed in Step 1. The key observation is that even though this procedure does not discover all the edges, the approximated distances lead to an approximate solution for edit distance.

### 30.1.3.2 LCS

Our algorithm for LCS mimics the same guideline. In addition to this, Steps 0 and 2 of our algorithm are LCS analogues of the ones used by Chakraborty et al. [CDG $\left.{ }^{+} 18\right]$. The main novelty of our algorithm is Step 1 which is a replacement for triangle inequality. Recall that unlike edit distance, triangle inequality does not hold for LCS.

Challenge 30.1.5. How can we introduce a notion similar to triangle inequality to a nonmetric setting such as LCS?

We introduce the notion of birthday triangle inequality to overcome the above difficulty. Given windows $w_{1}$, $w_{2}$, and $w_{3}$ of size $d$ such that $\left\|\operatorname{lcs}\left(w_{1}, w_{2}\right)\right\| \geq \lambda$ and $\left\|\operatorname{lcs}\left(w_{2}, w_{3}\right)\right\| \geq$ $\lambda$ hold, what can we say about the LCS of $w_{1}$ and $w_{3}$ ? In general, nothing! $\left\|\operatorname{lcs}\left(w_{1}, w_{3}\right)\right\|$ could be as small as 0 . However, let us add some randomness to the setting. Think of the LCS of $w_{1}$ and $w_{2}$ as a matching from the characters of $w_{1}$ to $w_{2}$ and similarly the LCS for $w_{2}$ and $w_{3}$ as another matching between characters of $w_{2}$ and $w_{3}$. Assume (for the sake of the thought experiment) that the characters of $w_{2}$ appear randomly in each matching. Since


Figure 30.1: Birthday paradox for triangle inequality: let $w_{1}, w_{2}, w_{3}$ be three windows of length $d=8$ and assume $\lambda=1 / 2$. The LCS between $w_{1}$ and $w_{2}$ is $\lambda d=4$ and the LCS between $w_{2}$ and $w_{3}$ is $\lambda d=4$. Finally due to birthday paradox, we expect that the LCS between $w_{1}$ and $w_{3}$ is $\lambda^{2} d=2$.
$\left\|\operatorname{lcs}\left(w_{1}, w_{2}\right)\right\| \geq \lambda$, each character of $w_{2}$ appears with probability at least $\lambda$ in the matching between $w_{1}$ and $w_{2}$. A similar argument implies that each character of $w_{2}$ appears with probability $\lambda$ in the matching of $w_{2}$ and $w_{3}$. Thus, (assuming independence), each character of $w_{2}$ appears in both matchings with probability $\lambda^{2}$. This means that in expectation, there are $\lambda^{2} d$ paths of length 2 between $w_{1}$ and $w_{3}$ which suggests $\left\|\operatorname{lcs}\left(w_{1}, w_{3}\right)\right\| \geq \lambda^{2}$ as shown in Figure 30.1. This is basically birthday paradox used for the sake of triangle inequality.

Replacing triangle inequality by birthday triangle inequality is particularly challenging since birthday triangle inequality only holds on average. In contrast, triangle inequality holds for any tuple of arbitrary strings. Most of our technical discussions is dedicated to proving that we can algorithmically use birthday triangle inequality to obtain a solution for the worst-case scenarios. The most inconvenient step of our analysis is to show that our algorithm estimates the LCS of most window pairs in the sparsification phase. While this is straightforward for edit distance, birthday triangle inequality requires a deeper analysis
of the underlying graph. In particular, we need to prove that if the undiscovered edges are too many, then birthday triangle inequality can be applied to certain neighborhoods of the graph.

There are two difficulties that we face here. On one hand, in order to apply birthday triangle inequality to a subgraph, we need to have enough structure for that subgraph to show the implication can be made. On the other hand, our assumptions cannot be too strong, otherwise such neighborhoods may not cover the edges of the graph. Therefore, the first challenge that we need to overcome is characterizing subgraphs in which birthday triangle inequality is guaranteed to be applicable. Our suggestion is the bi-cliques structure. Although combinatorial techniques seem unlikely to prove this, we use the Blakley-Roy inequality to show that in a large enough bi-clique, we can use birthday triangle inequality to imply a bound on the LCS of certain pairs. The second challenge is to prove that if the underlying graph is dense enough, the graph contains many bi-cliques that cover almost all the edges that we plan to discover. This is again a challenging graph theoretic problem. We leverage extremal graph theory tools such as Turan's theorem for cliques and bi-cliques to obtain this bound.

Similar to edit distance, we construct a set $W=W_{A} \cup W_{B}$ of $k$ windows in Step 0 and aim to sparsify the edges of the Ics-graph in Step 1 . Our construction ensures that $k d \simeq \Theta(n)$ and that knowing the LCS of the window pairs suffices to approximate the LCS of the two strings. For a threshold $0 \leq \lambda \leq 1$, define a matrix $\mathrm{O}:[k] \times[k] \rightarrow\{0,1\}$ to be a matrix which identifies whether $\left\|\operatorname{Ics}\left(w_{i}, w_{j}\right)\right\| \geq \lambda$. In other words, $\mathrm{O}[i][j]=1 \Longleftrightarrow\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \lambda$. For an $0<\alpha \leq 1$, we call a matrix $\mathrm{O}_{\alpha}$ an $\alpha$ approximation of O if it meets the following two conditions:

$$
\mathrm{O}_{\alpha}[i][j]=0 \Longrightarrow\left\|\operatorname{lss}\left(w_{i}, w_{j}\right)\right\|<\lambda \quad \text { and } \quad \mathrm{O}_{\alpha}[i][j]=1 \Longrightarrow\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \alpha \cdot \lambda
$$

Notice that $\mathrm{O}_{\alpha}$ gives more flexibility than O for the cases that $\lambda \alpha \leq \operatorname{lcs}\left(w_{i}, w_{j}\right)<\lambda$. That is, both 0 and 1 are acceptable in these cases. Indeed an $\alpha$ approximation algorithm for the above problem is enough to obtain an $\alpha$ approximation algorithm for LCS. However, this is not necessary as Step 2 allows for incorrect approximation for up to $k^{2-\Omega(1)}$ many window pairs. Therefore, the problem of approximating LCS essentially boils down to approximating O for a given basket of windows $W=W_{a} \cup W_{b}$ and a fixed $\lambda$ by allowing sufficiently small error in the output. A naive solution is to iterate over all pairs $w_{i}$ and $w_{j}$ and compute $\operatorname{Ics}\left(w_{i}, w_{j}\right)$ in time $O\left(d^{2}\right)$ and determine O accordingly. However, this amounts to a total running time of $O\left(k^{2} d^{2}\right)$ which is quadratic and not desirable. In order to save time, we need to compute the LCS of fewer than $k^{2}$ pairs of windows. To make this possible, we allow our algorithm to miss up to $O\left(k^{2-\Omega(1)}\right)$ edges of the graph. Step 2 ensures that this does not hurt the approximation factor significantly.

We construct a graph from the windows wherein each vertex corresponds to a window and each edge identifies a pair with a large LCS (in terms of $\lambda$ ). Let us call this graph the Ics-graph and denote it by $G_{\lambda}$. The goal is to detect the edges of the graph by allowing false-positive. As we discussed earlier, the hard instances of the problem are the cases where the Ics-graph is dense for which we need a sparsifier. Roughly speaking, in our sparsification technique, our algorithm constructs another graph $\widehat{G}_{\lambda}$ such that $\widehat{G}_{\lambda}$ is valid in the sense that the edges of $\widehat{\mathrm{G}}_{\lambda}$ correspond to pairs of windows with large enough LCS. In addition to this, our algorithm guarantees that after the removal of the edges of $\widehat{G}_{\lambda}$ from $G_{\lambda}$ the remainder is
sparse. In other words, $\left|E\left(\mathrm{G}_{\lambda}\right) \backslash E\left(\widehat{\mathrm{G}}_{\lambda}\right)\right|=O\left(\left|V\left(\mathrm{G}_{\lambda}\right)\right|^{2-\Omega(1)}\right)$. Of course, if the overall running time of the sparsification phase is truly subquadratic, the error of undiscovered edges can be addressed by the techniques of $\left[\mathrm{CDG}^{+} 18\right]$ in Step 2.

Below, we bring a formal definition for sparsification.

$$
\begin{aligned}
& \text { sparsification } \\
& \text { input: Windows } w_{1}, w_{2}, \ldots, w_{k} \text {, parameters } \lambda \text {, and } \alpha \text {. } \\
& \text { solution: A matrix } \widehat{\mathrm{O}}_{\alpha} \in\{0,1\}^{k \times k} \text { such that: } \\
& \text { - } \widehat{\mathrm{O}}_{\alpha}[i][j]=1 \Longrightarrow\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \alpha \cdot \lambda \\
& \text { - } \mid\left\{(i, j) \mid\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \lambda \text { and } \widehat{\mathrm{O}}_{\alpha}[i][j]=0\right\} \mid=k^{2-\Omega(1)}
\end{aligned}
$$

We present two sparsification techniques for LCS. The first one (Section 30.3.1), has an approximation factor of $(1-\epsilon) \cdot \lambda^{2}$. In Section 30.3 .2 we present another sparsification technique that has a worse approximation factor $O\left(\lambda^{3}\right)$ but leaves fewer edges behind. Although the second sparsification technique has a worse approximation factor, it has the advantage that the number of edges that remain in the sparse graph is truly subquadratic regardless of the value of $\lambda$ and therefore it extends our solution to the case that $\lambda=o(1)$ (see Section 30.3.2 for a detailed discussion).

Let us note one last algorithmic challenge to keep in mind before we begin to describe our sparsification techniques. For edit distance, if window pairs $\left(w_{1}, w_{2}\right)$ and $\left(w_{2}, w_{3}\right)$ are close, we are guaranteed that $w_{1}$ and $w_{3}$ are also close; for longest common subsequence, we will argue that $\left(w_{1}, w_{3}\right)$ are likely to be have a long LCS (for a "random" choice of $\left(w_{1}, w_{3}\right)$ ). Nonetheless, in order to add $\left(w_{1}, w_{3}\right)$ as an edge to our graph we have to verify that their

LCS is indeed long. If we were to verify an edge naively, we would need as much time as computing the LCS between $\left(w_{1}, w_{3}\right)$ from scratch!

Sparsification 1, ( $1-\epsilon) \lambda^{2}$-approximation: Similar to edit distance, applying birthday triangle inequality to paths of length 2 for LCS does not improve the running time significantly. Therefore, we need to use birthday triangle inequality for paths of length 3 . To this end, we define the notion of constructive tuples as follows: a tuple $\left\langle w_{i}, w_{a}, w_{b}, w_{j}\right\rangle$ is an $(\epsilon, \lambda)$ constructive tuple, if we have $\left\|\operatorname{lcs}\left(w_{i}, w_{a}\right)\right\| \geq \lambda,\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \lambda,\left\|\operatorname{lcs}\left(w_{b}, w_{j}\right)\right\| \geq \lambda$ and by taking the intersection of the three LCS matchings, we are able to imply $\left\|\operatorname{lcs}\left(w_{a}, w_{b}\right)\right\| \geq$ $(1-\epsilon) \lambda^{3}$ (see Figure 30.2 for an example). Taking the intersection of the matchings can be done in linear time which is faster than computing the LCS.

Our sparsification technique here is simple but the analysis is very intricate. We subsample a set $S$ of windows and compute the LCS of every window in $S$ and all other windows. We set $|S|=k^{\gamma} \log k$, where $\gamma \in(0,1)$. At this point, for some pairs, we already know their LCS. However, if neither $w_{i}$ nor $w_{j}$ is in $S$, we do not know if $\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \lambda$ or not. Therefore, for such pairs, we try to find windows $w_{a}, w_{b} \in S$ such that $\left\langle w_{i}, w_{a}, w_{b}, w_{j}\right\rangle$ is constructive. If such a constructive tuple is found for a pair of windows, then we conclude that their normalized LCS is at least $(1-\epsilon) \lambda^{3}$.

All that remains is to argue that this method discovers almost all the edges of the Ics-graph $\mathrm{G}_{\lambda}$ and the number of undiscovered edges is $k^{2-\Omega(1)}$. This is the most difficult part of the analysis. We note that proving the existence of only one constructive tuple is already non-trivial even when $G_{\lambda}$ is complete. However, our goal is to show almost all the edges are discovered via constructive tuples when $G_{\lambda}$ is dense (and of course not necessarily


Figure 30.2: Let $w_{i}, w_{a}, w_{b}$ and $w_{j}$ denote four windows and each of them has length $d=8$. This figure shows how the intersection of the edges of three windows are taken in order to construct a solution for the LCS of $w_{i}$ and $w_{j}$. If the size of the intersection is large, then such a tuple is called constructive. The solid lines represent LCS between two strings, and the dashed line represents the intersection of the three LCSs.
complete).
Define a pair of windows $w_{i}$ and $w_{j}$ to be well-connected, if there are at least $k^{2-\gamma}$ different $\left(w_{a}, w_{b}\right)$ pairs such that $\left\langle w_{i}, w_{a}, w_{b}, w_{j}\right\rangle$ is $(\epsilon, \lambda)$-constructive. Since each window appears in $S$ with probability $k^{\gamma-1} \log k$, for each well-connected pair we find one constructive tuple via our algorithm with high probablity. Therefore, we need to prove that the total number of pairs $\left(w_{i}, w_{j}\right)$ such that $\left(w_{i}, w_{j}\right)$ is not well-connected but $\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \lambda$ is subquadratic. Let us put these edges in a new graph $N G_{\lambda}$ whose vertices are all the windows. We first leverage the Blakley-Roy inequality and a double counting technique to prove that if $N G_{\lambda}$ has a large complete bipartite subgraph, then there is one constructive tuple which includes only the vertices of this subgraph (Lemma 30.3.2). Next, we apply the Turan's theorem to show that if $\mathrm{NG}_{\lambda}$ is dense, then it has a lot of large complete bipartite subgraphs. Finally, we use a probabilistic method to conclude that $N G_{\lambda}$ cannot be too dense otherwise
there are a lot of constructive tuples in the graph which implies that at least one edge ( $w_{i}, w_{j}$ ) in $N G_{\lambda}$ is well-connected. This is not possible since all the well-connected pairs are detected in our sparsification algorithm with high probability.

The above argument proves that if we sparsify our graph using our sparsification algorithm, the remainder graph would have a subquadratic number of edges. Therefore after plugging Step 2 into the algorithm, the running time remains subquadratic. However, since Turan theorem gives us a weak bound, the running time of the algorithm using this sparsification is $O\left(n^{2-\Omega(\lambda)}\right)$ and is only truly subquadratic when $1 / \lambda$ is constant.

The above method also gives a nice insight into computing edit distance. As we show in Section 30.3, if our Ics-graph or similarly ED-graph is very dense, it contains bipartite complete subgraphs that lead to strong implications. For LCS, we obtain a bound of ( $1-$ є) $\lambda^{2}$ for the approximation factor of the sparsifier. For edit distance, this may lead to an approximation factor better than 3. In particular, if one gives a positive answer to Question 30.1.6 for some $\alpha$, this technique improves the approximation factor of estimating edit distance in subquadratic time to $2+\alpha+\epsilon$. The authors are not aware of any counterexample for the question even when $\alpha=0$ and it is quite possible that this technique improves the approximation factor for edit distance down to $2+\epsilon$.

Question 30.1.6. Let $X$ and $Y$ be two sets of strings of length $n$ such that for each string $A \in X$ and $B \in Y$ we have $\|\operatorname{ED}(A, B)\| \leq \delta$. If $1 / \delta$ is constant and the sizes of $X$ and $Y$ are large enough constants (much larger than $1 / \delta$ ) can we imply that there exist two strings in $X$ or two strings in $Y$ whose normalized distance is bounded by $(1+\alpha) \delta$ ?

We remark that birthday triangle inequality proves Question 30.1.6 for $\alpha=1-O(\delta)$
even when $|X|=1$. While we do not provide a formal proof, we believe that this already gives a $3-o(1)$ approximation algorithm for edit distance in truly subquadratic time which slightly breaks the $3+\epsilon$ barrier suggested in previous work $\left[\mathrm{BEG}^{+} 18, \mathrm{CDG}^{+} 18\right]$.

Sparsification 2, $O\left(\lambda^{3}\right)$-approximation: In Section 30.3.1, we present another sparsification method that although gives us a slightly worse approximation factor $O\left(\lambda^{3}\right)$ it always leaves a truly subquadratic number of edges behind and therefore the running time of the algorithm would be truly subquadratic regardless of the parameter $\lambda$. This sparsification is based on a novel data structure.

Let opt ${ }_{i, a}$ denote the longest common subsequence of $w_{i}$ and $w_{a}$ (with some fixed tiebreaking rule, e.g. lexicographically first). Define $\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)$ to be the size of the longest common subsequence between opt ${ }_{i, a}$ and $w_{j}$. Notice that this definition is no longer symmetric. Let $\left\|\mid \operatorname{lss}_{w_{a}}\left(w_{i}, w_{j}\right)\right\|$ denote the relative value, i.e., $\left\|\mid \boldsymbol{\operatorname { c s s }}_{w_{a}}\left(w_{i}, w_{j}\right)\right\|=\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right) / \sqrt{\left|w_{i}\right| \cdot\left|w_{j}\right|}$. The first ingredient of the algorithm is a data-structure, namely Ics-cmp. After a preprocess of time $O\left(\left|w_{a}\right| \sum_{i \in S}\left|w_{i}\right|\right)$, Ics-cmp is able to answer queries of the following type in time $O\left(\left|w_{i}\right|+\left|w_{j}\right|\right):$

- "for a $0 \leq \tilde{\lambda} \leq 1$ either certify that $\left\|\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\| \geq \Omega\left(\widetilde{\lambda^{2}}\right)$ or report that $\left\|\mid \operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\|<$ $O(\widetilde{\lambda})^{\prime \prime}$.

In our sparsification, we repeat the following procedure $k^{\gamma}$ times, where $\gamma \in(0,1)$. We sample a window $w_{a}$ uniformly at random and construct $\operatorname{Ics-cmp}\left(w_{a}, S\right)$ for $S=\left\{w_{i} \mid i \neq\right.$ $a$ and $\left.\left|w_{i}\right| \geq\left|w_{a}\right|\right\}$. After the preprocessing step, we make a query for every pair of windows $\left(w_{i}, w_{j}\right)$ such that $w_{i}, w_{j} \in S$ and determine if $\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)$ is at least $\Omega\left(\lambda^{4}\right)$ or upper bounded
by $O\left(\lambda^{2}\right)$ (here $\widetilde{\lambda}=\lambda^{2}$ ). If their LCS is at least $\Omega\left(\lambda^{4}\right)$ we report this pair as an edge in our Ics-graph. Finally, we use the Turan theorem to prove that the number of remaining edges in our graph is small.

To be more precise, we first construct a graph $N G_{\lambda}$ that reflects the edges that are not detected via our sparsification. We give directions to the edges based on the length of the windows. If $\mathrm{NG}_{\lambda}$ is dense enough, then there is one vertex $v$ in $\mathrm{NG}_{\lambda}$ with a large enough outgoing degree. We use the neighbors of $v$ to construct another graph $\mathrm{NF}_{\lambda}$ with vertex set $N(v)$. An edge exists in $\mathrm{NF}_{\lambda}$ if $\max \left\{\left\|\operatorname{lcs}_{w_{v}\left(w_{i}, w_{j}\right)}\right\|,\left\|\operatorname{lcs}_{w_{v}\left(w_{j}, w_{i}\right)}\right\|\right\} \geq \Omega\left(\lambda^{2}\right)$. Edges of $\mathrm{NF}_{\lambda}$ are directed the same way as $N G_{\lambda}$. We prove that $\mathrm{NF}_{\lambda}$ has no large independent set. In other words, if we select a large enough set of vertices in $\mathrm{NF}_{\lambda}$, then there is at least one edges between them. Next, we apply the Turan theorem to prove that $\mathrm{NF}_{\lambda}$ is dense. Finally, we imply that since $\mathrm{NF}_{\lambda}$ is dense, there is one vertex $u$ in the neighbors of $v$ such that there are a lot of 2-paths between $v$ and $u$. This implies that the edge $(u, v)$ should have been detected in our sparsification and therefore must not exist in $N G_{\lambda}$. This contradiction implies that $N G_{\lambda}$ is sparse in the first place.

### 30.1.3.3 LIS

In this section, we present our result for longest increasing subsequence. More precisely, we show that when the solution size is lower bounded by $n \lambda(\lambda \in[0,1])$, one can approximate the solution within a factor $O\left(\lambda^{3}\right)$ in time $\widetilde{O}\left(\sqrt{n} / \lambda^{7}\right)$. This married with a simple sampling algorithm for the cases that $\lambda<n^{-\Omega(1)}$, provides an $O\left(\lambda^{3}\right)$-approximate algorithm with running time of $\widetilde{O}\left(n^{0.85}\right)$ (without further dependence on $\lambda$ ). We further extend this result to reduce the running time to $\widetilde{O}\left(n^{\kappa} \operatorname{poly}(1 / \lambda)\right)$ for any $\kappa>0$ by imposing
a multiplicative factor of $\operatorname{poly}(1 / \lambda)$ to the approximation.
Our algorithm heavily relies on sampling random elements of the array for which longest increasing subsequence is desired. Denote the input sequence by $A=\left\langle a_{1}, a_{2}, \ldots, a_{n}\right\rangle$. A naive approach to approximate the solution is to randomly subsample the elements of $A$ to obtain a smaller array $B$ and then compute the longest increasing subsequence of $B$ to estimate the solution size for $A$. Let us first show why this approach alone fails to provide a decent approximation factor. First, consider an array $A=\langle 1,2, \ldots, n\rangle$ which is strictly increasing. Based on $A$, we construct two inputs $A^{\prime}$ and $A^{\prime \prime}$ in the following way:

- $A^{\prime}$ is exactly equal to $A$ except that a $p$ fraction of the elements in $A^{\prime}$ are replaced by 0.
- $A^{\prime \prime}$ is exactly equal to $A$ except that every block of length $\sqrt{n}$ is reversed in $A^{\prime \prime}$. In other words, $A^{\prime \prime}=\langle\sqrt{n}, \sqrt{n}-1, \sqrt{n}-2, \ldots, 1,2 \sqrt{n}, 2 \sqrt{n}-1, \ldots, \sqrt{n}+1, \ldots, n, n-$ $1, n-2, \ldots, n-\sqrt{n}+1\rangle$.

We subsample the two arrays $A^{\prime}$ and $A^{\prime \prime}$ with a rate of $1 / \sqrt{n}$ to obtain two smaller arrays $B^{\prime}$ and $B^{\prime \prime}$ of size roughly $O(\sqrt{n})$. It is easy to prove that $\operatorname{lis}\left(B^{\prime}\right)=\Omega(\sqrt{n})$ and $\operatorname{lis}\left(B^{\prime \prime}\right)=\Omega(\sqrt{n})$ both hold even though $\operatorname{lis}\left(A^{\prime}\right)=\Omega(n)$ but $\operatorname{lis}\left(A^{\prime \prime}\right)=O(\sqrt{n})$. By setting $p=1 / e$ we can also make sure that $\operatorname{lis}\left(B^{\prime}\right)$ and $\operatorname{lis}\left(B^{\prime \prime}\right)$ are within a small multiplicative range even though the gap between $\operatorname{lis}\left(A^{\prime}\right)$ and $\operatorname{lis}\left(A^{\prime \prime}\right)$ is substantial.

The above observation shows that the problem is very elusive when random sampling is involved. We bring a remedy to this issue in the following. Divide the input array into $\sqrt{n}$ subarrays of size $\sqrt{n}$. We denote the subarrays by $\mathrm{sa}_{1}, \mathrm{sa}_{2}, \ldots, \mathrm{sa}_{\sqrt{n}}$ and fix an
optimal solution opt for the longest increasing subsequence of $A$. Define $\mathrm{sm}\left(\mathrm{sa}_{i}\right)$ to be the smallest number in $s a_{i}$ that contributes to opt and $\lg \left(\mathrm{sa}_{i}\right)$ to be the largest number in $\mathrm{sa}_{i}$ that contributes to opt. Moreover, define lis ${ }^{[\ell, r]}$ to be the longest increasing subsequence of an array subject to the elements whose values lie within the interval $[\ell, r]$. This immediately implies

$$
\operatorname{lis}(A)=\sum_{i=1}^{\sqrt{n}} \operatorname{lis}{ }^{\left[\mathrm{sm}\left(\mathbf{s a}_{i}\right), \lg \left(\mathbf{s a}_{i}\right)\right]}\left(\mathrm{sa}_{i}\right)
$$

Another observation that we make here is that since we assume $\|\operatorname{lis}(A)\| \geq \lambda$ and the size of each subarray is bounded by $\sqrt{n}$, then we have

$$
\frac{\operatorname{lis}(A)}{\max _{i} \operatorname{lis}^{\left[\operatorname{sm}\left(\mathrm{sa}_{i}\right), \lg \left(\mathrm{sa}_{i}\right)\right]}\left(\mathrm{sa}_{i}\right)} \geq \sqrt{n} \lambda
$$

which essentially means that in order to approximate lis $(A)$ it suffices to compute lis ${ }^{\left[\operatorname{sm}\left(s \mathbf{s a}_{i}\right), \lg \left(\mathbf{s a}_{i}\right)\right]}\left(\operatorname{sa}_{i}\right)$ for $\widetilde{O}(1 / \lambda)$ many randomly sampled subarrays. This is quite helpful since this shows that we only need to sample $\widetilde{O}(1 / \lambda)$ many subarrays and solve the problem for them. However, we do not know the values of $\mathrm{sm}\left(\mathrm{sa}_{i}\right)$ and $\lg \left(\mathrm{sa}_{i}\right)$ in advance. Therefore, the main challenge is to predict the values of $\mathrm{sm}\left(\mathrm{sa}_{i}\right)$ and $\lg \left(\mathrm{sa}_{i}\right)$ before we sample the subarrays.

Indeed, one needs to read the entire array to correctly compute $\mathrm{sm}\left(\mathrm{sa}_{i}\right)$ and $\lg \left(\mathrm{sa}_{i}\right)$ for each of the subarrays. However, we devise a method to approximately guess these values without losing too much in the size of the solution. Roughly speaking, we show that if we sample $k=O(1 /(\lambda \epsilon))$ different elements from a subarray sa ${ }_{i}$ for some constant $\epsilon$ and denote


Figure 30.3: Red rectangles show the elements of $\mathrm{sa}_{i}$ that contribute to lis $(A)$ and gray circles show the elements of sa that are sampled via our algorithm.
them by $a_{j_{1}}, a_{j_{2}}, \ldots, a_{j_{k}}$, then for at least one pair $(\alpha, \beta),\left[a_{j_{\alpha}}, a_{j_{\beta}}\right]$ is approximately close to $\left[\mathrm{sm}\left(\mathrm{sa}_{i}\right), \lg \left(\mathrm{sa}_{i}\right)\right]$ up to a $(1-\epsilon)$ factor.

The above argument provides $O\left((1 /(\lambda \epsilon))^{2}\right)$ candidate domain intervals for each sa ${ }_{i}$. However, this does not provide a solution since we do not know which candidate domain interval approximates $\left[\operatorname{sm}\left(\mathrm{sa}_{i}\right), \lg \left(\mathrm{sa}_{i}\right)\right]$ for each $\mathrm{sa}_{i}$. Of course, if we were to randomly choose one candidate interval for every subarray, we would make a correct guess for at least $O\left(\sqrt{n}(\lambda \epsilon)^{2}\right)$ subarrays which provides an approximation guarantee of $O\left(\lambda^{2}\right)$ for our algorithm. However, our assignments have to be monotone too. More precisely, let $\left[\widetilde{\operatorname{sm}}\left(\mathrm{sa}_{i}\right), \widetilde{\lg }\left(\mathrm{sa}_{i}\right)\right]$ be the guesses that our algorithm makes, then we should have

$$
\widetilde{\mathrm{sm}}\left(\mathrm{sa}_{1}\right) \leq \widetilde{\lg }\left(\mathrm{sa}_{1}\right) \leq \widetilde{\operatorname{sm}}\left(\mathrm{sa}_{2}\right) \leq \widetilde{\lg }\left(\mathrm{sa}_{2}\right) \leq \ldots \leq \widetilde{\mathrm{sm}}\left(\mathrm{sa}_{\sqrt{n}}\right) \leq \widetilde{\lg }\left(\mathrm{sa}_{\sqrt{n}}\right)
$$

Random sampling does not guarantee that the sampled intervals are monotone. To address this issue, we introduce the notion of pseudo-solutions. A pseudo-solution is an as-
signment of monotone intervals to subarrays in order to approximate $\mathrm{sm}\left(\mathrm{sa}_{i}\right)$ and $\lg \left(\mathrm{sa}_{i}\right)$. The quality of a pseudo solution with intervals $\left[\ell_{1}, r_{1}\right],\left[\ell_{2}, r_{2}\right], \ldots,\left[\ell_{\sqrt{n}}, r_{\sqrt{n}}\right]$ is equal to $\sum_{i} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathbf{s a}_{i}\right)$. For a fixed pseudo-solution, this can be easily approximated via random sampling. Thus, our goal is to construct a pseudo-solution whose quality is at least an $O\left(\lambda^{3}\right)$ approximation of the size of the optimal solution. To this end, we present a greedy method in Section 30.4.2 to construct the desired pseudo-solution.

Finally, in Section 30.4.4, we show how the above ideas can be generalized to improve the running time down to $\widetilde{O}\left(n^{\kappa}\right.$ poly $\left.(1 / \lambda)\right)$ for any arbitrarily small $\kappa>0$ by imposing a factor $\operatorname{poly}(1 / \lambda)$ to the approximation guarantee.

### 30.2 Organization of the Paper

Our algorithm for LCS is explained in Section 30.6 (Step 0), Section 30.3 (Step 1), and Section 30.7 (Step 2). Since Steps 0 and 2 follow from previous work, we only bring Step 1 in the main body of the paper and defer the rest to the appendices.

In both our results for LCS and LIS, we assume that the goal is to find approximate solutions, provided that the solution size is at least $\lambda_{0} n$. After the algorithms terminate, if the output is smaller than what we expect, we realize that the solution is smaller than $\lambda_{0} n$. Therefore, we begin by setting $\lambda_{0}=1$ and iteratively multiply $\lambda_{0}$ by a $1-\epsilon$ factor until we obtain a solution. This only adds a multiplicative factor of $\log 1 / \lambda$ to the running time and a multiplicative factor of $1-\epsilon$ to the approximation. Since we present two different sparsification techniques, we obtain two theorems: one is Theorem 30.2.1 and the other is Theorem 30.2.2.

Theorem 30.2.1. Given strings $A, B$ of length $|A|=|B|=n$ with $\|\operatorname{LCS}(A, B)\|=\lambda$, we can approximate the length of the LCS between the two strings within a factor $O\left(\lambda^{3}\right)$ in time $\widetilde{O}\left(n^{39 / 20}\right)$.

Proof. If $\lambda \leq n^{-20}$ we run the classic $O\left(n^{2} \lambda\right)$ time algorithm and get an exact solution in the desired time. Otherwise, we begin by setting $\lambda_{0}=1$ and iteratively multiply $\lambda_{0}$ by a factor $1-\epsilon$ until a solution is found with size at least $\Omega\left(\lambda^{3}\right) n$. This adds an overhead of $O(\log 1 / \lambda)$ to the running time. By Choosing $d=\sqrt{n}$ we can bound the total running time of Steps 0 , 1 , and 2 by $\widetilde{O}\left(n^{13 / 10} \lambda^{-13}\right) \leq \widetilde{O}\left(n^{39 / 20}\right)$.

Theorem 30.2.2. Given strings $A, B$ of length $|A|=|B|=n$ with $\|\operatorname{LCS}(A, B)\|=\lambda$, we can approximate the length of the LCS between the two strings within a factor $(1-\epsilon) \lambda^{2}$ in
$n^{2-\Omega(\epsilon \lambda)}$ time for any $\epsilon>0$.

Proof. The proof is identical to Theorem 30.2.1, we omit the details here.

As an immediate corollary of Theorem 30.2.1, we present an algorithm that beats the $1 /|\Sigma|$ approximation factor in truly subquadratic time, when the strings are balanced.

Corollary 30.2.3. Given a pair of strings $(A, B)$ of length $n$ over alphabet $\Sigma$ that satisfy the balance condition, we can approximate their LCS within an $O\left(|\Sigma|^{3 / 4}\right)$ factor in time $O\left(n^{39 / 20}\right)$.

Proof. Since $A$ and $B$ are balanced, there is a character $\sigma \in \Sigma$ that appears at least $n /|\Sigma|$ times in both strings. Indeed, finding a solution of size $n /|\Sigma|$ by restricting our attention to only character $\sigma$ can be done in time $O(n)$. If $\operatorname{lcs}(A, B) \leq n /|\Sigma|^{1 / 4}$ this already gives us an $O\left(|\Sigma|^{3 / 4}\right)$ approximate solution. Otherwise, $\|\operatorname{lcs}(A, B)\|>1 /|\Sigma|^{1 / 4}$ and the approximation factor of our $O\left(\lambda^{3}\right)$-approximation algorithm would be bounded by $O\left(|\Sigma|^{3 / 4}\right)$.

Finally, we bring our results for LIS in Section 30.4. We show that

Theorem 30.2.4. Given a length-n sequence $A$ with $\operatorname{lis}(A)=n \lambda$. We can approximate the length of the LIS within a factor of $O\left(\lambda^{3}\right)$ in time $\widetilde{O}\left(n^{17 / 20}\right)$.

Proof. If $\lambda<n^{-1 / 20}$ we sample the array with a rate of $n^{-3 / 20}$ and compute the LIS for the sampled array. The running time of the algorithm is $\widetilde{O}\left(n^{17 / 20}\right)$. The approximation factor is $O\left(n^{-3 / 20}\right) \geq O\left(\lambda^{3}\right)$. Otherwise, by Theorem 30.4.8, we estimate the size of LIS up to an $O\left(\lambda^{3}\right)$ approximation factor in time $\widetilde{O}\left(\lambda^{-7} \sqrt{n}\right) \leq \widetilde{O}\left(n^{17 / 20}\right)$.

### 30.3 LCS Step 1: Sparsification via Birthday Triangle Inequality

Recall that we are given two sets of windows $W_{A}$ and $W_{B}$ for the strings and our goal is to approximate the LCS of every pair of windows between $W_{A}$ and $W_{B}$. For simplicity, we put all the windows in the same basket $W=W_{1} \cup W_{2}$ and denote the windows by $w_{1}, w_{2}, \ldots, w_{k}$ where $k$ is the total number of windows. Since the windows have different lengths, we define $\mathrm{w}_{\text {max }}=\max _{i \in[k]}\left|w_{i}\right|$ to be the maximum length of the windows. Similarly, we also define $\mathrm{w}_{\text {min }}=\min _{i \in[k]}\left|w_{i}\right|$ to be the minimum length of the windows. Let $\mathrm{w}_{\text {gap }}=\mathrm{w}_{\max } / \mathrm{w}_{\text {min }}$. Let $\mathbf{w}_{\text {layers }}$ denote the number of different window sizes. Notations $w_{\text {gap }}$ and $w_{\text {layers }}$ will be used in the later analysis.

In order to approximate the LCS's we fix a $\lambda \in\left\{\epsilon \lambda_{0},(1+\epsilon) \epsilon \lambda_{0},(1+\epsilon)^{2} \epsilon \lambda_{0}, \ldots, 1\right\}$ and sparsify graph $G_{\lambda}$. In Section 30.3.1, we present a sparsification algorithm (Algorithm 30.1) which provides $\lambda^{2}$-approximation when $\lambda$ is constant. The formal guarantee of the algorithm is provided in Theorem 30.3.6. In Section 30.3.2, we present a sparsification which provides $O\left(\lambda^{3}\right)$-approximation for any (potentially super-constant) $\lambda$.

### 30.3.1 Sparsification for $\mathrm{O}_{\lambda^{2}}$

In our solution, we fix an arbitrary LCS for every pair of windows and refer to that as opt ${ }_{i, j}$ for two windows $w_{i}$ and $w_{j}$. Note that we do not explicitly compute opt $t_{i, j}$ in our algorithm. Let us for simplicity, think of each opt ${ }_{i, j}$ as a matching between the characters of the two windows. Also, denote by ( $\mathrm{opt}_{i, a} \cap \mathrm{opt}_{a, b} \cap \mathrm{opt}_{b, j}$ ) a solution which is constructed for windows $w_{i}$ and $w_{j}$ by taking the intersection of ( $\left.\mathrm{opt}_{i, a} \cap \mathrm{opt}_{a, b} \cap \mathrm{opt}_{b, j}\right)$. More precisely,
we have

$$
\begin{aligned}
(x, y) \in\left(\mathrm{opt}_{i, a} \cap \mathrm{opt}_{a, b} \cap \mathrm{opt}_{b, j}\right) \Longleftrightarrow & \exists \alpha, \beta \text { such that } \\
& (x, \alpha) \in \mathrm{opt}_{i, a} \text { and } \\
& (\alpha, \beta) \in \mathrm{opt}_{a, b} \text { and } \\
& (\beta, y) \in \mathrm{opt}_{b, j} .
\end{aligned}
$$

Definition 30.3.1 $((\epsilon, \lambda)$-constructive $)$. We call a tuple $\left\langle w_{i}, w_{a}, w_{b}, w_{j}\right\rangle\left(w_{i} \neq w_{a} \neq w_{b} \neq w_{j}\right)$ a constructive tuple, if $\left\|\left(\mathrm{opt}_{i, a} \cap \mathrm{opt}_{a, b} \cap \mathrm{opt}_{b, j}\right)\right\| \geq(1-\epsilon) \lambda^{3}$.

The advantage of a constructive tuple is that if opt ${ }_{i, a}$, opt $_{a, b}$, and opt $t_{b, j}$ are provided, one can construct a desirable solution for opt ${ }_{i, j}$ in linear time by taking the intersection of the given matchings. Our algorithm is actually based on the above observation. We bring our algorithm in the following:

We parameterize our algorithm by a value $0<\gamma<1$ to be set later. One may optimize the runtime of the algorithm by setting the value of $\gamma$ in terms of the number of windows and the length of the windows. We first sample a set $S$ of $O\left(k^{\gamma} \log k\right)$ windows. Next, we compute opt ${ }_{i, j}$ of every window $w_{i} \in S$ and every other window $w_{j}$ (not necessarily in $S$ ). Finally, we find all the constructive tuples $\left\langle w_{i}, w_{a}, w_{b}, w_{j}\right\rangle$ such that $w_{a}, w_{b} \in S$ and update $\widehat{\mathrm{O}}_{\lambda^{2}}$ accordingly. This is shown in Algorithm 30.1.

The running time of our algorithm is equal to $O\left(k|S| \mathbf{w}_{\max }^{2}+k^{2}|S|^{2} \mathbf{w}_{\max }\right)$. The rest of this section is dedicated to proving that what remains in the Ics-graph is sparse.

We first introduce the notion of well-connected pairs.

```
Algorithm 30.1 Sparsification for \(\mathrm{O}_{\lambda^{2}}\)
    procedure QuadraticSparsification \(\left(w_{1}, w_{2}, \ldots, w_{k}, \lambda, \epsilon\right) \quad \triangleright\) Theorem 30.3.6
        \(\gamma \leftarrow 0.1\)
        \(S \leftarrow 40 k^{\gamma} \log k\) i.i.d. samples of \([k]\)
        \(\widehat{\mathrm{O}}_{\lambda^{2}} \leftarrow\{0\}^{k \times k}\)
        for \(w_{i} \in S\) do \(\triangleright\) Takes \(|S| k \mathrm{w}_{\text {max }}^{2}\) time
            for \(j \leftarrow 1\) to \(k\) do
                opt \(_{i, j}\), opt \(_{j, i} \leftarrow \operatorname{lcs}\left(w_{i}, w_{j}\right)\)
            end for
        end for
        for \(w_{i} \in S\) do \(\quad \triangleright\) Takes \(|S| k \mathrm{w}_{\text {max }}\) time
            for \(j \leftarrow 1\) to \(k\) do
                if \(\mid\) opt \(_{i, j} \mid>\lambda\) then
                \(\widehat{\mathrm{O}}_{\lambda^{2}}[i][j] \leftarrow 1\)
            end if
            end for
        end for
        for \(i \leftarrow 1\) to \(k\) do \(\quad \triangleright\) Takes \(k^{2}|S|^{2} \mathbf{w}_{\text {max }}\) time
            for \(j \leftarrow 1\) to \(k\) do
                for \(w_{a} \in S\) do
                for \(w_{b} \in S\) do
                if \(\|\left(\right.\) opt \(\left._{i, a} \cap \mathrm{opt}_{a, b} \cap \mathrm{opt}_{b, j}\right) \| \geq(1-\epsilon) \lambda^{3}\) then
                        \(\widehat{\mathrm{O}}_{\lambda^{2}}[i][j] \leftarrow 1\)
                end if
                end for
                end for
            end for
        end for
        return \(\widehat{\mathrm{O}}_{\lambda^{2}}\)
    end procedure
```

Definition 30.3.2 (well-connected pair). We say a pair of windows $\left(w_{i}, w_{j}\right)$ is well-connected, if there are at least $k^{2-\gamma}$ pairs of windows $\left(w_{a}, w_{b}\right)$ such that $\left\langle w_{i}, w_{a}, w_{b}, w_{j}\right\rangle$ is $(\epsilon, \lambda)$ constructive.

It follows from the definition that well-connected pairs are detected in our algorithm with high probability.

Lemma 30.3.1. Let $\widehat{\mathrm{O}}_{\lambda^{2}} \in\{0,1\}^{k \times k}$ denote the output of Algorithm 30.1. With probability at least $1-1 / \operatorname{poly}(k)$, for all $(i, j) \in[k] \times[k]$ such that $\left(w_{i}, w_{j}\right)$ is a well-connected pair (Definition 30.3.2), we have $\widehat{\mathrm{O}}_{\lambda^{2}}[i][j]=1$.

Proof. We consider a fixed $(i, j)$ such that pair $\left(w_{i}, w_{j}\right)$ is well-connected. Let

$$
Q_{i, j}=\left\{(a, b) \mid\left\|\mathrm{opt}_{i, a} \cap \mathrm{opt}_{a, b} \cap \mathrm{opt}_{b, j}\right\| \geq(1-\epsilon) \lambda^{3}\right\}
$$

Conceptually, we divide the process of sampling $S$ into two phases: we sample $20 k^{\gamma} \log k$ windows in the first phase, and then we sample another $20 k^{\gamma} \log k$ windows in the second phase.

For each $a \in[k]$, let $Q_{i, j, a}=\left\{b:(a, b) \in Q_{i, j}\right\}$. Since $\sum_{a \in[k]}\left|Q_{i, j, a}\right|=\left|Q_{i, j}\right| \geq k^{2-\gamma}$, there are at least

$$
\frac{k^{2-\gamma}-k \cdot k^{1-\gamma} / 2}{k}=\frac{k^{1-\gamma}}{2}
$$

different number $a$ 's in $[k]$ such that $\left|Q_{i, j, a}\right| \geq \frac{k^{1-\gamma}}{2}$. Hence, in the first phase, there is a sampled number $q$ such that $\left|Q_{i, j, q}\right| \geq k^{1-\gamma} / 2$ with probability at least

$$
1-\left(1-\frac{k^{1-\gamma} / 2}{k}\right)^{20 k^{\gamma} \log k}<\frac{1}{k^{5}}
$$

We fix such a $q$. In the second phase, there is a sampled number $r$ such that $r \in Q_{i, j, q}$ with probability at least

$$
1-\left(1-\frac{k^{1-\gamma} / 2}{k}\right)^{20 k^{\gamma} \log k}<\frac{1}{k^{5}}
$$

The lemma is obtained by a union bound on all the well-connected pairs $\left(w_{i}, w_{j}\right)$.

In what follows, we bound $\mid\left\{(i, j) \mid\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \lambda\right.$ and $\left.\widehat{\mathrm{O}}_{\lambda^{2}}[i][j]=0\right\} \mid$. In other words, we prove an upper bound on the number of edges that remain in the graph.

Definition 30.3.3 $\left(\mathrm{NG}_{\lambda}\right)$. Define a graph $\mathrm{NG}_{\lambda}$ with $k$ vertices and the following edges:

$$
E\left(\mathrm{NG}_{\lambda}\right)=\left\{(i, j) \mid\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \lambda \quad \text { and } \quad \widehat{\mathrm{O}}_{\lambda^{2}}[i][j]=0\right\}
$$

We first prove that every $K_{\Omega\left(\mathrm{w}_{\text {gap }} /\left(\epsilon \lambda^{3}\right)\right), \Omega\left(\mathrm{w}_{\text {gap }} /\left(\epsilon \lambda^{3}\right)\right)}$ subgraph of $\mathrm{NG}_{\lambda}$ corresponds to at least one constructive tuple.

Lemma 30.3.2. Let $X$ and $Y$ be two sets of windows such that for every $w_{i} \in X$ and $w_{j} \in Y$ there is an $(i, j)$ edge in $E\left(\mathbf{N G}_{\lambda}\right)$, for every $w_{i}, w_{i^{\prime}} \in X,\left|w_{i}\right|=\left|w_{i^{\prime}}\right|$ and for every $w_{j}, w_{j^{\prime}} \in Y,\left|w_{j}\right|=\left|w_{j^{\prime}}\right|$. If $|X| \geq \Omega\left(\mathrm{w}_{\text {gap }} /\left(\epsilon \lambda^{3}\right)\right)$ and $|Y| \geq \Omega\left(\mathrm{w}_{\text {gap }} /\left(\epsilon \lambda^{3}\right)\right)$, then there exist $w_{i}, w_{a}, w_{b}, w_{j} \in X \cup Y$ such that $\left\langle w_{i}, w_{a}, w_{b}, w_{j}\right\rangle$ is $(\epsilon, \lambda)$-constructive, i.e.,

$$
\left\|\mathrm{opt}_{i, a} \cap \mathrm{opt}_{a, b} \cap \mathrm{opt}_{b, j}\right\| \geq(1-\epsilon) \lambda^{3} .
$$

Proof. By assumption, we know that $|X|,|Y| \geq \Omega\left(\mathrm{w}_{\text {gap }} /\left(\epsilon \lambda^{3}\right)\right)$. Let $t_{1}$ denote the window size for each window in $X$, and $t_{2}$ denote the window size for each window in $Y$.

We carefully select $X^{\prime} \subseteq X$ and $Y^{\prime} \subseteq Y$ such that

1. $\left|X^{\prime}\right| \geq \Omega\left(1 /\left(\epsilon \lambda^{3}\right)\right)$.
2. $\left|Y^{\prime}\right| \geq \Omega\left(1 /\left(\epsilon \lambda^{3}\right)\right)$.
3. $\left|X^{\prime}\right| t_{1}=(1 \pm O(\epsilon))\left|Y^{\prime}\right| t_{2}$.


Figure 30.4: The graph on the left is an example of the string-based graph and the graph on the right is an example of the character-based graph.

To do this, if $|X| t_{1}>(1+O(\epsilon))|Y| t_{2}$, then we set $Y^{\prime}=Y$ and select $X^{\prime}$ as an arbitrary subset of $X$ with size $\left\lceil\frac{|Y| t_{2}}{t_{1}}\right\rceil$. if $|X| t_{1}<(1-O(\epsilon))|Y| t_{2}$, then we set $X^{\prime}=X$ and select $Y^{\prime}$ an arbitrary subset of $Y$ with size $\left\lceil\frac{|X| t_{1}}{t_{2}}\right\rceil$.

We define a window-based (bipartite) graph $G_{W}=\left(V_{W}, E_{W}\right)$ to be the subgraph of $\delta$ on $V_{W}=X^{\prime} \cup Y^{\prime}$ removing all the edges within $X^{\prime}$ and all the edges within $Y^{\prime}$. Let $l_{1}=\left|X^{\prime}\right|$ and $l_{2}=\left|Y^{\prime}\right|$. The total number of nodes in window-based graph is $l_{1}+l_{2}$, and the window-based graph is a bi-clique.

We also define a character-based (bipartite) graph $G_{C}=\left(V_{C}, E_{C}\right)$ as follows: Each window of $X^{\prime}$ has $t_{1}$ nodes in the character-based graph such that each node represents a character of the window. Similarly, each window of $Y^{\prime}$ has $t_{2}$ nodes in the character-based graph. Two nodes $x, y$ in the character-based graph are adjacent iff $(x, y) \in \operatorname{opt}_{i, j}$ where $w_{i}$ is the window containing character $x$ and $w_{j}$ is the window containing character $y$. Hence, the total number of nodes in character-based graph is $l_{1} t_{1}+l_{2} t_{2}$. The total number of edges in character-based graph satisfy

$$
l_{1} l_{2} \lambda \sqrt{t_{1} t_{2}} \leq\left|E\left(G_{C}\right)\right| \leq l_{1} l_{2} \min \left\{t_{1}, t_{2}\right\} .
$$

The average degree of the character-based graph is at least $2 l_{1} l_{2} \lambda \sqrt{t_{1} t_{2}} /\left(l_{1} t_{1}+l_{2} t_{2}\right)$. By Blakley-Roy inequality (Lemma 30.5.3), the number of walks of length 3 in the characterbased graph is at least

$$
\begin{aligned}
\left|V\left(G_{C}\right)\right| \cdot\left(2 \frac{l_{1} l_{2} \lambda \sqrt{t_{1} t_{2}}}{l_{1} t_{1}+l_{2} t_{2}}\right)^{3} & =\left(l_{1} t_{1}+l_{2} t_{2}\right) \cdot\left(2 \frac{l_{1} l_{2} \lambda \sqrt{t_{1} t_{2}}}{l_{1} t_{1}+l_{2} t_{2}}\right)^{3} \\
& =8 \lambda^{3} \frac{\left(l_{1} l_{2} \sqrt{t_{1} t_{2}}\right)^{3}}{\left(l_{1} t_{1}+l_{2} t_{2}\right)^{2}}
\end{aligned}
$$

But many of the 3 -walks are degenerate. For the number of 2 -walks, we can upper bound it by
$\# 2$-walks $\leq\left(\max\right.$ degree of $\left.G_{W}\right) \cdot 2\left|E\left(G_{C}\right)\right| \leq \max \left\{l_{1}, l_{2}\right\} \cdot 2\left(l_{1} l_{2} \min \left\{t_{1}, t_{2}\right\}\right)=2(1+O(\epsilon)) l_{1}^{2} l_{2} t_{1}$.

Note that the $\# 1$-walks is $2\left|E\left(G_{C}\right)\right|$. We have

$$
\begin{aligned}
\frac{\# 2 \text {-walks }+\# 1 \text {-walks }}{\# 3 \text {-walks }} & \leq \frac{2(1+O(\epsilon)) l_{1}^{2} l_{2} t_{1}+2 l_{1} l_{2} \min \left\{t_{1}, t_{2}\right\}}{\left(8 \lambda^{3} \frac{\left(l_{1} l_{2} \sqrt{t_{1} t_{2}}\right)}{\left(l_{1} t_{1}+l_{2} t_{2}\right)^{2}}\right)} \\
& <\frac{\left(l_{1} t_{1}+l_{2} t_{2}\right)^{2} l_{1}^{2} l_{2} t_{1}}{2 \lambda^{3} l_{1}^{3} l_{2}^{3} t_{1}^{1.5} t_{2}^{1.5}} \\
& \leq \frac{(1+O(\epsilon)) l_{1}^{4} l_{2} t_{1}^{3}}{2 \lambda^{3} l_{1}^{4.5} l_{2}^{1.5} t_{1}^{3}} \\
& =\frac{1+O(\epsilon)}{2 \lambda^{3} l_{1}^{0.5} l_{2}^{0.5}} \\
& \leq O(\epsilon),
\end{aligned}
$$

where the second step and third step follow from $l_{1} t_{1}=(1 \pm O(\epsilon)) l_{2} t_{2}$, and the last step follows from $l_{1}=\Omega\left(1 /\left(\epsilon \lambda^{3}\right)\right), l_{2}=\Omega\left(1 /\left(\epsilon \lambda^{3}\right)\right)$. Hence, the total number of 3-paths (3-walks that are not degenerate) is at least

$$
8(1-O(\epsilon)) \lambda^{3} \frac{\left(l_{1} l_{2} \sqrt{t_{1} t_{2}}\right)^{3}}{\left(l_{1} t_{1}+l_{2} t_{2}\right)^{2}} .
$$

On the other hand, the number of 4 -tuples in window-based graph is $8\binom{l_{1}}{2}\binom{l_{2}}{2} \leq 2 l_{1}^{2} l_{2}^{2}$. Thus, there must exist a 4 -tuple containing 3 -walks with size at least,

$$
\begin{aligned}
8(1-O(\epsilon)) \lambda^{3} \frac{\left(l_{1} l_{2} \sqrt{t_{1} t_{2}}\right)^{3}}{\left(l_{1} t_{1}+l_{2} t_{2}\right)^{2}} \cdot \frac{1}{2 l_{1}^{2} l_{2}^{2}} & =4(1-O(\epsilon)) \lambda^{3} \sqrt{t_{1} t_{2}} \cdot \frac{l_{1} l_{2} t_{1} t_{2}}{\left(l_{1} t_{1}+l_{2} t_{2}\right)^{2}} \\
& =4(1-O(\epsilon)) \lambda^{3} \sqrt{t_{1} t_{2}} \cdot \frac{1}{\frac{l_{1} t_{1}}{l_{2} t_{2}}+2+\frac{l_{2} t_{2}}{l_{1} t_{1}}} \\
& \geq 4(1-O(\epsilon)) \lambda^{3} \sqrt{t_{1} t_{2}} \cdot \frac{1}{4+O(\epsilon)} \\
& \geq(1-O(\epsilon)) \lambda^{3} \sqrt{t_{1} t_{2}},
\end{aligned}
$$

where the third step follows from $l_{1} t_{1}=(1 \pm O(\epsilon)) l_{2} t_{2}$. Rescaling the factor of $\epsilon$ for $|X|$ and $|Y|$ completes the proof.

Definition 30.3.4 (reconstructive tuple). We say a tuple $\left\langle w_{i}, w_{a}, w_{b}, w_{j}\right\rangle$ is reconstructive if it is $(\epsilon, \lambda)$-constructive and also $(i, a),(a, b),(b, j),(i, j) \in E\left(\mathrm{NG}_{\lambda}\right)$.

Before proving Lemma 30.3.5, we need to define a new graph $\mathrm{NF}_{\lambda}$,
Definition 30.3.5 $\left(N F_{\lambda}\right)$. We construct a graph $N F_{\lambda}$ in the following way: for each node $v \in \mathrm{NG}_{\lambda}$ we keep $v$ in $\mathrm{NF}_{\lambda}$ with probability $p=k^{-\alpha} / 4$ where $\alpha=1-\gamma / 4$. For each $u, v \in \mathrm{NF}_{\lambda}$, if $(u, v) \in \mathrm{NG}_{\lambda}$ then we also draw an edge $u, v$ in $\mathrm{NF}_{\lambda}$.

It is obvious that $V\left(\mathrm{NF}_{\lambda}\right) \subseteq V\left(\mathrm{NG}_{\lambda}\right)$ and $E\left(\mathrm{NF}_{\lambda}\right) \subseteq E\left(\mathrm{NG}_{\lambda}\right)$.

Based on definition of $\mathrm{NF}_{\lambda}$, we are able to show that if $N G_{\lambda}$ is dense, so is $N F_{\lambda}$.
Claim 30.3.3 $\left(\mathrm{NF}_{\lambda}\right.$ is a dense graph). Let $\mathrm{w}_{\text {layers }}$ be the number of different window sizes. Let $\mathrm{NG}_{\lambda}$ denote a graph such that $\left|E\left(\mathrm{NG}_{\lambda}\right)\right| \geq k^{2-\beta} \cdot \mathrm{w}_{\text {layers }}^{2}$ for some $\beta$ and $\mathrm{NF}_{\lambda}$ be defined as Definition 30.3.5. If $\gamma / 4-\beta=\Omega(1)$, then with probability at least 0.98 , we have

$$
\left|E\left(\mathrm{NF}_{\lambda}\right)\right|=\Omega\left(\left|V\left(\mathrm{NF}_{\lambda}\right)\right|^{2-4 \beta / \gamma} \cdot \mathrm{w}_{\text {layers }}^{2}\right)
$$

Proof. Based on the sampling rate, we know that the following hold in expectation:

$$
\mathbb{E}\left[\left|V\left(\mathrm{NF}_{\lambda}\right)\right|\right]=\frac{1}{4 k^{\alpha}}\left|V\left(\mathrm{NG}_{\lambda}\right)\right|, \quad \text { and } \quad \mathbb{E}\left[\left|E\left(\mathrm{NF}_{\lambda}\right)\right|\right]=\frac{1}{4^{2} k^{2 \alpha}}\left|E\left(\mathrm{NG}_{\lambda}\right)\right|
$$

Using standard Chernoff bound, we have with probability 0.99,

$$
\left|V\left(\mathrm{NF}_{\lambda}\right)\right| \in[1 / 2,2] \cdot \frac{1}{4 k^{\alpha}}\left|V\left(\mathrm{NG}_{\lambda}\right)\right|=[1 / 2,2] \cdot \frac{1}{4} k^{1-\alpha}=[1 / 2,2] \cdot \frac{1}{4} k^{\gamma / 4}
$$

To show the concentration of $\left|E\left(\mathrm{NF}_{\lambda}\right)\right|$, we define a random variable $x_{u, v}$ such that $x_{u, v}=1$ if $(u, v) \in E\left(\mathrm{NF}_{\lambda}\right), 0$ otherwise. Then we have

$$
\begin{aligned}
& \mathbb{E}\left[\left|E\left(\mathrm{NF}_{\lambda}\right)\right|^{2}\right] \\
= & \mathbb{E}\left[\left(\sum_{(u, v) \in E\left(\mathrm{NG}_{\lambda}\right)} x_{u, v}\right)^{2}\right] \\
= & \mathbb{E}\left[\sum_{(u, v) \in E\left(\mathrm{NG}_{\lambda}\right)} x_{u, v}^{2}\right]+\mathbb{E}\left[\sum_{(u, v),\left(u, v^{\prime}\right) \in E\left(\mathrm{NG}_{\lambda}\right)} \mathbf{1}_{v \neq v^{\prime}} x_{u, v} x_{u, v^{\prime}}\right] \\
& +\mathbb{E}\left[\sum_{(u, v) \in E\left(\mathrm{NG}_{\lambda}\right)} \sum_{\left(u^{\prime}, v^{\prime}\right) \in E\left(\mathrm{NG}_{\lambda}\right)} \mathbf{1}_{u \neq v \neq u^{\prime} \neq v^{\prime}} x_{u, v} x_{u^{\prime}, v^{\prime}}\right] \\
\leq & p^{2}\left|E\left(\mathrm{NG}_{\lambda}\right)\right|+p^{3} \cdot \# 2 \text {-walks }+p^{4} \cdot\left(\left|E\left(\mathrm{NG}_{\lambda}\right)\right|^{2}-\# 2 \text {-walks }-\left|E\left(\mathrm{NG}_{\lambda}\right)\right|\right) .
\end{aligned}
$$

Since $\mathbb{E}\left[\left|E\left(\mathrm{NF}_{\lambda}\right)\right|\right]=p^{2}\left|E\left(\mathrm{NG}_{\lambda}\right)\right|$. Thus,

$$
\begin{array}{rlr}
\mathbb{V}\left[\left|E\left(\mathrm{NF}_{\lambda}\right)\right|\right] & =\mathbb{E}\left[\left(\left|E\left(\mathrm{NF}_{\lambda}\right)\right|\right)^{2}\right]-\left(\mathbb{E}\left[\left|E\left(\mathrm{NF}_{\lambda}\right)\right|\right]\right)^{2} & \\
& \leq p^{2}\left|E\left(\mathrm{NG}_{\lambda}\right)\right|+p^{3} \cdot \# 2 \text {-walks } \\
& \leq p^{2}\left|E\left(\mathrm{NG}_{\lambda}\right)\right|+p^{3}\left|V\left(\mathrm{NG}_{\lambda}\right)\right| \cdot\left|E\left(\mathrm{NG}_{\lambda}\right)\right| & \\
& \leq 2 p^{3}\left|V\left(\mathrm{NG}_{\lambda}\right)\right| \cdot\left|E\left(\mathrm{NG}_{\lambda}\right)\right| . & \text { by } p\left|V\left(\mathrm{NG}_{\lambda}\right)\right| \geq 1
\end{array}
$$

In order to use Chebyshev's inequality, we need to make sure

$$
p^{2}\left|E\left(\mathrm{NG}_{\lambda}\right)\right|>1, \quad p^{2}\left|E\left(\mathrm{NG}_{\lambda}\right)\right| \geq 100\left(2 p^{3}\left|V\left(\mathrm{NG}_{\lambda}\right)\right|\left|E\left(\mathrm{NG}_{\lambda}\right)\right|\right)^{1 / 2}
$$

By $p=k^{\gamma / 4-1} / 4$ and $\left|E\left(\mathrm{NG}_{\lambda}\right)\right| \geq k^{2-\beta}$, we need

$$
k^{\gamma / 2-\beta} / 16>1,
$$

and the second condition is equivalent to

$$
p\left|E\left(\mathrm{NG}_{\lambda}\right)\right| \geq 20000\left|V\left(\mathrm{NG}_{\lambda}\right)\right| .
$$

We just need $k^{\gamma / 4-\beta} \geq 80000$.
Using Chebyshev's inequality, we have

$$
\operatorname{Pr}\left[\left|\left|E\left(\mathrm{NF}_{\lambda}\right)\right|-\mathbb{E}\left[\left|E\left(\mathrm{NF}_{\lambda}\right)\right|\right]\right| \geq \tau \sqrt{\mathbb{V}\left[\left|E\left(\mathrm{NF}_{\lambda}\right)\right|\right]}\right] \leq \frac{1}{\tau^{2}}
$$

combining with $\sqrt{\mathbb{V}\left[\left|E\left(\mathrm{NF}_{\lambda}\right)\right|\right]} \leq \frac{1}{100} \mathbb{E}\left[\left|E\left(\mathrm{NF}_{\lambda}\right)\right|\right]$ together implies

$$
\operatorname{Pr}\left[\left|E\left(\mathrm{NF}_{\lambda}\right)\right|-\mathbb{E}\left[\left|E\left(\mathrm{NF}_{\lambda}\right)\right|\right] \left\lvert\, \geq \tau \frac{1}{100} \mathbb{E}\left[\left|E\left(\mathrm{NF}_{\lambda}\right)\right|\right]\right.\right] \leq \frac{1}{\tau^{2}}
$$

Choosing $\tau=10$, we have with probability 0.99 that $\left|E\left(\mathrm{NF}_{\lambda}\right)\right|>0.9 \mathbb{E}\left[\left|E\left(\mathrm{NF}_{\lambda}\right)\right|\right]$, and consequently

$$
\begin{aligned}
\left|E\left(\mathrm{NF}_{\lambda}\right)\right| & >0.9 \mathbb{E}\left[\left|E\left(\mathrm{NF}_{\lambda}\right)\right|\right] \\
& \geq 0.9 \frac{1}{4^{2} k^{2 \alpha}}\left|E\left(\mathrm{NG}_{\lambda}\right)\right| \\
& \geq 0.9 \frac{1}{4^{2}} k^{\gamma / 2-2} k^{2-\beta} \cdot \mathrm{w}_{\text {layers }}^{2} \\
& =0.9 \frac{1}{4^{2}} k^{\gamma / 2-\beta} \cdot \mathrm{w}_{\text {layers }}^{2} \\
& \geq 0.9 \frac{1}{4^{2}} k^{\frac{\gamma}{4}(2-4 \beta / \gamma)} \cdot \mathrm{w}_{\text {layers }}^{2} \\
& \geq \Omega\left(\left|V\left(\mathrm{NF}_{\lambda}\right)\right|^{2-4 \beta / \gamma} \cdot \mathrm{w}_{\text {layers }}^{2}\right) .
\end{aligned}
$$

Thus, we complete the proof.

Next, we show

Claim 30.3.4 $\left(\mathrm{NF}_{\lambda}\right.$ has no reconstructive tuple). Let $\mathrm{NF}_{\lambda}$ be defined as Definition 30.3.5. With probability at least 0.99, there is no reconstructive tuple in $\mathrm{NF}_{\lambda}$.

Proof. By Lemma 30.3.1 and the definition of $\mathbf{N G}_{\lambda}$, there is no pair of windows $\left(w_{i}, w_{j}\right)$ in $\mathrm{NG}_{\lambda}$ such that pair $(i, j)$ is well-connected. Formally speaking, for all pairs of windows $\left(w_{i}, w_{j}\right)$, there are at most $k^{2-\gamma}$ pairs of windows $\left(w_{a}, w_{b}\right)$ such that $\left\langle w_{i}, w_{a}, w_{b}, w_{j}\right\rangle$ is $(\epsilon, \lambda)$ constructive.

For a fixed $(i, a, b, j)$, the probability that we keep it in $\mathrm{NF}_{\lambda}$ is $\left(k^{-\alpha} / 4\right)^{4}$. We can compute the expected number of constructive tuples in $\mathrm{NF}_{\lambda}$,

$$
\mathbb{E}[\# \text { constructive tuple }] \leq k^{2} k^{2-\gamma} \cdot\left(k^{-\alpha} / 4\right)^{4}=k^{4-\gamma-4 \alpha} / 256=1 / 256
$$

where the last step follows from $\alpha=1-\gamma / 4$.
By Markov's inequality, we have

$$
\operatorname{Pr}[\# \text { constructive tuple } \geq 1 / 2] \leq \frac{\mathbb{E}[\# \text { constructive tuple }]}{1 / 2} \leq 1 / 100
$$

Thus, with probability 0.99 , there is no reconstructive tuple in $\mathrm{NF}_{\lambda}$.

Finally, we use Lemma 30.3.2 to prove an upper bound on the number of edges in $N G_{\lambda}$.

Lemma 30.3.5. Let $\mathrm{NG}_{\lambda}$ be as defined in Definition 30.3.3. Then with probability at least $0.9, \mathrm{NG}_{\lambda}$ has at most $k^{2-\Omega\left(\gamma \lambda \epsilon / \mathrm{w}_{\text {gap }}\right)} \cdot \mathrm{w}_{\text {layers }}^{2}$ edges.

Proof. We start with assuming

$$
\left|E\left(\mathrm{NG}_{\lambda}\right)\right| \geq k^{2-\beta} \cdot \mathrm{w}_{\text {layers }}^{2},
$$

and will get the contradiction at the end.
We construct a graph $\mathrm{NF}_{\lambda}$ based on Definition 30.3.5. Using Claim 30.3.4, we know that $\mathrm{NF}_{\lambda}$ has no reconstructive tuple. Next, we are going to show two facts, and they are contradicting to each other by some choice of $\beta$.
(Fact A). Note that $\mathrm{NF}_{\lambda}$ does not satisfy the assumption in Lemma 30.3.2, thus we cannot apply Lemma 30.3.2 to $\mathrm{NF}_{\lambda}$ directly. Let $\mathrm{w}_{\text {layers }}$ denote the number of different window sizes in total, we can decompose $\mathrm{NF}_{\lambda}$ into $w_{\text {layers }}^{2}$ graphs such that $N F_{\lambda}^{i, j}$ only involves size $i$ and $j, \forall i, j \in\left[\mathrm{w}_{\text {layers }}\right] \times\left[\mathrm{w}_{\text {layers }}\right]$. Now, applying Lemma 30.3.2, for every $i, j$, we imply that if $\mathrm{NF}_{\lambda}^{i, j}$ has no reconstructive tuple, then $\mathrm{NF}_{\lambda}^{i, j}$ has no bipartite graph of certain size, i.e., $K_{\Omega\left(\mathrm{w}_{\text {gap }} /\left(\epsilon \lambda^{3}\right)\right), \Omega\left(\mathrm{w}_{\text {gap }} /\left(\epsilon \lambda^{3}\right)\right)}$.
(Fact B). Using the Turán's Theorem (Lemma 30.5.4), we know for any integer $s \geq 2$, a graph $G$ with $n$ vertices and $\Omega\left(n^{2-1 / s}\right)$ edges has at least one $K_{s, s}$ subgraph. We consider graph $\mathrm{NF}_{\lambda}$. Since

$$
\left|E\left(\mathrm{NF}_{\lambda}\right)\right|=\Omega\left(\left|V\left(\mathrm{NF}_{\lambda}\right)\right|^{2-4 \beta / \gamma} \cdot \mathrm{w}_{\text {layers }}^{2}\right)
$$

by Claim 30.3.3, $\left|E\left(\mathrm{NF}_{\lambda}\right)\right|=\Omega\left(\left|V\left(\mathrm{NF}_{\lambda}\right)\right|^{2-4 \beta / \gamma} \cdot \mathrm{w}_{\text {layers }}^{2}\right)$. By the pigeonhole principle, there exist $i$ and $j$ such that $\left|E\left(\mathrm{NF}_{\lambda}^{i, j}\right)\right|=\Omega\left(\left|V\left(\mathrm{NF}_{\lambda}\right)\right|^{2-4 \beta / \gamma}\right)$. By Turán's Theorem, such an $\mathrm{NF}_{\lambda}^{i, j}$ contains a large complete bipartite subgraph $K_{\gamma /(4 \beta), \gamma /(4 \beta)}$.

In order to get a contraction, we just need

$$
\gamma /(4 \beta) \geq c \cdot \mathrm{w}_{\mathrm{gap}} /\left(\epsilon \lambda^{3}\right)
$$

for some sufficiently large constant $c>1$ related to Lemma 30.3.2. Thus, we have

$$
\beta \leq \gamma \epsilon \lambda^{3} /\left(4 c \cdot \mathrm{w}_{\mathrm{gap}}\right)
$$

Theorem 30.3.6 (quadratic sparsification for constant $\lambda$ ). Given $k$ windows $w_{1}, \cdots, w_{k}$. Let $\mathrm{w}_{\max }=\max _{i \in[k]}\left|w_{i}\right|$, $\mathrm{w}_{\text {min }}=\min _{i \in[k]}\left|w_{i}\right|$ and $\mathrm{w}_{\text {gap }}=\mathrm{w}_{\max } / \mathrm{w}_{\text {min }}$. Let the number of different window sizes be $\mathrm{w}_{\text {layers }}$. For any $\lambda \in(0,1)$ and $\epsilon \in(0,1 / 10)$, there is a randomized algorithm (Algorithm 30.1) that runs in time

$$
O\left(\mathrm{w}_{\max }^{2} k^{1.1} \log k+\mathrm{w}_{\max } k^{2.2} \log ^{2} k\right),
$$

outputs a table $\widehat{\mathrm{O}}_{\lambda^{2}} \in\{0,1\}^{k \times k}$ such that

$$
\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq(1-\epsilon) \lambda^{3}, \text { if } \widehat{\mathrm{O}}_{\lambda^{2}}[i][j]=1
$$

and

$$
\mid\left\{(i, j) \mid\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \lambda, \quad \text { and } \quad \widehat{\mathrm{O}}_{\lambda^{2}}[i][j]=0\right\} \mid=O\left(k^{2-\Omega\left(\lambda \epsilon / \mathrm{w}_{\mathrm{gap}}\right)} \cdot \mathrm{w}_{\text {layers }}^{2}\right) .
$$

The algorithm has success probability $1 / \operatorname{poly}(k)$.

Proof. The overall running time is

$$
\begin{aligned}
& O\left(k|S| \mathrm{w}_{\max }^{2}+k^{2}|S|^{2} \mathrm{w}_{\max }\right) \\
= & O\left(k \cdot k^{\gamma} \log k \cdot \mathrm{w}_{\max }^{2}+k^{2} \cdot k^{2 \gamma} \log ^{2} k \cdot \mathrm{w}_{\max }\right) \\
= & O\left(k^{1.1} \log k \cdot \mathrm{w}_{\max }^{2}+k^{2.2} \log ^{2} k \cdot \mathrm{w}_{\max }\right)
\end{aligned}
$$

where the first step follows from $|S|=O\left(k^{\gamma} \log k\right)$, the second step follows from $\gamma=0.1$.
The guarantee of table $\widehat{\mathrm{O}}_{\lambda^{2}}$ follows from properties of graph $\mathrm{NG}_{\lambda}$ (Algorithm 30.1 provides the first property of table in Theorem statement, Lemma 30.3.5 provides the second property of table in Theorem statement).

### 30.3.2 Sparsification for $\mathrm{O}_{\lambda^{3}}$

One shortcoming of Lemma 30.3.5 is that the number of remaining edges is only truly subquadratic if $\lambda$ is constant. As we discuss in Section 30.3.1, the overall running time of the algorithm depends on the number of edges in the remaining graph and in order for the running time to be truly subquadratic, we need to reduce the number of edges to truly subquadratic. In this section, we show how one can obtain this bound even when $\lambda$ is super-constant. However, instead of losing a factor $\lambda^{2}$ in the approximation, our technique loses a factor of $O\left(\lambda^{3}\right)$.

Definition 30.3.6 $\left(\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right)$. For two windows $w_{i}$ and $w_{j}$, and a window $w_{a}$, define $\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)$ as the length of LCS of opt ${ }_{i, a}$ and $w_{j}$, where opt ${ }_{i, a}$ denotes the LCS of $w_{i}$ and $w_{a}$.

Notice that unlike Ics, this new definition is not symmetric. Similar to Ics, we also normalize the size of $\mathrm{Ics}_{s}$ by the geometric mean of the lengths of the two windows. That is, we divide the size of the common string by $\sqrt{\left|w_{i}\right|\left|w_{j}\right|}$. In what follows, we first give an algorithm for detecting close pairs of windows, and then prove that the number of remaining pairs whose Ics is at least $\lambda$ is truly subquadratic.

Our algorithm is based on a data structure which we call $\operatorname{Ics-cmp}\left(w_{a}, S, \tilde{\lambda}\right)$. Roughly
speaking, we will choose $\widetilde{\lambda}=\lambda^{2}$ when we use it. Let us fix a threshold $\widetilde{\lambda}$ and a window $w_{a}$. $\operatorname{Ics}-\operatorname{cmp}\left(w_{a}, S, \widetilde{\lambda}\right)$ receives a set $S$ of windows as input such that the size of each window of $S$ is at least $\left|w_{a}\right|$. Upon receiving the windows, it makes a preprocess in time $O\left(\widetilde{\lambda}^{-2}\left|w_{a}\right| \sum_{w_{i} \in S}\left|w_{i}\right|\right)$. Next, $\operatorname{Ics-cmp}\left(w_{a}, S, \widetilde{\lambda}\right)$ would be able to answer each query of the following form in almost linear time:

- For two windows $w_{i}, w_{j} \in S$, either certify that $\left\|\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\|<\widetilde{\lambda} / 2$ or find a solution for $\left\|\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\|$ with a size of at least $\widetilde{\lambda}^{2} / 8$.

We first show in Section 30.3.2.1, how Ics-cmp gives us a sparsification in truly sublinear time and then discuss the algorithm for Ics-cmp in Section 30.3.2.2.

### 30.3.2.1 $\lambda^{3}$ Sparsification using Ics-cmp

Similar to what we did in Section 30.3.1, we again use a parameter $\gamma$ in our algorithm and in the end, we adjust $\gamma$ to minimize the total running time. In our algorithm, we repeat the following procedure $k^{\gamma} \log k$ times: sample a window $w_{a}$ uniformly at random and let $S$ be the set of all the other windows whose length is not smaller than $\left|w_{a}\right|$. Next, we obtain Ics-cmp $\left(w_{a}, S, \widetilde{\lambda}\right)$ via running the preprocessing step. Finally, for each pair of windows $w_{i}, w_{j} \in S$, we make a query to Ics-cmp to verify one of the following two possibilities:

- $\left\|\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\|<\lambda^{2} / 2 ;$
- $\left\|\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\| \geq \lambda^{4} / 8$.

If the latter is verified we set $\widehat{\mathrm{O}}_{\lambda^{3} / 8}[i][j]$ to 1 otherwise we take no action. In what follows, we prove that after the above sparsification, the number of edges in the remaining graph is
truly sublinear.

```
Algorithm 30.2 Sparsification for \(\mathrm{O}_{\lambda^{3}}\)
    procedure CubicSparsification \(\left(w_{1}, w_{2}, \ldots, w_{k}, \lambda\right) \quad \triangleright\) Theorem 30.3.10
        \(\underset{\lambda}{x} \leftarrow 0.1\)
        \(\tilde{\lambda} \leftarrow \lambda^{2}\)
        for counter \(=1 \rightarrow k^{\gamma} \log k\) do
            Sample \(a \sim[k]\) uniformly at random
        \(S \leftarrow \emptyset\)
        for \(i=1 \rightarrow k\) do
            if \(i \neq a\) and \(\left|w_{i}\right| \geq\left|w_{a}\right|\) then
                \(S \leftarrow S \cup\left\{w_{i}\right\}\)
                end if
            end for
            Ics-cmp.Initial \(\left(w_{a}, S, \widetilde{\lambda}\right) \quad \triangleright\) Algorithm 30.3, Lemma 30.3.11
            for \(w_{i} \in S\) do
                for \(w_{j} \in S\) do
                if Ics-cmp. Query \(\left(w_{i}, w_{j}\right)\) outputs accept then \(\triangleright\) Algorithm 30.3,
    Lemma 30.3.12
                \(\widehat{\mathrm{O}}_{\lambda^{3} / 8}[i][j] \leftarrow 1\)
                                    \(\triangleright\left\|\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\| \geq \widetilde{\lambda}^{2} / 8\)
                end if
                end for
            end for
        end for
        return \(\widehat{\mathrm{O}}_{\lambda^{3} / 8}\)
    end procedure
```

In the rest of this section, we bound $\mid\left\{(i, j)| | \mid \operatorname{lcs}\left(w_{i}, w_{j}\right) \| \geq \lambda\right.$ and $\left.\widehat{\mathrm{O}}_{\lambda^{3} / 8}[i][j]=0\right\} \mid$. In other words, we prove an upper bound on the number of edges that remain in the graph. Define a graph $N G_{\lambda}$ such that each vertex of $N G_{\lambda}$ corresponds to a window and an edge $(i, j)$ means that $\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \lambda$ but $\widehat{\mathrm{O}}_{\lambda^{3} / 8}[i][j]=0$. The goal is to prove an upper bound on the number of edges of $N G_{\lambda}$. We formally prove this in Lemma 30.3.9.

We define a notation called "close" which is similar to "well-connected" in Section 30.3.1.

In order to avoid confusion, we use "close" instead of well-connected.
Definition 30.3.7 (close). Let $\gamma \in(0,1)$ and $\lambda \in(0,1)$. We say a pair $\left(w_{i}, w_{j}\right)$ of windows is close, if there are at least $k^{1-\gamma}$ windows $w_{a}$ such that $\left|w_{a}\right| \leq \min \left\{\left|w_{i}\right|,\left|w_{j}\right|\right\}$ and $\left\|\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\| \geq \lambda^{2}$.

Our first observation is that Algorithm 30.2 detects all the close pairs with high probability.

Lemma 30.3.7. Let $\widehat{\mathrm{O}}_{\lambda^{3} / 8} \in\{0,1\}^{k \times k}$ be the output of Algorithm 30.2. For each $(i, j) \in$ $[k] \times[k]$, if $\left(w_{i}, w_{j}\right)$ is close (Definition 30.3.7) and $\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \lambda$, then $\widehat{\mathrm{O}}_{\lambda^{3} / 8}[i][j]=1$ holds with probability at least $1-1 / \operatorname{poly}(k)$.

Proof. We consider a fixed $(i, j)$ such that $\left(w_{i}, w_{j}\right)$ is close. By Definition 30.3.7, there are at least $k^{1-\gamma} w_{a}$ such that $\left|w_{a}\right| \leq \min \left\{\left|w_{i}\right|,\left|w_{j}\right|\right\}$ and $\left\|\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\| \geq \lambda^{2}$. Then the probability that none of these windows is sampled is at most

$$
\left(1-\frac{k^{1-\gamma}}{k}\right)^{t}=\left(1-\frac{1}{k^{\gamma}}\right)^{t}=\left(1-\frac{1}{k^{\gamma}}\right)^{10 k^{\gamma} \log k} \leq 1 / \operatorname{poly}(k)
$$

Taking a union over at most $k^{2}$ pairs completes the proof.

Before we proceed to Lemma 30.3.9, we bring Lemma 30.3.8 as an auxiliary observation.

Lemma 30.3.8 (existence of a correlated pair). Let $\epsilon \in(0,1), \lambda \in(0,1)$. Given a window $w_{a}$ and a set of windows $T$. Let $\mathrm{w}_{\text {gap }}$ denote the maximum size of windows divided by the minimum size of windows. Let $T$ contain at most $\mathrm{w}_{\text {layers }}$ different sizes and $|T|=\Omega\left(\mathrm{w}_{\text {layers }}\right.$. $\left.\mathrm{w}_{\text {gap }} /(\epsilon \lambda)\right)$. If for each $w_{i} \in T$ we have $\left\|\operatorname{lcs}\left(w_{a}, w_{i}\right)\right\| \geq \lambda$, then there exist two windows $w_{i}, w_{j} \in T$ such that $\left\|\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\| \geq(1-\epsilon) \lambda^{2}$.

Proof. Since $T$ has size $\Omega\left(\mathrm{w}_{\text {layers }} \cdot \mathrm{w}_{\text {gap }} /(\epsilon \lambda)\right)$ and the windows of $T$ have at most $\mathbf{w}_{\text {layers }}$ different sizes, there must exists a subset $T^{\prime} \subseteq T$ such that $\left|T^{\prime}\right|=\Omega\left(\mathrm{w}_{\text {gap }} /(\epsilon \lambda)\right)$ and all the windows in $T^{\prime}$ have the same size. In the rest of the proof, we will focus on $w_{a}$ and $T^{\prime}$.

Let $t_{0}$ be the size of $w_{a}$ and $t$ be the size of all the windows in $T^{\prime}$. We consider a window-based bipartite graph, on one side it has one node $w_{a}$, and on the other side it has $\left|T^{\prime}\right|$ nodes. Then we can expand the window-based bipartite graph into a character-based bipartite graph, on one side it has $t_{0}$ nodes, and on the other side it has $t\left|T^{\prime}\right|$ nodes. Two nodes $x, y$ in the character-based graph are adjacent iff $(x, y) \in \mathrm{opt}_{a, i}$ where $w_{a}$ is the window containing character $x$ and $w_{i}$ is the window containing character $y$. Since $\left\|\operatorname{lcs}\left(w_{a}, w_{i}\right)\right\| \geq \lambda$ for every $w_{i} \in T^{\prime}$, the number of edges in character-based bipartite graph is at least $\lambda\left|T^{\prime}\right| \sqrt{t_{0} t}$.

For $i$-th character in window $w_{a}$, we use $D_{i}$ to denote the degree of the corresponding node in the character-based bipartite graph. The number of 2 -walks is at least

$$
\begin{array}{rlr}
\sum_{i=1}^{t_{0}} D_{i}\left(D_{i}-1\right) & =\left(\sum_{i=1}^{t_{0}} D_{i}^{2}-\sum_{i=1}^{t_{0}} D_{i}\right) & \\
& \geq\left(\frac{1}{t_{0}}\left(\sum_{i=1}^{t_{0}} D_{i}\right)^{2}-\sum_{i=1}^{t_{0}} D_{i}\right) & \\
& \geq\left(\frac{1}{t_{0}}\left(\lambda\left|T^{\prime}\right| \sqrt{t_{0} t}\right)^{2}-\left(\lambda\left|T^{\prime}\right| \sqrt{t_{0} t}\right)\right) & \text { by Cauchy-Schwarz inequality } \\
& =\left(\lambda^{2} t\left|T^{\prime}\right|^{2}-\lambda\left|T^{\prime}\right| \sqrt{t_{0} t}\right) & \text { by Eq. (30.1) } \\
& \geq(1-\epsilon) \lambda^{2} t\left|T^{\prime}\right|^{2} & \text { by }\left|T^{\prime}\right|=\Omega\left(\mathrm{w}_{\text {gap }} /(\epsilon \lambda)\right)
\end{array}
$$

It remains to show Eq. (30.1). Since the number of edges in the character-based bipartite graph is at least $\lambda\left|T^{\prime}\right| \sqrt{t_{0} t}$, we have

$$
\sum_{i=1}^{t_{0}} D_{i} \geq \lambda\left|T^{\prime}\right| \sqrt{t_{0} t}
$$

By $\left|T^{\prime}\right|=\Omega\left(\mathrm{w}_{\text {gap }} /(\epsilon \lambda)\right)$, we have

$$
\lambda\left|T^{\prime}\right| \sqrt{t_{0} t} \geq t_{0}
$$

Due to a simple fact : $x(x-z)>y(y-z)$ as long as $x>y>z>0$, we have

$$
\begin{equation*}
\frac{1}{t_{0}}\left(\sum_{i=1}^{t_{0}} D_{i}\right)^{2}-\sum_{i=1}^{t_{0}} D_{i} \geq \frac{1}{t_{0}}\left(\lambda\left|T^{\prime}\right| \sqrt{t_{0} t}\right)^{2}-\left(\lambda\left|T^{\prime}\right| \sqrt{t_{0} t}\right) . \tag{30.1}
\end{equation*}
$$

The number of 3-tuple $\left(w_{i}, w_{a}, w_{j}\right)$ is at most $\left|T^{\prime}\right|^{2}$. Thus, there must exists a pair such that

$$
\left\|\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\| \geq(1-\epsilon) \lambda^{2}
$$

We define graph $N G_{\lambda}$ and $\mathrm{NF}_{\lambda}$ as follows:

Definition 30.3.8 $\left(\mathrm{NG}_{\lambda}\right)$. We assume that the edges of $N G_{\lambda}$ are directed in the following way: For an edge $(i, j)$ if $\left|w_{i}\right| \neq\left|w_{j}\right|$ the starting point of the edge would be the vertex corresponding to the longer window and the ending point of the edge would be the one corresponding to the shorter window. If both corresponding windows have the same lengths, then we use an arbitrary direction for $(i, j)$.

Definition 30.3.9 $\left(\mathrm{NF}_{\lambda}\right)$. We construct graph $\mathrm{NF}_{\lambda}$ in the following way: let $a$ denote the node in $V\left(\mathrm{NG}_{\lambda}\right)$ that has the highest outgoing degree, let $V\left(\mathrm{NF}_{\lambda}\right)=N(a)$. We add an edge between vertices $(i, j)$ in $\mathrm{NF}_{\lambda}$ if $\left\|\mid \operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\| \geq(1-\epsilon) \lambda^{2}$. Give directions to the edges of $\mathrm{NF}_{\lambda}$ the same way we did it for $\mathrm{NG}_{\lambda}$.

Now, we are ready to prove the main observation of this section.

Lemma 30.3.9 (upper bound on $E\left(\mathrm{NG}_{\lambda}\right)$ ). Let $\mathrm{NG}_{\lambda}$ be defined as Definition 30.3.3. Then

$$
\left|E\left(\mathrm{NG}_{\lambda}\right)\right|=O\left(k^{2-\gamma} \cdot \mathrm{w}_{\text {layers }} \cdot \mathrm{w}_{\text {gap }} /(\epsilon \lambda)\right) .
$$

holds with probability $1-1 / \operatorname{poly}(k)$.

Proof. We start with assuming

$$
\left|E\left(\mathrm{NG}_{\lambda}\right)\right| \geq \Omega\left(k^{2-\beta} \cdot \mathrm{w}_{\text {layers }} \cdot \mathrm{w}_{\text {gap }} /(\epsilon \lambda)\right) .
$$

We construct a graph $\mathrm{NF}_{\lambda}$ based on Definition 30.3.9. By Definition 30.3.9, we know that $\left|V\left(\mathrm{NF}_{\lambda}\right)\right| \geq \Omega\left(k^{1-\beta} \cdot \mathrm{w}_{\text {layers }} \cdot \mathrm{w}_{\text {gap }} /(\epsilon \lambda)\right)$. By choosing some $\beta=\gamma$ we will make a contradiction.

Using Lemma 30.3.8, we have for each set $T \subseteq V\left(\mathrm{NF}_{\lambda}\right)$ with $|T|=\Omega\left(\mathrm{w}_{\text {layers }} \cdot \mathrm{w}_{\text {gap }} /(\epsilon \lambda)\right)$, there exist two nodes $u$ and $v$ in $T$ such that edge $(u, v) \in \mathrm{NF}_{\lambda}$.

If we look at the complement of graph $\mathrm{NF}_{\lambda}$, we know there is no clique $K_{r+1}$ where $r=O\left(\mathrm{w}_{\text {layers }} \cdot \mathrm{w}_{\text {gap }} /(\epsilon \lambda)\right)$. Using Turan's theorem (Lemma 30.5.5) we know that complement of graph $\mathrm{NF}_{\lambda}$ has at most $\left(1-\frac{1}{r}\right) \frac{q^{2}}{2}$ edges, where $q=\left|V\left(\mathrm{NF}_{\lambda}\right)\right|$. Then we have

$$
\left|E\left(\mathrm{NF}_{\lambda}\right)\right| \geq \frac{q(q-1)}{2}-\left(1-\frac{1}{r}\right) \frac{q^{2}}{2}=\frac{q^{2}}{2 r}-\frac{q}{2} .
$$

This implies that there is one vertex $j$ whose outgoing degree in $\mathrm{NF}_{\lambda}$ is at least at least

$$
\frac{\left|V\left(\mathrm{NF}_{\lambda}\right)\right|}{2 r}-\frac{1}{2} \geq\left|V\left(\mathrm{NF}_{\lambda}\right)\right| \cdot \Omega\left(\frac{\epsilon \lambda}{\mathrm{w}_{\text {layers }} \cdot \mathrm{w}_{\text {gap }}}\right) \geq k^{1-\beta} \geq k^{1-\gamma}
$$

Thus, there is a node in $N G_{\lambda}$ that has a lot of outgoing edges, which means $N G_{\lambda}$ has a close pair.

On the other hand, Using definition of $\mathrm{NG}_{\lambda}$ and Lemma 30.3.7, we know that $\mathrm{NG}_{\lambda}$ should not contain any close pair. Thus, we get a contradiction.

Theorem 30.3.10 (cubic sparsification for arbitrary $\lambda$ ). Given $k$ windows $w_{1}, \cdots, w_{k}$. Let $\mathrm{w}_{\text {layers }}$ denote the number of different sizes for windows. Let $\mathrm{w}_{\max }=\max _{i \in[k]}\left|w_{i}\right|, \mathrm{w}_{\min }=$ $\min _{i \in[k]}\left|w_{i}\right|$ and $\mathrm{w}_{\text {gap }}=\mathrm{w}_{\max } / \mathrm{w}_{\min }$. For any $\lambda \in(0,1), \gamma \in(0,1)$, there is a randomized algorithm (Algorithm 30.2) that runs in time

$$
O\left(\lambda^{-4} k^{1+\gamma} \mathbf{w}_{\max }^{2} \log k+k^{2+\gamma} \mathbf{w}_{\max } \log k\right)
$$

and outputs a table $\widehat{\mathrm{O}}_{\lambda^{3}} \in\{0,1\}^{k \times k}$ such that

$$
\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \lambda^{4} / 8, \text { if } \widehat{\mathrm{O}}_{\lambda^{3}}[i][j]=1
$$

and

$$
\mid\left\{(i, j) \mid\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \lambda, \quad \text { and } \quad \widehat{\mathrm{O}}_{\lambda^{3}}[i][j]=0\right\} \mid=O\left(k^{2-\gamma} \mathbf{w}_{\text {layers }} \mathbf{w}_{\text {gap }} / \lambda\right)
$$

The algorithm has success probability $1-1 / \operatorname{poly}(k)$.

Proof. The running time of each round of Algorithm 30.2 is

$$
\begin{aligned}
& =O\left(\text { constructing } S \text { time }+ \text { Ics-cmp.Initial time }+|S|^{2} \cdot(\text { Ics-cmp.QUERY time })\right) \\
& =O\left(k+\widetilde{\lambda}^{-2}|S| \mathrm{w}_{\max }^{2}+|S|^{2} \mathrm{w}_{\max }\right) \\
& =O\left(\lambda^{-4} k \mathrm{w}_{\max }^{2}+k^{2} \mathrm{w}_{\max }\right)
\end{aligned}
$$

where the last step follows by $\widetilde{\lambda}=\lambda^{2}$.
Since it repeats $O\left(k^{\gamma} \log k\right)$ rounds, thus the overall running time is

$$
O\left(k^{\gamma} \log k\right) \cdot O\left(\lambda^{-4} k \mathrm{w}_{\max }^{2}+k^{2} \mathbf{w}_{\max }\right)=O\left(\lambda^{-4} k^{1+\gamma} \mathrm{w}_{\max }^{2} \log k+k^{2+\gamma} \mathbf{w}_{\max } \log k\right) .
$$



Figure 30.5: computing a set of common subsequences between $w_{a}$ and $w_{i}$ for each $i$ such that every character of $w_{a}$ appears at most once among the sequences for a fixed $i$.

### 30.3.2.2 Implementation of Ics-cmp

Given $\widetilde{\lambda}, w_{a}$ and $w_{1}, w_{2}, \ldots, w_{s}$, we present an $O\left(\widetilde{\lambda}^{-1}\left|w_{a}\right| \sum_{i=1}^{s}\left|w_{i}\right|\right)$ time preprocessing algorithm and $O\left(\left|w_{a}\right|\right)$ time query algorithm for Ics-cmp such that for any $i, j \in[s]$

1. if $\left\|\operatorname{lss}_{w_{a}}\left(w_{i}, w_{j}\right)\right\|>\widetilde{\lambda} / 2$, then the query algorithm outputs accept;
2. if $\left\|\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\| \leq \widetilde{\lambda}^{2} / 8$, then the query algorithm outputs reject.

Algorithm 30.3 is our Ics-cmp preprocessing and query algorithm. The high level idea is to compute opt ${ }_{a, i}$ for every $i \in[s]$ and a set of at most $2 / \lambda$ common subsequences between $w_{a}$ and $w_{i}$ for each $i$ such that every character of $w_{a}$ appears at most once among the sequences for a fixed $i$. $\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)$ is approximated by taking the longest intersection of opt ${ }_{a, i}$ and some sequence in the set of window $w_{j}$. Also see Figure 30.5 for the intuition.

Now we prove that Algorithm 30.3 implements Ics-cmp.
Lemma 30.3.11 (Initial). Given parameter $\widetilde{\lambda}$, a window $w_{a}$ and a set of windows $w_{1}, \cdots, w_{s}$. The Initial of data structure Ics-cmp (Algorithm 30.3) takes $O\left(\widetilde{\lambda}^{-1}\left|w_{a}\right| \sum_{i=1}^{s}\left|w_{i}\right|\right)$ time, and outputs $\left\{X_{a, i}\right\}_{i \in[s]}$ and $\left\{Y_{a, i}\right\}_{i \in[s]}$ such that

1. $X_{a, i}$ corresponds to indices of a longest common subsequence between $w_{a}$ and $w_{i}$ for every $i \in[s]$.
2. $Y_{a, i}$ corresponds to indices of at most $2 / \tilde{\lambda}$ common subsequence between $w_{a}$ and $w_{i}$ for every $i \in[s]$.
3. If none of the element indices of a common subsequence between $w_{a}$ and $w_{i}$ is in $Y_{a, i}$, then the length of this common subsequence is less than $\tilde{\lambda}\left|w_{a}\right|$.

Proof. The first property follows from the definition of the algorithm.
For the second and third property, if an opt ${ }_{a, i}^{\prime}$ obtained in Line 9 has length less than $\widetilde{\lambda}\left|w_{a}\right| / 2$, then the element indices of $\mathrm{opt}_{a, i}^{\prime}$ are not in $Y_{a, i}$. Hence, the third property holds. Also, it means that every time that Line 13 is executed, the size of $Y_{a, i}$ increases by at least $\widetilde{\lambda}\left|w_{a}\right| / 2$. For a fixed $i \in[s]$, Line 13 is executed at most $2 / \widetilde{\lambda}$ times. Thus, the second property holds.

The running time is also obtained by the fact that Line 13 is executed at most $2 / \tilde{\lambda}$ times for every fixed $i \in[s]$.

Lemma 30.3.12 (Query). For any $(i, j) \in[s] \times[s]$, the QuERY of data structure Ics-cmp (Algorithm 30.3) runs in time $O\left(\left|w_{a}\right|\right)$ with the following properties:

1. If $\left\|\operatorname{lss}_{w_{a}}\left(w_{i}, w_{j}\right)\right\|>\tilde{\lambda} / 2$, then it outputs accept.
2. If $\left\|\mid \operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\| \leq \widetilde{\lambda}^{2} / 8$, then it outputs reject.

Proof. Let opt ${ }_{a, i, j}$ be the LCS between opt ${ }_{a, i}$ and $w_{j}$. By the definition, we have

$$
\left\|\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\|=\frac{\left|\mathrm{opt}_{a, i, j}\right|}{\sqrt{\left|w_{i}\right| \cdot\left|w_{j}\right|}}
$$

Consider the case of $\left\|\mid \operatorname{ls}_{w_{a}}\left(w_{i}, w_{j}\right)\right\| \geq \tilde{\lambda} / 2$. By the third property of Lemma 30.3.11, there are less than $\tilde{\lambda}\left|w_{a}\right| / 4$ elements of opt ${ }_{a, i, j}$ with indices not in $Y_{a, j}$. So we have

$$
\left|X_{a, i} \cap Y_{a, j}\right|>\frac{1}{2} \widetilde{\lambda} \sqrt{\left|w_{i}\right|\left|w_{j}\right|}-\frac{1}{4} \widetilde{\lambda}\left|w_{a}\right| \geq \frac{1}{4} \widetilde{\lambda} \sqrt{\left|w_{i}\right|\left|w_{j}\right|}
$$

where the last inequality follows from the fact that $\left|w_{a}\right|$ is smaller than or equal to $\left|w_{i}\right|$ for every $i \in[s]$. Hence the algorithm accepts.

Consider the case of $\left\|\mid \operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\| \leq \widetilde{\lambda}^{2} / 8$. By the second property of Lemma 30.3.11, $Y_{a, j}$ corresponds to at most $2 / \tilde{\lambda}$ mutually non-overlapping common subsequence between $w_{a}$ and $w_{j}$. Since $\left\|\operatorname{lcs}_{w_{a}}\left(w_{i}, w_{j}\right)\right\| \leq \widetilde{\lambda}^{2} / 8$, the set of indices of every such a common subsequence has an intersection with $X_{a, i}$ of size at most $\widetilde{\lambda}^{2} \sqrt{\left|w_{i}\right|\left|w_{j}\right|} / 8$. Thus, we have

$$
\left|X_{a, i} \cap Y_{a, j}\right| \leq \frac{1}{8} \widetilde{\lambda}^{2} \sqrt{\left|w_{i}\right|\left|w_{j}\right|} \cdot \frac{2}{\widetilde{\lambda}} \leq \frac{1}{4} \widetilde{\lambda} \sqrt{\left|w_{i}\right|\left|w_{j}\right|} .
$$

Hence the algorithm rejects for this case.

### 30.4 Longest Increasing Subsequence

We outlined the algorithm in Section 30.1.3.3. Here, we bring the details for each step of the algorithm. In Section 30.4.1 we discuss the solution domains and show how we construct them. Next, in Section 30.4.2 we discuss the details of constructing pseudosolutions and finally in Section 30.4.3 we show how we can obtain an approximate solution from the pseudo-solutions. Also, in Section 30.4.4 we bring an improvement to the running time at the expense of having a larger approximation factor for the algorithm.

### 30.4.1 Solution Domains

We assume from now on that $\operatorname{lis}(A)>n \lambda$ holds. As mentioned earlier, we divide the input array into $\sqrt{n}$ subarrays of size $\sqrt{n}$ and denote them by $\mathrm{sa}_{1}, \mathrm{sa}_{2}, \ldots$, sa ${ }_{\sqrt{n}}$. For a fixed optimal solution opt, our goal is to approximate the smallest and the largest number of each subarray that contributes to opt. Let us refer to these numbers as the domain of each subarray. Let $\epsilon=1 / 1000$ be accuracy. For a subarray sa ${ }_{i}$, we sample $k$ (will be decided later) different elements and refer to them by $a_{j_{1}}, a_{j_{2}}, \ldots, a_{j_{k}}$.

We first prove that,

Lemma 30.4.1 (constructing candidate domains). Let $\lambda \in(0,1), \epsilon \in(0,1 / 2)$ and $\delta \in$ $(0,1 / 10)$. Let $\mathrm{sa}_{i}$ be a length $-\sqrt{n}$ subarray whose contribution to the optimal solution is at least $\epsilon \sqrt{n} \lambda$, i.e., $\operatorname{lis}{ }^{\left[\operatorname{sm}\left(\operatorname{sa}_{i}\right), \lg \left(\operatorname{sa}_{i}\right)\right]}\left(\operatorname{sa}_{i}\right) \geq \epsilon \sqrt{n} \lambda$. If we uniformly sample $k=20 \log (1 / \delta) /\left(\lambda \epsilon^{2}\right)$ elements $a_{j_{1}}, a_{j_{2}}, \ldots, a_{j_{k}}$ from sa $_{i}$, then with probability at least $1-\delta$, there exists a pair $(\alpha, \beta) \in[k] \times[k]$ such that the following two conditions hold

1. $\operatorname{sm}\left(\operatorname{sa}_{i}\right) \leq a_{j_{\alpha}} \leq a_{j_{\beta}} \leq \lg \left(\mathrm{sa}_{i}\right)$,
2. $\operatorname{lis}{ }^{\left[a_{j_{\alpha}}, a_{j_{\beta}}\right]}\left(\operatorname{sa}_{i}\right) \geq(1-\epsilon) \operatorname{lis}{ }^{\left[\operatorname{sm}\left(\mathbf{s a}_{i}\right), \lg \left(\mathrm{sa}_{i}\right)\right]}\left(\mathrm{sa}_{i}\right)$.

Proof. At least $\epsilon \sqrt{n} \lambda$ elements of $s_{i}$ appear in opt. Let us put all these elements in an array b in the same order that they appear in $\mathrm{sa}_{i}$. Then it is obvious that b has at least $\epsilon \sqrt{n} \lambda$ elements. To prove the lemma, we bound the probability that none of the first $\epsilon / 2$ fraction of the elements of $b$ is sampled in our algorithm.
$\operatorname{Pr}[$ none of the elements in the first $\epsilon / 2$ fraction of $b$ is sampled $]$

$$
\begin{aligned}
& \leq\left(1-\frac{\epsilon}{2} \cdot \epsilon \sqrt{n} \lambda \cdot \frac{1}{\sqrt{n}}\right)^{k} \\
& =\left(1-\frac{\epsilon^{2} \lambda}{2}\right)^{\frac{2}{\epsilon^{2} \lambda} \cdot 10 \log (1 / \delta)} \\
& \leq e^{-10 \log (1 / \delta)} \\
& \leq \delta / 2
\end{aligned}
$$

where the first step follows from the fact that b contains at least $\epsilon \lambda \sqrt{n}$ elements, the second step follows from $k=20 \log (1 / \delta) /\left(\lambda \epsilon^{2}\right)$, and the third step follows from the fact that the $(1-1 / x)^{x} \leq 1 / e$ for $\forall x \geq 4$. Hence, with probability at least $1-\delta / 2$, at least one of the elements in the first $\epsilon / 2$ fraction of $\mathbf{b}$ are sampled.

With the same analysis, one can prove that with probability at least $1-\delta / 2$ at least one of the elements in the last $\epsilon / 2$ fraction of $\mathbf{b}$ are also sampled.

Taking a union bound of two events, with probability at least $1-\delta$, at least one of the elements in the first and at least one of the elements in the last $\epsilon / 2$ fraction $\mathbf{o f} \mathbf{b}$ are sampled.

Therefore, the lis of sa ${ }_{i}$ subject to this interval is at least a $1-\epsilon$ fraction of lis ${ }^{\left[\operatorname{sm}\left(\operatorname{sa}_{i}\right), \lg \left(\operatorname{sa}_{i}\right)\right]}\left(\operatorname{sa}_{i}\right)$.

Notice that the average contribution of each subarray to opt is $\sqrt{n} \lambda$ and Lemma 30.4.1 applies to a subarray if its contribution to opt is at least an $\epsilon$ fraction of this value. Therefore Lemma 30.4.1 implies that a considerable fraction of the solution is covered by the candidate domains.

Corollary 30.4.2 (existence of a desirable solution). Let $\lambda \in(0,1)$ such that $\operatorname{lis}(A) \geq n \lambda$ and $\epsilon \in(0,1 / 4)$. If we run Algorithm 30.4 with parameter $\delta=\epsilon$ on every subarray independently, then with probability at least $1-\exp \left(-\Omega\left(\epsilon^{2} \sqrt{n} \lambda\right)\right)$, there exist a set $T \subseteq[\sqrt{n}]$ and elements $\alpha_{i}$ and $\beta_{i}$ sampled from $\mathbf{s a}_{i}$ for each $i \in T$ such that the following conditions hold:

1. For any $i \in T, \alpha_{i} \leq \beta_{i}$.
2. For any $i, j \in T$ satisfying $i<j, \beta_{i}<\alpha_{j}$.
3. $\sum_{i \in T} \operatorname{lis}^{\left[\alpha_{i}, \beta_{i}\right]}\left(\operatorname{sa}_{i}\right) \geq(1-4 \epsilon) \operatorname{lis}(A)$.

Proof. Lemma 30.4.1 holds for all subarrays whose contribution to opt is at least $\epsilon \sqrt{n} \lambda$. Let $S \subseteq[\sqrt{n}]$ denote the set of coordinates such that for each $i \in S$

$$
\operatorname{lis}^{\left[\operatorname{sm}\left(\mathrm{sa}_{i}\right), \lg \left(\mathbf{s a}_{i}\right)\right]}\left(\mathrm{sa}_{i}\right) \geq \epsilon \sqrt{n} \lambda .
$$

Since $\sum_{i=1}^{\sqrt{n}} \operatorname{lis}{ }^{\left[s m\left(\mathbf{s a}_{i}\right), \lg \left(\mathbf{s a}_{i}\right)\right]}\left(\mathrm{sa}_{i}\right) \geq n \lambda$ and $\operatorname{lis}{ }^{\left[\mathrm{sm}\left(\mathrm{sa}_{i}\right), \lg \left(\mathbf{s a}_{i}\right)\right]}\left(\mathrm{sa}_{i}\right) \leq \sqrt{n}$, we have

$$
\begin{equation*}
\sum_{i \in S} \operatorname{lis}^{\left[\operatorname{sm}\left(\mathbf{s a}_{i}\right), \lg \left(\mathbf{s a}_{i}\right)\right]}\left(\operatorname{sa}_{i}\right) \geq \sum_{i=1}^{\sqrt{n}} \operatorname{lis}{ }^{\left[\operatorname{sm}\left(\mathbf{s a}_{i}\right), \lg \left(\mathbf{s a}_{i}\right)\right]}\left(\operatorname{sa}_{i}\right)-\sqrt{n} \cdot \epsilon \sqrt{n} \lambda \geq \operatorname{lis}(A)-\epsilon n \lambda . \tag{30.2}
\end{equation*}
$$

Let $T \subseteq S$ denote the set of coordinates such that for each $i \in T$,

$$
\operatorname{lis}{ }^{\left[\alpha_{i}, \beta_{i}\right]}\left(\mathbf{s a}_{i}\right) \geq(1-\epsilon) \operatorname{lis}{ }^{\left[\mathrm{sm}\left(\mathbf{s a}_{i}\right), \lg \left(\mathbf{s a}_{i}\right)\right]}\left(\mathrm{sa}_{i}\right), \quad \text { and } \quad \operatorname{lis}{ }^{\left[\mathrm{sm}\left(\mathbf{s a}_{i}\right), \lg \left(\mathbf{s a}_{i}\right)\right]}\left(\mathrm{sa}_{i}\right) \geq \epsilon \sqrt{n} \lambda .
$$

Now we show that with probability at least $1-\exp \left(-\Omega\left(\epsilon^{2} \sqrt{n} \lambda\right)\right)$,

$$
\begin{equation*}
\sum_{i \in T} \operatorname{lis}^{\left[s m\left(\operatorname{sa}_{i}\right), \lg \left(\mathrm{sa}_{i}\right)\right]}\left(\operatorname{sa}_{i}\right) \geq(1-2 \epsilon) \sum_{i \in S} \operatorname{lis}^{\left[\mathrm{sm}\left(\mathrm{sa}_{i}\right), \lg \left(\mathrm{sa}_{i}\right)\right]}\left(\mathrm{sa}_{i}\right) . \tag{30.3}
\end{equation*}
$$

For each $i \in S$, let $X_{i}$ denote a random variable such that

$$
X_{i}= \begin{cases}\operatorname{lis}{ }^{\left[\operatorname{sm}\left(\mathrm{sa}_{i}\right), \lg \left(\mathrm{sa}_{i}\right)\right]}\left(\mathrm{sa}_{i}\right), & \text { with probability of } i \in T \\ 0, & \text { with probability of } i \notin T\end{cases}
$$

and $X=\sum_{i \in S} X_{i}$. By Lemma 30.4.1 (with $\delta=\epsilon$ ), We have $\mathbb{E}[X] \geq(1-\epsilon) \sum_{i \in S} \operatorname{lis}{ }^{\left[\operatorname{sm}\left(\mathbf{s a}_{i}\right), \lg \left(\mathbf{s a}_{i}\right)\right]}\left(\operatorname{sa}_{i}\right)$. By Hoeffding bound (Theorem 30.5.2),

$$
\begin{aligned}
\operatorname{Pr}[X-\mathbb{E}[X] \geq \epsilon \mathbb{E}[X]] & \leq 2 \exp \left(-\frac{2 \epsilon^{2}(\mathbb{E}[X])^{2}}{\sum_{i \in S}\left(\operatorname{lis}^{\left[\operatorname{sm}\left(\mathbf{s a}_{i}\right), \mathrm{lg}\left(\mathbf{s a}_{i}\right)\right]}\left(\mathbf{s a}_{i}\right)\right)^{2}}\right) \\
& \leq 2 \exp \left(-\frac{2 \epsilon^{2}(1-\epsilon)^{2}\left(\sum_{i \in S} \operatorname{lis}^{\left[\operatorname{sm}\left(\mathbf{s a}_{i}\right), \mathrm{lg}\left(\mathbf{s a}_{i}\right)\right]}\left(\mathbf{s a}_{i}\right)\right)^{2}}{n^{3 / 2}}\right) \\
& \leq 2 \exp \left(-2 \epsilon^{2}(1-\epsilon)^{4} \sqrt{n} \lambda\right) \\
& \leq \exp \left(-\Omega\left(\epsilon^{2} \sqrt{n} \lambda\right)\right) .
\end{aligned}
$$

Hence, Equation (30.3) holds with probability at least $1-\exp \left(-\Omega\left(\epsilon^{2} \sqrt{n} \lambda\right)\right)$.

Conditioned on Equation (30.3), we have

$$
\begin{align*}
& \sum_{i \in T} \operatorname{lis}^{\left[\alpha_{i}, a_{\beta_{i}}\right]}\left(\mathrm{sa}_{i}\right) \\
\geq & \sum_{i \in T}(1-\epsilon) \operatorname{lis}^{\left[\mathrm{sm}\left(\mathrm{sa}_{i}\right), \lg \left(\mathrm{sa}_{i}\right)\right]}\left(\mathrm{sa}_{i}\right) \\
\geq & (1-\epsilon)(1-2 \epsilon) \sum_{i \in S} \operatorname{lis}^{\left[\operatorname{sm}\left(\mathbf{s a}_{i}\right), \lg \left(\mathbf{s a}_{i}\right)\right]}\left(\mathrm{sa}_{i}\right)  \tag{30.4}\\
\geq & (1-\epsilon)(1-2 \epsilon)(\operatorname{lis}(A)-\epsilon n \lambda) \\
\geq & (1-4 \epsilon) \operatorname{lis}(A)
\end{align*}
$$

where the first inequality follows from the definition of $T$, the second inequality follows from Equation (30.3), the third inequality follows from Equation (30.2) and the last inequality follows from $\operatorname{lis}(A) \geq n \lambda$.

Finally, by Equation (30.3) and (30.4), we have

$$
\operatorname{Pr}\left[\sum_{i \in T} \operatorname{lis}^{\left[\alpha_{i}, a_{\beta_{i}}\right]}\left(\operatorname{sa}_{i}\right) \geq(1-4 \epsilon) \operatorname{lis}(A)\right] \geq 1-\exp \left(-\Omega\left(\epsilon^{2} \sqrt{n} \lambda\right)\right) .
$$

### 30.4.2 Constructing Approximately Optimal Pseudo-solutions

We call a sequence of $\sqrt{n}$ intervals $\left[\ell_{1}, r_{1}\right],\left[\ell_{2}, r_{2}\right], \ldots,\left[\ell_{\sqrt{n}}, r_{\sqrt{n}}\right]$ a pseudo-solution if all of the intervals are monotone. That is $\ell_{1} \leq r_{1}<\ell_{2} \leq r_{2}<\ell_{3} \leq r_{3} \leq \ldots \leq \ell_{\sqrt{n}} \leq r_{\sqrt{n}}$. These intervals denote solution-domains for the subarrays. We also may decide not assign any solution domain to a subarray in which case we show the corresponding interval by $\emptyset$. Indeed, $\emptyset$ does not break the monotonicity of a pseudo-solution. The quality of a pseudosolution is defined as $\sum_{i} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right)$. We denote the quality of a pseudo-solution ps by $\mathrm{q}(\mathrm{ps})$.

Another way to interpret Corollary 30.4.2 is that one can construct a pseudo-solution using the sampled elements whose quality is at least a $1-\epsilon$ fraction of $\operatorname{lis}(A)$. In this section, we present an algorithm to construct a small set of pseudo-solutions with the promise that at least one of them has a quality of at least $\operatorname{lis}(A) / t$, where $t$ is the number of pseudosolutions. Finally, in Section 30.4.3, we present a method to approximate the size of the optimal solution using pseudo-solutions.

We construct the pseudo-solutions via Algorithm 30.5. The input of Algorithm 30.5 is the set of candidate domain intervals obtained by Algorithm 30.4 on every subarray. We first find an assignment of candidate solution domains to the subarrays which is monotone and has the largest number of candidate domain intervals. (This step can be implemented by dynamic programming.) We make a pseudo-solution out of this assignment and update the set of candidate intervals by removing the ones which are used in our pseudo-solution. We then repeat the same procedure to construct the second pseudo-solution and update the candidate solution domains accordingly. We continue on, until the number of solution domains used in our pseudo-solution drops below $\epsilon \sqrt{n}$ in which case we stop.

We first prove in Lemma 30.4.3 that the number of pseudo-solutions constructed in Algorithm 30.5 is bounded by $O\left(k^{2} /(\lambda \epsilon)\right)$. Next, we show in Lemma 30.4.4 that at least one of the pseudo-solutions constructed by Algorithm 30.5 has a quality of at least $\Omega(\operatorname{lis}(A) / t)$ where $t$ is the number of pseudo-solutions. Finally we prove in Lemma 30.4.5 that the running time of Algorithm 30.5 is $O\left(t k^{2} \sqrt{n} \log n\right)$.

Lemma 30.4.3 (number of pseudo-solutions). For each $i \in[\sqrt{n}]$, let $\mathrm{cdi}_{i}$ be a set of at most $k^{2}$ candidate domain intervals. Let $t$ denote the number of pseudo-solutions constructed in Algorithm 30.5. Then, we have $t \leq k^{2} /(\lambda \epsilon)$.

Proof. Note that for each subarray we have at most $k^{2}$ candidate domain intervals. Since there are $\sqrt{n}$ subarrays, then in total we have $\sqrt{n} k^{2}$ candidate domain intervals. Each time we construct a pseudo-solution, the total number of the candidate domain intervals is decreased by at least $\epsilon \sqrt{n} \lambda$. Thus, the total number of pseudo-solutions $t$ can be upper bounded,

$$
t \leq \frac{\sqrt{n} k^{2}}{\epsilon \sqrt{n} \lambda} \leq \frac{k^{2}}{\lambda \epsilon}
$$

Lemma 30.4.4 (quality of pseudo-solutions). Let $\mathrm{ps}_{1}, \mathrm{ps}_{2}, \ldots, \mathrm{ps}_{t}$ be the pseudo-solutions constructed by Algorithm 30.5. If $\operatorname{lis}(A) \geq n \lambda$ holds, then with probability at least $1-$ $\exp (-\Omega(\sqrt{n} \lambda))$, there exists an $i \in[t]$ such that

$$
\mathrm{q}\left(\mathrm{ps}_{i}\right) \geq \frac{\operatorname{lis}(A)}{2 t}
$$

Proof. Let us focus again on the actual solution domains

$$
\left[\operatorname{sm}\left(s a_{1}\right), \lg \left(\mathrm{sa}_{1}\right)\right],\left[\mathrm{sm}\left(\mathrm{sa}_{2}\right), \lg \left(\mathrm{sa}_{2}\right)\right], \ldots,\left[\mathrm{sm}(\mathrm{sa} \sqrt{n}), \lg \left(\mathrm{sa}_{\sqrt{n}}\right)\right] .
$$

We define set $S \subseteq[\sqrt{n}]$ such that

$$
\operatorname{lis}^{\left[s m\left(\operatorname{sa}_{i}\right), \lg \left(\mathrm{sa}_{i}\right)\right]} \geq \epsilon \sqrt{n} \lambda, \forall i \in S
$$

Using Corollary 30.4.2 with $\epsilon=1 / 10$, we know that there is a monotone pseudosolution $\left[\alpha_{i}, \beta_{i}\right]_{i \in T}(T \subseteq S)$ such that $\left[\alpha_{i}, \beta_{i}\right]$ are candidate domain intervals and

$$
\sum_{i \in T} \operatorname{lis}^{\left[\alpha_{i}, \beta_{i}\right]}\left(\mathbf{s a}_{i}\right) \geq(1-4 \epsilon) \operatorname{lis}(A) .
$$

Denote this pseudo-solution as sol.
At the time we terminate Algorithm 30.5, there are at most $\epsilon \sqrt{n} \lambda$ candidate domain intervals of sol does not belongs to any pseudo-solution of the pseudo-solution set. Also, since each candidate domain interval contributes at most $\sqrt{n}$ to the quality of the pseudo-solution containing the interval, we have

$$
\sum_{i=1}^{t} \mathrm{q}\left(\mathrm{ps}_{i}\right) \geq(1-4 \epsilon) \operatorname{lis}(A)-\epsilon \sqrt{n} \lambda \cdot \sqrt{n} \geq(1-4 \epsilon) \operatorname{lis}(A)-\epsilon \operatorname{lis}(A)=(1-5 \epsilon) \operatorname{lis}(A) .
$$

Thus, there exists an $i \in[t]$ such that

$$
\mathrm{q}\left(\mathrm{ps}_{i}\right) \geq(1-5 \epsilon) \operatorname{lis}(A) / t=\operatorname{lis}(A) /(2 t)
$$

Lemma 30.4.5 (running time). For each $i \in[\sqrt{n}]$, let $\operatorname{cdi}_{i}$ be a set of at most $k^{2}$ candidate domain intervals. Let $t$ denote the number of pseudo-solutions. The running time of Algorithm 30.5 is bounded by $O\left(t k^{2} \sqrt{n} \log n\right)$.

Proof. Lemma 30.4.3 states that Algorithm 30.5 terminates after constructing $t$ pseudosolutions.

Now we show that constructing each pseudo-solution takes time $O\left(k^{2} \sqrt{n} \log n\right)$. Our solution is based on a dynamic programming technique. Let $D:[\sqrt{n}] \times\left[k^{2}\right] \rightarrow \mathbb{N}$ be an array such that $D[i][j]$ stores the size of the largest monotone pseudo-solution for the first $i$ subarrays which ends with the $j$ 'th candidate domain interval of sa ${ }_{i}$. Using classic segmenttree data structure (this data structure can be found in many textbooks, e.g. [CLRS09]), one can compute the value of $D[i][j]$ in time $O(\log n)$ from the previously computed elements of the array.

Thus, the total running is bounded by $O\left(t k^{2} \sqrt{n} \log n\right)$.

### 30.4.3 Evaluating the Pseudo-solutions

We finally use a concentration bound to show that the quality of a pseudo-solution can be approximated well by sampling a small number of subarrays. Since a pseudo-solution specifies the range of the numbers used in every subarray, the quality of a pseudo-solution, or in other words, the size of the corresponding increasing subsequence of a pseudo-solution can be formulated as

$$
\mathrm{q}(\mathrm{ps})=\sum_{i=1}^{\sqrt{n}} \mathrm{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right)
$$

where $\left[\ell_{i}, r_{i}\right]$ denotes the corresponding solution domain of ps for $\mathrm{sa}_{i}$.
In Lemma 30.4.6, we prove that by sampling $O\left(\log ^{O(1)} n / \lambda^{4}\right)$ many subarrays and computing lis ${ }^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right)$ for them, one can approximate the quality of a pseudo-solution pretty accurately.

Lemma 30.4.6 (the quality of pseudo-solution). Let $\lambda \in(0,1)$ and $\epsilon$ be a constant in $(0,1 / 100)$. Let $\mathrm{ps}_{1}, \mathrm{ps}_{2}, \cdots, \mathrm{ps}_{t}$ be a set of $t$ pseudo-solutions. With probability at least $1-\exp \left(-\Omega\left(\log ^{2} n\right)\right)$, Algorithm 30.6 runs in time $O\left(t^{2} \sqrt{n} \log ^{O(1)} n / \lambda\right)$ such that,

1. If there exists an $i \in[t], \mathrm{q}\left(\mathrm{ps}_{i}\right) \geq \frac{\lambda n}{2 t}$, then the algorithm outputs an estimation at least $\frac{\lambda n}{4 t}$.
2. If $\mathrm{q}\left(\mathrm{ps}_{i}\right)<\frac{\lambda n}{8 t}$ for al the $i \in[t]$, then the algorithm outputs an estimation smaller than $\frac{\lambda n}{4 t}$.

Proof. By Chernoff bound, with probability $1-\exp \left(-\Omega\left(\log ^{3} n\right)\right)$ most $2 p \sqrt{n}=\frac{2000 t \log ^{4} n}{\epsilon^{4} \lambda}$ subarrays are sampled. For each pseudo-solution, Algorithm 30.6 runs the algorithm for longest increasing subsequence once for every sampled subarray. Hence, the running time of the algorithm is $O\left(t^{2} \sqrt{n} \log ^{O(1)} n / \lambda\right)$.

Consider an arbitrary pseudo-solution $\mathrm{ps} \in\left\{\mathrm{ps}_{1}, \ldots, \mathrm{ps}_{t}\right\}$ for now. We show that

$$
\operatorname{Pr}\left[\left(1-\frac{\epsilon}{4}\right)^{2}\left(\mathrm{q}(\mathrm{ps})-\frac{\epsilon \lambda^{4} n}{100 t}\right) \widetilde{\mathrm{q}}(\mathrm{ps}) \leq\left(1+\frac{\epsilon}{4}\right)^{2} \mathrm{q}(\mathrm{ps})+\frac{\epsilon \lambda^{4} n}{100 t}\right] \leq 1-\exp \left(-\Omega\left(\log ^{3} n\right)\right)
$$

Then the lemma holds by a union bound on all the pseudo-solutions.
Let $T$ be the set of subarray indices such that $i \in T$ iff there is a non-empty interval [ $\left.\ell_{i}, r_{i}\right]$ of ps corresponding to subarray sa ${ }_{i}$. Let $p=\frac{1000 t \log ^{4} n}{\epsilon^{4} \lambda \sqrt{n}}$ be the probability of sampling a subarray. For each $i \in T$, let $X_{i}$ denote a random variable such that

$$
X_{i}= \begin{cases}1, & \text { with prob. } p \\ 0, & \text { with prob. } 1-p\end{cases}
$$

and

$$
X=\sum_{i \in T} \frac{1}{p} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) X_{i} .
$$

We have

$$
\mathbb{E}[X]=\mathbb{E}\left[\sum_{i \in T} \frac{1}{p} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) X_{i}\right]=\sum_{i \in T} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right)=\mathrm{q}(\mathrm{ps}) .
$$

Let

$$
T_{j}=\left\{i \in T:(1+\epsilon / 4)^{j-1} \leq \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right)<(1+\epsilon / 4)^{j}\right\}
$$

for integer $1 \leq j \in\left\lceil\log _{1+\epsilon} \sqrt{n}\right\rceil$.

Let $\Delta=\frac{\epsilon^{2} \lambda \sqrt{n}}{1000 t \log n}$. Consider $T_{j}$ 's such that $\left|T_{j}\right| \geq \Delta$. The contribution of $\mathrm{q}(\mathrm{ps})$ mostly comes from $T_{j}$ in the following sense

$$
\begin{equation*}
\sum_{j:\left|T_{j}\right| \geq \Delta} \sum_{i \in T_{j}} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) \leq \mathrm{q}(\mathrm{ps}) \tag{30.5}
\end{equation*}
$$

and

$$
\begin{align*}
\sum_{j:\left|T_{j}\right| \geq \Delta} \sum_{i \in T_{j}} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) & =\mathrm{q}(\mathrm{ps})-\sum_{j:\left|T_{j}\right|<\Delta} \sum_{i \in T_{j}} \mathrm{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) \\
& \geq \mathrm{q}(\mathrm{ps})-\Delta \sqrt{n} \cdot\left\lceil\log _{1+\epsilon} \sqrt{n}\right\rceil  \tag{30.6}\\
& \geq \mathrm{q}(\mathrm{ps})-\frac{\epsilon \lambda n}{100 t} .
\end{align*}
$$

Now we bound the random variable $\sum_{j:\left|T_{j}\right| \geq \Delta} \sum_{i \in T_{j}}$ lis $^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) X_{i} / p$. By Chernoff bound, for each $j$ such that $\left|T_{j}\right| \geq \Delta$, we have

$$
\begin{equation*}
\operatorname{Pr}\left[\left(1-\frac{\epsilon}{4}\right) p\left|T_{j}\right| \leq \sum_{i \in T_{j}} X_{i} \leq\left(1+\frac{\epsilon}{4}\right) p\left|T_{j}\right|\right] \geq 1-\exp \left(-\Omega\left(\epsilon^{2} p\left|T_{j}\right|\right)\right)=1-\exp \left(-\Omega\left(\log ^{3} n\right)\right) \tag{30.7}
\end{equation*}
$$

Notice that for set $T_{j}$,

$$
\begin{equation*}
\frac{(1+\epsilon / 4)^{j-1}}{p} \sum_{i \in T_{j}} X_{i} \leq \sum_{i \in T_{j}} \frac{1}{p} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) X_{i} \leq \frac{(1+\epsilon / 4)^{j}}{p} \sum_{i \in T_{j}} X_{i} . \tag{30.8}
\end{equation*}
$$

We have

$$
\begin{aligned}
& \operatorname{Pr}\left[\left(1-\frac{\epsilon}{4}\right)^{2} \sum_{i \in T_{j}} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathbf{s a}_{i}\right) \leq \sum_{i \in T_{j}} \frac{1}{p} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathbf{s a}_{i}\right) X_{i} \leq\left(1+\frac{\epsilon}{4}\right)^{2} \sum_{i \in T_{j}} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathbf{s a}_{i}\right)\right] \\
& \geq \operatorname{Pr}\left[\left(1-\frac{\epsilon}{4}\right)^{2} \sum_{i \in T_{j}} \mathrm{is}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) \leq \frac{(1+\epsilon / 4)^{j-1}}{p} \sum_{i \in T_{j}} X_{i}\right. \\
& \left.\bigwedge \frac{(1+\epsilon / 4)^{j}}{p} \sum_{i \in T_{j}} X_{i} \leq\left(1+\frac{\epsilon}{4}\right)^{2} \sum_{i \in T_{j}} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathbf{s a}_{i}\right)\right] \\
& \geq \operatorname{Pr}\left[\left(1-\frac{\epsilon}{4}\right)^{2}\left(1+\frac{\epsilon}{4}\right)^{j}\left|T_{j}\right| \leq \frac{(1+\epsilon / 4)^{j-1}}{p} \sum_{i \in T_{j}} X_{i}\right. \\
& \left.\bigwedge \frac{(1+\epsilon / 4)^{j}}{p} \sum_{i \in T_{j}} X_{i} \leq\left(1+\frac{\epsilon}{4}\right)^{j+1}\left|T_{j}\right|\right] \\
& \geq \operatorname{Pr}\left[\left(1-\frac{\epsilon}{4}\right) p\left|T_{j}\right| \leq \sum_{i \in T_{j}} X_{i} \quad \bigwedge \sum_{i \in T_{j}} X_{i} \leq\left(1+\frac{\epsilon}{4}\right) p\left|T_{j}\right|\right] \\
& =1-\exp \left(-\Omega\left(\log ^{3} n\right)\right),
\end{aligned}
$$

where the first inequality follows from Equation (30.8), the second inequality follows from the definition of $T_{j}$ and the last inequality follows from Equation (30.7). By Equation (30.5),
(30.6), (30.10) and union bound, we have

$$
\begin{aligned}
& \quad \operatorname{Pr}\left[\left(1-\frac{\epsilon}{4}\right)^{2}\left(\mathrm{q}(\mathrm{ps})-\frac{\epsilon \lambda n}{100 t}\right) \leq \sum_{j:\left|T_{j}\right| \geq \Delta} \sum_{i \in T_{j}} \frac{1}{p} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) X_{i} \leq\left(1+\frac{\epsilon}{4}\right)^{2} \mathrm{q}(\mathrm{ps})\right] \\
& \geq \operatorname{Pr}\left[\left(1-\frac{\epsilon}{4}\right)^{2} \sum_{j:\left|T_{j}\right| \geq \Delta} \sum_{i \in T_{j}} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) \leq \sum_{j:\left|T_{j}\right| \geq \Delta} \sum_{i \in T_{j}} \frac{1}{p} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) X_{i}\right. \\
& \left.\quad \leq \sum_{j:\left|T_{j}\right| \geq \Delta}\left(1+\frac{\epsilon}{4}\right)^{2} \sum_{i \in T_{j}} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right)\right] \\
& \geq \\
& \\
& 1-O\left(\frac{\log n}{\epsilon}\right) \exp \left(-\Omega\left(\log ^{3} n\right)\right) \\
& = \\
& 1-\exp \left(-\Omega\left(\log ^{3} n\right)\right) .
\end{aligned}
$$

Now we bound random variable $\sum_{j:\left|T_{j}\right|<\Delta} \sum_{i \in T_{j}} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) X_{i} / p$. If $\left|T_{j}\right|<\Delta$, then by Chernoff bound we have

$$
\operatorname{Pr}\left[\sum_{i \in T_{j}} X_{i} \leq \frac{2 \log ^{3} n}{\epsilon^{2}}\right] \geq 1-\exp \left(-\Omega\left(\log ^{3} n\right)\right)
$$

By union bound we have
$\left.\operatorname{Pr}\left[\sum_{j:\left|T_{j}\right|<\Delta} \sum_{i \in T_{j}} \frac{1}{p} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\operatorname{sa}_{i}\right) X_{i} \leq \frac{2 \log ^{3} n}{\epsilon^{2}} \cdot \frac{\sqrt{n}}{p} \cdot\left\lceil\log _{1+\epsilon} \sqrt{n}\right\rceil<\frac{\epsilon \lambda n}{100 t}\right] \geq 1-\exp \left(-\Omega\left(\log ^{3} n\right)\right)\right)$.

If $\mathrm{q}(\mathrm{ps}) \geq \frac{\lambda n}{2 t}$, then then by Equation (30.10)

$$
\begin{aligned}
& \operatorname{Pr}\left[X>\frac{\lambda n}{4 t}\right] \\
\geq & \operatorname{Pr}\left[\left(1-\frac{\epsilon}{4}\right)^{2}\left(\mathrm{q}(\mathrm{ps})-\frac{\epsilon \lambda n}{100 t}\right) \leq \sum_{j:\left|T_{j}\right| \geq \Delta} \sum_{i \in T_{j}} \frac{1}{p} \mathrm{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) X_{i}\right] \\
\geq & 1-\exp \left(-\Omega\left(\log ^{3} n\right)\right) .
\end{aligned}
$$

If $\mathrm{q}(\mathrm{ps})<\frac{\lambda n}{8 t}$, then by Equation (30.10) and Equation (30.11),

$$
\begin{aligned}
& \operatorname{Pr}\left[X<\frac{\lambda n}{4 t}\right] \\
\geq & \operatorname{Pr}\left[\sum_{j:\left|T_{j}\right| \geq \Delta} \sum_{i \in T_{j}} \frac{1}{p} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) X_{i} \leq\left(1+\frac{\epsilon}{4}\right)^{2} \mathrm{q}(\mathrm{ps}) \quad \bigwedge \sum_{j:\left|T_{j}\right|<\Delta} \sum_{i \in T_{j}} \frac{1}{p} \operatorname{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) X_{i} \leq \frac{\epsilon \lambda n}{100 t}\right] \\
\geq & 1-\exp \left(-\Omega\left(\log ^{3} n\right)\right) .
\end{aligned}
$$

Putting all the previous Lemmas together gives the following result:
Corollary 30.4.7 (algorithm for LIS decision problem). Given a length-n sequence A, let $\lambda \in[1 / n, 1]$. There is a randomized algorithm that runs in time $O\left(\lambda^{-7} \sqrt{n} \log ^{O(1)} n\right)$ such that with probability $1-1 / \operatorname{poly}(n)$

- The algorithm accepts if $\operatorname{lis}(A)>n \lambda$.
- The algorithm rejects if $\operatorname{lis}(A)=O\left(n \lambda^{4}\right)$.

Proof. The correctness follows from Lemma 30.4.6, Lemma 30.4.4, and Lemma 30.4.5.
Running time: The running time is

$$
\begin{aligned}
\text { time } & =\underbrace{O\left(t k^{2} \sqrt{n} \log ^{O(1)} n\right)}_{\text {Lemma } 30.4 .5}+\underbrace{O\left(t^{2} \lambda^{-1} \sqrt{n} \log ^{O(1)} n\right)}_{\text {Lemma 30.4.6 }} & & \\
& =O\left(t^{2} \lambda^{-1} \sqrt{n} \log ^{O(1)} n\right) & & \text { by Lemma 30.4.3 } \\
& =O\left(k^{4} \lambda^{-3} \sqrt{n} \log ^{O(1)} n\right) & & \text { by Lemma 30.4.3 } \\
& =O\left(\lambda^{-7} \sqrt{n} \log ^{O(1)} n\right) & &
\end{aligned}
$$

Thus, we complete the proof.

Finally, by starting with $\lambda=1$ and iteratively multiplying $\lambda$ by a $1 /(1+\epsilon)$ factor until a solution is found, we can approximate lis $(A)$ within an approximation factor of $O\left(\lambda^{3}\right)$.

Theorem 30.4.8 (polynomial approximation for LIS). Given a length-n sequence $A$ such that $\operatorname{lis}(A)=n \lambda$ where $\lambda \in[1 / n, 1]$ is unknown to the algorithm. There is an algorithm that runs in time $\widetilde{O}\left(\lambda^{-7} \sqrt{n}\right)$ and outputs a number est such that

$$
\Omega\left(\operatorname{lis}(A) \lambda^{3}\right) \leq \text { est } \leq O(\operatorname{lis}(A))
$$

with probability at least $1-1 / \operatorname{poly}(n)$.

We remark that one can turn Theorem 30.4.8 into an algorithm with running time $\widetilde{O}\left(n^{17 / 20}\right)$ by considering two cases separately. If $\lambda<n^{-1 / 20}$ we sample the array with a rate of $n^{-3 / 20}$ and compute the LIS for the sampled array. Otherwise, the running time of the algorithm is already bounded by $\widetilde{O}\left(n^{17 / 20}\right)$.

### 30.4.4 An $O\left(n^{\kappa}\right)$ Time Algorithm via Bootstrapping

Let us move a step backward and analyze the previous algorithm for obtaining an $O\left(\lambda^{3}\right)$ approximate solution. We first divide the input array into $\sqrt{n}$ subarrays of size $\sqrt{n}$ and after constructing the pseudo-solutions, we sample $O\left(\lambda^{4}\right)$ subarrays to estimate the size of the solution for pseudo-solutions. The reason we set the size of the subarrays to $\sqrt{n}$ is that there is a trade-off between the first and the last steps of the algorithm. More precisely, if we have more than $\sqrt{n}$ subarrays then the number of samples we draw in the beginning would exceed $O_{\lambda}(\sqrt{n})$. On the other hand, having fewer than $\sqrt{n}$ subarrays results in larger subarrays which would be costly in the last step.

If one favors the running time over the approximation factor, the following improvement can be applied to the algorithm: In the last step of the algorithm, instead of sampling the entire subarrays and computing lis for every pseudo-solution, we recursively call the same procedure to approximate the size of the solution for each subarray. This way, having large subarrays would no longer be an issue and therefore we can have fewer subarrays to improve the number of samples we draw in the first step of the algorithm.

More formally, in order to obtain a running time of $O\left(\operatorname{poly}(\lambda) n^{\kappa}\right)$ we set the size of each subarray to $n^{1-\kappa}$ and therefore after constructing the pseudo-solutions, the problem boils down to approximating the solution for poly $(\lambda)$ many subarrays of length $n^{1-\kappa}$. By running the same algorithm, we would have $n^{\kappa}$ subarrays of length $n^{1-2 \kappa}$ in the second iteration. After $1 / \kappa-1$ iterations, the subarrays are small enough and we can access all their elements in time $O\left(\operatorname{poly}(\lambda) n^{\kappa}\right)$. Of course, this imposes a factor of poly $(\lambda)$ to the approximation.

By generalizing the ideas from previous subsections, we show that if there is an algorithm for LIS with approximation factor $f(\lambda)$, then we can get a $\left(f\left(\frac{\lambda^{4}}{2^{32}}\right) \cdot \frac{\lambda^{4}}{2^{33}}\right)$-approximate LIS algorithm with better running time using the $f(\lambda)$-approximate algorithm as a subroutine.

Lemma 30.4.9. Assume we partition the sequence into $\zeta$ subarrays, where $\zeta$ is polynomially related to the length of the input sequence. For parameter $\lambda \in(0,1)$, let Oracle be af( $\lambda$ )approximate LIS algorithm (with respect to a domain interval) with running time $g(n, \lambda)$ and success probability $1-\exp \left(-\Omega\left(\log ^{2} n\right)\right)$ where $n$ is the length of the input sequence. Then Algorithm 30.7 using ORACLE as a subroutine is a $\left(f\left(\frac{\lambda^{4}}{2^{32}}\right) \cdot \frac{\lambda^{4}}{2^{33}}\right)$-approximate LIS algorithm


Figure 30.6: The flowchart of the $O_{\lambda}\left(n^{\epsilon}\right)$ time algorithm is shown.
with

$$
O\left(\lambda^{-4} g\left(\frac{n}{\zeta}, \frac{\lambda^{4}}{2^{32}}\right) \log ^{O(1)} n+\lambda^{-7} \zeta \log ^{O(1)} n\right)
$$

running time and success probability $1-\exp \left(-\Omega\left(\log ^{2} n\right)\right)$, where $\zeta$ is the number of subarrays.

Proof. We first prove the correctness of the algorithm. Let $A$ be a sequence of length $n$, and $\mathrm{sa}_{1}, \ldots, \mathrm{sa}_{\zeta}$ be the subarrays.

Consider the case of $\operatorname{lis}^{[\ell, r]}(A) \geq \lambda n$. By Corollary 30.4.2 and Lemma 30.4.4 with $\epsilon=\delta=1 / 10$, with probability $1-\exp (-\Omega(\zeta \lambda))$, there exists a pseudo-solution $\mathrm{ps}_{j}$ within interval $[\ell, r]$ satisfying

$$
\mathrm{q}\left(\mathrm{ps}_{j}\right) \geq \frac{\operatorname{lis}^{[\ell, r]}(A)}{2 t} \geq \frac{\operatorname{lis}^{[\ell, r]}(A) \lambda \epsilon}{2 k^{2}} \geq \frac{\operatorname{lis}^{[\ell, r]}(A) \lambda \epsilon}{2 \cdot 20^{2} \log ^{2}(1 / \delta) /\left(\lambda^{2} \epsilon^{4}\right)} \geq \frac{\epsilon^{5} \lambda^{4}}{800 \log ^{2}(1 / \delta)} n \geq \frac{\lambda^{4}}{2^{31}} n .
$$

Let $\alpha$ denote the number of subarrays $\operatorname{sa}_{i}$ such that $\operatorname{lis}{ }^{\left[\ell_{i}, r_{i}\right]}\left(\operatorname{sa}_{i}\right) \geq \lambda_{0} n / \zeta$ where $\left[\ell_{i}, r_{i}\right]$ is the
interval for subarray $\mathrm{sa}_{i}$ in $\mathrm{ps}_{j}$. We have

$$
\alpha \geq \frac{\mathrm{q}\left(\mathrm{ps}_{j}\right)-\lambda^{4} n / 2^{32}}{n / \zeta} \geq \frac{\lambda^{4} \zeta}{2^{32}} .
$$

By Chernoff bound, Step 14 to Step 22 of Algorithm 30.7 accepts on $\mathrm{ps}_{j}$ with probability at least $1-\exp \left(-\Omega\left(\log ^{2} n\right)\right)$.

Consider the case of

$$
\operatorname{lis}^{[\ell, r]}(A) \leq f\left(\frac{\lambda^{4}}{2^{32}}\right) \frac{\lambda^{4}}{2^{33}} n .
$$

Then for any pseudo-solution $\mathrm{ps}_{j}$, we have

$$
\mathrm{q}\left(\mathrm{ps}_{j}\right) \leq f\left(\frac{\lambda^{4}}{2^{32}}\right) \frac{\lambda^{4}}{2^{33}} n .
$$

Let $\beta$ denote the number of subarrays sa ${ }_{i}$ such that lis ${ }^{\left[\ell_{i}, r_{i}\right]}\left(\right.$ sa $\left._{i}\right) \geq f\left(\lambda^{4} / 2^{32}\right) n / \zeta$ where $\left[\ell_{i}, r_{i}\right]$ is the interval for subarray $\mathrm{sa}_{i}$ in $\mathrm{ps}_{j}$. We have

$$
\beta \leq \frac{\mathrm{q}\left(\mathrm{ps}_{j}\right)}{f\left(\lambda^{4} / 2^{32}\right) n / \zeta} \leq \frac{\lambda^{4}}{2^{33}} \zeta .
$$

By Chernoff bound, Step 14 to Step 22 of Algorithm 30.7 do not accept on $\mathrm{ps}_{j}$ with probability at least $1-\exp \left(-\Omega\left(\log ^{2} n\right)\right)$. By union bound, Algorithm 30.7 rejects with probability at least $1-\exp \left(-\Omega\left(\log ^{2} n\right)\right)$.

Hence, Algorithm 30.7 is a $\left(f\left(\frac{\lambda^{4}}{2^{32}}\right) \cdot \frac{\lambda^{4}}{2^{33}}\right)$-approximate algorithm for LIS of length $n$ with success probability at least $1-\exp \left(-\Omega\left(\log ^{2} n\right)\right)$.

Finally, we discuss the running time of Algorithm 30.7. By the definition of procedures ConstructCandidateDomains and ConstructPseudoSolutions, the running time of Step 5 to Step 9 of Algorithm 30.7 is $O\left(\lambda^{-9} \zeta \log ^{O(1)} n\right)$. By Lemma 30.4.3, $t=O\left(\lambda^{-3}\right)$, and by Chernoff bound with probability $1-\exp \left(-\Omega\left(\log ^{2} n\right)\right)$ the size of $Q$ is
at most $O\left(\lambda^{-4} \log ^{4} n\right)$. Hence, the running time of Step 12 to Step 24 of Algorithm 30.7 is $O\left(\lambda^{-7} g\left(n / \zeta, \lambda^{4} / 2^{32}\right) \log ^{4} n\right)$.

By definition, we have the following basic fact about approximate ratio.

Fact 30.4.10. Let $f$ and $f^{\prime}$ be two functions mapping $(0,1)$ to $(0,1)$ such that $f(\lambda) \geq f^{\prime}(\lambda)$ for any $\lambda \in(0,1)$. If there is a $f(\lambda)$-approximate LIS algorithm, then the algorithm is also $f^{\prime}(\lambda)$-approximate.

Now we present algorithm to approximate LIS using $\widetilde{O}\left(n^{\kappa} \operatorname{poly}\left(\lambda^{-1}\right)\right)$ space by applying the pseudo-solution construction-evaluation framework recursively. In particular, we use the same algorithm on subarrays as an oracle and apply Lemma 30.4.9 recursively to approximate the entire sequence with slightly worse approximation ratio (compared with approximation ratio of the oracle).

Lemma 30.4.11. Let $\kappa$ be a constant of $(0,1)$ and $\lambda \in(0,1)$. Algorithm 30.8 approximates LIS with approximation ratio

$$
\frac{\lambda^{2 \cdot 4^{([1 / \kappa]-1)}}}{256^{3 \cdot 4([1 / k]-1)}}
$$

and running time $O\left(n^{\kappa} \cdot \lambda^{-4^{O(1 / \kappa)}} \log ^{O(1)} n\right)$ and success probability $1-\exp \left(-\Omega\left(\log ^{3} n\right)\right)$.

Proof. We first prove the correctness of the algorithm by induction. Without loss of generality, we assume $1 / \kappa$ is an integer.

For $i \in\{2,3, \ldots, 1 / \kappa\}$, denote

$$
h_{i}(\lambda)=\frac{\lambda^{2 \cdot 4^{(i-1)}-4}}{256^{2 \cdot 4^{(i-1)}+3 \cdot 4^{(i-2)}-7}} .
$$

We show that if the length of the input sequence is $n^{i \cdot \kappa}$ then Algorithm 30.8 is $h_{i}(\lambda)$ approximate.

If the length of input sequence is $n^{2 \kappa}$, then $h_{2}(\lambda)=\frac{\lambda^{4}}{2^{32}}$. By Corollary 30.4.2, Lemma 30.4.4, and Lemma 30.4.6, Algorithm 30.8 is $h_{2}(\lambda)$-approximate.

In the induction step, for an integer $2 \leq i<1 / \kappa$, we assume Algorithm 30.8 is $h_{i}(\lambda)$-approximate for input instance of length $n^{i \cdot \kappa}$. By Lemma 30.4.9, Algorithm 30.8 is $\left(h_{i}\left(\frac{\lambda^{4}}{2^{32}}\right) \frac{\lambda^{4}}{2^{33}}\right)$-approximate for input instance of length $n^{(i+1) \cdot \kappa}$. Since

$$
\begin{aligned}
h_{i}\left(\frac{\lambda^{4}}{2^{32}}\right) \frac{\lambda^{4}}{2^{33}} & =\frac{\lambda^{4 \cdot\left(2 \cdot 4^{(i-1)}-4\right)} \cdot \lambda^{4}}{256^{4 \cdot\left(2 \cdot 4^{(i-1)}-4\right)} \cdot 256^{2 \cdot 4^{(i-1)}+3 \cdot 4^{(i-2)}-7} \cdot 2^{33}} \\
& =\frac{\lambda^{2 \cdot 4^{i}-12}}{256^{2 \cdot 4^{i}+2 \cdot 4^{(i-1)}+3 \cdot 4^{(i-2)}-18.875}} \\
& >\frac{\lambda^{2 \cdot 4^{i}-4}}{256^{2 \cdot 4^{i}+3 \cdot 4^{(i-1)}-7}} \\
& =h_{i+1}(\lambda),
\end{aligned}
$$

by Fact 30.4.10, Algorithm 30.8 is $h_{i+1}(\lambda)$-approximate for input instance of length $n^{(i+1) \cdot \kappa}$. Since

$$
h_{1 / \kappa}(\lambda)>\frac{\lambda^{2 \cdot 4^{(1 / \kappa)-1)}}}{256^{3 \cdot 4^{(1 / \kappa)-1)}}},
$$

by Fact 30.4.10, Algorithm 30.8 is $\left(\frac{\lambda^{2 \cdot 4}\left(\frac{4(1 / \kappa)-1)}{\left.256^{3 \cdot 4}(1 / \kappa)-1\right)}\right) \text {-approximate for input instance of length }}{\text { a }}\right.$ $n$.

By Corollary 30.4.2, Lemma 30.4.4, Lemma 30.4.6 and Lemma 30.4.9, we have the desired running time. The success probability is obtained by same corollaries/lemmas and union bound.

Finally, by starting with $\lambda=1$ and iteratively multiplying $\lambda$ by a $1 /(1+\epsilon)$ factor until a solution is found, we can approximate $\operatorname{lis}(A)$ within an approximation factor of $\lambda^{O\left(4^{1 / \kappa}\right)}$. Theorem 30.4.12. Let $\kappa$ be a constant of $(0,1)$ and $\lambda \in(0,1)$. There exists a $\widetilde{O}\left(n^{\kappa}\right.$. $\left.\lambda^{-4^{O(1 / \kappa)}}\right)$ time algorithm for lis with approximation factor $\lambda^{4^{O(1 / k)}}$ and success probability $1-\exp \left(-\Omega\left(\log ^{3} n\right)\right)$.

### 30.5 Probability and Graph Tools

In this section, we restate probability and graph tools that we use throughout this paper. All these theorems are proven in previous work.

Theorem 30.5.1 (Chernoff Bounds). Let $X=\sum_{i=1}^{n} X_{i}$, where $X_{i}=1$ with probability $p_{i}$ and $X_{i}=0$ with probability $1-p_{i}$, and all $X_{i}$ are independent. Let $\mu=\mathbb{E}[X]=\sum_{i=1}^{n} p_{i}$. Then

1. $\operatorname{Pr}[X \geq(1+\delta) \mu] \leq \exp \left(-\delta^{2} \mu / 3\right), \forall \delta>0$;
2. $\operatorname{Pr}[X \leq(1-\delta) \mu] \leq \exp \left(-\delta^{2} \mu / 2\right), \forall 0<\delta<1$.

Theorem 30.5.2 (Hoeffding bound). Let $X_{1}, \cdots, X_{n}$ denote $n$ independent bounded variables in $\left[a_{i}, b_{i}\right]$. Let $X=\sum_{i=1}^{n} X_{i}$, then we have

$$
\operatorname{Pr}[|X-\mathbb{E}[X]| \geq t] \leq 2 \exp \left(-\frac{2 t^{2}}{\sum_{i=1}^{n}\left(b_{i}-a_{i}\right)^{2}}\right)
$$

Theorem 30.5.3 (Blakley-Roy inequality, [BR65], see also Proposition 3.1 in [KSV13]). Let $G$ denote a graph that has $n$ vertices and average degree $d$. The number of walks of length $k$ in graph $G$ is at least $n d^{k}$.

Theorem 30.5.4 (Turán theorem for bipartite graphs, [KST54], see also [BBK13]). For a graph $G$ the Turán number $\operatorname{ex}(G, n)$ is the maximum number of edges that a graph on $n$ vertices can have without containing a copy of $G$. For any $s \leq t$, $\operatorname{ex}\left(K_{s, t}, n\right) \leq \frac{1}{2}(t-$ $1)^{1 / s} n^{2-1 / s}+o\left(n^{2-1 / s}\right)$

Theorem 30.5.5 (Turán theorem for cliques [Tur41]). Let $G$ be any graph with $n$ vertices, such that $G$ is $K_{r+1}$-free. Then the number of edges in $G$ is at most

$$
\left(1-\frac{1}{r}\right) \cdot \frac{n^{2}}{2} .
$$

Corollary 30.5.6 (of Theorem 30.5.5). Let $G$ be any graph with $n$ vertices, such that $G$ has no independent set of size $r+1$. Then the number of edges in $G$ is at least

$$
\binom{n}{2}-\left(1-\frac{1}{r}\right) \cdot \frac{n^{2}}{2}=\frac{n r+n^{2}}{2 r} .
$$

### 30.6 LCS Step 0: Window-compatible Solutions

The goal of this section is to show window-compatible solutions provide very accurate estimates for LCS. Roughly speaking, in this section, we construct a set of windows for $A$ and a set of windows for $B$ and we show that it suffices to only compute the LCS of the windows. If that information is available, then we can estimate the LCS of two strings very accurately.

We use the same definition of window-compatible transformation as $\left[\mathrm{BEG}^{+} 18\right]$, but we will use a richer set of windows. Every window is essentially a substring of each string. Notice that windows may not have the same length. We use $W_{A}$ to denote the set of windows constructed for $A$ and $W_{B}$ for the set of the windows that we construct for $B$.

Let us first state our definition of window-compatible solutions for LCS.

Definition 30.6.1 (Window-compatible LCS). Let $S=\left\langle w_{1}, w_{2}, \cdots, w_{\alpha}\right\rangle$ and $S^{\prime}=\left\langle w_{1}^{\prime}, w_{2}^{\prime}, \cdots w_{\alpha}^{\prime}\right\rangle$ be two sequences of non-overlapping windows from $W_{1}$ and $W_{2}$, respectively.

- We call a common substring of $(A, B)$ window-compatible with respect to $S$ and $S^{\prime}$, if it is a union of $k$ common substrings of $\left(w_{1}, w_{1}^{\prime}\right), \ldots,\left(w_{\alpha}, w_{\alpha}^{\prime}\right)$

We call a common substring window-compatible, if it is window-compatible with respect to at least a pair of sequence of non-overlapping windows from $W_{1}$ and $W_{2}$, respectively.

### 30.6.1 Construction of the Windows

This section presents a definition (see Definition 30.6.2) of a multiple layers window set for strings $A$ and $B$. The construction is the same for both strings, therefore, we only
state it for $A$. As mentioned earlier, each window is a substring of the strings. For clarity, we use $A^{(j, l)}$ to denote a length- $l$ substring of $A$ which starts at index $j$ and ends at index $j+l-1$. Our construction has multiple layers, each of which contains windows with equal lengths.

Given two length $n$ strings $A, B$, and two parameters $1 \leq d \leq n$ and $\epsilon_{0} \in(0,1)$, we define window set $W_{A}$ to be a set of continuous subsequence of string $A$. The sizes of the windows depend on parameter $d$ and the accuracy of our construction depends on parameter $\epsilon$. Setting $d=n^{x}$ leads to a truly subquadratic time solution for any $0<x<1$, however, one can play with the value of $d$ to optimize the running time.

Our construction has multiple layers, where in each layer the lengths of the windows are the same. Let $f=\Theta\left(\frac{1}{\epsilon_{0}} \log \frac{1}{\epsilon_{0}}\right)$ and $f+1$ denote the number of different layers of windows in $W_{A}$.

For each $i \in\{0,1, \cdots, f\}$, we define $d_{i}=d\left(1+\epsilon_{0}\right)^{i}, g_{i}=\epsilon_{0} d_{i}$, and $t_{i}=n / d_{i}$. For each $i$, let $W_{A}^{i}$ denote the set of all the windows in this layer. We consider window set $W_{A}^{i}$, in $i$-th layer, all the window have the same size which is $d_{i} ; g_{i}$ is the shift size, it means for each window (except the leftmost and rightmost one), if we shift either to the left or to the right, it becomes another window in $W_{A}^{i}$.

Formally speaking, we define $W_{A}^{i}$ as follows

Definition 30.6.2. Let $W_{A}$ be the set of windows we construct for string $A$. Then we have $W_{A}^{i}=\left\{A^{(\text {left,len })} \mid\right.$ left $=x \cdot g_{i}+y \cdot d_{i}$, len $\left.=d_{i}, \forall x \in\left\{0,1, \cdots, d_{i} / g_{i}-1\right\}, \forall y \in\left\{0,1, \cdots, t_{i}-1\right\}\right\}$
where we use $A^{(\text {left,len })}$ denotes a continuous subsequence of string $A$ that starts at left and
has length len. Finally $W_{A}$ is $\cup_{i \in\{0,1, \cdots, f\}} W_{A}^{i}$. Similarly, for each $i \in\{0,1, \cdots, f\}$, we can obtain $W_{B}^{i}$ and $W_{B}$.

Indeed, the running time of constructing the windows is truly subquadratic.
Lemma 30.6.1 (Generating ( $d, \epsilon_{0}$ )-multiple layers window set). Given two strings $A, B$ of $n$ symbols over $\Sigma$ symbols, and parameters $d \geq 1, \epsilon_{0} \in(0,1)$. Let set $W_{A}$ and $W_{B}$ be generated by procedure GeneratemultipleLevels (Algorithm 30.9). The algorithm takes $n \cdot \widetilde{O}\left(1 / \epsilon_{0}^{2}\right)$ time and $\left|W_{A}\right|=\left|W_{B}\right|=O\left(n /\left(d \epsilon_{0}^{2}\right)\right)$.

Proof. Given two sets of blocks $W_{A}$ and $W_{B}$, for each block $A^{\left(j_{1}, l_{1}\right)} \in W_{A}, A^{\left(j_{1}, l_{1}\right)}$ starts at $j_{1}$ and has length $l_{1}$. Similarly, for each block $B^{\left(j_{2}, l_{2}\right)} \in W_{B}, B^{\left(j_{2}, l_{2}\right)}$ starts at $j_{2}$ and has length $l_{2}$.

We can compute

$$
\begin{equation*}
\sum_{i=0}^{f} t_{i}=n \sum_{i=0}^{f} 1 / d_{i}=(n / d) \sum_{i=0}^{f}\left(1+\epsilon_{0}\right)^{-i} \leq 2 n /\left(d \epsilon_{0}\right) \tag{30.12}
\end{equation*}
$$

Then we calculate the size of window set $W_{A}$,

$$
\begin{aligned}
\left|W_{A}\right| & =\sum_{i=0}^{f}(\# \text { shifts }) \cdot(\# \text { windows per shift }) \\
& =\sum_{i=0}^{f}\left(d_{i} / g_{i}\right) \cdot t_{i} \\
& =\sum_{i=0}^{f} \frac{1}{\epsilon_{0}} \cdot t_{i} \\
& =\frac{1}{\epsilon_{0}} \sum_{i=0}^{f} t_{i} \\
& =2 n /\left(d \epsilon_{0}^{2}\right)
\end{aligned}
$$

where the second step follows from definition of window set in each layer $i$, the third step follows from $d_{i} / g_{i}=1 / \epsilon_{0}$, and last step follows from Eq. (30.12).

The running time is

$$
\sum_{i=0}^{f}\left(d_{i} / g_{i}\right) \cdot t_{i} \cdot d_{i}=n \sum_{i=0}^{f}\left(d_{i} / g_{i}\right)=n \sum_{i=0}^{f} 1 / \epsilon_{0}=\widetilde{O}\left(n / \epsilon_{0}^{2}\right) .
$$

Remark 30.6.1. In order to make sure the loss in the approximation is negligible, we use $\epsilon_{0}=\epsilon \lambda$ and therefore we have

- The total number of windows is equal to $k=\left|W_{A}\right|+\left|W_{B}\right|=O\left((1 / \lambda)^{2} n / d\right)$.
- The maximum size of the windows is equal to $\mathrm{w}_{\max }=O(d / \lambda)$.
- The minimum size of the windows is equal to $\mathrm{w}_{\text {min }}=d$.
- The ratio of the maximum window size over the minimum window size is bounded by $\mathrm{w}_{\text {gap }}=\mathrm{w}_{\text {max }} / \mathrm{w}_{\text {min }} \leq O(1 / \lambda)$.
- The number of different layers which is equal to the number of different window sizes is equal to $\mathrm{w}_{\text {layers }}=O((1 / \lambda) \log (1 / \lambda))$.


### 30.6.2 Optimality of the Windows

We present a structural lemma for the windows in the following. This essentially, shows that the problem of computing LCS between the two strings reduces to the problem of computing the LCS's between the windows.

Lemma 30.6.2 (Improved Structural Lemma). For any two strings $A, B$ of length $n$ :

- If $\|\operatorname{LCS}(A, B)\| \geq \lambda$, there exists a window-compatible common substring of $A, B$ of length at least $\lambda n-8 \epsilon_{0} n$ via $\left(d, \epsilon_{0}\right)$-Multiple layers window (Definition 30.6.2).

Proof. We fix $d$ and $\epsilon_{0}$.
We pick the first block $\widetilde{A}_{1}$ to be length $d$, and consider the optimal matching. Suppose the optimal matching block for $\widetilde{A}_{1}$ is $\widetilde{B}_{1}$, there are three possibilities.

First, $\widetilde{B}_{1}$ has length between $\left[\epsilon_{0} d, d / \epsilon_{0}\right]$. That's good and we keep this pair.
Second, $\widetilde{B}_{1}$ has length at least $d / \epsilon_{0}$, then we throw out this pair. The characters we throw out from $2 n$ is at least $d / \epsilon_{0}$, and we decrease the entire LCS by at most $d$.

Third, the length of $\widetilde{B}_{1}$ has length at most $d \epsilon_{0}$, then we also throw out this pair. The characters we throw out from $2 n$ is at least $d$, and we decrease the entire LCS by at most $\epsilon_{0} d$.

The above procedure can be recursively applied until we found all the matching blocks. Let's say there are $k$ matching pairs, $\left(\widetilde{A}_{1}, \widetilde{B}_{1}\right),\left(\widetilde{A}_{2}, \widetilde{B}_{2}\right), \cdots\left(\widetilde{A}_{\alpha}, \widetilde{B}_{\alpha}\right)$. Note that $\widetilde{A}_{1}, \widetilde{A}_{2}, \cdots, \widetilde{A}_{\alpha}$ is a partition of $A$. Similarly, $\widetilde{B}_{1}, \widetilde{B}_{2}, \cdots, \widetilde{B}_{\alpha}$ is a partition of $B$. Let $z_{i}$ denote the variable such that $z_{i}=0$ indicates we throw out the pair, and $z_{i}=1$ otherwise. Since whenever we throw out a pair, the amount of LCS got decreased is always at most $\epsilon_{0}$ fraction of the characters we throw out. Since there are $2 n$ characters in total. Then we LCS we decrease is at most $2 \epsilon_{0} n$, i.e.,

$$
\begin{equation*}
\sum_{i \in Z} \operatorname{LCS}\left(\widetilde{A}_{i}, \widetilde{B}_{i}\right) \geq \operatorname{LCS}(A, B)-2 \epsilon_{0} n \tag{30.13}
\end{equation*}
$$

where $Z=\left\{i \mid z_{i}=1, i \in[k]\right\}$.
We consider a fixed $i \in Z$. For each pair $\widetilde{A}_{i}, \widetilde{B}_{i}$, let $p$ and $q$ denote the integer that

$$
\left|\widetilde{A}_{i}\right| \in\left(\left(1+\epsilon_{0}\right)^{p},\left(1+\epsilon_{0}\right)^{p+1}\right], \text { and }\left|\widetilde{B}_{i}\right| \in\left(\left(1+\epsilon_{0}\right)^{q},\left(1+\epsilon_{0}\right)^{q+1}\right] .
$$

Let $\widehat{A}_{i}$ and $\widehat{B}_{i}$ denote the two pairs from $W_{A}$ and $W_{B}$ that has the largest LCS and satisfies the following

$$
\left|\widehat{A}_{i}\right|=\left(1+\epsilon_{0}\right)^{p-1}, \widehat{A}_{i} \subseteq \widetilde{A}_{i},\left|\widehat{B}_{i}\right|=\left(1+\epsilon_{0}\right)^{q-1}, \widehat{B}_{i} \subseteq \widetilde{B}_{i} .
$$

Since $\left|\widetilde{A}_{i}\right|-\left|\widehat{A}_{i}\right| \geq\left(1+\epsilon_{0}\right)^{p}-\left(1+\epsilon_{0}\right)^{p-1}=\epsilon_{0}\left(1+\epsilon_{0}\right)^{p-1}$. Recall the construction of $W_{A}$, we places a length $\left(1+\epsilon_{0}\right)^{p-1}$ block for every $\epsilon_{0}\left(1+\epsilon_{0}\right)^{p-1}$ shifts. Thus, there must exist such block $\widehat{A}_{i}$ from $W_{A}$ and $\widehat{A}_{i} \subseteq \widetilde{A}_{i}$. Similarly, the same argument holds for $\widehat{B}_{i}$.

According to the definition of $W_{A}$ and $W_{B}$, we know that

$$
\begin{aligned}
\left|\widetilde{A}_{i} \backslash \widehat{A}_{i}\right| & \leq\left(1+\epsilon_{0}\right)^{p+1}-\left(1+\epsilon_{0}\right)^{p-1} \\
& \leq\left(1+\epsilon_{0}\right)^{p-1} 3 \epsilon_{0} \\
& \leq 3 \epsilon_{0}\left|\widetilde{A}_{i}\right|
\end{aligned}
$$

Similarly, $\left|\widetilde{B}_{i} \backslash \widehat{B}_{i}\right| \leq 3 \epsilon_{0}\left|\widetilde{B}_{i}\right|$. The characters we ignore is at most $3 \epsilon_{0} \cdot \max \left(\left|\widetilde{A}_{i}\right|,\left|\widetilde{B}_{i}\right|\right)$. Let $\tau_{i}=3 \epsilon_{0} \cdot \max \left(\left|\widetilde{A}_{i}\right|,\left|\widetilde{B}_{i}\right|\right)$, then $\tau_{i} \leq 3 \epsilon_{0}\left(\left|\widetilde{A}_{i}\right|+\left|\widetilde{B}_{i}\right|\right)$.

If we consider $\sum_{i \in Z}\left|\operatorname{LCS}\left(\widehat{A}_{i}, \widehat{B}_{i}\right)\right|$,

$$
\begin{aligned}
\sum_{i \in Z}\left|\operatorname{LCS}\left(\widehat{A}_{i}, \widehat{B}_{i}\right)\right| & \geq \sum_{i \in Z}\left(\left|\operatorname{LCS}\left(\widetilde{A}_{i}, \widetilde{B}_{i}\right)\right|-\tau_{i}\right) \\
& \geq \sum_{i \in Z}\left|\operatorname{LCS}\left(\widetilde{A}_{i}, \widetilde{B}_{i}\right)\right|-6 \epsilon_{0} n \\
& \geq|\operatorname{LCS}(A, B)|-8 \epsilon_{0} n
\end{aligned}
$$

where the second step follows from $\sum_{i \in Z} \tau_{i} \leq 6 \epsilon_{0} n$, the third step follows from Eq. (30.13).

### 30.6.3 Dynamic Programming for Block-based LCS

We showed in Section 30.6.2 that there always exists a window-compatible solution whose size is very close to that of the optimal solution. In this section, we show how one can find such a solution via dynamic programming. Therefore, we assume that a matrix $M$ is provided as input that stores the LCS of every pair of windows.

Lemma 30.6.3 (Dynamic Programming for Longest Common Sequence). Given two strings A, $B$ of length $n$. Let $\epsilon \in(0,1), d \geq 1$. Let $W_{A}, W_{B}$ be generated by procedure CreateMulTiPLELEVELS $(A, B, n, d, \Theta(\epsilon \lambda))$ (Algorithm 30.9). Let table $M$ be generated by procedure NonMetricEstimation $\left(W_{A}, W_{B}\right)$.

Then there is an algorithm (Algorithm 30.10) that runs in $(n / d)^{2} \cdot \operatorname{poly}(\lambda, 1 / \epsilon)$ time to output a monotone matching between two sets $W_{A}$ and $W_{B}$ which gives a string $C$ such that

$$
|C| \geq(1-o(1)) \cdot(1-\epsilon) \cdot|\operatorname{LCS}(A, B)| .
$$

Proof. We choose $\epsilon_{0}=\Theta(\epsilon / \lambda)$.
Using Lemma 30.6.2, we know there exits a batch of window-pairs $\left(\widehat{A}_{i_{1}}, \widehat{B}_{j_{1}}\right),\left(\widehat{A}_{i_{2}}, \widehat{B}_{j_{2}}\right), \ldots$ from $W_{A}, W_{B}$ such that

$$
\begin{equation*}
\sum_{l}\left|\operatorname{LCS}\left(\widehat{A}_{i_{l}}, \widehat{B}_{i_{l}}\right)\right| \geq|\operatorname{LCS}(A, B)|-8 \epsilon_{0} n \tag{30.14}
\end{equation*}
$$

By assumption on $M$, we know that $M$ gives a good estimation for every pair in $W_{A}, W_{B}$ :

$$
\delta_{i_{l}, j_{l}} d_{i_{l}, j_{l}} \geq M_{i_{l}, j_{l}} \geq(1-o(1)) \delta_{i_{l}, j_{l}}^{3} \cdot d_{i_{l}, j_{l}} .
$$

where $d_{i_{l}, j_{l}}=\max \left\{\left|\widehat{A}_{i_{l}}\right|,\left|\widehat{B}_{j_{l}}\right|\right\}, \delta_{i_{l}, j_{l}}=\left|\operatorname{LCS}\left(\widehat{A}_{i_{l}}, \widehat{B}_{j_{l}}\right)\right| / d_{i_{l}, j_{l}}$. It implies that

$$
\left|\operatorname{LCS}\left(\widehat{A}_{i_{l}}, \widehat{B}_{j_{l}}\right)\right| \geq M_{i_{l}, j_{l}} \geq(1-o(1)) \delta_{i_{l}, j_{l}}^{2}\left|\operatorname{LCS}\left(\widehat{A}_{i_{l}}, \widehat{B}_{j_{l}}\right)\right|
$$

Now, if we consider this cost $\sum_{l} M_{i_{l}, j_{l}}$,

$$
\begin{aligned}
\sum_{l} M_{i_{l}, j_{l}} & \geq(1-o(1)) \cdot \sum_{l} \delta_{i_{l}, j_{l}}^{2}\left|\operatorname{LCS}\left(\widehat{A}_{i_{l}}, \widehat{B}_{j_{l}}\right)\right| \\
& \geq(1-o(1)) \cdot \lambda^{-2} \sum_{l}\left|\operatorname{LCS}\left(\widehat{A}_{i_{l}}, \widehat{B}_{j_{l}}\right)\right| \\
& \geq(1-o(1)) \cdot\left(|\operatorname{LCS}(A, B)|-8 \epsilon_{0} n\right) \\
& =(1-o(1)) \cdot(|\operatorname{LCS}(A, B)|-\epsilon n / \lambda) \\
& \geq(1-o(1))(1-\epsilon)|\operatorname{LCS}(A, B)|
\end{aligned}
$$

where the second step follows from $\delta_{i_{l}, j_{l}} \geq \lambda^{-1}$, the third step follows from Eq. (30.14), the fourth step follows from $\epsilon_{0}=\Theta(\epsilon / \lambda)$, the last step follows from $\operatorname{LCS}(A, B) \geq \lambda n$.

Next, we prove the correctness of running time. Using Lemma 30.6.1, we know that $\left|W_{A}\right|=\left|W_{B}\right|=O\left(n /\left(d \epsilon_{0}^{2}\right)\right)$. Then $\left|S_{1}\right|=\left|S_{2}\right|=O\left(n /\left(d \epsilon_{0}^{2}\right)\right)$. The running time of sorting takes $O\left(\left|S_{1}\right| \log \left|S_{1}\right|\right)$.

The dominated part of the running time is three for-loops. For the third one, if we consider a fixed $i_{1}$ and $i_{2}$, there are $O\left(\left(\frac{1}{\epsilon_{0}} \log \left(1 / \epsilon_{0}\right)\right)^{2}\right)$ pairs could have right index at $i_{1}, i_{2}$. Thus the running time of dynamic programming takes

$$
\left(n /\left(d \epsilon_{0}^{2}\right)\right) \cdot\left(n /\left(d \epsilon_{0}^{2}\right)\right) \cdot O\left(\left(\left(1 / \epsilon_{0}\right) \log \left(1 / \epsilon_{0}\right)\right)^{2}\right)=(n / d)^{2} \cdot \widetilde{O}\left(1 / \epsilon_{0}^{6}\right)
$$

For the correctness of dynamic programming, it follows from three cases cover all the possibilities. For each $i_{1}, i_{2}$, we define $C\left[i_{1}\right]\left[i_{2}\right]$ to be optimal solution that only uses blocks from $W_{A}$ with ending index no later than $i_{1}$ and blocks from $W_{B}$ with ending index no later than $i_{2}$. Suppose we are at $i_{1}, i_{2}$. There are three cases, the first case is, the last pair of the optimal solution for $C\left[i_{1}\right]\left[i_{2}\right]$ is ending at $i_{1}, i_{2}$. The second case is, the last pair of the optimal solution for $C\left[i_{1}\right]\left[i_{2}\right]$ is ending at $\operatorname{prev}\left(i_{1}\right), i_{2}$. The third case is, the last pair of the optimal solution for $C\left[i_{1}\right]\left[i_{2}\right]$ is ending at $i_{1}, \operatorname{prev}\left(i_{2}\right)$.

Remark 30.6.2. If in the above algorithm, $M$ instead of the correct LCS's has a $(1-\epsilon) \lambda^{2}$ or $O\left(\lambda^{3}\right)$ approximation of the solutions, then the overall approximation factor of the algorithm would be $(1-\epsilon) \lambda^{2}$ or $O\left(\lambda^{3}\right)$, respectively.

### 30.7 LCS Step 2: Discovering the Sparse Graph

The arguments of this section follow from the work of $\left[\mathrm{CDG}^{+} 18\right]$. However, for the sake of completeness, we restate the overall ideas here. This section is dedicated to presenting a method NonMetricEstimation $\left(W_{A}, W_{B}\right)$ to be used for estimating the longest common subsequence. The output of $\operatorname{NonMEtricEstimation}\left(W_{A}, W_{B}\right)$ is a matrix $M:\left[\left|W_{A}\right|\right] \times$ $\left[\left|W_{B}\right|\right] \rightarrow \mathbb{Z}^{+}$where $M[i][j]$ estimates the LCS of windows $w_{i} \in W_{A}$ and $w_{j} \in W_{B}$. We allow for errors in a few elements of $M$ but we prove at the end of this section that with high probability, the error does not affect the solution significantly.

Similar to Section 30.3, we define $W=W_{A} \cup W_{B}$ and denote the windows inside $W$ by $w_{1}, w_{2}, \ldots, w_{k}$. For a threshold $\lambda$ the edges of graph $\mathrm{G}_{\lambda}$ reflect pairs of windows $\left(w_{i}, w_{j}\right)$ such that $\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \lambda$. In order to estimate $M$, we need to discover the edges of $\mathrm{G}_{\lambda}$ for $\lambda \in\left\{\epsilon \lambda_{0}, \epsilon \lambda_{0}(1+\epsilon), \epsilon \lambda_{0}(1+\epsilon)^{2}, \ldots, 1\right\}$ and build $M$ based on that. In other words, for each $w_{i}$ and $w_{j}$, we find the largest $\lambda$ such that $\mathrm{G}_{\lambda}$ contains edge $(i, j)$ and set $M[i][j]=\lambda$.

Let us for simplicity fix the approximation factor to be $(1-\epsilon) \lambda^{2}$. As mentioned earlier, we represent the edges of $\mathrm{G}_{\lambda}$ by a matrix $\mathrm{O}:[k] \times[k] \rightarrow\{0,1\}$ where $\mathrm{O}[i][j]=1$ implies that edge $(i, j)$ exists in $\mathrm{G}_{\lambda}$. The $(1-\epsilon) \lambda^{2}$ sparsification gives us an output $\widehat{\mathrm{O}}_{(1-\epsilon) \lambda^{2}}$ where except for a few pairs the rest of the edges of $G_{\lambda}$ are detected within an approximation factor $(1-\epsilon) \lambda^{2}$. If all the edges were detected with this algorithm, this would provide us with a matrix $M$ such that $(1-\epsilon)\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\|^{3} \leq M[i][j] \leq\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\|$ holds if $\left\|\operatorname{lcs}\left(w_{i}, w_{j}\right)\right\| \geq \epsilon \lambda$. If such a matrix is given to Algorithm 30.11, it would result in a solution with approximation factor $(1-\epsilon) \lambda^{2}$.

However, not all the edges of $G_{\lambda}$ are detected in Step 1. Therefore, if we use the
discovered edges to construct a matrix $\widehat{M}$ for the LCS of the windows, $\widehat{M}$ may not meet the guarantee of $M$ for at most $\widetilde{O}\left(k^{2-\Omega(1)}\right)$ pairs. In what follows, we present an algorithm that resolves this issue.

To simplify the notation, define a graph NG in which every window is a vertex and there is an edge between pair $(i, j)$ if the LCS of windows $w_{i}$ and $w_{j}$ is not estimated correctly in the sparsification step. Our algorithm is not aware of the edges of NG but since the number of remaining edges is truly subquadratic, we have $|E(\mathrm{NG})|=\widetilde{O}\left(k^{2-\Omega(1)}\right)$.

Let us assume that NG contains at most $\widetilde{O}\left(k^{2-\eta}\right)$ edges. This implies that the average degrees of the vertices of NG is upper bounded by $\widetilde{O}\left(k^{1-\eta}\right)$. Although we are not given the edges of NG explicitly, we can verify for each pair $(i, j)$ whether there is an edge between them in NG by computing the LCS of $w_{i}$ and $w_{j}$ and comparing it to the value that has been obtained in $\widehat{M}$.

Let $\mathrm{w}_{\max }=\max _{i}\left|w_{i}\right|$ be the maximum length of the windows. Define two windows $\left(w_{i}, w_{j}\right)$ of $W_{A}$ or two windows $\left(w_{i}, w_{j}\right)$ of $W_{B}$ to be nearby if their distance is bounded by at most $\mathbf{w}_{\max } k^{\eta / 2}$. That is $w_{i}$ covers an index $a$ and $w_{j}$ covers an index $b$ such that $|a-b| \leq \mathrm{w}_{\max } k^{\eta / 2}$. We subsample a set $S$ of $W_{A}$ with a rate of $\epsilon^{-1} k^{-\eta / 2}$. That is, each window of $W_{A}$ appears in $S$ independently, with probability $\epsilon^{-1} k^{-\eta / 2}$. Next, for each window $w_{i} \in S$, we discover all of its incident edges in NG by computing the LCS of $w_{i}$ and every window $w_{j} \in W_{B}$. Every time we discover one edge of NG we add this edge to $\mathrm{NG}^{\prime}$. Notice that $\mathrm{NG}^{\prime}$ contains a subset of edges of NG however, the edges of $\mathrm{NG}^{\prime}$ are known to the algorithm and are available. Finally, we update $M$ by computing the LCS of pairs of windows ( $w_{i}, w_{j}$ ) such that there exist windows $w_{i^{\prime}} \in W_{A}$ and $W_{j^{\prime}} \in W_{B}$ such that 1) $w_{i}$ and $w_{i^{\prime}}$ are nearby, 2) $w_{j}$ and $w_{j^{\prime}}$ are nearby, and 3) vertices $i$ and $j$ are neighbors in $\mathrm{NG}^{\prime}$. This is shown in Algorithm
30.12 .

It follows from the analysis of $\left[\mathrm{CDG}^{+} 18\right]$ that if we use $M$ as an estimation of the windows, then Algorithm 30.11 finds the desired solution even though $M$ does not correctly estimate all the values.

Lemma 30.7.1 ([CDG+18]). Let $\widehat{M}$ be the output of the sparsification algorithm and $M$ be constructed by Algorithm 30.12 from $\widehat{M}$. If NG has at most $\widetilde{O}\left(k^{2-\eta}\right)$ edges, then the runtime of Algorithm 30.12 is bounded by $\widetilde{O}\left(k^{2-\eta / 2} \mathrm{w}_{\max }^{2} / \lambda^{4}\right)$. Also using $M$ as an estimate for the LCS of the windows hurts the solution by a factor of at most $1+\epsilon$.

Proof. The proof follow from the arguments of $\left[\mathrm{CDG}^{+} 18\right]$. Here we just mention the intuition and the overall ideas.

Runtime: Since every window of $W_{A}$ is sampled with probability $k^{-\eta / 2}$ the size of $S$ is bounded by $O\left(k^{1-\eta / 2}\right)$ with high probability. In order to discover the incident edges of set $S$ in NG, we compute the LCS of every element in $S$ and all other windows. Therefore, the total running time at this point is $\widetilde{O}\left(k^{2-\eta / 2} w_{\text {max }}^{2}\right)$.

NG contains at most $k^{2-\eta}$ edges and therefore the expected number of edges that we put in $\mathrm{NG}^{\prime}$ is $\widetilde{O}\left(k^{2-3 \eta / 2}\right)$. For every edge that we discover in $\mathrm{NG}^{\prime}$ we compute the LCS of $\widetilde{O}\left(\left(2 k^{\eta / 2} / \lambda^{2}\right)^{2}\right)=\widetilde{O}\left(k^{\eta} / \lambda^{4}\right)$ pairs of windows. Since computing the LCS of every pair of windows takes time $O\left(\mathrm{w}_{\max }^{2}\right)$ then the expected overall running time is

$$
\widetilde{O}\left(k^{2-\eta / 2} \mathbf{w}_{\max }^{2}+k^{2-3 \eta / 2} \cdot k^{\eta} / \lambda^{4} \cdot \mathrm{w}_{\max }^{2}\right)=\widetilde{O}\left(k^{2-\eta / 2} \mathrm{w}_{\max }^{2} / \lambda^{4}\right) .
$$

Notice that $k^{2} \mathbf{w}_{\max }^{2} \simeq n^{2}$ and thus the above running time is truly subquadratic.

Sketch of the correctness: We refer the reader to $\left[\mathrm{CDG}^{+} 18\right]$ for a complete proof. The high-level idea is based on the following: Originally, the window size is defined based on a parameter $d$. Now, suppose we define some larger windows by setting $d^{\prime}=d k^{\eta / 2}$. Let $W_{A}^{\prime}$ and $W_{B}^{\prime}$ be the two sets of windows we obtain by using $d^{\prime}$ as the base parameter for constructing windows. Our window-compatibility lemma holds for any window size and therefore there is also a desired window-compatible solution when using $W_{A}^{\prime}$ and $W_{B}^{\prime}$.

The intuition behind the proof of $\left[\mathrm{CDG}^{+} 18\right]$ is the following: Suppose we know a window $w_{x}^{\prime}$ of $W_{A}^{\prime}$ is mapped to a window $w_{y}^{\prime}$ of $W_{B}^{\prime}$ in the optimal window-compatible solution for $W_{A}^{\prime}$ and $W_{B}^{\prime}$. There are at least $k^{\eta / 2}$ smaller windows of $W_{A}$ that lie inside $w_{x}^{\prime}$ and similarly $w_{y}^{\prime}$ contains at least $k^{\eta / 2}$ windows of $W_{B}$. Let the two sets be $X$ and $Y$. There is indeed a window-compatible solution for the windows of $X$ and $Y$. Now, the question that we need to answer, is how well such a window-compatible solution between $w_{x}^{\prime}$ and $w_{y}^{\prime}$ is approximated via $\widehat{M}$. If fewer than $\epsilon|X|$ pairs of this mapping are among the edges of NG then we can be sure that the approximation factor does not hurt by much. On the other hand, if more than $\epsilon|X|$ pairs are among the edges of NG then we can imply that one of such edges is discovered via Algorithm 30.12 in $\mathrm{NG}^{\prime}$ and thus $M$ would have the correct value of LCS between all windows of $X$ and all windows of $Y$. Thus, when we run Algorithm 30.10 on $M$, the output has an error of at most $\epsilon$ even though $M$ does not have the correct estimates for all pairs.

```
Algorithm 30.3 Ics-cmp data structure
    data structure Ics-cmp
    members
        \(X_{a, 1}, \cdots, X_{a, s}\)
        \(Y_{a, 1}, \cdots, Y_{a, s}\)
        \(\widetilde{\lambda} \in(0,1)\)
    end members
    procedure \(\operatorname{Initial}\left(w_{a},\left\{w_{i}\right\}_{i \in[s]}, \widetilde{\lambda}\right)\)
                            \(\triangleright\) Lemma 30.3.11
        for \(i \in[s]\) do
            compute opt \({ }_{a, i}\), a LCS between \(w_{a}\) and \(w_{i}\)
            let \(X_{a, i}\) be the set of indices of all the element of opt \({ }_{a, i}\) with respect to \(w_{a}\)
        end for
        for \(i \in[s]\) do
            let \(Y_{a, i}\) be an empty set
            while true do
                compute opt \({ }_{a, i}^{\prime}\), a LCS between \(w_{a}\) removing the elements with index in \(Y_{a, i}\)
    and \(w_{i}\)
                if the length of opt \({ }_{a, i}^{\prime}\) is less than \(\widetilde{\lambda}\left|w_{a}\right| / 2\) then
                    break
                else
                    put the indices of all the elements of opt \(_{a, i}^{\prime}\) with respect to \(w_{a}\) into \(Y_{a, i}\)
                end if
            end while
        end for
        return \(\left\{X_{a, i}\right\}_{i \in[s]},\left\{Y_{a, i}\right\}_{i \in[s]}\)
    end procedure
    procedure \(\operatorname{QUERY}\left(w_{i}, w_{j}\right) \quad \triangleright\) Lemma 30.3.12
        if \(\left|X_{a, i} \cap Y_{a, j}\right|>\widetilde{\lambda} \sqrt{\left|w_{i}\right| \cdot\left|w_{j}\right|} / 4\) then
        return accept
        else
            return reject
        end if
    end procedure
    end data structure
```

```
Algorithm 30.4 Constructing the candidate domains
    procedure ConstructCandidateDomains \(\left(\mathrm{sa}_{i}\right) \quad \triangleright\) Lemma 30.4.1
                                    \(\triangleright\) Given random access to a subarray sa \({ }_{i}\)
        \(k \leftarrow 20 \log (1 / \delta) /\left(\lambda \epsilon^{2}\right)\)
        Sample \(k\) elements from sa \({ }_{i}\), and denote the sampled elements by \(a_{j_{1}}, a_{j_{2}}, \ldots, a_{j_{k}}\)
        for \(\alpha\) in \([k]\) do
            for \(\beta\) in \([k]\) do
                    If \(a_{j_{\alpha}} \leq a_{j_{\beta}}\), then construct a candidate domain \(\left[a_{j_{\alpha}}, a_{j_{\beta}}\right]\)
            end for
        end for
        return all the constructed candidate domains
    end procedure
```

```
Algorithm 30.5 Constructing the pseudo solutions
    procedure ConstructPseudoSolutions \(\left(\operatorname{cdi}_{1}, \ldots, \operatorname{cdi}_{\sqrt{n}}\right) \quad \triangleright\)
    Lemma 30.4.3,30.4.4,30.4.5
                                    \(\triangleright\left\{\mathrm{cdi}_{i}\right\}_{i \in[\sqrt{n}]}\) is \(\sqrt{n}\) sets of candidate domain intervals
        pseudo-solutions \(\leftarrow \emptyset\)
        while true do
            assg \(\leftarrow\) largest assigment of candidate domain intervals to subarrays which is
    monotone
            if assg contains less than \(\epsilon \sqrt{n} \lambda\) non-empty candidate domain intervals then
                break
            else
                Add assg to pseudo-solutions
                for \(i \leftarrow 1\) to \(\sqrt{n}\) do
                if assg contains a candidate domain interval for subarray sa \({ }_{i}\) then
                    remove the corresponding candidate domain interval from \(\mathrm{cdi}_{i}\)
                end if
                end for
            end if
        end while
        return pseudo-solutions \(\mathrm{ps}_{1}, \mathrm{ps}_{2}, \ldots, \mathrm{ps}_{t}\)
    end procedure
```

```
Algorithm 30.6 Evaluate the pseudo solutions
    procedure EvaluatePseudoSolutions \(\left(\mathrm{ps}_{1}, \ldots, \mathrm{ps}_{t}\right) \quad \triangleright\) Lemma 30.4.6
                                    \(\triangleright\left\{\mathrm{ps}_{i}\right\}_{i \in[t]}\) is a set of pseudo solutions
        \(p \leftarrow \frac{1000 t \log ^{4} n}{\epsilon^{4} \lambda \sqrt{n}}\)
            Randomly sample each \(i \in[\sqrt{n}]\) with probability p , and put all the samples in a set
    W
        for each \(\mathrm{ps}_{j}\) do
            \(\widetilde{\mathrm{q}}\left(\mathrm{ps}_{j}\right) \leftarrow 0\)
            for each interval \(\left[\ell_{i}, r_{i}\right]\) in \(\mathrm{ps}_{j}\) do
                if \(i \in W\) then
                \(\widetilde{\mathrm{q}}\left(\mathrm{ps}_{j}\right) \leftarrow \widetilde{\mathrm{q}}\left(\mathrm{ps}_{j}\right)+\mathrm{lis}^{\left[\ell_{i}, r_{i}\right]}\left(\mathrm{sa}_{i}\right) / p\)
                end if
            end for
            end for
            return largest \(\widetilde{\mathrm{q}}\left(\mathrm{ps}_{j}\right)\) for all \(j \in[t]\)
    end procedure
```

```
Algorithm 30.7 Recursive Estimate LIS with Oracle
    procedure RecursiveEstimationWithOracle(Oracle, \(A, \lambda, \ell, r)\)
    Lemma 30.4.9
    :
    .
    :
        for \(i \in[\zeta]\) do
            \(\operatorname{cdi}_{i} \leftarrow\) ConstructCandidateDomains \(\left(\mathrm{sa}_{i}\right)\)
            discard all the intervals which are not in \([\ell, r]\) from \(\mathrm{cdi}_{i}\)
        end for
        \(\left\{\mathrm{ps}_{1}, \ldots, \mathrm{ps}_{t}\right\} \leftarrow{\text { ConstructPseudoSolutions }\left(\mathrm{cdi}_{1}, \ldots, \mathrm{cdi}_{\zeta}\right)}\) )
        \(\lambda_{0} \leftarrow\left(\frac{\lambda}{2^{8}}\right)^{4}\)
        \(p \leftarrow \frac{20 \log ^{4} n}{\lambda_{0} \zeta}\)
        randomly sample each \(i \in[\zeta]\) with probability \(p\), and put all the samples in a set \(Q\)
        for \(j \in[t]\) do
            \(c \leftarrow 0\)
            for \(i \in W\) do
                if \(\exists\left[\ell_{i}, r_{i}\right] \in \mathrm{ps}_{j}\) and \(\operatorname{Oracle}\left(\mathrm{sa}_{i}, \lambda_{0}, \ell_{i}, r_{i}\right)\) accepts then
                    \(c \leftarrow c+1\)
                end if
            end for
            if \(c \geq 3 \lambda_{0} p \zeta / 4\) then
                return accept
            end if
        end for
        return reject
    end procedure
```

```
Algorithm 30.8 Recursive Estimate LIS
    procedure RecursiveLiS \((A, \lambda, \ell, r) \quad \triangleright\) Lemma 30.4.11
        \(\triangleright\) input: sequence \(A\), parameter \(\lambda\), domain interval \([\ell, r]\)
                    \(\triangleright\) assume \(\mathrm{sa}_{1}, \mathrm{sa}_{2}, \ldots, \mathrm{sa}_{n^{\kappa}}\) are subarrays of \(A\)
        if the length of \(A\) is greater than \(n^{2 \kappa}\) then
            return RecursiveLiSWithOracle(RecursiveLis, \(A, \lambda, \ell, r)\) with \(\zeta=n^{\kappa}\)
        else
            for \(i \in\left[n^{\kappa}\right]\) do
                \(\operatorname{cdi}_{i} \leftarrow\) CONSTRUCtCANDIDATEDOMAINS \(\left(\mathrm{sa}_{i}\right)\)
                discard all the intervals which are not in \([\ell, r]\) from \(\mathrm{cdi}_{i}\)
            end for
            \(\left\{\mathrm{ps}_{1}, \ldots, \mathrm{ps}_{t}\right\} \leftarrow\) ConstructPseudosolutions \(\left(\mathrm{cdi}_{1}, \ldots, \mathrm{cdi}_{n^{\kappa}}\right)\)
            if EvaluatePseudoSolutions \(\left(\mathrm{ps}_{1}, \ldots, \mathrm{ps}_{t}\right) \geq \lambda|A|\) then
                    return accept
            else
                return reject
            end if
        end if
    end procedure
```

```
Algorithm 30.9 An algorithm that is able to generate multiple layers
    procedure GenerateMultipleLevels \(\left(A, B, n, d, \epsilon_{0}\right) \quad \triangleright\) Lemma 30.6.1
        \(f \leftarrow \Theta\left(\frac{1}{\epsilon_{0}} \log \frac{1}{\epsilon_{0}}\right)\)
        \(\triangleright d \cdot\left(1+\epsilon_{0}\right)^{f}=d / \epsilon_{0}\)
        \(W_{A} \leftarrow \emptyset, W_{B} \leftarrow \emptyset\)
        for \(i=0 \rightarrow f\) do
            \(W_{A}^{i} \leftarrow \emptyset, W_{B}^{i} \leftarrow \emptyset\)
            \(d_{i} \leftarrow d\left(1+\epsilon_{0}\right)^{i}, t_{i} \leftarrow n / d_{i}, g_{i} \leftarrow \epsilon_{0} d_{i} \quad \triangleright d_{i}\) is the length, \(g_{i}\) is the shift
            for \(x=0 \rightarrow d_{i} / g_{i}-1\) do
                for \(y=0 \rightarrow t_{i}-1\) do
                    left \(\leftarrow x \cdot g_{i}+y \cdot d_{i}\)
                    len \(\leftarrow d_{i}\)
                    \(W_{A}^{i} \leftarrow W_{A}^{i} \cup A^{(\text {left,len })}, W_{B}^{i} \leftarrow W_{B}^{i} \cup B^{(\text {left,len })}\)
                end for
            end for
            \(W_{A} \leftarrow W_{A} \cup W_{A}^{i}, W_{B} \leftarrow W_{B} \cup W_{B}^{i}\)
        end for
        return \(W_{A}, W_{B}, f\)
    end procedure
```

```
Algorithm 30.10 Dynamic programming algorithm for block-based LCS problem
    procedure DP-LCS \(\left(W_{A}, W_{B}, M\right) \quad \triangleright\) Lemma 30.6.3
        Note that \(M\) is table that for each \(A_{k_{1}}\) and \(B_{k_{2}}, M\left(A_{k_{1}}, B_{k_{2}}\right)\) gives a cost
        Note that \(W_{A}\) and \(W_{B}\) are sets produce by Algorithm 30.9
        Let \(S_{1}\) denote the sorted list of the set(without duplicate) of right indices of each \(A_{k_{1}}\)
        Let \(S_{2}\) denote the sorted list of the set(without duplicate) of right indices of each \(B_{k_{2}}\)
        for \(i_{1} \in S_{1}\) do
            for \(i_{2} \in S_{2}\) do
                \(C\left[i_{1}\right]\left[i_{2}\right] \leftarrow 0\)
                \(x_{1} \leftarrow 0\)
                for \(A_{k_{1}}, B_{k_{2}}\) have right index \(i_{1}, i_{2}\) do \(\triangleright\) Case 1
                    left \(_{1} \leftarrow\) left index of \(A_{k_{1}}\)
                left \(_{2} \leftarrow\) left index of \(B_{k_{2}}\)
                \(\mathrm{tmp} \leftarrow C\left[\mathrm{left}_{1}\right]\left[\mathrm{left}_{2}\right]+M\left(A_{k_{1}}, B_{k_{2}}\right)\)
                if \(\mathrm{tmp}>x_{1}\) then
                    \(x_{1} \leftarrow \mathrm{tmp}\)
                end if
                end for
                Let \(\operatorname{prev}\left(i_{1}\right)\) denote the first index \(\in S_{1}\) that is early than \(i_{1}\)
                Let \(\operatorname{prev}\left(i_{2}\right)\) denote the first index \(\in S_{2}\) that is early than \(i_{2}\)
                \(x_{2} \leftarrow C\left[\operatorname{prev}\left(i_{1}\right)\right]\left[i_{2}\right] \quad \triangleright\) Case 2
                \(x_{3} \leftarrow C\left[i_{1}\right]\left[\operatorname{prev}\left(i_{2}\right)\right] \quad \triangleright\) Case 3
                \(C\left[i_{1}\right]\left[i_{2}\right] \leftarrow \max \left(x_{1}, x_{2}, x_{3}\right)\)
            end for
        end for
    end procedure
```

```
Algorithm 30.11
    procedure ApproxLCS \((A, B, n, \lambda)\)
        \(d \leftarrow n^{0.2}, \beta \leftarrow \epsilon \leftarrow \Theta(1), \epsilon_{0} \leftarrow \Theta(\epsilon / \lambda)\)
        \(W_{A}, W_{B}, L \leftarrow \operatorname{CreateMultipleLevels}\left(A, B, n, d, \epsilon_{0}\right) \quad \triangleright\) Lemma 30.6.1,
    Algorithm 30.9
        \(M \leftarrow\) NonMetricEstimation \(\left(W_{A}, W_{B}\right)\)
        \(C \leftarrow \operatorname{DPLCS}\left(W_{A}, W_{B}, M\right) \quad \triangleright\) Lemma 30.6.3, Algorithm 30.10
        return \(C\)
    end procedure
    procedure NonMEtricEstimation \(\left(W_{A}, W_{B}\right) \quad \triangleright\) Section 30.3 and Section 30.7
        for \(i=1 \rightarrow L\) do
            for \(j=1 \rightarrow L\) do
                Make strings in \(W_{A}^{i}\) and \(W_{B}^{j}\) have the same length by appending \(\perp\) to the
    short one
            \(W_{A}^{i}\) has \(t_{i}\) blocks and each of it has length \(d_{i}\)
            \(W_{B}^{j}\) has \(t_{j}\) blocks and each of it has length \(d_{j}\)
            \(d \leftarrow \max \left(d_{i}, d_{j}\right), t \leftarrow \max \left(t_{i}, t_{j}\right)\)
            Create \(S^{(i, j)}\) by Step 1 (Section 30.3) and Step 2 (Section 30.7)
                end for
        end for
        Form cost table \(M\) by putting all \(S^{(i, j)}\) together
        return \(M\)
    end procedure
```

```
Algorithm 30.12 Constructing \(M\) based on \(\widehat{M}\)
    procedure NonMEtricEstimation \(\left(W_{A}, W_{B}\right) \quad \triangleright\) Lemma 30.7.1
        \(S \leftarrow \emptyset\)
        for \(w_{i} \in W_{A}\) do
            Add \(w_{i}\) to \(S\) with probability \(\epsilon^{-1} k^{-\eta / 2}\)
        end for
        for \(w_{i} \in S\) do
            for \(w_{j} \in W_{B}\) do
            Compute \(\operatorname{lcs}\left(w_{i}, w_{j}\right)\) and put an edge in \(\mathbf{N G}^{\prime}\) if \(\widehat{M}[i][j]\) does not correctly esti-
    mate \(\operatorname{lcs}\left(w_{i}, w_{j}\right)\).
        end for
        end for
        \(M \leftarrow \widehat{M}\)
        for \(w_{i} \in W_{a}\) do
            for \(w_{j} \in W_{B}\) do
            if there exist nearby windows \(w_{i^{\prime}}, w_{j^{\prime}}\) such that \(\left(i^{\prime}, j^{\prime}\right) \in E\left(\mathrm{NG}^{\prime}\right)\) then
                    \(M[i][j] \leftarrow \operatorname{lcs}\left(w_{i}, w_{j}\right)\)
            end if
            end for
        end for
        return \(M\)
    end procedure
```


## Chapter 31

## Fourier Transform

### 31.1 One dimensional Fourier transform

The result in this section is based on the unpublished manuscript by Jerry Li, Ruoqi Shen and Zhao Song [LSS19].

Lemma 31.1.1 (Claim 5.2 in [CKPS16]). For any $k$, there exists $m=O\left(k^{2} \log k\right)$ such that for any $k$-Fourier-sparse signal $x(t)$, any $t_{0} \geq 0$ and $\tau>0$, there always exists $C_{1}, \ldots, C_{m} \in \mathbb{C}$ such that $\left|C_{j}\right| \leq 11$ for all $j \in[m]$ and

$$
x\left(t_{0}\right)=\sum_{j \in[m]} C_{j} \cdot x\left(t_{0}+j \cdot \tau\right) .
$$

Lemma 31.1.2 (Bounding absolute by discrete sum on grids). For any $k$-Fourier-sparse signal $x(t)=\sum_{j=1}^{k} v_{j} e^{2 \pi \mathrm{i} f_{j} t}$, we have

$$
\max _{t \in[0, T]}|x(t)|^{2} \leq O\left(k^{6} \log ^{3} k\right) \cdot \frac{1}{T} \sum_{t=0}^{T}|x(t)|^{2}
$$

Proof. Without loss of generality, we fix $T=1$ and assume $|x(0)|^{2}=\max _{t \in[0, T]}|x(t)|^{2}$. If $t^{*}=\arg \max _{t \in[0, T]}|x(t)|^{2}$ is not 0 or $T=1$, we can rescale the two intervals $\left[0, t^{*}\right]$ and $\left[t^{*}, T\right]$ to $[0,1]$. By Lemma 31.1.1, we choose $t_{0}=0$ such that $\forall \tau>0$, there exists $C_{1}, \ldots, C_{m} \in \mathbb{C}$ and

$$
x(0)=\sum_{j \in[m]} C_{j} \cdot x(j \cdot \tau) .
$$

By the Cauchy-Schwarz inequality, it implies that for any $\tau$,

$$
\begin{aligned}
|x(0)|^{2} & \leq m \sum_{j \in[m]}\left|C_{j}\right|^{2}|x(j \cdot \tau)|^{2} \\
& \lesssim m \sum_{j \in[m]}|x(j \cdot \tau)|^{2} .
\end{aligned}
$$

Then,

$$
\begin{aligned}
|x(0)|^{2} & =\frac{1}{\lfloor T / m\rfloor} \sum_{\tau=0}^{\lfloor T / m\rfloor}|x(0)|^{2} \\
& \lesssim \frac{m}{T} \sum_{\tau=0}^{\lfloor T / m\rfloor}\left(m \sum_{j \in[m]}|x(j \cdot \tau)|^{2}\right) \\
& =\frac{m^{2}}{T} \sum_{j \in[m]} \sum_{\tau=0}^{\lfloor T / m\rfloor}|x(j \cdot \tau)|^{2} \\
& \leq \frac{m^{3}}{T} \sum_{t=0}^{T}|x(t)|^{2} .
\end{aligned}
$$

From $m=O\left(k^{2} \log k\right)$, we obtain $|x(0)|^{2}=O\left(k^{6} \log ^{3} k\|x\|_{[T]}\right)$.

### 31.2 High dimensional Fourier transform

The result in this section is based on the unpublished result by Zhao Song [Son17].

Lemma 31.2.1 ( $d$ dimensional version of Lemma 5.1 in [CKPS16]). For any d-dimensional $k$-Fourier-sparse signal $x(t): \mathbb{R}^{d} \rightarrow \mathbb{C}$ and any duration $T$, we have

$$
\max _{t \in[0, T]^{d}}|x(t)|^{2} \leq k^{O(d)}\|x\|_{T}^{2}
$$

where $\|x\|_{T}^{2}=\frac{1}{T^{d}} \int_{[0, T]^{d}}|x(t)|^{2} \mathrm{~d} t$.

Proof. Without loss of generality, we fix $T=1$. Then $\|x\|_{T}^{2}=\int_{[0,1]^{d}}|x(t)| \mathrm{d} t$. Because $\|x\|_{T}^{2}$ is the average over the interval $[0, T]^{d}$, if $t^{*}=\arg \max _{t \in[0, T]^{d}}|x(t)|^{2}$ is not $0^{d}$ or $T=1$, we can rescale the two intervals $\left[0^{d}, t^{*}\right]$ and $\left[t^{*}, T^{d}\right]$ to $[0,1]$ and prove the desired property separately. Hence we assume that $|x(0)|^{2}=\max _{t \in[0, T]}|x(t)|^{2}$ in the proof.

The proof of the following Claim can be found in Section 31.3 (is similarly to the proof of Claim 5.2 in [CKPS16])

Claim 31.2.2 ( $d$-variables version of Claim 5.2 in [CKPS16]). For any $k$ and $d$, there exists $m=O\left(k^{2} \log k\right)$ such that for any d-dimensional $k$-Fourier-sparse signal $x(t)$, any $t_{0} \in \mathbb{R}_{\geq 0}^{d}$ and $\tau \in \mathbb{R}_{>0}^{d}$, there always exist $C_{1}, C_{2}, \cdots, C_{m} \in \mathbb{C}$ such that the following properties hold,

$$
\begin{aligned}
\text { Property I } & \left|C_{j}\right| \leq 11 \text { for all } j \in[m], \\
\text { Property II } & x\left(t_{0}\right)=\sum_{j \in[m]} C_{j} \cdot x\left(t_{0}+j \cdot \tau\right) .
\end{aligned}
$$

In the next a few paragraph, we show how to use the above Claim to prove Lemma 31.2.1.
We use $0^{d}$ to denote a length- $d$ vector with 0 everywhere. We choose $t_{0}=0^{d}$ such that $\forall \tau \in \mathbb{R}_{>0}^{d}$ there always exist $C_{1}, \cdots, C_{m} \in \mathbb{C}$, and

$$
x\left(0^{d}\right)=\sum_{j \in[m]} C_{j} \cdot x(j \cdot \tau) .
$$

By the Cauchy-Schwarz inequality, it implies that for any $\tau$,

$$
\left|x\left(0^{d}\right)\right|^{2} \leq m \sum_{j \in[m]}\left|C_{j}\right|^{2}|x(j \cdot \tau)|^{2}
$$

At last, we obtain

$$
\begin{align*}
\left|x\left(0^{d}\right)\right|^{2} & =m^{d} \int_{[0,1 / m]^{d}}\left|x\left(0^{d}\right)\right|^{2} \mathrm{~d} \tau \\
& \lesssim m^{d} \cdot \int_{[0,1 / m]^{d}}\left(m \sum_{j=1}^{m}|x(j \cdot \tau)|^{2}\right) \mathrm{d} \tau \\
& =m^{d+1} \cdot \sum_{j=1}^{m} \int_{[0,1 / m]^{d}}|x(j \cdot \tau)|^{2} \mathrm{~d} \tau \\
& =m^{d+1} \cdot \sum_{j=1}^{m} \frac{1}{j^{d}} \int_{[0, j / m]^{d}}|x(\tau)|^{2} \mathrm{~d} \tau \\
& =m^{d+1} \cdot \sum_{j=1}^{m} \frac{1}{j^{d}} \int_{[0,1]^{d}}|x(\tau)|^{2} \mathrm{~d} \tau \\
& \lesssim m^{d+1} \log m \cdot\|x\|_{T}^{2} \\
& \leq k^{O(m)}\|x\|_{T}^{2}, \tag{31.1}
\end{align*}
$$

where the third step follows by moving $m$ outside of the integral and swapping the integration and the summation, the fourth step follows by replacing $j \tau$ by $\tau$, the fifth step follows by $j / m \leq 1$, the sixth step follows by $\sum_{j=1}^{m} 1 / j=O(\log m)$ and the definition of $\|x\|_{T}^{2}$, and the last step follows by $m=\operatorname{poly}(k)$.

Thus, we have the desired bound.

### 31.3 Proof of Claim 31.2.2

For $x(t)=\sum_{i=1}^{k} v_{i} e^{2 \pi \mathbf{i} f_{i}^{\top} t}$, where $f_{i} \in \mathbb{R}^{d}$ and $t \in \mathbb{R}^{d}$. We fix $t_{0} \in \mathbb{R}^{d}$ and $\tau \in \mathbb{R}^{d}$ then rewrite $x\left(t_{0}+j \cdot \tau\right)$ as a polynomial of $b_{i}=v_{i} \cdot e^{2 \pi \mathbf{i} f_{i}^{\top} t_{0}}$ and $z_{i}=e^{2 \pi \mathbf{i} f_{i}^{\top} \tau}$ for each $i \in[k]$.

$$
\begin{aligned}
x\left(t_{0}+j \cdot \tau\right) & =\sum_{i=1}^{k} v_{i} e^{2 \pi \mathbf{i} f_{i}^{\top}\left(t_{0}+j \tau\right)} \\
& =\sum_{i=1}^{k} v_{i} e^{2 \pi \mathbf{i} f_{i}^{\top} t_{0}} e^{2 \pi \mathbf{i} f_{i}^{\top} j \tau} \\
& =\sum_{i=1}^{k} b_{i} \cdot z_{i}^{j}
\end{aligned}
$$

Given $k$ and $z_{1}, \cdots, z_{k}$, let $P(z)=\sum_{j=0}^{m} c_{j} z^{j}$ be the degree $m$ polynomial in Lemma 5.4 ([CKPS16]).

$$
\begin{aligned}
\sum_{j=0}^{m} c_{j} x\left(t_{0}+j \tau\right) & =\sum_{j=0}^{m} c_{j} \sum_{i=1}^{k} b_{i} \cdot z_{i}^{j} \\
& =\sum_{i=1}^{k} b_{i} \sum_{j=0}^{m} c_{j} \cdot z_{i}^{j} \\
& =\sum_{i=1}^{k} b_{i} P\left(z_{i}\right) \\
& =0
\end{aligned}
$$

where the last step follows by Property I of $\mathrm{P}(\mathrm{z})$ in Lemma 5.4 ([CKPS16]). From the Property II and III of $P(z)$ (Lemma 5.4 [CKPS16]), we obtain $x\left(t_{0}\right)=-\sum_{j=1}^{m} c_{j} x\left(t_{0}+j \tau\right)$.

## Chapter 32

## Map-Reduce

Many modern parallel systems, such as MapReduce, Hadoop and Spark, can be modeled well by the MPC model. The MPC model captures well coarse-grained computation on large data - data is distributed to processors, each of which has a sublinear (in the input data) amount of memory and we alternate between rounds of computation and rounds of communication, where each machine can communicate an amount of data as large as the size of its memory. This model is stronger than the classical PRAM model, and it is an intriguing question to design algorithms whose running time is smaller than in the PRAM model.

One fundamental graph problem is connectivity. On an undirected graph with $n$ nodes and $m$ edges, $O(\log n)$ round connectivity algorithms have been known for over 35 years. However, no algorithms with better complexity bounds were known. In this work, we give fully scalable, faster algorithms for the connectivity problem, by parameterizing the time complexity as a function of the diameter of the graph. Our main result is a $O\left(\log D \log \log _{m / n} n\right)$ time connectivity algorithm for diameter- $D$ graphs, using $\Theta(m)$ total memory. If our algorithm can use more memory, it can terminate in fewer rounds, and there is no lower bound on the memory per processor.

We extend our results to related graph problems such as spanning forest, finding a

DFS sequence, exact/approximate minimum spanning forest, and bottleneck spanning forest. We also show that achieving similar bounds for reachability in directed graphs would imply faster boolean matrix multiplication algorithms.

We introduce several new algorithmic ideas. We describe a general technique called double exponential speed problem size reduction which roughly means that if we can use total memory $N$ to reduce a problem from size $n$ to $n / k$, for $k=(N / n)^{\Theta(1)}$ in one phase, then we can solve the problem in $O\left(\log \log _{N / n} n\right)$ phases. In order to achieve this fast reduction for graph connectivity, we use a multistep algorithm. One key step is a carefully constructed truncated broadcasting scheme where each node broadcasts neighbor sets to its neighbors in a way that limits the size of the resulting neighbor sets. Another key step is random leader contraction, where we choose a smaller set of leaders than many previous works do.

This part is based upon the following previous publication

- Alexandr Andoni, Zhao Song, Clifford Stein, Zhengyu Wang, Peilin Zhong

Parallel graph connectivity in log diameter rounds.
FOCS 2018 [ASS $\left.{ }^{+} 18\right]$

### 32.1 Introduction

Recently, several parallel systems, including MapReduce [DG04, DG08], Hadoop [Whi12], Dryad [IBY $\left.{ }^{+} 07\right]$, Spark [ZCF ${ }^{+}$10], and others, have become successful in practice. This success has sparked a renewed interest in algorithmic ideas for these parallel systems.

One important theoretical direction has been to develop good models of these modern systems and to relate them to classic models such as PRAM. The work of $\left[F M S^{+} 10, \mathrm{KSV} 10\right.$, GSZ11, BKS13, ANOY14] have led to the model of Massive Parallel Computing (MPC) that balances accurate modeling with theoretical elegance. MPC is a variant of the Bulk Synchronous Parallel (BSP) model [Val90]. In particular, MPC allows $N^{\delta}$ space per machine (processor), where $\delta \in(0,1)$ and $N$ is the input size, with alternating rounds of unlimited local computation, and communication of up to $N^{\delta}$ data per processor. An MPC algorithm can equivalently be seen as a small circuit, with arbitrary, $N^{\delta}$-fan-in gates; the depth of the circuit is the parallel time. Any PRAM algorithm can be simulated on MPC in the same parallel time [KSV10, GSZ11]. However, MPC is in fact more powerful than the PRAM: even computing the XOR of $N$ bits requires near-logarithmic parallel-time on the most powerful CRCW PRAMs [BH89], whereas it takes constant, $O(1 / \delta)$, parallel time on the MPC model.

The main algorithmic question of this area is then: for which problems can we design MPC algorithms that are faster than the best PRAM algorithms? Indeed, this question has been the focus of several recent papers, see, e.g., [KSV10, LMSV11, EIM11, ANOY14, AG18, AK17, IMS17b, CLM ${ }^{+}$18]. Graph problems have been particularly well studied and one fundamental problem is connectivity in a graph. While this problem has a standard logarithmic time PRAM algorithm [SV82], we do not know whether we can solve it faster in the MPC model.

While we would like fully scalable algorithms - which work for any value of $\delta>0$ there have been graph algorithms that use space close to the number of vertices $n$ of the graph. In particular, the result of [LMSV11] showed a faster algorithm for the setting when the space per machine is polynomially larger than the number of vertices, i.e., $s \geq n^{1+\Omega(1)}$, and hence the number of edges is necessarily $m \geq n^{1+\Omega(1)}$. In fact, similar space restrictions are pervasive for all known sub-logarithmic time graph algorithms, which require $s=\Omega\left(\frac{n}{\log ^{O(1)} n}\right)$ [LMSV11, AG18, AK17, $\mathrm{CLM}^{+}$18] (the only exception is [ANOY14] who consider geometric graphs). We highlight the work of $\left[\mathrm{CLM}^{+} 18\right]$, who manage to obtain slightly sublinear space of $n / \log ^{\Omega(1)} n$ in $\log ^{O(1)} \log n$ parallel time, for the approximate matching problem and $\left[\mathrm{ABB}^{+} 17\right]$ who obtain slightly sublinear space of $n / \log ^{\Omega(1)} n$ in $O(\log \log n)$ parallel time. We note that the space of $\sim n$ also coincides with the space barrier of the semi-streaming model: essentially no graph problems are solvable in less than $n$ space in the streaming model, unless we have many more passes; see e.g. the survey [McG09].

It remains a major open question whether there exist fully scalable connectivity MPC algorithms with sub-logarithmic time (e.g., for sparse graphs). There are strong indications that such algorithms do not exist: [BKS13] show logarithmic lower bounds for restricted algorithms. Alas, showing an unconditional lower bound may be hard to prove, as that would imply circuit lower bounds [RVW16].

In this work, we show faster, fully scalable algorithms for the connectivity problem, by parameterizing the time complexity as a function of the diameter of the graph. The diameter of the graph is the largest diameter of its connected components. Our main result is an $O\left(\log D \log \log _{m / n} n\right)$ time connectivity algorithm for diameter- $D$ graphs with $m$ edges. Parameterizing as a function of $D$ is standard, say, in the distributed computing
literature [PRS16, HHW18]. In fact, some previous MPC algorithms for connectivity in the applied communities have been conjectured to obtain $O(\log D)$ time [RMCS13]; alas, we show in Section 32.11 the algorithm of [RMCS13] has a lower bound of $\Omega(\log n)$ time.

Our algorithms exhibit a tradeoff between the total amount of memory available and the number of rounds of computation needed. For example, if the total space is $\Omega\left(n^{1+\gamma^{\prime}}\right)$ for some constant $\gamma^{\prime}>0$, then our algorithms run in $O(\log D)$ rounds only.

### 32.1.1 The MPC model

Before stating our full results, we briefly recall the MPC model [BKS13]. A detailed discussion appears in Section 32.7, along with some core primitives implementable in the MPC model.

Definition 32.1.1 $((\gamma, \delta)$-MPC model). Fix parameters $\gamma, \delta>0$, and suppose $N \geq 1$ is the input size. There are $p \geq 1$ machines (processors) each with local memory size $s=\Theta\left(N^{\delta}\right)$, such that $p \cdot s=O\left(N^{1+\gamma}\right)$. The space size is measured by words, each of $\Theta(\log (s \cdot p))$ bits. The input is distributed on the local memory of $\Theta(N / s)$ input machines. The computation proceeds in rounds. In each round, each machine performs computation on the data in its local memory, and sends messages to other machines at the end of the round. The total size of messages sent or received by a machine in a round is bounded by $s$. In the next round, each machine only holds the received messages in its local memory. At the end of the computation, the output is distributed on the output machines. Input/output machines and other machines are identical except that input/output machine can hold a part of the input/output. The parallel time of an algorithm is the number of rounds needed to finish the computation.

In this model, the space per machine is sublinear in $N$, and the total space is only an $O\left(N^{\gamma}\right)$ factor more than the input size $N$. In this paper, we consider the case when $\delta$ is an arbitrary constant in $(0,1)$. Our results are for both the most restrictive case of $\gamma=0$ (total space is linear in the input size), as well as $\gamma>0$ (for which our algorithms are a bit faster). The model from Definition 32.1.1 matches the model MPC $(\epsilon)$ from [BKS13] with $\epsilon=\gamma /(1+\gamma-\delta)$ and the number of machines $p=O\left(N^{1+\gamma-\delta}\right)$.

### 32.1.2 Our Results

While our main result is a $\sim \log D$ time connectivity MPC algorithm, our techniques extend to related graph problems, such as spanning forest, finding a DFS sequence, and exact/approximate minimum spanning forest. We also prove a lower bound showing that, achieving similar bounds for reachability in directed graphs would imply faster boolean matrix multiplication algorithms.

We now state our results formally. For all results below, consider an input graph $G=(V, E)$, with $n=|V|, N=|V|+|E|$, and $D$ being the upper bound on the diameter of any connected component of $G$.

Connectivity: In the connectivity problem, the goal is to output the connected components of an input graph $G$, i.e. at the end of the computation, $\forall v \in V$, there is a unique tuple $(x, y)$ with $x=v$ stored on an output machine, where $y$ is called the color of $v$. Any two vertices $u, v$ have the same color if and only if they are in the same connected component.

Theorem 32.1.1 (Connectivity in MPC, restatement of Theorem 32.8.4). For any $\gamma \in$ $[0,2]$ and any constant $\delta \in(0,1)$, there is a randomized $(\gamma, \delta)-$ MPC algorithm (see Algorithm 32.3) which outputs the connected components of the graph $G$ in $O(\min (\log D$. $\left.\left.\log \frac{\log n}{\log \left(N^{1+\gamma} / n\right)}, \log n\right)\right)$ parallel time. The success probability is at least 0.98. In addition, if the algorithm fails, then it returns FAIL.

Notice that in the most restrictive case of $\gamma=0$ and $m=n$, we obtain $O(\min (\log D$. $\log \log n, \log n))$ time. When the total space is slightly larger, or the graph is slightly denseri.e. $\gamma>c$ or $\log _{n} m>1+c$, where $c>0$ is an arbitrarily small constant-then we obtain $O(\log D)$ time.

Remark 32.1.1. We note the concurrent and independent work of [ASW18], who also give a connectivity algorithm in the MPC model but with different guarantees. In particular, their runtime is parameterized as a function of $\lambda$, which is a lower bound on the spectral gap ${ }^{1}$ of the connected components of $G$. For a graph $G$ with $n$ vertices and $m=\widetilde{O}(n)$ edges, their algorithm runs in $O(\log \log n+\log (1 / \lambda))$ parallel time and uses $\widetilde{O}\left(n / \lambda^{2}\right)$ total space. In contrast, our algorithm has a runtime of $O\left(\log D \cdot \log \log _{N / n} n\right)$, where $D$ is the largest diameter of a connected component of $G$, and $N=\Omega(m)$ is the total space available. To compare the two runtimes, we note that: 1) $D \leq O\left(\frac{\log n}{\lambda}\right)$ for any undirected graph $G$; and 2) there exist sparse graphs $G^{2}$ with $n$ vertices and $O(n)$ edges such that $\frac{1}{\lambda} \geq D \cdot n^{\Omega(1)}$ and $D \leq O(\log n)$. Thus, our results subsume [ASW18] in the case when total space is $N=n^{1+\Omega(1)}$, but are incomparable otherwise.

Spanning forest problem: In the spanning forest problem, the goal is to output a subset of edges of an input graph $G$ such that the output edges together with the vertices of $G$ form a spanning forest of the graph $G$. In the rooted spanning forest problem, in addition to the edges of the spanning forest, we are also required to orient the edge from child to parent, so that the parent-child pairs form a rooted spanning forest of the input graph $G$.

Theorem 32.1.2 (Spanning Forest, restatement of Theorem 32.8.14). For any $\gamma \in[0,2]$ and any constant $\delta \in(0,1)$, there is a randomized $(\gamma, \delta)$ - MPC algorithm (see Algorithm 32.11 and Algorithm 32.12) which outputs the rooted spanning forest of the graph $G$ in $O(\min (\log D$. $\left.\left.\log \frac{\log n}{\log \left(N^{1+\gamma} / n\right)}, \log n\right)\right)$ parallel time. The success probability is at least 0.98. In addition, if

[^109]the algorithm fails, then it returns FAIL.

Our spanning forest algorithm can also output an approximation to the diameter, as follows.

Theorem 32.1.3 (Diameter Estimator, restatement of Theorem 32.8.15). For any $\gamma \in[0,2]$ and any constant $\delta \in(0,1)$, there is a randomized $(\gamma, \delta)-$ MPC algorithm which outputs a diameter estimator $D^{\prime}$ of the input graph $G$ in $O\left(\min \left(\log D \cdot \log \frac{\log n}{\log \left(N^{1+\gamma / n)}\right.}, \log n\right)\right)$ parallel time such that $D \leq D^{\prime} \leq D^{O\left(\log \left(1 / \gamma^{\prime}\right)\right)}$, where $\gamma^{\prime}=\frac{\log \left(N^{1+\gamma} / n\right)}{\log n}$. The success probability is at least 0.98. In addition, if the algorithm fails, then it returns FAIL.

Depth-First-Search sequence: If the input graph $G$ is a tree, then we are able to output a Depth-First-Search sequence of that tree in $O(\log D)+T$ parallel time, where $T$ is parallel time to compute a rooted tree (see Theorem 32.1.2 for our upper bound of $T$ ) for $G$. (See Section 32.7.2 for a discussion how to represent a sequence in the MPC model.)

Theorem 32.1.4 (DFS Sequence of a Tree, restatement of Theorem 32.8.21). Suppose the graph $G$ is a tree. For any $\gamma \in[\beta, 2]$ and any constant $\delta \in(0,1)$, there is a randomized $(\gamma, \delta)$ - MPC algorithm (Algorithm 32.17) that outputs a Depth-First-Search sequence for the input graph $G$ in $O(\min (\log D \cdot \log (1 / \gamma), \log n))$ parallel time, where $\beta=\Theta(\log \log n / \log n)$. The success probability is at least 0.98. In addition, if the algorithm fails, then it returns FAIL.

Applications of DFS sequence of a tree include lowest common ancestor, tree distance oracle, the size of every subtree, and others. See Section 32.6.4 for a more detailed discussion of the DFS sequence of a tree.

Minimum Spanning Forest: In the minimum spanning forest problem, the goal is to compute the minimum spanning forest of a weighted graph $G$.

Theorem 32.1.5 (Minimum Spanning Forest, restatement of Theorem 32.9.3). Consider a weighted graph $G$ with weights $w: E \rightarrow \mathbb{Z}$ such that $\forall e \in E,|w(e)| \leq \operatorname{poly}(n)$. For any $\gamma \in[0,2]$ and any constant $\delta \in(0,1)$, there is a randomized $(\gamma, \delta)-$ MPC algorithm which outputs a minimum spanning forest of $G$ in $O\left(\min \left(\log D_{M S F} \cdot \log \left(\frac{\log n}{1+\gamma \log n}\right), \log n\right) \cdot \frac{\log n}{1+\gamma \log n}\right)$ parallel time, where $D_{M S F}$ is the diameter (with respect to the number of edges/hops) of a minimum spanning forest of $G$. The success probability is at least 0.98 . In addition, if the algorithm fails, then it returns FAIL.

We note that we require the bounded weights condition merely to ensure that each weight is described by one word.

Theorem 32.1.6 (Approximate Minimum Spanning Forest, restatement of Theorem 32.9.4). Consider a weighted graph $G$ with weights $w: E \rightarrow \mathbb{Z}_{\geq 0}$ such that $\forall e \in E,|w(e)| \leq$ poly $(n)$. For any $\epsilon \in(0,1), \gamma \in[\beta, 2]$ and any constant $\delta \in(0,1)$, there is a randomized $(\gamma, \delta)$ - MPC algorithm which can output a $(1+\epsilon)$ approximate minimum spanning forest for $G$ in $O\left(\min \left(\log D_{M S F} \cdot \log \left(\frac{\log n}{\log \left(N^{1+\gamma} /\left(\epsilon^{-1} n \log n\right)\right)}\right), \log n\right)\right)$ parallel time, where $\beta=$ $\Theta\left(\log \left(\epsilon^{-1} \log n\right) / \log n\right)$, and $D_{M S F}$ is the diameter (with respect to the number of edges/hops) of a minimum spanning forest of $G$. The success probability is at least 0.98. In addition, if the algorithm fails, then it returns FAIL.

Theorem 32.1.7 (Bottleneck Spanning Forest, restatement of Theorem 32.9.5). Consider a weighted graph $G$ with weights $w: E \rightarrow \mathbb{Z}$ such that $\forall e \in E,|w(e)| \leq \operatorname{poly}(n)$. For any $\gamma \in$
$[0,2]$ and any constant $\delta \in(0,1)$, there is a randomized $(\gamma, \delta)-$ MPC algorithm which can output a bottleneck spanning forest for $G$ in $O\left(\min \left(\log D_{M S F} \cdot \log \left(\frac{\log n}{1+\gamma \log n}\right), \log n\right) \cdot \log \left(\frac{\log n}{1+\gamma \log n}\right)\right)$ parallel time, where $D_{M S F}$ is the diameter (with respect to the number of edges/hops) of a minimum spanning forest of $G$. The success probability is at least 0.98. In addition, if the algorithm fails, then it returns FAIL.

Conditional hardness for directed reachability. We also consider the reachability question in the directed graphs, for which we show similar to the above results are unlikely. In particular, we show that if there is a fully scalable multi-query directed reachability $(0, \delta)$ - MPC algorithm with $n^{o(1)}$ parallel time and polynomial local running time, then we can compute the Boolean Matrix Multiplication in $n^{2+\epsilon+o(1)}$ time for arbitrarily small constant $\epsilon>0$. We note that the equivalent problem for undirected graphs can be solved in $O(\log D \log \log n)$ parallel time via Theorem 32.1.1.

Theorem 32.1.8 (Directed Reachability vs. Boolean Matrix Multiplication, restatement of Theorem 32.10.1). Consider a directed graph $G=(V, E)$. If there is a polynomial local running time, fully scalable $(\gamma, \delta)-\mathrm{MPC}$ algorithm that can answer $|V|+|E|$ pairs of reachability queries simultaneously for $G$ in $O\left(|V|^{\alpha}\right)$ parallel time, then there is a sequential algorithm which can compute the multiplication of two $n \times n$ boolean matrices in $O\left(n^{2} \cdot n^{2 \gamma+\alpha+\epsilon}\right)$ time, where $\epsilon>0$ is a constant which can be arbitrarily small.

Finally, in Section 32.11 we show hard instances for the algorithm [RMCS13].

### 32.1.3 Our Techniques

In this section, we give an overview of the various techniques that we use in our algorithms. More details, as well as some of the low level details of the implementation in the MPC model, are defered to later sections.

Before getting into our techniques, we mention two standard tools to help us build our MPC subroutines. The first one is sorting: while in the PRAM model it takes $\sim \log N$ parallel time, sorting takes only constant parallel time in the MPC model [Goo99, GSZ11]. The second tool is indexing/predecessor search [GSZ11], which also has a constant parallel time in MPC. Furthermore, these two tools are fully scalable, and hence all the subroutines built on these two tools are also fully scalable. See Section 32.7 for how to use these two tools to implement the MPC operations needed for our algorithms.

Graph Connectivity: A natural approach to the graph connectivity problem is via the classic primitive of contracting to leaders: select a number of leader verteces, and contract every vertex (or most vertices) to a leader from its connected component (this is usually implemented by labeling the vertex by the corresponding leader). Indeed, many previous works (see e.g. [KSV10, RMCS13, $\left.\mathrm{KLM}^{+} 14 \mathrm{~b}\right]$ ) are based on this approach. There are two general questions to address in this approach: 1) how to choose leader vertices, and 2) how to label each vertex by its leader. For example, the algorithm in [KSV10] randomly chooses half of the vertices as leaders, and then contracts each non-leader vertex to one of its neighbor leader vertex. Thus, in each round of their algorithm, the number of vertices drops by a constant fraction. At the same time, half of the vertices are leaders, and hence their algorithm still needs at least $\Omega(\log n)$ rounds to contract all the vertices to one leader. Note
that a constant fraction of leaders is needed to ensure that there is a constant fraction of non-leader vertices who are adjacent to at least one leader vertex and hence are contracted. This leader selection method appears optimal for some graphs, e.g. path graphs.

To improve the runtime to $\ll \log n$, one would have to choose a much smaller fraction of the vertices to be leaders. Indeed, for a graph where every vertex has a large degree, say at least $d \gg \log n$, we can choose fewer leaders: namely, we can choose each vertex to be a leader with probability $p=\Theta((\log n) / d)$. Then the number of leaders will be about $\widetilde{O}(n / d)$, while each non-leader vertex has at least one leader neighbor with high probability. After contracting non-leader vertices to leader vertices, the number of remaining vertices is only a $1 / d$ fraction of original number of vertices.

By the above discussion, the goal would now be to modify our input graph $G$ so that every vertex has a uniformly large degree, without affecting the connectivity of the graph. An obvious such modification is to add edges between pairs of vertices that are already in the same connected component. In particular, if a vertex $v$ learns of a large number of vertices which are in the same connected component as $v$, then we can add edges between $v$ and those vertices to increase the degree of $v$. A naïve way to implement the latter is via broadcasting: each vertex $v$ first initializes a set $S_{v}$ which contains all the neighbors of $v$, and then, in each round, every vertex $v$ updates the set $S_{v}$ by adding the union of the sets $S_{u}$ over all neighbors $u$ of $v$ (old and new). This approach takes log-diameter number of rounds, and each vertex learns all vertices which are in the same connected component at the end of the procedure. However, in a single round, the total communication needed may be as huge as $\Omega\left(n^{3}\right)$ since each of $n$ vertices may have $\Omega(n)$ neighbors, each with a set of size $\Omega(n)$.

Since our goal of each vertex $v$ is to learn only $d$ vertices in the same component (not
necessarily the entire component), we can therefore implement a "truncated" version of the above broadcasting procedure:

1. If $S_{v}$ already had size $d$, then we do not need any further operation for $S_{v}$.
2. If $u$ is in $S_{v}$, and $S_{u}$ already has $d$ vertices, then we can just put all the elements from $S_{u}$ into $S_{v}$ and thus $S_{v}$ becomes of size $d$.
3. If $\left|S_{v}\right|<d$, and for every $u \in S_{v}$, the set $S_{u}$ is also smaller than $d$, then we can implement one step of the broadcasting - add the union of $S_{u}$ 's, for all neighbors $u \in S_{v}$, to $S_{v}$.

In the above procedure, if the number of vertices in $S_{v}$ is smaller than $d$ after the $i^{\text {th }}$ round, then we expect $S_{v}$ to contain all the vertices whose distance to $v$ is at most $2^{i}$. Thus, the above procedure also takes at most log-diameter rounds. Furthermore, the total communication needed is at most $O\left(n \cdot d^{2}\right)$.

Our full graph connectivity algorithm implements the above "truncated broadcasting" procedure iteratively, for values $d$ that follow a certain "schedule", depending on the available space. At the beginning of the algorithm, we have an $n$ vertex graph $G$ with diameter $D$, and a total of $\Omega(m)$ space. The algorithm proceeds in phases, where each phase takes $O(\log D)$ rounds of communication. In the first phase, the starting number of vertices is $n_{1}=n$. We implement a truncated broadcasting procedure where the target degree $d$ is $d_{1}=\left(m / n_{1}\right)^{1 / 2}$, using $O(\log D)$ rounds and $O(m)$ total space. Then we can randomly select $\widetilde{O}\left(n_{1} / d_{1}\right)$ leaders, and contract all the non-leader vertices to leader vertices. At the end of the first phase, the total number of remaining vertices is at most $n_{2}=\widetilde{O}\left(n_{1} / d_{1}\right)=\widetilde{O}\left(n_{1}^{1.5} / m^{0.5}\right)$. In general,
suppose, at the beginning of the $i^{\text {th }}$ phase, the number of remaining vertices is $n_{i}$. Then we use the truncated broadcasting procedure for value $d$ set to $d_{i}=\left(m / n_{i}\right)^{1 / 2}$, thus making each vertex have degree at least $d_{i}=\left(m / n_{i}\right)^{1 / 2}$ in $O(\log D)$ number of communication rounds and $O(m)$ total space. Then we choose $\widetilde{O}\left(n_{i} / d_{i}\right)$ leaders, and, after contracting non-leaders, the number $n_{i+1}$ of remaining vertices is at most $\widetilde{O}\left(n_{i}^{1.5} / m^{0.5}\right)$. Let us look at the progress of the value $d_{i}$. We have that $d_{i+1}=\widetilde{\Omega}\left(\left(m / n_{i+1}\right)^{1 / 2}\right)=\widetilde{\Omega}\left(\left(m^{1.5} / n_{i}^{1.5}\right)^{1 / 2}\right)=\widetilde{\Omega}\left(d_{i}^{1.5}\right)$. Thus, we are making double exponential progress on $d_{i}$, which implies that the total number of phases needed is at most $O\left(\log \log _{m / n} n\right)$, and the total parallel time is thus $O\left(\log D \cdot \log \log _{m / n} n\right)$.

This technique of double-exponential progress is more general and extends to other problems beyond connectivity. In particular, for a problem, suppose its size is characterized by a parameter $n$ (not necessarily the input size - e.g. in connectivity problem, $n$ is the number of vertices). When $n$ is a constant, the problem can be solved in $O(1)$ parallel time. If there is a procedure that uses total space $\Theta(m)$ to reduce the problem size to at most $n / k$ for $k=(m / n)^{c}, c=\Omega(1)$, then we can repeat the procedure $O\left(\log \log _{m / n} n\right)$ times to solve the overall problem. In particular, after repeating the procedure $i$ times, the problem size is $n_{i} \leq n_{i-1} /\left(m / n_{i-1}\right)^{c} \leq n \cdot(n / m)^{(1+c)^{i}-1}$. We call this technique double-exponential speed problem size reduction.

Remark 32.1.2. For any problem characterized by a size parameter $n$, if we can use parallel time $T$ and total space $\Theta(m)$ to reduce the problem size such that the reduced problem size is $n / k$ for $k=(m / n)^{\Omega(1)}$, then we can solve the problem in $O(m)$ total space and $O\left(T \cdot \log \log _{m / n} n\right)$ parallel time.

Spanning Forest and Diameter Estimator: Extending a connectivity algorithm to a spanning forest algorithm is usually straightforward. For example, in [KSV10], they only contract a non-leader vertex to an adjacent leader vertex, thus their algorithm can also give a spanning forest, using the contracted edges. Here however, extending our connectivity algorithm to a spanning forest algorithm requires several new ideas. In our connectivity algorithm, because of the added edges, we only ensure that when a vertex $u$ is contracted to a vertex $v, u$ and $v$ must be in the same connected component; but $u$ and $v$ may not be adjacent in the original graph. Thus, we need to record more information to help us build a spanning forest.

We can represent a forest as a collection of parent pointers part $(v)$, one for each vertex $v \in V$. If $v$ is a root in the forest, then we let $\operatorname{part}(v)=v$. We use $\operatorname{dep}_{\mathrm{par}}(v)$ to denote the depth of $v$ in the forest, i.e. $\operatorname{dep}_{\mathrm{par}}(v)$ is the distance from $v$ to its root. Let $\operatorname{dist}_{G}(u, v)$ denote the distance between two vertices $u$ and $v$ in a graph $G$.

Our connectivity algorithm uses the "neighbor increment" procedure described above. We observed that if the set $S_{v}$ has fewer than $d$ vertices after the $i^{\text {th }}$ round, then $S_{v}$ should contain all the vertices with distance at most $2^{i}$ to $v$. This motivates us to maintain a shortest path tree for $S_{v}$, with root $v$. In the $i^{\text {th }}$ round, if we need to update $S_{v}$ to be $\bigcup_{u \in S_{v}} S_{u}$, then we can update the shortest path tree of $S_{v}$ in the following way:

1. For each $x \in S_{u}$ for some $u \in S_{v}$, we can create a tuple $(x, u)$.
2. Then, for each $x \in\left(\bigcup_{u \in S_{v}} S_{u}\right) \backslash S_{v}$, we can sort all the tuples $\left(x, u_{1}\right),\left(x, u_{2}\right), \cdots,\left(x, u_{k}\right)$ such that $u_{1}$ minimizes $\min _{u \in S_{v}} \operatorname{dist}_{G}(v, u)+\operatorname{dist}_{G}(u, x)$. Since $u$ is in $S_{v}, x$ is in $S_{u}$, it is easy to get the value of $\operatorname{dist}_{G}(v, u), \operatorname{dist}_{G}(u, x)$ by the information of shortest path
tree for $S_{v}$ and $S_{u}$. Then we set the new parent of $x$ in the shortest path tree for $S_{v}$ to be the parent of $x$ in the shortest path tree for $S_{u_{1}}$.

Since $S_{v}$ before the update contains all the vertices which have distance to $v$ at most $2^{i-1}$, the union of the shortest path from $x$ to $u_{1}$ and the shortest path from $u_{1}$ to $v$ must be the shortest path from $x$ to $v$. Then by induction, we can show that the parent of $x$ in the shortest path tree for $S_{u_{1}}$ is also the parent of $x$ in the shortest path tree for updated $S_{v}$. Thus, this modified "neighbor increment" procedure can find $n$ local shortest path trees where there is a tree with root $v$ for each vertex $v$. Furthermore, the procedure still takes $O(\log D)$ rounds. And we can still use $O\left(n d^{2}\right)$ total space to make each shortest path tree have size at least $d$. Next, we show how to use these $n$ local shortest path trees to construct a forest with the roots in the forest being the leaders.

As discussed in the connectivity algorithm, if every local shortest path tree has size at least $d$, we can choose each vertex as a leader with probability $p=\Theta((\log n) / d)$ and then every tree will contain at least one leader with high probability. Let $L$ be the set of sampled leaders, and let $\operatorname{dist}_{G}(v, L)$ be defined as $\min _{u \in L} \operatorname{dist}_{G}(v, u)$. Let $v$ be a non-leader vertex, i.e. $v \in V \backslash L$. According to the shortest path tree for $S_{v}{ }^{3}$, since $L \cap S_{v} \neq \emptyset$, we can find a child $u$ of the root $v$ such that $\operatorname{dist}_{G}(v, L)>\operatorname{dist}_{G}(u, L)$; in this case we set part $(v)=u$. For vertex $v \in L$, we can set $\operatorname{part}(v)=v$. We can see now that par denotes a rooted forest where the roots are sampled leaders. Furthermore, since $\forall v \notin L,(v, \operatorname{part}(v))$ is from the shortest path tree for $S_{v}$, we know that $v$ and $\operatorname{part}(v)$ are adjacent in the original graph $G$. After doing

[^110]the above for all nodes $v$, the forest denoted by the resulting vector par must be a subgraph of the spanning forest of $G$. We then apply the standard doubling algorithm to contract all the vertices to their leaders (roots), in $O(\log D)$ rounds. Therefore, the problem is reduced to finding a spanning forest in the contracted graph. The number of vertices remaining in the contracted graph is at most $\widetilde{O}(n / d)$, where $d=(m / n)^{\Theta(1)}$. By Remark 32.1.2, we can output a spanning forest in $O\left(\log D \cdot \log \log _{m / n} n\right)$ parallel time.

Although the above algorithm can output the edges of a spanning forest, it cannot output a rooted spanning forest. To output a rooted spanning forest, we follow a top-down construction. Suppose now we have a rooted spanning forest of the contracted graph. Since we have all the information of how vertices were contracted, we know the contraction trees in the original graph. To merge these contraction trees into the rooted spanning forest of the contracted graph, we only need to change the root of each contraction tree to a proper vertex in that tree. This changing root operation can be implemented by the doubling algorithm via a divide-and-conquer approach.

Since the spanning forest algorithm needs $O\left(\log \log _{m / n} n\right)$ phases to contract all vertices to a single vertex, the total parallel time to compute a rooted spanning forest is $O\left(\log D \cdot \log \log _{m / n} n\right)$. Furthermore, the depth of the rooted spanning forest will be at most $O\left(D^{O\left(\log \log _{m / n} n\right)}\right)$. Thus, we can use the doubling algorithm to calculate the depth of the tree, and output this depth as an estimator of the diameter of the input graph.

Depth-First-Search Sequence: Here, when the input graph $G$ is a tree, our goal is to output a DFS sequence for this tree. Once we have this sequence, it is easy to output a rooted tree. Thus, computing a DFS sequence is at least as hard as computing a rooted tree,
and all the previous algorithms need $\Omega(\log n)$ parallel time to do so.

First of all, we use our spanning forest algorithm to compute a rooted tree, reducing the problem to computing a DFS sequence for a rooted tree. The idea is motivated by TeraSort [O'M08]. If the size of the tree is small enough such that it can be handled by a single machine, then we can just use a single machine to generate its DFS sequence. Otherwise, our algorithm can be roughly described as follows. (Recall that $\delta$ is the parameter such that each machine has $\Theta\left(n^{\delta}\right)$ local memory.)

1. Sample $n^{\delta / 2}$ leaves $l_{1}, l_{2}, \cdots, l_{s}$.
2. Determine the order of sampled leaves in the DFS sequence.
3. Compute the DFS sequence $\widetilde{A}$ of the tree which only consists of sampled leaves and their ancestors.
4. Compute the DFS sequence $A_{v}$ of every root- $v$ subtree which does not contain any sampled leaf.
5. Merge $\widetilde{A}$ and all the $A_{v}$.

The first and second steps go as follows. Since we only sample $n^{\delta / 2}$ leaves, we can send them to a single machine. We generate queries for every pair of sampled leaves where each query $\left(l_{i}, l_{j}\right)$ queries the lowest common ancestor of $\left(l_{i}, l_{j}\right)$. We have $n^{\delta}$ such queries in total. Since the input tree is rooted, we can use a doubling algorithm to preprocess a data structure in $O(\log D)$ parallel time and answer all the queries simultaneously in $O(\log D)$ parallel time. Thus, we know the lowest common ancestor of any pair of sampled leaves,
and we can store this all on a single machine. Based on the information of lowest common ancestors of each pair of sampled leaves, we are able to determine the order of the leaves.

For the third step, suppose the sampled leaves have order $l_{1}, l_{2}, \cdots, l_{s}$. Let $v$ be the root of the tree. Then the DFS sequence $\widetilde{A}$ should be: the path from $v$ to $l_{1}$, the path from $l_{1}$ to the lowest common ancestor of $\left(l_{1}, l_{2}\right)$, the path from the lowest common ancestor of $\left(l_{1}, l_{2}\right)$ to $l_{2}$, the path from $l_{2}$ to the lowest common ancestor of $\left(l_{2}, l_{3}\right), \ldots$, the path from $l_{s}$ to $v$. We can find these paths simultaneously by a doubling algorithm together with a divide-and-conquer algorithm in $O(\log D)$ parallel time.

In the fourth step, we apply the procedure recursively. Suppose the total number of leaves in the tree is $q \leq n$. Since we randomly sampled $n^{\delta / 2}$ number of leaves, with high probability, each subtree which does not contain a sampled leaf will have at most $O\left(q / n^{\delta / 2}\right)$ number of leaves. Thus, the depth of the recursion will be at most a constant, $O(1 / \delta)$.

Minimum Spanning Forest and Bottleneck Spanning Forest. Recall that the input is a graph $G=\left(V, E=\left(e_{1}, e_{2}, \cdots, e_{m}\right)\right)$ together with a weight function $w$ on $E$. Without loss of generality, we only consider the case when all the weights of edges are different, i.e. $w\left(e_{1}\right)<w\left(e_{2}\right)<\cdots<w\left(e_{m}\right)$. Since the weights of edges are different, the minimum spanning forest of the graph is unique. By Kruskal's algorithm, the diameter of the graph induced by the first $i$ edges for any $i \in[m]$ is at most the depth of the minimum spanning forest. Now, let us use $D$ to denote the depth of the minimum spanning forest.

We first discuss the minimum spanning forest algorithm. A crucial observation of Kruskal's algorithm is: if we want to determine which edges in $e_{i}, e_{i+1}, \cdots, e_{j}$ are in the minimum spanning forest, we can always contract the first $i-1$ edges to obtain a graph $G^{\prime}$,
run a minimum spanning forest algorithm on the contracted graph $G^{\prime}$, and observe whether an edge is included in the spanning forest of $G^{\prime}$. Thus, if the total space is $\Theta\left(m^{1+\gamma}\right)$, we can have $m^{\gamma}$ copies of the graph, where the $i^{\text {th }}$ copy contracts the first $(i-1) \cdot m^{1-\gamma}$ edges. Thus, we are able to divide the edges into $m^{\gamma}$ groups where each group has $m^{1-\gamma}$ number of edges. We only need to solve the minimum spanning forest problem for each group. Then in the second phase, we can divide the edges into $m^{2 \gamma}$ groups where each group has $m^{1-2 \gamma}$ number of edges. Thus, the total number of phases needed is at most $O(1 / \gamma)$. In each phase, we just need to run our connectivity algorithm to contract the graph.

For the approximate minimum spanning forest algorithm, we use a similar idea. If we want a $(1+\epsilon)$ approximation, then we round each weight to the closest value $(1+\epsilon)^{i}$ for some integer $i$. After rounding, there are only $O(1 / \epsilon \cdot \log n)$ edge groups. Since our total space is at least $\Omega(m \log (n) / \epsilon)$, we can make $O(1 / \epsilon \cdot \log n)$ copies of the graph. The $i^{\text {th }}$ copy of the graph contracts all the edges in group $1,2, \cdots, i-1$. Then, we only need to run our spanning forest algorithm on each copy to determine which edges should be chosen in each group.

Another application of our double exponential speed problem size reduction technique is bottleneck spanning forest. For the bottleneck spanning forest, suppose we have $\Theta(k m)$ total space. We can have $k$ copies of the graph where the $i^{\text {th }}$ copy contracts the first $(i-1) \cdot m / k$ number of edges. We can determine the group of $O(m / k)$ edges which contains the bottleneck edge. Thus, we reduce the problem to $O(m / k)$. According to Remark 32.1.2, the number of phases is at most $O\left(\log _{\log }^{k} 2 m\right.$, and each phase needs $T$ parallel time, where $T$ is the parallel time for spanning forest.

Directed Reachability vs. Boolean Matrix Multiplication If there is a fully scalable multi-query directed reachability MPC algorithm with almost linear total space, we can simulate the algorithm in sequential model. Thus, it will imply a good sequential multi-query directed reachability algorithm which implies a good sequential Boolean Matrix Multiplication algorithm.

### 32.1.4 Roadmap

The rest of the paper contains the technical details of our algorithms. In Section 32.2, we described a simplified connectivity algorithm. In Section 32.3, we described the notations. In Sections 32.4, 32.5, and 32.6, we give the details of our main algorithms for connectivity, spanning forest and depth first search sequence. In these sections, we focus on the design of the algorithms and the analysis of the number of rounds. In Section 32.7, we describe the MPC model in detail and discuss some known primitives in that model. In Section 32.8, we discuss how to implement the details of our algorithms in the MPC model to achieve the bounds claimed in the previous sections. In Section 32.9, we show how to apply our connectivity and spanning forest algorithm in minimum spanning forest and bottleneck spanning forest problems. In Section 32.11, we show hard instances for the algorithm [RMCS13]. In Section 32.12, we show an alternative approach for random leader selection.

### 32.2 A Simplified Batch Algorithm for Connectivity

In this section, we show a simplified version of our connectivity algorithm.
Firstly, let us describe the simplified version of truncated broadcasting procedure in the following. Since $G^{\prime}$ is obtained by adding edges between the vertices in the same component of $G$. $G^{\prime}$ will preserve the connectivity of $G$. The parallel time needed is at most $O(\log D)$ where $D$ is the diameter of $G$. The procedure takes at most $O\left(n d^{2}+m\right)$ total space.

Truncated Broadcasting for Neighbor Increment:

- Input:
- A graph $G=(V, E)$ with $n=|V|$ vertices and $m=|E|$ number of edges.
- A parameter $d$.
- Output:
- A graph $G^{\prime}=\left(V, E^{\prime}\right)$ such that $\forall v \in V,|\Gamma(v)| \geq d$. $\triangleright \Gamma(v)$ denotes the neighbors of $v$.
- While $\exists x \in V$ such that $|\Gamma(x)|<d$ :
- For each $v \in V$ with $|\Gamma(v)|<d$ :
* If $\exists u \in \Gamma(v)$ which has $|\Gamma(u)| \geq d$, then $\Gamma(v) \leftarrow \Gamma(v) \cup \Gamma(u)$.
* Otherwise, $\Gamma(v) \leftarrow \Gamma(v) \cup \bigcup_{u \in \Gamma(v)} \Gamma(u)$.

We can apply the above procedure to make each vertex have a large degree. Next, let us briefly describe how to choose the leaders and implement the contraction operation for the graph where each vertex has a large degree. The following procedure just needs $O(n+m)$ total space and $O(1)$ parallel time. If every vertex has degree at least $d$, then in the following
procedure we can reduce the number of vertices to $\widetilde{O}(n / d)$ by contracting all the vertices to $\widetilde{O}(n / d)$ number of leaders.

Random Leader Contraction:

- Input:
- A graph $G=(V, E)$ with $n=|V|$ vertices where each vertex has degree at least $d$.


## - Output:

- A graph $G^{\prime}=\left(V^{\prime}, E^{\prime}\right)$ with $\widetilde{O}(n / d)$ vertices.
- A mapping par : $V \rightarrow V^{\prime}$, such that part $(v)$ is the vertex that $v$ contracts to.


## - Leader Selection:

- Let $L$ denote the set of leaders.
- For each $v \in V$, with probability at least $\widetilde{\Omega}(1 / d)$, choose $v$ as a leader, i.e. $L \leftarrow$ $L \cup\{v\}$.
- Contraction:
- For each $v \in L$, let $\operatorname{part}(v)=v$, and put $v$ into $V^{\prime}$.
- For each $v \in V \backslash L$, choose $u \in \Gamma(v) \cap L$, and set $\operatorname{part}(u)=v$.
- For each $(u, v) \in E$, if $\operatorname{part}(u) \neq \operatorname{part}(v)$, put the edge $(\operatorname{part}(u), \operatorname{part}(v))$ into $E^{\prime}$.

Finally, we describe the simplified version of our connectivity algorithm in the following.

## Connectivity:

- Input:
- A graph $G=(V, E)$ with $n=|V|$ vertices and $m=|E|$ edges.
- Total space $N$ which is $\Theta(m)$.
- Output:
- A mapping col : $V \rightarrow V$ satisfies $\forall u, v \in V, \operatorname{col}(u)=\operatorname{col}(v)$ if and only if $u$ and $v$ are connected.
- Initialization:
- Let $G_{0} \leftarrow G, n_{0} \leftarrow n$.
- In phase $i$ :
- Compute $G_{i-1}^{\prime}$ : Increase the degree of every vertex in $G_{i-1}$ to at least $d_{i}=$ $\Theta\left(\left(N / n_{i-1}\right)^{1 / 2}\right)$.
- Compute $G_{i}$ : Select $n_{i}=\widetilde{O}\left(n_{i-1} / d_{i}\right)$ leaders in $G_{i-1}^{\prime}$ and contract all the vertices to the leaders.
- If $v$ is contracted to $u$, record $\operatorname{part}(v)=u$.
- If $G_{i}$ does not have any edges, then for every vertex $v$ in $G_{i}$, set $\operatorname{part}(v)=v$, and exit the loop.
- Finding the root leader:
- For each $v \in V$, find the root of $v$ in par, i.e. find $u=\operatorname{part}(\operatorname{part}(\cdots \operatorname{part}(v)))$ such that $\operatorname{part}(u)=u$.
$-\operatorname{Set} \operatorname{col}(v)=u$.

After phase $i$, the number of vertices survived is at most $\widetilde{O}\left(n_{i-1} /\left(N / n_{i-1}\right)^{1 / 2}\right)$. By Remark 32.1.2, there will be at most $O\left(\log \log _{N / n} n\right)$ phases. For phase $i$, we need $O(\log D)$ parallel time to increase the neighbors of every vertex in $G_{i-1}$. The total parallel time is thus
$O\left(\log D \cdot \log \log _{N / n} n\right)$. The total space used in phase $i$ is at most $O\left(m+\left(N / n_{i-1}\right)^{1 / 2} \cdot n_{i-1}\right)=$ $O(N)$.

### 32.3 Notations

$[n]$ denotes the set $\{1,2, \cdots, n\}$. Let $G$ be an undirected graph with vertex set $V$ and edge set $E$. For $v \in V, \Gamma_{G}(v)$ denotes the set of neighbors of $v$ in $G$, i.e. $\Gamma_{G}(v)=\{u \in V \mid$ $(v, u) \in E\}$. For any $u, v \in V, \operatorname{dist}_{G}(u, v)$ denotes the distance between $u, v$ in graph $G$. If $u, v$ are not in the same connected component, then $\operatorname{dist}_{G}(u, v)=\infty$. If $u, v$ are in the same connected component, then $\operatorname{dist}_{G}(u, v)<\infty$. For $v \in V,\left\{u \in V \mid \operatorname{dist}_{G}(u, v)<\infty\right\}$ is the set of all the vertices in the same connected component as $v$. The diameter $\operatorname{diam}(G)$ of $G$ is the largest diameter of its components, i.e. $\operatorname{diam}(G)=\max _{u, v \in V: \operatorname{dist}_{G}(u, v)<\infty} \operatorname{dist}_{G}(u, v)$.

### 32.4 Graph Connectivity

### 32.4.1 Neighbor Increment Operation

In this section, we describe a procedure which can increase the number of neighbors of every vertex and preserve the connectivity at the same time. The input of the procedure is an undirected graph $G=(V, E)$ and a parameter $m$ which is larger than $|V|$. The output is a graph $G^{\prime}=\left(V, E^{\prime}\right)$ such that for each vertex $v$, either the connected component which contains $v$ is a clique or $v$ has at least $\left\lceil(m /|V|)^{1 / 2}\right\rceil-1$ neighbors. Furthermore, $\left|E^{\prime}\right| \leq|E|+m$. We use $\Gamma_{G}(v)$ to denote the neighbors of $v$ in graph $G$, i.e. $\Gamma_{G}(v)=\{u \in V \mid(u, v) \in E\}$. Similarly, we let $\Gamma_{G^{\prime}}(v)$ be the neighbors of $v$ in $G^{\prime}$, i.e. $\Gamma_{G^{\prime}}(v)=\left\{u \in V \mid(u, v) \in E^{\prime}\right\}$.

Lemma 32.4.1. Let $G=(V, E)$ be an undirected graph, $m \in \mathbb{Z}_{\geq 0}$ which has $m \geq 4|V|$. Let $n=|V|$. Let $r$ be the value at the end of the procedure $\operatorname{Neighborincrement~}(m, G)$ (Algorithm 32.1.) Then $\forall i \in\{0,1, \cdots, r\}, v \in V, S_{v}^{(i)}$ satisfies the following properties:

1. $v \in S_{v}^{(i)}$.
2. $\forall u \in S_{v}^{(i)}, \operatorname{dist}_{G}(u, v)<\infty$.
3. $\left|S_{v}^{(i)}\right|<\left\lceil(m / n)^{1 / 2}\right\rceil \Rightarrow S_{v}^{(i)}=\left\{u \in V \mid \operatorname{dist}_{G}(u, v) \leq 2^{i}\right\}$.
4. $\left|S_{v}^{(i)}\right| \leq m / n$.

Proof. For property 1, we can prove it by induction. When $i=0$, due to line 3, we know $v \in S_{v}^{(0)}$. Suppose property 1 holds for $S_{v}^{(i-1)}$ for all $v \in V$. If $S_{v}^{(i)}$ is updated by line 17 , there are two cases: 1. if $u=v$, then $v \in S_{u}^{(i-1)}$, and the condition of line 17 does not hold, thus $v \in S_{v}^{(i)} ; 2$. if $u \neq v$, then after implementing line 17, $v$ will not be removed, thus $v \in S_{v}^{(i)}$.

If $S_{v}^{(i)}$ is updated by line 20, then since $v \in S_{v}^{(i-1)}, v$ is also in the set $S_{v}^{(i)}$. Thus, property 1 holds for every $S_{v}^{(i)}$.

For property 2, we can prove it by induction. When $i=0$, it is easy to see $S_{v}^{(0)} \subseteq$ $\Gamma_{G}(v) \cup\{v\}$, thus property 2 holds for it. Suppose property 2 holds for $S_{v}^{(i-1)}$ for all $v \in V$. If $S_{v}^{(i)}$ is updated by line 17 , then since $u \in S_{v}^{(i-1)}$ and $S_{v}^{(i)} \subseteq S_{u}^{(i-1)} \cup\{v\}$, all the vertices from $S_{v}^{(i)}$ are in the same connected component as $u$ and $u$ is in the same connected component as $v$. Thus, property 2 holds in this case. If $S_{v}^{(i)}$ is updated by the line 20 , then $\forall p \in S_{v}^{(i)}$, there exists $u \in S_{v}^{(i-1)}$ such that $p \in S_{u}^{(i-1)}$. We have $p$ is in the same connected component as $u$, and $u$ is in the same connected component as $v$. Thus, property 2 also holds in this case.

For property 3, we can prove it by induction. When $i=0$, due to line 7 , we have $\left|S_{v}^{(0)}\right|<\left\lceil(m / n)^{1 / 2}\right\rceil \rightarrow S_{v}^{(0)}=\Gamma_{G}(v) \cup\{v\}=\left\{u \in V \mid \operatorname{dist}_{G}(u, v) \leq 1\right\}$. Suppose property 3 holds for $S_{v}^{(i-1)}$ for all $v \in V$. Since if $\left|S_{v}^{(i)}\right|<\left\lceil(m / n)^{1 / 2}\right\rceil$, then $S_{v}^{(i)}$ can only be updated by line 20, and $\forall u \in S_{v}^{(i-1)}$, it has $\left|S_{u}^{(i-1)}\right|<\left\lceil(m / n)^{1 / 2}\right\rceil$. Thus, $S_{v}^{(i)}=\bigcup_{u \in S_{v}^{(i-1)}} S_{u}^{(i-1)}=$ $\bigcup_{u \in V, \operatorname{dist}_{G}(u, v) \leq 2^{i-1}}\left\{p \in V \mid \operatorname{dist}_{G}(p, u) \leq 2^{i-1}\right\}=\left\{u \in V \mid \operatorname{dist}_{G}(u, v) \leq 2^{i}\right\}$. Thus, property 3 holds.

For property 4 , we can prove it by induction. When $i=0$, due to line $7, \forall v \in V$, we have $\left|S_{v}^{(0)}\right| \leq\left\lceil(m / n)^{1 / 2}\right\rceil \leq m / n$, where the last inequality follows by $m / n \geq 4$. Now suppose property 4 holds for $S_{v}^{(i-1)}$ for all $v \in V$. If $S_{v}^{(i)}$ is updated by line 17 , then $\left|S_{v}^{(i)}\right|=\left|S_{u}^{(i-1)}\right| \leq$ $m / n$. If $S_{v}^{(i)}$ is updated by line 20, we know $\forall u \in S_{v}^{(i-1)},\left|S_{u}^{(i-1)}\right|<\left\lceil(m / n)^{1 / 2}\right\rceil$. Notice that by property $1, v \in S_{v}^{(i-1)}$, so $\left|S_{v}^{(i-1)}\right|<\left\lceil(m / n)^{1 / 2}\right\rceil$. Thus, $\left|S_{v}^{(i)}\right|=\left|\bigcup_{u \in S_{v}^{(i-1)}} S_{u}^{(i-1)}\right| \leq$ $(m / n)^{1 / 2} \cdot(m / n)^{1 / 2} \leq m / n$.

```
Algorithm 32.1 Neighbor Increment Operation
    procedure NeighborIncrement \((m, G=(V, E)) \quad \triangleright\) Lemma 32.4.1, Lemma 32.4.2
                                    \(\triangleright\) Output: \(G^{\prime}=\left(V, E^{\prime}\right)\)
        Initially, \(n=|V|, E^{\prime}=\emptyset\) and let \(S_{v}^{(0)}=\{v\}\) for all \(v \in V\).
        for \(v \in V\) do \(\quad \triangleright\) Initially, let \(S_{v}^{(0)}\) be the set (or subset) of direct neighbors
            for \(u \in \Gamma_{G}(v)\) do
                if \(\left|S_{v}^{(0)}\right|<\left\lceil(m / n)^{1 / 2}\right\rceil\) then
                    \(S_{v}^{(0)} \leftarrow S_{v}^{(0)} \cup\{u\}\).
                end if
            end for
        end for
        \(r \leftarrow 1\).
        for true do
            for \(v \in V\) do
                if \(\exists u \in S_{v}^{(r-1)},\left|S_{u}^{(r-1)}\right| \geq\left\lceil(m / n)^{1 / 2}\right\rceil\) then \(\triangleright\) neighbor \(u\) has many neighbors
                \(S_{v}^{(r)}=S_{u}^{(r-1)} \cup\{v\}\).
                if \(\left|S_{v}^{(r)}\right|>\left|S_{u}^{(r-1)}\right|\) then
                    \(S_{v}^{(r)} \leftarrow S_{v}^{(r)} \backslash\{u\}\).
                end if
                else
                \(S_{v}^{(r)}=\bigcup_{u \in S_{v}^{(r-1)}} S_{u}^{(r-1)} . \quad \triangleright\) neighbors of \(v\) 's neighbors are \(v\) 's new
    neighbors.
                end if
            end for
                                    \(\triangleright S_{v}^{(r)}\) is large or is a component
            if \(\forall v \in V\), either \(\left|S_{v}^{(r)}\right| \geq\left\lceil(m / n)^{1 / 2}\right\rceil\) or \(\left|S_{v}^{(r)}\right|=\left|S_{v}^{(r-1)}\right|\) then
                Let \(E^{\prime}=E \cup \bigcup_{v \in V}\left\{(v, u) \mid v \in S_{u}^{(r)}\right.\) or \(\left.u \in S_{v}^{(r)}, u \neq v\right\}\).
                return \(G^{\prime}=\left(V, E^{\prime}\right)\)
            else
                \(r \leftarrow r+1\).
            end if
        end for
    end procedure
```

The following definition defines the number of iterations of Algorithm 32.1.

Definition 32.4.1. Given an undirected graph $G=(V, E)$ and a parameter $m \in \mathbb{Z}_{\geq 0}, m \geq$ $4|V|$, the number of iterations of NeighborIncrement $(m, G)$ (Algorithm 32.1) is the value of $r$ at the end of the procedure.

In the following lemma, we characterize the properties of Algorithm 32.1.

Lemma 32.4.2. Let $G=(V, E)$ be an undirected graph, $m \in \mathbb{Z}_{\geq 0}$ which has $m \geq 4|V|$. Let $G^{\prime}=\left(V, E^{\prime}\right)$ be the output of Neighborincrement $(m, G)$. We have:

1. The number of iterations (Definition 32.4.1), $r \leq \min (\lceil\log (\operatorname{diam}(G))\rceil,\lceil\log (m / n)\rceil)+1$.
2. For all $u, v \in V$, $\operatorname{dist}_{G}(u, v)<\infty \Leftrightarrow \operatorname{dist}_{G^{\prime}}(u, v)<\infty$.
3. $\forall v \in V$, if $\left|\Gamma_{G^{\prime}}(v)\right|<\left\lceil(m / n)^{1 / 2}\right\rceil-1$, then the connected component in $G^{\prime}$ which contains $v$ is a clique. It also implies that $\forall u, v \in V$, if $\left|\Gamma_{G^{\prime}}(v)\right|<\left\lceil(m / n)^{1 / 2}\right\rceil-1$ and $\left|\Gamma_{G^{\prime}}(u)\right| \geq\left\lceil(m / n)^{1 / 2}\right\rceil-1$, then $\operatorname{dist}_{G^{\prime}}(u, v)=\infty$.
4. $E \subseteq E^{\prime},\left|E^{\prime}\right| \leq|E|+m$.

Proof. For property 1, if $r>\lceil\log (\operatorname{diam}(G))\rceil+1$, then let $i=\lceil\log (\operatorname{diam}(G))\rceil+1$. Let $v \in V$. By property 3 of Lemma 32.4.1, if $\left|S_{v}^{(i)}\right|<\left\lceil(m / n)^{1 / 2}\right\rceil$, then $S_{v}^{(i)}=\left\{u \in V \mid \operatorname{dist}_{G}(u, v) \leq\right.$ $\left.2^{i}\right\}=\left\{u \in V \mid \operatorname{dist}_{G}(u, v) \leq 2 \cdot 2^{\lceil\log (\operatorname{diam}(G))\rceil}\right\}=\left\{u \in V \mid \operatorname{dist}_{G}(u, v)<\infty\right\}$. Furthermore, if $\left|S_{v}^{(i)}\right|<\left\lceil(m / n)^{1 / 2}\right\rceil$, then $\left|S_{v}^{(i-1)}\right|<\left\lceil(m / n)^{1 / 2}\right\rceil$, which means that $S_{v}^{(i-1)}=\{u \in V \mid$ $\left.\operatorname{dist}_{G}(u, v) \leq 2^{i-1}\right\}=\left\{u \in V \mid \operatorname{dist}_{G}(u, v) \leq 2^{[\log (\operatorname{diam}(G))\rceil}\right\}=\left\{u \in V \mid \operatorname{dist}_{G}(u, v)<\right.$
$\infty\}=S_{v}^{(i)}$. Then due to the condition in line 24 , it will end the procedure in this round, which contradicts to $r>\lceil\log (\operatorname{diam}(G))\rceil+1=i$. Similarly, if $r>\lceil\log (m / n)\rceil+1$, then let $i=\lceil\log (m / n)\rceil+1$. Furthermore, if $\left|S_{v}^{(i)}\right|<\left\lceil(m / n)^{1 / 2}\right\rceil$, then $\left|S_{v}^{(i-1)}\right|<\left\lceil(m / n)^{1 / 2}\right\rceil$, which means that $S_{v}^{(i-1)}=\left\{u \in V \mid \operatorname{dist}_{G}(u, v) \leq 2^{i-1}\right\} \neq\left\{u \in V \mid \operatorname{dist}_{G}(u, v) \leq 2^{i}\right\}=S_{v}^{(i)}$. Thus, there exists $u \in S_{v}^{(i)}$ such that $\left|S_{v}^{(i)}\right|>\operatorname{dist}_{G}(u, v)>2^{i-1}>m / n \geq\left\lceil(m / n)^{1 / 2}\right\rceil$ which leads to a contradiction.

For property 2 , if $u, v$ are in the same connected component in $G$, then since $E \subseteq$ $E^{\prime}, u, v$ are in the same connected component in $G^{\prime}$. If $u, v$ are in the same connected component in $G^{\prime}$, then there should be a path $u=u_{1} \rightarrow u_{2} \rightarrow \cdots \rightarrow u_{p}=v$ in $G^{\prime}$, i.e. $\forall j \in[p-1],\left(u_{j}, u_{j+1}\right) \in E^{\prime} .\left(u_{j}, u_{j+1}\right) \in E^{\prime}$ implies that either $\left(u_{j}, u_{j+1}\right) \in E$ or $u_{j} \in S_{u_{j+1}}^{(r)}$ or $u_{j+1} \in S_{u_{j}}^{(r)}$. By property 2 of Lemma 32.4.1, we know that $\forall j \in[p-1], u_{j}$ and $u_{j+1}$ are in the same connected component in $G$. Thus, $u$ and $v$ are in the same connected component in $G$.

For property 3, due to line 25, if $\left|\Gamma_{G^{\prime}}(v)\right|<\left\lceil(m / n)^{1 / 2}\right\rceil-1$, then we have $\left|S_{v}^{(r)}\right|<$ $\left\lceil(m / n)^{1 / 2}\right\rceil$. By property 3 of Lemma 32.4.1, and the condition in line 24, we know $\{u \in V \mid$ $\left.\operatorname{dist}_{G}(u, v) \leq 2^{r}\right\}=S_{v}^{(r)}=S_{v}^{(r-1)}=\left\{u \in V \mid \operatorname{dist}_{G}(u, v) \leq 2^{r-1}\right\}$. Thus, $S_{v}^{(r)}=\{u \in V \mid$ $\left.\operatorname{dist}_{G}(u, v)<\infty\right\}$. Due to property 2, we have $\Gamma_{G^{\prime}}(v) \cup\{v\} \subseteq\left\{u \in V \mid \operatorname{dist}_{G}(u, v)<\infty\right\}$. Notice that $S_{v}^{(r)} \subseteq \Gamma_{G^{\prime}}(v) \cup\{v\}$, thus, we have $\Gamma_{G^{\prime}}(v) \cup\{v\}=\left\{u \in V \mid \operatorname{dist}_{G}(u, v)<\infty\right\}$. Let $v^{\prime} \in\left\{u \in V \mid \operatorname{dist}_{G}(u, v)<\infty\right\}$, then due to property $2, \Gamma_{G^{\prime}}\left(v^{\prime}\right) \cup\left\{v^{\prime}\right\} \subseteq\{u \in V \mid$ $\left.\operatorname{dist}_{G}\left(u, v^{\prime}\right)<\infty\right\}=\left\{u \in V \mid \operatorname{dist}_{G}(u, v)<\infty\right\}$, then we have $\left|\Gamma_{G^{\prime}}\left(v^{\prime}\right) \cup\left\{v^{\prime}\right\}\right|<\left\lceil(m / n)^{1 / 2}\right\rceil$. Thus, $\left|S_{v^{\prime}}^{(r)}\right|<\left\lceil(m / n)^{1 / 2}\right\rceil$. By property 3 of Lemma 32.4.1, and the condition in line 24, we know $\left\{u \in V \mid \operatorname{dist}_{G}\left(u, v^{\prime}\right) \leq 2^{r}\right\}=S_{v^{\prime}}^{(r)}=S_{v^{\prime}}^{(r-1)}=\left\{u \in V \mid \operatorname{dist}_{G}\left(u, v^{\prime}\right) \leq 2^{r-1}\right\}$. Thus, $S_{v^{\prime}}^{(r)}=\left\{u \in V \mid \operatorname{dist}_{G}\left(u, v^{\prime}\right)<\infty\right\}$. Thus, $\Gamma_{G^{\prime}}\left(v^{\prime}\right) \cup\left\{v^{\prime}\right\}=\left\{u \in V \mid \operatorname{dist}_{G}\left(u, v^{\prime}\right)<\infty\right\}=$
$\left\{u \in V \mid \operatorname{dist}_{G}(u, v)<\infty\right\}$. Thus, $\forall p, q \in\left\{u \in V \mid \operatorname{dist}_{G}(u, v)<\infty\right\}$, we have $(p, q) \in E^{\prime}$, which means that $\left\{u \in V \mid \operatorname{dist}_{G}(u, v)<\infty\right\}$ is a clique in $G^{\prime}$.

Now consider two vertices $u, v \in V$. Suppose $\left|\Gamma_{G^{\prime}}(v)\right|<\left\lceil(m / n)^{1 / 2}\right\rceil-1$, then we have that $\left\{p \in V \mid \operatorname{dist}_{G}(p, v)<\infty\right\}$ is a clique in $G^{\prime}$. Thus, $\forall q \in\left\{p \in V \mid \operatorname{dist}_{G}(p, v)<\infty\right\}$, we have $\left|\Gamma_{G^{\prime}}(q)\right|=\left|\Gamma_{G^{\prime}}(v)\right|<\left\lceil(m / n)^{1 / 2}\right\rceil-1$. If $\left|\Gamma_{G^{\prime}}(u)\right| \geq\left\lceil(m / n)^{1 / 2}\right\rceil-1$, then $\operatorname{dist}_{G^{\prime}}(u, v)=\infty$.

For property 4 , by line 25 , we have $E \subseteq E^{\prime}$ and $\left|E^{\prime}\right| \leq|E|+\sum_{v \in V}\left|S_{v}^{(r)}\right| \leq|E|+n$. $m / n=|E|+m$ where the last inequality follows by the property 4 of Lemma 32.4.1.

### 32.4.2 Random Leader Selection

Given an undirected graph $G=(V, E)$, to design a connected component algorithm, a natural way is constantly contracting the vertices in the same component. One way to do the contraction is that we randomly choose some vertices as leaders, then contract non-leader vertices to the neighbor leader vertices.

In this section, we show that if $\forall v \in V$, the number of neighbors of $v$ is large enough, then we can just sample a small number of leaders such that for each non-leader vertex $v \in V$, there is at least one neighbor of $v$ which is chosen as a leader. A more generalized statement is stated in the following lemma.

Lemma 32.4.3. Let $V$ be a vertex set with $n$ vertices. Let $0<\gamma \leq n, \delta \in(0,1)$. For each $v \in V$, let $S_{v}$ be a subset of $V \backslash\{v\}$ with size at least $\gamma-1$. Let $l: V \rightarrow\{0,1\}$ be a random hash function such that $\forall v \in V, l(v)$ are i.i.d. Bernoulli random variables, i.e.

$$
l(v)= \begin{cases}1 & \text { with probability } p \\ 0 & \text { otherwise }\end{cases}
$$

If $p \geq \min ((10 \log (2 n / \delta)) / \gamma, 1)$, then, with probability at least $1-\delta$,

1. $\sum_{v \in V} l(v) \leq \frac{3}{2} p n$;
2. $\forall v \in V, \exists u \in S_{v} \cup\{v\}$ such that $l(u)=1$.

Proof. For a fixed vertex $v \in V$, we have

$$
\begin{aligned}
& \operatorname{Pr}\left(\sum_{u \in S_{v} \cup\{v\}}(\mathbb{E}(l(u))-l(u))>\frac{1}{2} \sum_{u \in S_{v} \cup\{v\}} \mathbb{E}(l(u))\right) \\
\leq & \exp \left(-\frac{\frac{1}{2}\left(\frac{1}{2} \sum_{u \in S_{v} \cup\{v\}} \mathbb{E}(l(u))\right)^{2}}{\sum_{u \in S_{v} \cup\{v\}} \mathbb{V}(l(u))+\frac{1}{3} \cdot 1 \cdot \frac{1}{2} \sum_{u \in S_{v} \cup\{v\}} \mathbb{E}(l(u))}\right) \\
\leq & \exp \left(-\frac{\frac{1}{2}\left(\frac{1}{2} \sum_{u \in S_{v} \cup\{v\}} \mathbb{E}(l(u))\right)^{2}}{\sum_{u \in S_{v} \cup\{v\}} \mathbb{E}(l(u))+\frac{1}{3} \cdot 1 \cdot \frac{1}{2} \sum_{u \in S_{v} \cup\{v\}} \mathbb{E}(l(u))}\right) \\
= & \exp \left(-\frac{3}{28} \cdot \sum_{u \in S_{v} \cup\{v\}} \mathbb{E}(l(u))\right)=\exp \left(-\frac{3}{28} \cdot p \cdot\left|S_{v} \cup\{v\}\right|\right) \leq \frac{\delta}{2 n},
\end{aligned}
$$

where the first inequality follows by Bernstein inequality and $|l(u)-E(l(u))| \leq 1$, the second inequality follows by $\mathbb{V}(l(u)) \leq \mathbb{E}\left(l^{2}(u)\right)=\mathbb{E}(l(u))$. The last inequality follows by $\left|S_{v} \cup\{v\}\right| \geq \gamma$, and $p \geq \min ((10 \log (2 n / \delta)) / \gamma, 1)$. Since $\frac{1}{2} \sum_{u \in S_{v} \cup\{v\}} \mathbb{E}(l(u)) \geq 1$, with probability at least $1-\delta /(2 n), \sum_{u \in S_{v} \cup\{v\}} l(v) \geq 1$. By taking union bound over all $S_{v}$, with probability at least $1-\delta / 2, \forall v \in V, \exists u \in S_{v} \cup\{v\}, l(u)=1$.

Similarly, we have

$$
\begin{aligned}
& \operatorname{Pr}\left(\sum_{u \in V}(l(u)-\mathbb{E}(l(u)))>\frac{1}{2} \sum_{u \in V} \mathbb{E}(l(u))\right) \\
\leq & \exp \left(-\frac{\frac{1}{2}\left(\frac{1}{2} \sum_{u \in V} \mathbb{E}(l(u))\right)^{2}}{\sum_{u \in V} \mathbb{V}(l(u))+\frac{1}{3} \cdot 1 \cdot \frac{1}{2} \sum_{u \in V} \mathbb{E}(l(u))}\right) \\
\leq & \exp \left(-\frac{\frac{1}{2}\left(\frac{1}{2} \sum_{u \in V} \mathbb{E}(l(u))\right)^{2}}{\sum_{u \in V} \mathbb{E}(l(u))+\frac{1}{3} \cdot 1 \cdot \frac{1}{2} \sum_{u \in V} \mathbb{E}(l(u))}\right) \\
= & \exp \left(-\frac{3}{28} \cdot \sum_{u \in V} \mathbb{E}(l(u))\right) \\
= & \exp \left(-\frac{3}{28} \cdot p \cdot|V|\right) \leq \frac{\delta}{2 n} \leq \frac{\delta}{2} .
\end{aligned}
$$

Since $\sum_{u \in V} \mathbb{E}(l(u))=p \cdot n$, with probability at least $1-\delta / 2, \sum_{u \in V} l(u) \leq 1.5 p n$.
By taking union bound, with probability at least $1-\delta, \sum_{u \in V} l(u) \leq 1.5 p n$ and $\forall v \in$ $V, \exists u \in S_{v} \cup\{v\}, l(u)=1$.

If the number of neighbors of each vertex is not large, then we can still have a constant fraction of vertices which can contract to a leader.

Lemma 32.4.4. Let $V$ be a vertex set with $n$ vertices. Let $S_{v}$ be a subset of $V \backslash\{v\}$ with size at least 1. Let $l: V \rightarrow\{0,1\}$ be a random hash function such that $\forall v \in V, l(v)$ are i.i.d. Bernoulli random variables, i.e.

$$
l(v)= \begin{cases}1 & \text { with probability } \frac{1}{2} \\ 0 & \text { otherwise }\end{cases}
$$

Let $L=\{v \in V \mid l(v)=1\} \cup\left\{v \in V \mid \forall u \in S_{v} \cup\{v\}, l(u)=0\right\} . \mathbb{E}(L) \leq 0.75 n$.

Proof. For $v \in V, \operatorname{Pr}(l(v)=1)=\frac{1}{2}$. Let $u \in S_{v}$. Then $\operatorname{Pr}\left(\forall x \in S_{v} \cup\{v\}, l(x)=0\right) \leq$ $\operatorname{Pr}(l(v)=0, l(u)=0)=0.25 . \mathbb{E}(|L|)=\sum_{v \in V} \operatorname{Pr}(v \in L) \leq 0.75 n$.

### 32.4.3 Tree Contraction Operation

In this section, we introduce the contraction operation. Firstly, let us introduce the concept of the parent pointers which can define a rooted forest.

Definition 32.4.2. Given a set of vertices $V$, let par : $V \rightarrow V$ satisfy that $\forall v \in V, \exists i>0$ such that $\operatorname{par}^{(i)}(v)=\operatorname{par}^{(i+1)}(v)$, where $\forall v \in V, j>0, \operatorname{par}^{(j)}(v)$ is defined as part $\left(\operatorname{par}^{(j-1)}(v)\right)$, and $\operatorname{par}^{(0)}(v)=v$. Then, we call such par a set of parent pointers on $V$. For $v \in V$, if $\operatorname{part}(v)=v$, then we say $v$ is a root of par. par can have more than one root. The depth of $v \in V, \operatorname{dep}_{\mathrm{par}}(v)$ is the smallest $i \in \mathbb{Z}_{\geq 0}$ such that $\operatorname{par}^{(i)}(v)=\operatorname{par}^{(i+1)}(v)$. The root of $v \in V$, $\operatorname{par}^{(\infty)}(v)$ is defined as $\operatorname{par}^{\left(\operatorname{dep}_{\operatorname{par}}(v)\right)}(v)$. The depth of par, dep(par) is defined as $\max _{v \in V} \operatorname{dep}_{\mathrm{par}}(v)$.

It is easy to see that a set of parent pointers par on $V$ formed a rooted forest on $V$. For a vertex $v \in V$, if $\operatorname{part}(v)=v$, then $v$ is a root in the forest. Otherwise $\operatorname{part}(v)$ is the parent of $v$ in the forest.

In the following, we define the union operation of several sets of parent pointers.

Definition 32.4.3. Let $\operatorname{par}_{1}: V_{1} \rightarrow V_{1}, \operatorname{par}_{2}: V_{2} \rightarrow V_{2}, \cdots, \operatorname{par}_{k}: V_{k} \rightarrow V_{k}$ be $k$ sets of parent pointers on vertex sets $V_{1}, V_{2}, \cdots, V_{k}$ respectively, where $\forall i \neq j \in[k], V_{i} \cap V_{j}=\emptyset$. Then par $=\operatorname{par}_{1} \cup \operatorname{par}_{2} \cup \cdots \cup \operatorname{par}_{k}$ is a set of parent pointers on the vertex set $V_{1} \cup V_{2} \cup \cdots \cup V_{k}$ such that $\forall i \in[k], v \in V_{i}, \operatorname{part}(v)=\operatorname{par}_{i}(v)$.

Now we focus on the parent pointers which can preserve the connectivity of the graph.

Definition 32.4.4. Given a graph $G=(V, E)$ and a set of parent pointers par on $V$, if $\forall v \in V$, we have $\operatorname{dist}_{G}(v, \operatorname{part}(v))<\infty$, then par is compatible with $G$.

It is easy to show the following fact:
Fact 32.4.5. Given a graph $G=(V, E)$ and a set of parent pointers par which is compatible with $G$, then $\forall u, v \in V$ with $\operatorname{par}^{(\infty)}(u)=\operatorname{par}^{(\infty)}(v)$, we have $\operatorname{dist}_{G}(u, v)<\infty$.

Proof. By the definition of compatible, $\forall v \in V$, $\operatorname{dist}_{G}(v, \operatorname{part}(v))<\infty$. By induction, $\forall l \in$ $\mathbb{Z}_{>0}, v \in V$, we have $\operatorname{dist}_{G}\left(v, \operatorname{par}^{(l)}(v)\right) \leq \operatorname{dist}_{G}\left(v, \operatorname{par}^{(l-1)}(v)\right)+\operatorname{dist}_{G}\left(\operatorname{par}^{(l-1)}(v), \operatorname{par}^{(l)}(v)\right)<$ $\infty$. Thus, for any pair of vertices $u, v \in V$, if $\operatorname{par}^{(\infty)}(u)=\operatorname{par}^{(\infty)}(v)$, then $\operatorname{dist}_{G}(u, v) \leq$ $\operatorname{dist}_{G}\left(u, \operatorname{par}^{(\infty)}(u)\right)+\operatorname{dist}_{G}\left(\operatorname{par}^{(\infty)}(v), v\right)<\infty$.

In this section, we describe a procedure which can be used to reduce the number of vertices. The input of the procedure is an undirected graph $G=(V, E)$ and a set of parent pointers par : $V \rightarrow V$, where par is compatible with $G$. The output of the procedure will be the root of each vertex in $V$ and an undirected graph $G^{\prime}=\left(V^{\prime}, E^{\prime}\right)$ which satisfies $V^{\prime}=\{v \in$ $V \mid \operatorname{part}(v)=v\}, E^{\prime}=\left\{(u, v) \in V^{\prime} \times V^{\prime} \mid u \neq v, \exists(p, q) \in E, \operatorname{par}^{(\infty)}(p)=u, \operatorname{par}^{(\infty)}(q)=v\right\}$. Notice that $V^{\prime}$ only contains all the roots in the forest induced by par, and $\left|E^{\prime}\right| \leq|E|$.

Lemma 32.4.6. Let $G=(V, E)$ be an undirected graph, par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2). Then TreeContraction(G, par) (See Algorithm 32.2) will output $\left(G^{\prime}, g^{(r)}\right)$ with $r \leq\lceil\log \operatorname{dep}(\mathrm{par})\rceil$ satisfies the following properties:

1. $\forall v \in V, g^{(r)}(v)=\operatorname{par}^{(\infty)}(v)$.
2. $V^{\prime}=\{v \in V \mid \operatorname{part}(v)=v\}$.
3. $E^{\prime}=\left\{(u, v) \in V^{\prime} \times V^{\prime} \mid u \neq v, \exists(p, q) \in E, \operatorname{par}^{(\infty)}(p)=u, \operatorname{par}^{(\infty)}(q)=v\right\}$.

Proof. One crucial observation is the following claim.

```
Algorithm 32.2 Tree Contraction Operation
    procedure TreeContraction \((G=(V, E)\), par : \(V \rightarrow V) \quad \triangleright\) Lemma 32.4.6,
    Corollary 32.4.8
                            \(\triangleright\) Output: \(G^{\prime}=\left(V^{\prime}, E^{\prime}\right), \operatorname{par}^{(\infty)}(v)\) for all \(v \in V\)
        Initially, for each \(v \in V\) let \(g^{(0)}(v) \leftarrow \operatorname{part}(v)\). Let \(V^{\prime}=\emptyset, E^{\prime}=\emptyset\).
        \(l \leftarrow 0\).
        for \(\exists v \in V, \operatorname{part}\left(g^{(l)}(v)\right) \neq g^{(l)}(v)\) do
                \(l \leftarrow l+1\).
            For each \(v \in V\), compute \(g^{(l)}(v)=g^{(l-1)}\left(g^{(l-1)}(v)\right) . \quad \triangleright g^{(l)}\) is \(\operatorname{par}^{\left(2^{l}\right)}\)
        end for
        \(r \leftarrow l . \quad \triangleright r\) is the number of iterations, and is used in the analysis.
        For \(v \in V\), if \(\operatorname{part}(v)=v\), let \(V^{\prime} \leftarrow V^{\prime} \cup\{v\}\).
        For \((u, v) \in E\), if \(g^{(r)}(u) \neq g^{(r)}(v)\), let \(E^{\prime} \leftarrow E^{\prime} \cup\left\{\left(g^{(r)}(u), g^{(r)}(v)\right)\right\}\).
        \(\triangleright \forall v \in V\), contract \(v\) to \(\operatorname{par}^{(\infty)}(v)\)
        return \(g^{(r)}(v)\) as \(\operatorname{par}^{(\infty)}(v)\) for all \(v \in V\), and \(G^{\prime}=\left(V^{\prime}, E^{\prime}\right)\)
    end procedure
```

Claim 32.4.7. $\forall l \in\{0,1, \cdots, r\}, v \in V$, we have $g^{(l)}(v)=\operatorname{par}^{\left(2^{l}\right)}(v)$.

Proof. The proof is by induction. When $l=0, \forall v \in V, g^{(0)}(v)=\operatorname{part}(v)=\operatorname{par}^{(1)}(v)$, the claim is true. Suppose for $l-1$, we have $\forall v \in V, g^{(l-1)}(v)=\operatorname{par}^{\left(2^{l-1}\right)}(v)$, then $\forall v \in V, g^{(l)}(v)=$ $g^{(l-1)}\left(g^{(l-1)}(v)\right)=\operatorname{par}^{\left(2^{l-1}\right)}\left(\operatorname{par}^{\left(2^{l-1}\right)}(v)\right)=\operatorname{par}^{\left(2^{l}\right)}(v)$. So the claim is true.

If $r>\lceil\log \operatorname{dep}(\operatorname{par})\rceil$, then $r-1 \geq\lceil\log \operatorname{dep}(\operatorname{par})\rceil$. Due to claim 32.4.7, we have $\forall v \in V, g^{(r-1)}(v)=\operatorname{par}^{\left(2^{r-1}\right)}(v)=\operatorname{par}^{(\infty)}(v)$. Due to the condition in line 5 , the loop will stop when $l \leq r-1$ which leads to a contradiction to line 9 . Thus, at the end of the algorithm, $r$ should be at most $\lceil\log \operatorname{dep}($ par $)\rceil$.

Since we have $\forall v \in V, \operatorname{part}\left(g^{(r)}(v)\right)=g^{(r)}(v)$ at the end of the Algorithm 32.2, $\forall v \in V, g^{(r)}(v)$ must be $\operatorname{par}^{(\infty)}(v)$. Then due to line 10 and line 11, we have $V^{\prime}=\{v \in V \mid$ $\operatorname{part}(v)=v\}, E^{\prime}=\left\{(u, v) \in V^{\prime} \times V^{\prime} \mid u \neq v, \exists(p, q) \in E, \operatorname{par}^{(\infty)}(p)=u, \operatorname{par}^{(\infty)}(q)=v\right\}$.

Definition 32.4.5. Let $G=(V, E)$ be an undirected graph, par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2). Then the number of iteration of Treecontraction( $G$, par) is defined as the value of $r$ at the end of the procedure.

Corollary 32.4.8 (Preserved connectivity and diameter). Let $G=(V, E)$ be an undirected graph, par: $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) which is compatible (See Definition 32.4.4) with $G$. Then at the end of the Algorithm 32.2, $r \leq\lceil\log \operatorname{dep}(\operatorname{par})\rceil$ and the output $\left(G^{\prime}, g^{(r)}\right)$ will satisfy the following properties:

1. $\operatorname{diam}\left(G^{\prime}\right) \leq \operatorname{diam}(G)$.
2. $\forall u, v \in V, \operatorname{dist}_{G}(u, v)<\infty \Rightarrow \operatorname{dist}_{G^{\prime}}\left(\operatorname{par}^{(\infty)}(u), \operatorname{par}^{(\infty)}(v)\right)<\infty$.
3. $\forall u, v \in V, \operatorname{dist}_{G}(u, v)<\infty \Leftarrow \operatorname{dist}_{G^{\prime}}\left(\operatorname{par}^{(\infty)}(u), \operatorname{par}^{(\infty)}(v)\right)<\infty$.

Proof. By Lemma 32.4.6, we have $r \leq\lceil\log \operatorname{dep}(\operatorname{par})\rceil$, $V^{\prime}=\{v \in V \mid \operatorname{part}(v)=v\}$ and $E^{\prime}=\left\{(u, v) \in V^{\prime} \times V^{\prime} \mid u \neq v, \exists(p, q) \in E, \operatorname{par}^{(\infty)}(p)=u, \operatorname{par}^{(\infty)}(q)=v\right\}$.

For any two vertices $u, v \in V$ which are in the same connected component in $G$, then there should be a path $u=u_{1} \rightarrow u_{2} \rightarrow \cdots \rightarrow u_{p}=v$ in graph $G$. So $\forall i \in[p-1],\left(u_{i}, u_{i+1}\right) \in E$ which means that either $\operatorname{par}^{(\infty)}\left(u_{i}\right)=\operatorname{par}^{(\infty)}\left(u_{i+1}\right)$ or $\left(\operatorname{par}^{(\infty)}\left(u_{i}\right), \operatorname{par}^{(\infty)}\left(u_{i+1}\right)\right) \in E^{\prime}$. Thus, $\operatorname{par}^{(\infty)}\left(u_{1}\right) \rightarrow \operatorname{par}^{(\infty)}\left(u_{2}\right) \rightarrow \cdots \rightarrow \operatorname{par}^{(\infty)}\left(u_{p}\right)$ is a valid path in $G^{\prime}$, and the length of this path in $G^{\prime}$ is at most $p$. Thus, the properties 1 and 2 are true.

For any two vertices $u, v \in V$ which are not in the same connected component in $G$, but there is a path $\operatorname{par}^{(\infty)}(u)=u_{1}^{\prime} \rightarrow u_{2}^{\prime} \rightarrow \cdots \rightarrow u_{p}^{\prime}=\operatorname{par}^{(\infty)}(v)$ in $G^{\prime}$, then it means that there exists vertices $u_{1,1}, u_{1,2}, u_{2,1}, u_{2,2}, \cdots, u_{p, 1}, u_{p, 2} \in V$ which satisfies
(a) $\forall i \in[p-1],\left(u_{i, 2}, u_{i+1,1}\right) \in E, \operatorname{par}^{(\infty)}\left(u_{i, 2}\right)=u_{i}^{\prime}, \operatorname{par}^{(\infty)}\left(u_{i+1,1}\right)=u_{i+1}^{\prime}$.
(b) $u_{1,1}=u, u_{p, 2}=v$.
(c) $\forall i \in[p], \operatorname{par}^{(\infty)}\left(u_{i, 1}\right)=\operatorname{par}^{(\infty)}\left(u_{i, 2}\right)$. By Fact 32.4.5, we have $\operatorname{dist}_{G}\left(u_{i, 1}, u_{i, 2}\right)<\infty$.

Thus, there exists a path from $u$ to $v$. This contradicts to that $u, v$ are not in the same connected component. Therefore, property 3 is also true.

### 32.4.4 Connectivity Algorithm

In this section, we described a batch algorithm for graph connectivity/connected components problem. The input is an undirected graph $G=(V, E)$, a space/rounds tradeoff parameter $m$, and the rounds parameter $r \leq|V|$. The output is a function col : $V \rightarrow V$ such that $\forall u, v \in V, \operatorname{dist}_{G}(u, v)<\infty \Leftrightarrow \operatorname{col}(u)=\operatorname{col}(v)$.

The algorithm is described in Algorithm 32.3. The following theorem shows the correctness of Algorithm 32.3.

Theorem 32.4.9 (Correctness of Algorithm 32.3). Let $G=(V, E)$ be an undirected graph, $m \geq 4|V|$, and $r \leq|V|$ be the rounds parameter. If $\operatorname{Connectivity~}(G, m, r)$ (Algorithm 32.3) does not output FAIL, then $\forall u, v \in V$, we have $\operatorname{dist}_{G}(u, v)<\infty \Leftrightarrow \operatorname{col}(u)=$ $\operatorname{col}(v)$.

Proof. Firstly, we show that the input of line 18 is valid.

Claim 32.4.10. $\forall i \in[r], \operatorname{par}_{i}$ is a set of parent pointers on $V_{i}^{\prime \prime}$, (See Definition 32.4.2) and is compatible (See Definition 32.4.4) with $G_{i}^{\prime \prime}$.

Proof. $\forall v \in V_{i}^{\prime \prime}$, if $v \in L_{i}$, then $\operatorname{par}_{i}(v)=v$. For $v \in V_{i}^{\prime \prime} \backslash L_{i}$, due to property 3 of Lemma 32.4.2, we have $\operatorname{par}_{i}(v) \in V_{i}^{\prime \prime}$. Since $\operatorname{par}_{i}(v) \in L_{i}$, we have $\operatorname{par}_{i}\left(\operatorname{par}_{i}(v)\right)=\operatorname{par}_{i}(v)$. Thus, $\operatorname{par}_{i}: V_{i}^{\prime \prime} \rightarrow V_{i}^{\prime \prime}$ is a set of parent pointers on $V_{i}^{\prime \prime}$. Due to property 2 of Lemma 32.4.2 and $\operatorname{dist}_{G_{i}^{\prime}}\left(\operatorname{par}_{i}(v), v\right)<\infty$, we know that $\operatorname{dist}_{G_{i-1}}\left(\operatorname{par}_{i}(v), v\right)<\infty$. Thus, $\operatorname{dist}_{G_{i}^{\prime \prime}}\left(\operatorname{par}_{i}(v), v\right)<$ $\infty$. It implies that $\operatorname{par}_{i}$ is compatible with $G_{i}^{\prime \prime}$.

The following claim shows that the number of the remaining vertices cannot increase after each round.

Claim 32.4.11. If $\operatorname{Connectivity~}(G, m, r)$ does not output FAIL, then $\forall i \in[r], V_{i} \subseteq V_{i}^{\prime \prime} \subseteq$ $V_{i}^{\prime}=V_{i-1}$.

Proof. Let $i \in[r]$. Due to Claim 32.4.10, the input of line 18 is valid. Then, we can apply property 2 of Lemma 32.4 .6 to get $V_{i} \subseteq V_{i}^{\prime \prime}$. By the construction of $V_{i}^{\prime \prime}$ we have $V_{i}^{\prime \prime} \subseteq V_{i}^{\prime}$. Since the procedure Neighborincrement $\left(m, G_{i-1}\right)$ (Algorithm 32.1) does not change the vertex set, we have $V_{i}^{\prime}=V_{i-1}$.

Now, we show that $\forall u, v \in V_{i}, \operatorname{dist}_{G_{i}}(u, v)<\infty \Leftrightarrow \operatorname{dist}_{G}(u, v)<\infty$.
Claim 32.4.12. If $\operatorname{Connectivity~}(G, m, r)$ does not output FAIL, then $\forall i \in[r], \forall u, v \in V_{i}$, we have $\operatorname{dist}_{G_{i}}(u, v)<\infty \Leftrightarrow \operatorname{dist}_{G}(u, v)<\infty$.

Proof. The proof is by induction. Suppose $\forall u, v \in V_{i-1}, \operatorname{dist}_{G_{i-1}}(u, v)<\infty \Leftrightarrow \operatorname{dist}_{G}(u, v)<$ $\infty . \forall w, z \in V_{i}$, according to Claim 32.4.11, $w, z \in V_{i}^{\prime \prime}$. By property 2,3 of Lemma 32.4.8, and property 2 of Lemma 32.4.6, $\operatorname{dist}_{G_{i}}(w, z)<\infty \Leftrightarrow \operatorname{dist}_{G_{i}^{\prime \prime}}(w, z)<\infty$. Due to property 2,3 of Lemma 32.4.2, there is no edge in $E_{i-1}$ between $V_{i}^{\prime \prime}$ and $V_{i}^{\prime} \backslash V_{i}^{\prime \prime}$. According to Claim 32.4.11, $w, z \in V_{i-1}$. Thus, $\operatorname{dist}_{G_{i}^{\prime \prime}}(w, z)<\infty \Leftrightarrow \operatorname{dist}_{G_{i-1}}(w, z)<\infty$. By induction hypothesis, we have $\forall w, z \in V_{i}, \operatorname{dist}_{G_{i}}(w, z)<\infty \Leftrightarrow \operatorname{dist}_{G}(w, z)$.

The following claim states that once a vertex $v \in V$ is contracted to an another vertex, it will never be operated.

Claim 32.4.13. Suppose Connectivity $(G, m, r)$ does not output FAIL. $\forall i \in\{0,1, \cdots, r\}$, $v \in V$, we have $h_{i}(v)=$ null $\Leftrightarrow v \in V_{i}$. Furthermore, $\forall v \in V, \exists j \in[r]$ such that $h_{0}(v)=$ $h_{1}(v)=\cdots=h_{j-1}(v)=$ null and $h_{j}(v)=h_{j+1}(v)=\cdots=h_{r}(v) \neq{\operatorname{null}, \operatorname{dist}_{G}\left(v, h_{r}(v)\right)<\infty . ~}_{\text {. }}$.

Proof. When $i=0, \forall v \in V, h_{0}(v)=$ null, $v \in V_{0}=V$. Suppose it is true that $\forall v \in$ $V, h_{i-1}(v)=$ null $\Leftrightarrow v \in V_{i-1}$. If $v \notin V_{i}$, according to Claim 32.4.11, there are three cases: $v \in V_{i}^{\prime \prime} \backslash V_{i}, v \in V_{i}^{\prime} \backslash V_{i}^{\prime \prime}, v \notin V_{i-1}$. In the first case, due to line $22, h_{i}(v) \neq$ null. In the second case, due to line $21, h_{i}(v) \neq$ null, In the third case, due to line $23, h_{i}(v) \neq$ null. If $h_{i}(v)=$ null, then $h_{i}(v)$ cannot be updated by line 21 , line 22 or line 23 which implies that $v \in V_{i-1}, v \notin V_{i}^{\prime} \backslash V_{i}^{\prime \prime}, v \notin V_{i}^{\prime \prime} \backslash V_{i}$. Thus, $v \in V_{i}$.

Since the procedure does not FAIL, we have $n_{r}=0$ which means that $\forall v \in V, h_{r}(v) \neq$ null. Notice that by line 23 , if $h_{i-1}(v) \neq$ null, then $h_{i}(v)=h_{i-1}(v)$. Thus, $\forall v \in V, \exists j \in[r]$ such that $h_{0}(v)=h_{1}(v)=\cdots=h_{j-1}(v)=$ null and $h_{j}(v)=h_{j+1}(v)=\cdots=h_{r}(v) \neq$ null.

For $v \in V$, if $h_{j}(v) \neq$ null and $h_{j-1}(v)=$ null, then $h_{j}(v)$ can only be updated by 21 or line 22 . In both cases, $\operatorname{dist}_{G_{j-1}}\left(v, h_{j}(v)\right)<\infty$. By Claim 32.4.12, we have that $\operatorname{dist}_{G}\left(v, h_{j}(v)\right)<\infty$.

In the following, we show that $h_{r}$ is a rooted tree such that $\operatorname{dist}_{G}(u, v)<\infty \Leftrightarrow u, v$ have the same root. Due to Claim 32.4.13, if Connectivity $(G, m, r)$ does not output FAIL, then $n_{r}=0$ which implies that $\forall v \in V, h_{r}(v) \neq$ null. Thus, we can define $h_{r}^{(k)}(v)$ for $k \in \mathbb{Z}_{>0}$ as applying $h_{r}$ on $v k$ times. $\forall v \in V$, by Claim 32.4.13, let $j \in[r]$ satisfy that $h_{j}(v) \neq$ null and $h_{j-1}(v)=$ null. If $h_{j}(v)$ is updated by line 22 , then $h_{j}\left(h_{j}(v)\right)=$ null. If $h_{j}(v)$ is updated by line 21 , then $h_{j}\left(h_{j}(v)\right)=h_{j}(v)$. In both cases, $h_{j}$ cannot create a cycle. Thus, we can define $h_{r}^{(\infty)}(v)=h_{r}^{(k)}(v)$ for some $k$ which satisfies $h_{r}\left(h_{r}^{(k)}(v)\right)=h_{r}^{(k)}(v)$.

Claim 32.4.14. Suppose Connectivity $(G, m, r)$ does not output FAIL. Then $\forall u, v \in V$, we have $\operatorname{dist}_{G}(u, v)<\infty \Leftrightarrow h_{r}^{(\infty)}(u)=h_{r}^{(\infty)}(v)$.

Proof. Let $u, v \in V$. By Claim 32.4.13, if $h_{r}^{\infty}(u)=h_{r}^{\infty}(v)$ we have $\operatorname{dist}_{G}(u, v)<\infty$.

If $\operatorname{dist}_{G}(u, v)<\infty$, then let $u^{\prime}=h_{r}^{(\infty)}(u), v^{\prime}=h_{r}^{(\infty)}(v)$. By Claim 32.4.13, $\operatorname{dist}_{G}\left(u^{\prime}, v^{\prime}\right) \leq$ $\operatorname{dist}_{G}\left(u, u^{\prime}\right)+\operatorname{dist}_{G}(u, v)+\operatorname{dist}_{G}\left(v, v^{\prime}\right)<\infty$, and we can find $j \in[r]$ such that $h_{j}\left(u^{\prime}\right) \neq$ null, $h_{j-1}\left(u^{\prime}\right)=$ null. Without loss of generality, we can assume $h_{j-1}\left(v^{\prime}\right)=$ null (otherwise we can swap $u^{\prime}$ and $\left.v^{\prime}\right)$. Due to Claim 32.4.13, $u^{\prime}, v^{\prime} \in V_{i-1}$. Since $h_{j}\left(u^{\prime}\right)=h_{r}\left(u^{\prime}\right)=u^{\prime}, h_{j}\left(u^{\prime}\right)$ can be only updated by line 21, and $u^{\prime} \in V_{j}^{\prime} \backslash V_{j}^{\prime \prime}$. Then due to property 3 of Lemma 32.4.6, $v^{\prime}$ should be in $\Gamma_{G_{i}^{\prime}}(u) \cup\{u\}$. Since $h_{j}\left(v^{\prime}\right)=h_{r}\left(v^{\prime}\right)=v^{\prime}$, we can conclude that $u^{\prime}=v^{\prime}$.

If Connectivity $(G, m, r)$ does not output FAIL, then in line 26 , col is exactly $h_{r}^{(\infty)}$. By Claim 32.4.14, we have $\forall u, v \in V$, $\operatorname{dist}_{G}(u, v)<\infty \Leftrightarrow \operatorname{col}(u)=\operatorname{col}(v)$.

Now let us consider the number of iterations of Algorithm 32.3 and the success probability.

Definition 32.4.6 (Total iterations). Let $G=(V, E)$ be an undirected graph, $\operatorname{poly}(n) \geq$ $m>4 n$, and $r \leq n$ be the rounds parameter where $n$ is the number of vertices in $G$. The total number of iterations of Connectivity $(G, m, r)$ (Algorithm 32.3) is defined as $\sum_{i=1}^{r}\left(k_{i}+r_{i}^{\prime}\right)$, where $k_{i}$ denotes the number of iterations (See Definition 32.4.1) of Neighborincrement ( $m, G_{i-1}$ ) (see line 9), and $r_{i}^{\prime}$ denotes the number of iterations (See Definition 32.4.5) of TreeContraction $\left(G_{i}^{\prime \prime}, \operatorname{par}_{i}\right)$ (see line 18).

Theorem 32.4.15 (Success probability and total iterations). Let $G=(V, E)$ be an undirected graph, $\operatorname{poly}(n) \geq m>4 n$, and $r \leq n$ be the rounds parameter where $n=|V|$. Let $c>0$ be a sufficiently large constant. If $r \geq c \log \log _{m / n}(n)$, then with probability at least 0.98 , Connectivity $(G, m, r)$ (Algorithm 32.3) will not return FAIL. If $\operatorname{Connectivity~}(G, m, r)$
succeeds, let $k_{i}$ denote the number of iterations (See Definition 32.4.1) of NeighborIncrement $\left(m, G_{i-1}\right)$ (see line 9), and let $r_{i}^{\prime}$ denote the number of iterations of (See Definition 32.4.5) of TreeContraction( $G_{i}^{\prime \prime}$ (see line 18), then

1. $\forall i \in[r], r_{i}^{\prime}=0$.
2. $\forall i \in[r], k_{i}$ is at most $\lceil\log (\operatorname{diam}(G))\rceil+1$.
3. The number of iterations of line 26 is at most $\lceil\log r\rceil$.
4. $\sum_{i=1}^{r} k_{i} \leq O(r \log (\operatorname{diam}(G)))$.

Let $c_{1}>0$ be a sufficiently large constant. If $m \geq c_{1} n \log ^{4} n$, then with probability at least 0.99, $\sum_{i=1}^{r} k_{i} \leq O\left(\log (\operatorname{diam}(G)) \log \log _{\operatorname{diam}(G)}(n)\right)$. If $m<c_{1} n \log ^{4} n$, then with probability at least 0.98, $\sum_{i=1}^{r} k_{i} \leq O\left(\log (\operatorname{diam}(G)) \log \log _{\operatorname{diam}(G)}(n)+(\log \log (n))^{2}\right)$.

Proof. Suppose Connectivity $(G, m, r)$ succeeds. Property 1 follows by $\forall v \in V_{i}^{\prime \prime}, \operatorname{par}_{i}\left(\operatorname{par}_{i}(v)\right)=$ $\operatorname{par}_{i}(v)$ and Lemma 32.4.6. Property 2 follows by $\operatorname{diam}\left(G_{r}\right) \leq \operatorname{diam}\left(G_{r}^{\prime \prime}\right) \leq \operatorname{diam}\left(G_{r}^{\prime}\right) \leq$ $\operatorname{diam}\left(G_{r-1}\right) \leq \operatorname{diam}\left(G_{r-1}^{\prime \prime}\right) \leq \operatorname{diam}\left(G_{r-1}^{\prime}\right) \leq \cdots \leq \operatorname{diam}\left(G_{0}\right)=\operatorname{diam}(G)$ and property 1 of Lemma 32.4.2. Property 3 follows by the depth of $h_{r}$ is at most $r$ and Lemma 32.4.6. Property 4 follows by property 2 .

Now let us prove the success probability. Let $i \in[r]$. If $p_{i}<0.5$, then we can apply Lemma 32.4.3 on vertex set $V_{i}^{\prime \prime}$, parameter $\gamma_{i}$, and hash function $l_{i}$. Notice that the set $S_{v}$ in the statement of Lemma 32.4.3 is $\Gamma_{G_{i}^{\prime}}(v)$ in the algorithm. Notice that $\left|V_{i}^{\prime \prime}\right| \leq n$. Then in the $i^{\text {th }}$ round, if $p_{i}<0.5$, then with probability at most $1 /\left(100 n^{2}\right), L_{i}$ will be $\left\{v \in V_{i}^{\prime \prime} \mid l_{i}(v)=1\right\}$, and $n_{i}=\left|L_{i}\right| \leq 1.5 p_{i} n_{i-1}$. By taking union bound over all $i \in[r]$,
we have that with probability at least 0.99 , event $\mathcal{E}$ happens: for all $i \in[r]$, if $p_{i}<0.5$, then $n_{i} \leq 1.5 p_{i} n_{i-1} \leq 0.75 n_{i-1}$. Suppose $\mathcal{E}$ happens. For $i \in[r], p_{i}=0.5$, if we apply Lemma 32.4.4, then condition on $n_{i-1}$, we have $\mathbb{E}\left(n_{i}\right) \leq 0.75 n_{i-1}$. Thus, we know $\forall i \in$ $[r], \mathbb{E}\left(n_{i}\right) \leq 0.75 \mathbb{E}\left(n_{i-1}\right) \leq 0.75^{i} n$.

Next, we discuss the case for $p_{0}=0.5$ and the case for $p_{0}<0.5$ separately.
If $p_{0}=0.5$, then $m \leq n \cdot(600 \log n)^{4}$. By Markov's inequality, when $i^{*} \geq 4 \log _{4 / 3}(6000 \log n)$, with probability at least $0.99, n_{i^{*}} \leq n /(600 \log n)^{4}$ and thus $p_{i^{*}}<0.5$. Condition on this event and $\mathcal{E}$, we have

$$
\begin{aligned}
n_{r} & \leq \frac{\left(\frac{\left(\frac{n_{i i^{1.5}}^{m_{0} .5}}{}(45 \log n+150)\right)^{1.5}}{m^{0.5}}(45 \log n+150)\right)^{\cdots}}{\cdots} \quad \text { (Apply } r^{\prime}=r-i^{*} \text { times) } \\
& =\frac{n_{i^{*}}^{1.5 r^{\prime}}}{m^{1.5 r^{\prime}-1}}(45 \log n+150)^{2 \cdot\left(1.5^{r^{\prime}}-1\right)} \\
& =n_{i^{*}} /\left(m / n_{i^{*}}\right)^{1.5^{r^{\prime}}-1} \cdot(45 \log n+150)^{2 \cdot\left(1.5^{r^{\prime}}-1\right)} \\
& \leq n /\left(m /\left(n_{i^{*}}(45 \log n+150)^{2}\right)\right)^{1.5^{r^{\prime}}-1} \\
& \leq n /\left(m /\left(n_{i^{*}}(45 \log n+150)^{2}\right)\right)^{1.5^{r^{\prime} / 2}} \\
& \leq n /(m / n)^{1.5^{r^{\prime} / 2}} \leq \frac{1}{2}
\end{aligned}
$$

where the second inequality follows by $n_{i^{*}} \leq n$, the third inequality follows by $r^{\prime} \geq 5$, the forth inequality follows by $n_{i^{*}} \leq n /(600 \log n)^{4}$, and the last inequality follows by $r^{\prime} \geq \frac{2}{\log 1.5} \log \log _{m / n}(2 n)$. Since $4 n \leq m \leq n \cdot(600 \log n)^{4}, \log \log _{m / n} n=\Theta(\log \log n)$. Let $c>0$ be a sufficiently large constant. Thus, when $r \geq c \log \log _{m / n} n \geq i^{*}+r^{\prime}=$ $4 \log (6000 \log n) / \log (4 / 3)+\frac{2}{\log 1.5} \log \log _{m / n}(2 n)$, with probability at least 0.98 , CONNECTIVITY $(G, m, r)$ will not fail.

Since property 1 of Lemma 32.4.2, we have $k_{i} \leq O\left(\log \left(\min \left(m / n_{i-1}, \operatorname{diam}(G)\right)\right)\right)$.
Thus,

$$
\begin{aligned}
& \sum_{i=1}^{r} k_{i}=\sum_{i=1}^{i^{*}} k_{i}+\sum_{i=i^{*}+1}^{r} k_{i} \leq O\left((\log \log n)^{2}\right)+\sum_{i=i^{*}+1}^{r} k_{i} \\
\leq & O\left((\log \log n)^{2}\right)+\sum_{i: i \geq i^{*}+1, m / n_{i-1} \leq \operatorname{diam}(G)} k_{i}+\sum_{i: i \leq r, m / n_{i-1}>\operatorname{diam}(G)} k_{i} \\
\leq & O\left((\log \log n)^{2}\right)+O\left(\sum_{i=0}^{\left\lceil\log _{1.25} \log _{2}(\operatorname{diam}(G))\right\rceil} \log \left(2^{1.25^{i}}\right)\right)+O\left(\sum_{i=0}^{\left\lceil\log _{1.25} \log _{\operatorname{diam}(G)}(m)\right\rceil} \log (\operatorname{diam}(G))\right) \\
\leq & O\left((\log \log n)^{2}\right)+O(\log (\operatorname{diam}(G)))+O\left(\log (\operatorname{diam}(G)) \log \log _{\operatorname{diam}(G)}(n)\right) \\
\leq & O\left(\log (\operatorname{diam}(G)) \log \log _{\operatorname{diam}(G)}(n)+(\log \log (n))^{2}\right),
\end{aligned}
$$

where the first inequality follows by $i^{*}=O(\log \log n)$ and $\forall i \leq\left[i^{*}\right], m / n_{i-1} \leq \operatorname{poly}(\log n)$, the third inequality follows by $m / n_{i+1} \geq\left(m / n_{i}\right)^{1.5} /(45 \log n+150) \geq\left(m / n_{i}\right)^{1.25}$. If $m>n \cdot(600 \log n)^{4}$, then $\forall i \in\{0\} \cup[r-1]$, we have $p_{i}<0.5$. Since $\mathcal{E}$ happens. We have:

$$
\begin{aligned}
n_{r} & \leq \frac{\left(\frac{\left(\frac{n^{1.5}}{m^{0.5}}(45 \log n+150)\right)^{1.5}}{m^{0.5}}(45 \log n+150)\right)}{\cdots} \\
& =\frac{n^{1.5^{r}}}{m^{1.5^{r}-1}}(45 \log n+150)^{2 \cdot\left(1.5^{r}-1\right)} \\
& =n /(m / n)^{1.5^{r}-1} \cdot(45 \log n+150)^{2 \cdot\left(1.5^{r}-1\right)} \\
& =n /\left(m /\left(n(45 \log n+150)^{2}\right)\right)^{1.5^{r}-1} \\
& \leq n /\left(m /\left(n(45 \log n+150)^{2}\right)\right)^{1.5^{r / 2}} \\
& \leq n /\left(m /\left(n(200 \log n)^{2}\right)\right)^{1.5^{r / 2}} \\
& \leq \frac{1}{2}
\end{aligned}
$$

where the second inequality follows by $r \geq 5$, the third inequality follows by $45 \log n+150 \leq$ $200 \log n$, and the last inequality follows by

$$
r \geq c \log \log _{m / n} n \geq 2 \log _{1.5} \log _{(m / n)^{1 / 2}} 2 n \geq 2 \log _{1.5} \log _{m /\left(n(200 \log n)^{2}\right)} 2 n
$$

for a sufficiently large constant $c>0$.
By property 1 of Lemma 32.4.2, we have $k_{i} \leq O\left(\log \left(\min \left(m / n_{i-1}, \operatorname{diam}(G)\right)\right)\right)$. Thus,

$$
\begin{aligned}
\sum_{i=1}^{r} k_{i} & \leq \sum_{m / n_{i-1} \leq \operatorname{diam}(G)} k_{i}+\sum_{m / n_{i-1}>\operatorname{diam}(G)} k_{i} \\
& \leq O\left(\sum_{i=0}^{\left\lceil\log _{1.25} \log _{2}(\operatorname{diam}(G))\right\rceil} \log \left(2^{1.25^{i}}\right)\right)+O\left(\sum_{i=0}^{\left\lceil\log _{1.25} \log _{\operatorname{diam}(G)}(m)\right\rceil} \log (\operatorname{diam}(G))\right) \\
& \leq O(\log (\operatorname{diam}(G)))+O\left(\log (\operatorname{diam}(G)) \log \log _{\operatorname{diam}(G)}(n)\right),
\end{aligned}
$$

where the first inequality follows by $m / n_{i+1} \geq\left(m / n_{i}\right)^{1.5} /(45 \log n+150) \geq\left(m / n_{i}\right)^{1.25}$.
Since $n_{r}$ is an integer, $n_{r}$ must be 0 when $n_{r} \leq 1 / 2$. Let $c>0$ be a sufficiently large constant. For all $m \geq 4 n$, if $r \geq c \log \log _{m / n} n$ then $\operatorname{Connectivity~}(G, m, r)$ will succeed with probability at least 0.98 .

### 32.5 Spanning Forest

### 32.5.1 Local Shortest Path Tree

In this section, we introduce an important procedure which will be used in the spanning tree algorithm. Roughly speaking, our procedure can merge several local shortest path trees into a larger local shortest path tree. Before we describe the details of the procedure, let us look at some concepts.

Definition 32.5.1 (Local shortest path tree (LSPT)). Let $V^{\prime}$ be a set of vertices, $v$ be a vertex in $V^{\prime}$, and par : $V^{\prime} \rightarrow V^{\prime}$ be a set of parent pointers (See Definition 32.4.2) on $V^{\prime}$ which satisfies that $v$ is the only root of par. Let $T=\left(V^{\prime}\right.$, par $)$. Given an undirected graph $G=(V, E)$, if $V^{\prime} \subseteq V$ and $\forall u \in V^{\prime} \backslash\{v\},(u, \operatorname{part}(u)) \in E, \operatorname{dep}_{\mathrm{par}}(u)=\operatorname{dist}_{G}(u, v)$, then we say $T$ is a local shortest path tree (LSPT) in $G$, and $T$ has root $v$. The vertex set ( $V^{\prime}$ in the above) in $T$ is denoted as $V_{T}$. The set of parent pointers (par in the above) in $T$ is denoted as $\operatorname{par}_{T}$. For short, $\operatorname{dep}_{\operatorname{par}_{T}}$ is denoted as $\operatorname{dep}_{T}$, and $\operatorname{dep}(\operatorname{part}(T))$ is denoted as $\operatorname{dep}(T)$.

Definition 32.5.2. Given an undirected graph $G=(V, E)$, a vertex $v \in V$, and $s \in \mathbb{Z}_{\geq 0}$, we define the ball centered at $v$ with radius $s$ as the set $B_{G, s}(v)=\left\{u \in V \mid \operatorname{dist}_{G}(u, v) \leq s\right\}$.

If in the context graph $G$ is clear, then we use $B_{s}(v)$ to denote $B_{G, s}(v)$.
Definition 32.5.3 (Local complete shortest path tree (LCSPT)). Given an undirected graph $G=(V, E), s \in \mathbb{Z}_{\geq 0}$ and a local shortest path tree $T=\left(V_{T}, \operatorname{par}_{T}\right)$ in $G$ where $T$ has root $v \in V$. If $V_{T}=B_{G, s}(v)$, then we call $T$ a local complete shortest path tree (LCSPT) in $G$. The root of $T$ is $v$. The radius of $T$ is $s$.

Let $\widetilde{T}=\left(V_{\widetilde{T}}, \operatorname{par}_{\widetilde{T}}\right)$ with radius $s_{1} \in \mathbb{Z}_{\geq 0}$ and root $v$ be a local complete shortest path tree in some graph $G=(V, E)$. For $s_{2} \in \mathbb{Z}_{\geq 0}$, if for every $u \in V_{\widetilde{T}}$, we have a local complete
shortest path tree $T(u)=\left(V_{T(u)}, \operatorname{par}_{T(u)}\right)$ with root $u$ and radius $s_{2}$, then we can compute a larger local complete shortest path tree $\widehat{T}$ with root $v$ and radius $s_{1}+s_{2}$. The procedure is described in Algorithm 32.4.

Lemma 32.5.1. Let $G=(V, E)$ be an undirected graph, $s_{1}, s_{2} \in \mathbb{Z}_{\geq 0}$, and $v \in V$. Let $\widetilde{T}=\left(V_{\widetilde{T}}, \operatorname{par}_{\widetilde{T}}\right)$ with root $v$ and radius $s_{1}$ be a local complete shortest path tree in $G$, and $\operatorname{dep}_{\widetilde{T}}: V_{\widetilde{T}} \rightarrow \mathbb{Z}_{\geq 0}$ be the depth of every vertex in $\widetilde{T} . \forall u \in V_{\widetilde{T}}$, let $T(u)$ with root $u$ and radius $s_{2}$ be a local complete shortest path tree in $G$, and $\operatorname{dep}_{T(u)}: V_{T(u)} \rightarrow \mathbb{Z}_{\geq 0}$ be the depth of every vertex in $T(u)$. Let $\left(\widehat{T}=\left(V_{\widehat{T}}, \operatorname{par}_{\widehat{T}}\right), \operatorname{dep}_{\widehat{T}}\right)=\operatorname{TrEEEXPANSION}\left(\widetilde{T}, \operatorname{dep}_{\widetilde{T}},\{T(u) \mid u \in\right.$ $\left.V_{\widetilde{T}}\right\},\left\{\operatorname{dep}_{T(u)} \mid u \in V_{\widetilde{T}}\right\}$ ) (Algorithm 32.4), then $\widehat{T}$ is a local complete shortest path tree with root $v$ and radius $s_{1}+s_{2}$ in $G$. In addition, $\operatorname{dep}_{\widehat{T}}$ records the depth of every vertex in $\widehat{T}$.

Proof. If $x \in B_{s_{1}+s_{2}}(v)$, then there must exist $u \in V$ such that $\operatorname{dist}_{G}(v, u) \leq s_{1}$ and $\operatorname{dist}_{G}(u, x) \leq s_{2}$. Thus, $V_{\widehat{T}}=\bigcup_{u \in \widetilde{T}} V_{T(u)}=\bigcup_{u \in B_{s_{1}}(v)} B_{s_{2}}(u)=B_{s_{1}+s_{2}}(v)$.

Now we want to prove that $\operatorname{par}_{\widehat{T}}: V_{\widehat{T}} \rightarrow V_{\widehat{T}}$ also satisfies the condition that $\widehat{T}$ is a local shortest path tree. We can prove it by induction. If $\operatorname{dist}_{G}(u, v)=0$, then it means $u=v$. In this case, $\operatorname{par}_{\widehat{T}}(u)=\operatorname{par}_{\widetilde{T}}(u)=v$, and $h(u)=\operatorname{dep}_{\widetilde{T}}(u)=0$. Let $s \in\left[s_{1}+s_{2}\right]$. Suppose $\forall x \in B_{s-1}(v)$, we have $h(x)=\operatorname{dep}_{\widehat{T}}(x)=\operatorname{dist}_{G}(x, v)$. If $B_{s}(v)=B_{s-1}(v)$, then we are already done. Otherwise, let $x$ be the vertex which has $\operatorname{dist}_{G}(x, v)=s$. If $x \in B_{s_{1}}(v)$, then $h(x)=$ $\operatorname{dep}_{\widetilde{T}}(x)=\operatorname{dist}_{G}(x, v)$. Additionally, we have $\operatorname{par}_{\widehat{T}}(x)=\operatorname{par}_{\widetilde{T}}(x)$. Therefore, $\operatorname{dep}_{\widehat{T}}(x)=$ $\operatorname{dep}_{\widehat{T}}\left(\operatorname{par}_{\widetilde{T}}(x)\right)+1=\operatorname{dist}_{G}\left(v, \operatorname{par}_{\widetilde{T}}(x)\right)+1=\operatorname{dist}_{G}(v, x)$. If $x \in B_{s_{2}}(v) \backslash B_{s_{1}}(v)$, then $h(x)=$ $\min _{u: \operatorname{dist}_{G}(v, u) \leq s_{1}, \operatorname{dist}_{G}(u, x) \leq s_{2}} \operatorname{dep}_{\widetilde{T}}(u)+\operatorname{dep}_{T(u)}(x)=\min _{u: \operatorname{dist}_{G}(v, u) \leq s_{1}, \operatorname{dist}_{G}(u, x) \leq s_{2}} \operatorname{dist}_{G}(v, u)+$ $\operatorname{dist}_{G}(u, x)=\operatorname{dist}_{G}(v, x)=s$. And we have $\operatorname{dist}_{G}(v, x)=\operatorname{dist}_{G}\left(v, u_{x}\right)+\operatorname{dist}_{G}\left(u_{x}, x\right)$. Notice that $\operatorname{dist}_{G}\left(v, \operatorname{par}_{T\left(u_{x}\right)}(x)\right)=\operatorname{dist}_{G}\left(v, u_{x}\right)+\operatorname{dist}_{G}\left(u_{x}, \operatorname{par}_{T\left(u_{x}\right)}(x)\right)=\operatorname{dist}_{G}(v, x)-1=s-1$.

Thus,

$$
\operatorname{dep}_{\widehat{T}}(x)=\operatorname{dep}_{\widehat{T}}\left(\operatorname{par}_{\widehat{T}}(x)\right)+1=\operatorname{dep}_{\widehat{T}}\left(\operatorname{par}_{T\left(u_{x}\right)}(x)\right)+1=s
$$

To conclude, $\widehat{T}$ is a local complete shortest path tree with root $v$ and radius $s_{1}+s_{2}$ in $G$. In addition, $\operatorname{dep}_{\widehat{T}}$ records the depth of every vertex in $\widehat{T}$.

### 32.5.2 Multiple Local Shortest Path Trees

In this section, we show a procedure which is a generalization of neighbor increment procedure shown in Section 32.4.1. The input of the procedure is an undirected graph $G=(V, E)$ and a parameter $m$ which is larger than $|V|=n$. The output will be $n$ local shortest path trees (See Definition 32.5.1) such that $\forall v \in V$, there is a shortest path tree with root $v$. Furthermore, the size of each shortest path tree is at least $\left\lceil(m /|V|)^{1 / 4}\right\rceil$ and at most $\left[(m /|V|)^{1 / 2}\right\rceil$. The algorithm is described in Algorithm 32.6. The high level idea is that we firstly use doubling technique and the algorithm described in Section 32.5.1 to get local complete shortest path trees rooted at every vertex with multiple radius, and then use these LCSPTs to find large enough local shortest path trees rooted at every vertex. The doubling algorithm is described in Algorithm 32.5.

Definition 32.5.4. Given a graph $G=(V, E)$ and a parameter $m \in \mathbb{Z}_{\geq 0}, m \geq|V|$, the number of iterations of MultiRadiusLCSPT $(G, m)$ (Algorithm 32.5) is the value of $r$ at the end of the procedure.

Lemma 32.5.2. Let $G=(V, E)$ be an undirected graph, and $m$ be a parameter which is at least $|V|$. Let $\left(r,\left\{T_{i}(v) \mid i \in\{0\} \cup[r], v \in V\right\},\left\{\operatorname{dep}_{T_{i}(v)} \mid i \in\{0\} \cup[r], v \in V, T_{i}(v) \neq \operatorname{null}\right\}\right)=$ MultiRadiusLCSPT(G,m) (Algorithm 32.5). We have following properties.

1. $\forall i \in\{0\} \cup[r], v \in V$, if $T_{i}(v) \neq$ null, then $T_{i}(v)$ is a LCSPT (See Definition 32.5.3) with root $v$ and radius $2^{i}$ in $G$. Furthermore, $\operatorname{dep}_{T_{i}(v)}$ records the depth of every vertex in $T_{i}(v)$.
2. $\forall i \in\{0\} \cup[r], v \in V,\left|B_{G, 2^{i}}(v)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil \Leftrightarrow T_{i}(v)=$ null.
3. For $v \in V$, if $T_{r}(v) \neq$ null, then $V_{T_{r}(v)}=\left\{u \in V \mid \operatorname{dist}_{G}(u, v)<\infty\right\}$.
4. The number of iterations (see Definition 32.5.4) $r \leq \min (\lceil\log (\operatorname{diam}(G))\rceil,\lceil\log (m / n)\rceil)+$ 1.

Proof. For property 1, we can prove it by induction. If $i=0$, the property holds by line 4, line 5 and line 7. Now suppose $\forall v \in V$, if $T_{i-1}(v)$ is not null, then $T_{i-1}(v)$ is a LCSPT with root $v$ and radius $2^{i-1}$ in $G$, and $\operatorname{dep}_{T_{i}(v)}$ records the depth of every vertex in $T_{i}(v)$. For $v \in V$, notice that the only place that will make $T_{i}(v)$ not null is line 15 , and if the procedure run line 15, any of $T_{i-1}(v)$ and $T_{i-1}(u)$ with $u \in V_{T_{i-1}(v)}$ cannot be null. By Lemma 32.5.1, since the radius of $T_{i-1}(v)$ is $2^{i-1}$, and $\forall u \in V_{T_{i-1}(v)}, T_{i-1}(u)$ has radius $2^{i-1}, T_{i}(v)$ is a LCSPT with root $v$ and radius $2^{i}$. Furthermore $\operatorname{dep}_{T_{i}(v)}$ records the depth of every vertex in $T_{i}(v)$.

For property 2 , if $i=0$, then this property holds by line 4 to line 7 . For $i \in[r]$, our proof is by induction. Suppose the property holds for $i-1$. Now consider $T_{i}(v)$ for $v \in V$. The only way to make $T_{i}(v)$ not null is line 15 . If the procedure invokes line 15 , then any of $T_{i-1}(v)$ and $T_{i-1}(u)$ with $u \in V_{T_{i-1}(v)}$ cannot be null. By property 1 and Lemma 32.5.1, $T_{i}(v)$ will be a LCSPT with root $v$ and radius $2^{i}$ in line 15. If $\left|B_{G, 2^{i}}(v)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil$, then $T_{i}(v)$ is set to be null in line 16 . Thus, we already got $\left|B_{G, 2^{i}}(v)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil \Rightarrow T_{i}(v)=$ null. Now we want to show $\left|B_{G, 2^{i}}(v)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil \Leftarrow T_{i}(v)=$ null. If $T_{i}(v)=$ null, then there are three cases. The first case is that $T_{i}(v)$ is set at line 12 . In this case, $T_{i-1}(v)=$ null implies $\left|B_{G, 2^{i}}(v)\right| \geq\left|B_{G, 2^{i-1}}(v)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil$. The second case is that $T_{i}(v)$ is set at line 13. In this case, $\exists u \in V_{T_{i-1}(v)}=B_{G, 2^{i-1}}(v)$ such that $\left|B_{G, 2^{i-1}}(u)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil$ which implies $\left|B_{G, 2^{i}}(v)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil$. In the final case, $T_{i}(v)$ is set at line 16 , and thus, $\left|B_{G, 2^{i}}(v)\right| \geq$ $\left\lceil(m / n)^{1 / 4}\right\rceil$.

For property 3 , if $T_{r}(v) \neq$ null, then by property 1 , we know $V_{T_{r}(v)}=B_{G, 2^{r}}(v)$. By the condition in line 19, we know $V_{T_{r}(v)}=V_{T_{r-1}(v)}$ which implies $B_{G, 2^{r}}(v)=B_{G, 2^{r-1}}(v)$. Thus, $V_{T_{r}(v)}=\left\{u \in V \mid \operatorname{dist}_{G}(u, v)<\infty\right\}$.

For property 4, we can prove it by contradiction. If $r>\lceil\log (\operatorname{diam}(G))\rceil+1$, then let $i=\lceil\log (\operatorname{diam}(G))\rceil+1$. By the condition in line 19, we know there is a vertex $v \in$ $V$ such that $T_{i}(v) \neq$ null and $V_{T_{i}(v)} \neq V_{T_{i-1}(v)}$. It means that $B_{G, 2^{i}}(v) \neq B_{G, 2^{i-1}}(v)$, i.e. $\exists u \in V, \operatorname{dist}_{G}(v, u)>2^{i-1}$. But this contradicts to $i=\lceil\log (\operatorname{diam}(G))\rceil+1$. Similarly, if $r>\lceil\log (m / n)\rceil+1$, then let $i=\lceil\log (m / n)\rceil+1$. By the condition in line 19, we know there is a vertex $v \in V$ such that $T_{i}(v) \neq$ null and $V_{T_{i}(v)} \neq V_{T_{i-1}(v)}$. If $2^{i-1} \leq \operatorname{diam}(G)$, then we have $V_{T_{i}(v)} \neq V_{T_{i-1}(v)}$ which leads to a contradiction. If $2^{i-1} \geq \operatorname{diam}(G)$, then $\left|B_{G, 2^{i-1}}(v)\right| \geq 2^{i-1} \geq m / n \geq\left\lceil(m / n)^{1 / 4}\right\rceil$ which contradicts to property 2.

Next, we show how to use Algorithm 32.5 to design an algorithm which can output $|V|$ number of local shortest path trees rooted at every vertex in $V$. The details of the algorithm is described in Algorithm 32.6, and the guarantees of the algorithm is stated in the following lemma.

Lemma 32.5.3. Let $G=(V, E)$ be an undirected graph, and $m$ be a parameter which is at least $16|V|$. Let $\left(\{\widetilde{T}(v) \mid v \in V\}\right.$, $\left.\left\{\operatorname{dep}_{\widetilde{T}(v)} \mid v \in V\right\}\right)=\operatorname{MultipleLargeTrees}(G, m)$. (Algorithm 32.6) Then, the output satisfies the following properties.

1. $\forall v \in V, \widetilde{T}(v)$ is a LSPT (See Definition 32.5.1) with root $v$, and $\operatorname{dep}_{\widetilde{T}(v)}$ records the depth of every vertex in $\widetilde{T}(v)$.
2. $\forall v \in V, u \in V_{\widetilde{T}(v)}, w \in V \backslash V_{\widetilde{T}(v)}$, it satisfies $\operatorname{dist}_{G}(v, u) \leq \operatorname{dist}_{G}(v, w)$.

$$
\begin{aligned}
& \text { 3. } \forall v \in V \text {, either }\left|V_{\widetilde{T}(v)}\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil \text { or } V_{\widetilde{T}(v)}=\left\{u \in V \mid \operatorname{dist}_{G}(u, v)<\infty\right\} \text {. } \\
& \text { 4. } \forall v \in V,\left|V_{\widetilde{T}(v)}\right| \leq\left\lfloor(m / n)^{1 / 2}\right\rfloor
\end{aligned}
$$

Proof. Before we prove above properties, we first show some crucial observations.
Claim 32.5.4. $\forall v \in V, i \in\{0\} \cup[r]$, if $T_{i}(v) \neq$ null, then $T_{i}(v)$ is a LCSPT with root $v$ and radius $2^{i}$ in graph $G$. Furthermore, $\operatorname{dep}_{T_{i}(v)}: V_{T_{i}(v)} \rightarrow \mathbb{Z}_{\geq 0}$ records the depth of every vertex in $T_{i}(v)$. If $T_{i}(v)=$ null, then $\left|B_{G, 2^{i}}(v)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil$.

Proof. Follows by property 1 and property 2 of Lemma 32.5.2 directly.
Claim 32.5.5. Let $v \in V$ be a vertex with $T_{r}(v)=$ null. Then, $\forall i \in\{0\} \cup[r], \widetilde{T}_{i}(v)$ is a LCSPT (See Definition 32.5.3) with root $v$ and radius $s_{i}(v)$ in $G$, and $\operatorname{dep}_{\widetilde{T}_{i}(v)}$ records the depth of every vertex in $\widetilde{T}_{i}(v)$. Furthermore, we have $\left|B_{G, s_{i}(v)}(v)\right|<\left\lceil(m / n)^{1 / 4}\right\rceil,\left|B_{G, s_{i}(v)+2^{r-i}}(v)\right| \geq$ $\left\lceil(m / n)^{1 / 4}\right\rceil$.

Proof. Let $v \in V$ be a vertex with $T_{r}(v)=$ null. When $i=0$, then due to line 5 and line 6, $\widetilde{T}_{0}(v)$ is a LCSPT (See Definition 32.5.3) with root $v$ and radius $0=s_{0}(v)$ in $G$. According to property 2 of Lemma 32.5.2, since $T_{r}(v)=$ null, we know $\left|B_{G, 0+2^{r}}(v)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil$.

For $i \in[r]$, we prove it by induction. Suppose the claim is true for $i-1$. By Claim 32.5.4, Lemma 32.5.1 and the condition in line 9, if the procedure executes line 10 , then we know $\widetilde{T}_{i}(v)$ is a LCSPT with root $v$ and radius $s_{i-1}(v)+2^{r-i}$ in $G$ at the end of the execution of line 10 , and $\operatorname{dep}_{\widetilde{T}_{i}(v)}$ records the depth of every vertex in $\widetilde{T}_{i}(v)$. If $\left|V_{\widetilde{T}_{i}(v)}\right|<\left\lceil(m / n)^{1 / 4}\right\rceil$, then $\left|B_{G, s_{i-1}(v)+2^{r-i}}(v)\right|<\left\lceil(m / n)^{1 / 4}\right\rceil$. The procedure will execute line 11, and thus $s_{i}(v)$ is the radius of $\widetilde{T}_{i}(v)$. In addition, since $s_{i}(v)=s_{i-1}(v)+2^{r-i}$ and
$\left|B_{G, s_{i-1}(v)+2^{r-i+1}}(v)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil$, we have $\left|B_{G, s_{i}(v)+2^{r-i}}(v)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil$. If at the end of line 10, $\left|V_{\widetilde{T}_{i}(v)}\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil$, then we know $\left|B_{G, s_{i-1}(v)+2^{r-i}}(v)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil$. In this case, $\widetilde{T}_{i}(v), s_{i}(v)$ and $\operatorname{dep}_{\widetilde{T}_{i}(v)}$ will be set to be $\widetilde{T}_{i-1}(v), s_{i-1}(v)$ and $\operatorname{dep}_{\widetilde{T}_{i}(v)}$ respectively, and thus $\left|B_{G, s_{i}(v)}(v)\right|<\left\lceil(m / n)^{1 / 4}\right\rceil$. If the condition in line 9 does not hold, then we know $\left|B_{G, s_{i-1}(v)+2^{r-i}}(v)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil$ by claim 32.5.4. In this case, $\widetilde{T}_{i}(v), s_{i}(v)$ and $\operatorname{dep}_{\widetilde{T}_{i}(v)}$ will also be set to be $\widetilde{T}_{i-1}(v), s_{i-1}(v)$ and $\operatorname{dep}_{\widetilde{T}_{i}(v)}$ respectively, and thus $\left|B_{G, s_{i}(v)}(v)\right|<$ $\left\lceil(m / n)^{1 / 4}\right\rceil$.

Claim 32.5.5 shows that for each vertex $v \in V$, we know $\widetilde{T}_{r}(v)$ is a LCSPT with root $v$ and radius $s_{r}(v)$ in $G$ such that $\left|B_{G, s_{r}(v)}(v)\right|<\left\lceil(m / n)^{1 / 4}\right\rceil$ and $\left|B_{G, s_{r}(v)+1}(v)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil$.

Now, let us prove property 1 and property 2. For $v \in V$, if $T_{r}(v) \neq$ null, then $\widetilde{T}(v), \operatorname{dep}_{\widetilde{T}(v)}$ will be set to be $T_{r}(v), \operatorname{dep}_{T_{r}(v)}$ respectively. By Claim 32.5.4, the properties holds. Let $v$ be a vertex in $V$ with $T_{r}(v)=$ null. If $\widetilde{T}(v)$ is assigned at line 22 , then by Lemma 32.5.1, we know $\widetilde{T}(v)$ is a LCSPT with root $v$, and $\operatorname{dep}_{\widetilde{T}(v)}$ records the depth of every vertex in $\widetilde{T}(v)$. Thus, both properties hold. If $\widetilde{T}(v)$ is assigned at line 28 , then there are two cases for the vertices in $V_{\widetilde{T}(v)}$ :

1. If $x$ is in $V_{\widetilde{T}_{r}(v)}$, then since Claim 32.5 .5 shows $\widetilde{T}_{r}(v)$ is a LCSPT with root $v$, it is easy to show $\operatorname{dep}_{\widetilde{T}(v)}(x)=\operatorname{dep}_{\widetilde{T}_{r}(v)}(x)=\operatorname{dist}_{G}(v, x)$, and $\operatorname{par}_{\widetilde{T}(v)}(x)=\operatorname{par}_{\widetilde{T}_{r}(v)}(x) \in E$.
2. if $x$ is in $N\left(u_{v}\right)$ but not in $V_{\widetilde{T}_{r}(v)}$, then since $\widetilde{T}_{r}(v)$ is a LCSPT with root $v$ and radius $s_{r}(v), \operatorname{dist}_{G}(v, x) \geq s_{r}(v)+1$. Also notice that $\operatorname{dist}_{G}(v, x) \leq \operatorname{dist}_{G}\left(v, u_{v}\right)+\operatorname{dist}_{G}\left(u_{v}, x\right) \leq$ $s_{r}(v)+1$. Therefore, $\operatorname{dist}_{G}(v, x)=s_{r}(v)+1, \operatorname{dep}_{\widetilde{T}(v)}(x)=\operatorname{dep}_{\widetilde{T}(v)}\left(u_{v}\right)+1=\operatorname{dist}_{G}(v, x)=$ $s_{r}(v)+1$. Since $x \in N\left(u_{v}\right),\left(\operatorname{par}_{\widetilde{T}(v)}(x), x\right)=\left(u_{v}, x\right) \in E$.

Thus, $\widetilde{T}$ is a LSPT with root $v$, and it proves property 1. Due to above both cases, we know $B_{G, s_{r}(v)}(v) \subseteq V_{\widetilde{T}}$, and $\forall x \in V_{\widetilde{T}(v)}$, $\operatorname{dist}_{G}(v, x) \leq s_{r}(v)+1$. Thus, property 2 holds.

For property 3 and property 4, we have two cases. The first case is when $v$ satisfies $T_{r}(v) \neq$ null. In this case $\widetilde{T}(v)=T_{r}(v)$, due to property 3,4 of Lemma 32.5.2, we have $V_{\widetilde{T}(v)}=\left\{u \in V \mid \operatorname{dist}_{G}(v, u)<\infty\right\}$, and $\left|V_{\widetilde{T}(v)}\right|<\left\lceil(m / n)^{1 / 4}\right\rceil \leq\left\lfloor(m / n)^{1 / 2}\right\rfloor$. The second case is $T_{r}(v)=$ null. In this case, if $\widetilde{T}(v)$ is assigned at line 22 , then $V_{\widetilde{T}(v)}=B_{G, s_{r}(v)+1}(v)$. Then by Claim 32.5.5, we can get $\left|V_{\widetilde{T}(v)}\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil$. Because $\left|V_{\widetilde{T}_{r}(v)}\right|<\left\lceil(m / n)^{1 / 4}\right\rceil$ and $\forall u \in V_{\widetilde{T}_{r}(v)},\left|V_{T_{0}(u)}\right|<\left\lceil(m / n)^{1 / 4}\right\rceil$, we know $\left|V_{\widetilde{T}(v)}\right| \leq\left\lfloor(m / n)^{1 / 2}\right\rfloor$. If $\widetilde{T}(v)$ is assigned at line 28, then $\left|V_{\widetilde{T}(v)}\right| \geq\left|N\left(u_{v}\right)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil$, and $\left|V_{\widetilde{T}(v)}\right| \leq\left|V_{\widetilde{T}_{r}(v)}\right|+\left|N\left(u_{v}\right)\right|<2 \cdot\left\lceil(m / n)^{1 / 4}\right\rceil \leq$ $\left\lfloor(m / n)^{1 / 2}\right\rfloor$.

Definition 32.5.5. Let graph $G=(V, E)$, and let $m$ be a parameter which is at least $16|V|$. The number of iterations of $\left(\{\widetilde{T}(v) \mid v \in V\},\left\{\operatorname{dep}_{\widetilde{T}(v)} \mid v \in V\right\}\right)=\operatorname{MultipleLargeTrees}(G, m)$ (Algorithm 32.6) is defined as the value of $r$ in the procedure.

Lemma 32.5.6 (Number of iterations of Algorithm 32.6). Let $G=(V, E)$ be an undirected graph, and let $m$ be a parameter which is at least $16|V|$. The number of iterations (see Definition 32.5.5) of $\left(\{\widetilde{T}(v) \mid v \in V\},\left\{\operatorname{dep}_{\widetilde{T}(v)} \mid v \in V\right\}\right)=\operatorname{MultipleLargeTrees}(G, m)$ (Algorithm 32.6) is at most $\min (\lceil\log (\operatorname{diam}(G))\rceil,\lceil\log (m / n)\rceil)+1$.

Proof. It follows by property 4 of Lemma 32.5.2 directly.

### 32.5.3 Path Generation and Root Changing

In this section, we show a procedure which can output a path from a certain vertex to the root in a rooted tree. Then we show how to use the procedure to change the root of a rooted tree to a certain vertex in the tree. To output the vertex-root path, we have two stages. The first stage is using doubling method to compute the depth and the $2^{i}$ th (for all $i \in\{0,1, \cdots, \log (\operatorname{dep})\}$ ) ancestor of each vertex. The second stage is using divide-and-conquer technique to split the path into segments, and recursively find the path for each segment. Once we have the procedure to find the vertex-root path, then we can use it to implement root-changing. The idea is very simple, if we want to change the root to a certain vertex, we just need to find the path from that vertex to the root, and reverse the parent pointers of every vertex on the path. The path finding procedure is described in Algorithm 32.8. The root changing procedure is described in Algorithm 32.9.

Definition 32.5.6. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$. The number of iterations of FindAncestors(par) is defined as the value of $r$ at the end of the procedure.

Lemma 32.5.7. Let par : $V \rightarrow V$ be $a$ set of parent pointers (See Definition 32.4.2) on $a$ vertex set V. Let $\left(r, \operatorname{dep}_{\mathrm{par}},\left\{g_{i} \mid i \in\{0\} \cup[r]\right\}\right)=$ FindAnCestors(par) (Algorithm 32.7). Then the number of iterations (see Definition 32.5.6) r should be at most $\lceil\log (\operatorname{dep}(\operatorname{par})+1)\rceil$, $\operatorname{dep}_{\mathrm{par}}: V \rightarrow \mathbb{Z}_{\geq 0}$ records the depth of every vertex in $V$, and $\forall i \in\{0\} \cup[r], v \in V g_{i}(v)=$ $\operatorname{par}^{\left(2^{i}\right)}(v)$.

Proof. $h_{l}$ and $g_{l}$ will satisfies the properties in the following claim.

Claim 32.5.8. $\forall i \in\{0\} \cup[r], v \in V g_{i}(v)=\operatorname{par}^{\left(2^{i}\right)}(v)$, and if $\operatorname{dep}_{\mathrm{par}}(v) \leq 2^{i}-1$ then $h_{i}(v)=\operatorname{dep}_{\mathrm{par}}(v)$. Otherwise $\operatorname{dep}_{\mathrm{par}}(v)=$ null.

Proof. The proof is by induction. The claim is obviously true when $i=0$. Suppose the claim is true for $i-1$. We have $g_{i}(v)=g_{i-1}\left(g_{i-1}(v)\right)=\operatorname{par}^{\left(2^{i-1}\right)}\left(\operatorname{par}^{\left(2^{i-1}\right)}(v)\right)=\operatorname{par}^{\left(2^{i}\right)}(v)$. If $h_{i}(v) \neq$ null, then there are two cases. In the first case, we have $h_{i}(v)=h_{i-1}(v)$. By induction we know $h_{i}(v)=\operatorname{dep}_{\mathrm{par}}(v)$. In the second case, we have $h_{i}(v)=h_{i-1}\left(g_{i-1}(v)\right)+2^{i-1}=$ $\operatorname{dep}_{\mathrm{par}}\left(\operatorname{par}^{\left(2^{i-1}\right)}(v)\right)+2^{i-1}$. Notice that in this case $h_{i-1}(v)=$ null, thus by the induction, $\operatorname{dep}_{\mathrm{par}}(v) \geq 2^{i-1}$. Therefore, $\operatorname{dep}_{\mathrm{par}}(v)=\operatorname{dep}_{\mathrm{par}}\left(\operatorname{par}^{\left(2^{i-1}\right)}(v)\right)+2^{i-1}=h_{i}(v)$. If $h_{i}(v)=$ null, then it means that $h_{i-1}\left(\operatorname{par}^{\left(2^{i-1}\right)}(v)\right)=$ null which implies that $\operatorname{dep}_{\mathrm{par}}(v) \geq 2^{i}$.

Due to the above claim, we know that if $i \geq\left\lceil\log \left(\operatorname{dep}_{\mathrm{par}}(v)+1\right)\right\rceil$ then $h_{i}(v) \neq$ null. Thus, we have $r \leq\lceil\log (\operatorname{dep}(\operatorname{par})+1)\rceil$. Since the procedure returns $h_{r}$ as $\operatorname{dep}_{\mathrm{par}}$, the returned $\operatorname{dep}_{\mathrm{par}}$ is correct.

Lemma 32.5.9. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$. Let $q$ be a vertex in $V$. Let $\left(\operatorname{dep}_{\text {par }}, P, w\right)=\operatorname{FindPath(par,q)}$ (Algorithm 32.8). Then $\operatorname{dep}_{\mathrm{par}}: V \rightarrow \mathbb{Z}_{\geq 0}$ records the depth of every vertex in $V$ and $P \subseteq V$ is the set of all vertices on the path from $q$ to the root of $q$, i.e. $P=\left\{v \in V \mid \exists k \geq 0, v=\operatorname{par}^{(k)}(q)\right\}$. If $\operatorname{dep}_{\mathrm{par}}(q) \geq 1$, then $w=\operatorname{par}^{\left(\operatorname{dep}_{\mathrm{par}}(q)-1\right)}(q)$. Furthermore, $k$ should be at most $\lceil\log (\operatorname{dep}(\operatorname{par}))\rceil$.

Proof. By Lemma 32.5.7, since $\left(r, \operatorname{dep}_{\mathrm{par}},\left\{g_{i} \mid i \in\{0\} \cup[r]\right\}\right)=\operatorname{FindAncestors}($ par $)$, we know $r$ should be at most $\lceil\log (\operatorname{dep}($ par $)+1)\rceil$, $\operatorname{dep}_{p a r}: V \rightarrow \mathbb{Z}_{\geq 0}$ records the depth of every vertex in $V$, and $\forall i \in\{0\} \cup[r], v \in V g_{i}(v)=\operatorname{par}^{\left(2^{i}\right)}(v)$. Thus $k=\left\lceil\log \left(\operatorname{dep}_{\mathrm{par}}(q)\right)\right\rceil \leq$ $\lceil\log (\operatorname{dep}($ par $)+1)\rceil$

Now let us prove that $P$ is the vertex set of all the vertices on the path from $q$ to the root of $q$. We use divide-and-conquer to get $P$. The following claim shows that $S_{i}$ is a set of segments which is a partition of the path, and each segment has length at most $2^{k-i}$.

Claim 32.5.10. $\forall i \in\{0\} \cup[k], S_{i}$ satisfies the following properties:

1. $\exists(x, y) \in S_{i}$ such that $x=q$.
2. $\exists(x, y) \in S_{i}$ such that $y=g_{r}(q)$.
3. $\forall(x, y) \in S_{i}, \operatorname{dep}_{\mathrm{par}}(y)-\operatorname{dep}_{\mathrm{par}}(x) \leq 2^{k-i}$.
4. $\forall(x, y) \in S_{i}$, if $y \neq g_{r}(q)$, then $\exists\left(x^{\prime}, y^{\prime}\right) \in S_{i}, x^{\prime}=y$.
5. $\forall(x, y) \in S_{i}, \exists j \in \mathbb{Z}_{\geq 0}, \operatorname{par}^{(j)}(x)=y$.

Proof. Our proof is by induction. According to line 4, all the properties hold when $i=0$. Suppose all the properties hold for $i-1$. For property 1, by induction we know there exists $(x, y) \in S_{i-1}$ such that $x=q$. Then by line 8 and line 9 , there must be an $\left(x, y^{\prime}\right)$ in $S_{i}$. For property 2, by induction we know there exists $(x, y) \in S_{i-1}$ such that $y=g_{r}(q)$. Thus, there must be an $\left(x^{\prime}, y\right)$ in $S_{i}$. For property 3 , if $(x, y)$ is added into $S_{i}$ by line 9 , then $\operatorname{dep}_{\mathrm{par}}(x)-\operatorname{dep}_{\mathrm{par}}(y) \leq 2^{k-i}$. Otherwise, in line 8 , we have $\operatorname{dep}_{\mathrm{par}}(x)-\operatorname{dep}_{\mathrm{par}}\left(g_{k-i}(x)\right) \leq$ $2^{k-i}, \operatorname{dep}_{\mathrm{par}}\left(g_{k-i}(x)\right)-\operatorname{dep}_{\mathrm{par}}(y) \leq 2^{k-i+1}-2^{k-i}=2^{k-i}$. For property 4 , if $(x, y)$ is added into $S_{i}$ by line 9 , then by induction there is $\left(y, y^{\prime}\right) \in S_{i-1}$, and thus by line 9 and line 8 , there must be $\left(y, y^{\prime \prime}\right) \in S_{i}$. Otherwise, in line 8 will generate two pairs $\left(x, g_{k-i}(x)\right),\left(g_{k-i}(x), y\right)$. For $\left(x, g_{k-i}(x)\right)$, the property holds. For $\left(g_{k-i}(x), y\right)$, there must be $\left(y, y^{\prime}\right) \in S_{i-1}$ and thus there should be $\left(y, y^{\prime \prime}\right) \in S_{i}$. For property 5 , since $g_{k-i}(x)=\operatorname{par}^{(k-i)}(x)$, for all pairs generated by line 8 and line 9 , the property holds.

By Claim 32.5.10, we know
$S_{k}=\left\{(q, \operatorname{part}(q)),\left(\operatorname{part}(q), \operatorname{par}^{(2)}(q)\right),\left(\operatorname{par}^{(2)}(q), \operatorname{par}^{(3)}(q)\right), \cdots,\left(\operatorname{par}^{\left(\operatorname{dep}_{\operatorname{par}}(q)-1\right)}(q), \operatorname{par}^{\left(\operatorname{dep}_{\operatorname{par}}(q)\right)}(q)\right)\right\}$.
Thus, $P$ is the set of all the vertices on the path from $q$ to the root of $q$. And $w=$ $\operatorname{par}^{\left(\operatorname{dep}_{\mathrm{par}}(q)-1\right)}(q)$ when $\operatorname{dep}_{\mathrm{par}}(q) \geq 1$.

Lemma 32.5.11. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$. Let $q$ be a vertex in $V$. Let $\widehat{\text { par }}=$ RootChange(par, q) (Algorithm 32.9).
 if $\operatorname{par}^{(\infty)}(v)=\operatorname{par}^{(\infty)}(q)$ then $\widehat{\operatorname{par}}^{(\infty)}(v)=q$. Otherwise $\widehat{\operatorname{par}}^{(\infty)}(v)=\operatorname{par}^{(\infty)}(v) . \forall u \neq v \in$ $V, \operatorname{part}(v)=u \Leftrightarrow$ either $\widehat{\operatorname{par}}(v)=u$ or $\widehat{\operatorname{par}}(u)=v$. Furthermore, $\operatorname{dep}(\widehat{\operatorname{par}}) \leq 2 \operatorname{dep}($ par $)$.

Proof. For a vertex $v \in V$, if $\left\{u \mid i \in \mathbb{Z}_{\geq 0}, u=\operatorname{par}^{(i)}(v)\right\} \cap P=\emptyset$, then we have $\forall i \in$ $\mathbb{Z}_{\geq 0}, \operatorname{par}^{(i)}(v)=\widehat{\operatorname{par}^{(i)}}(v)$. According to Lemma 32.5.9, $P=\left\{u \in V \mid i \in \mathbb{Z}_{\geq 0}, \operatorname{par}^{(i)}(q)=u\right\}$. Then for all $v \in P \backslash\{q\}$, if $\operatorname{part}(u)=v$ then $\widehat{\operatorname{par}}(v)=u$. Thus, $\forall u \in P, \widehat{\operatorname{par}}^{(\infty)}(u)=q$. Let $i^{*}$ be the smallest number such that $\operatorname{par}^{\left(i^{*}\right)}(v) \in P$. Then $\widehat{\operatorname{par}}^{\left(i^{*}\right)}(v) \in P$. Thus, $\widehat{\operatorname{par}}^{(\infty)}(v)=$ $\widehat{\operatorname{par}}{ }^{(\infty)}\left(\widehat{\operatorname{par}}^{\left(i^{*}\right)}(v)\right)=q$. Furthermore, we have $\forall v \in V, \operatorname{dep}(\widehat{\operatorname{par}}) \leq \operatorname{dep}(\operatorname{par})+\operatorname{dep}_{\mathrm{par}}(q) \leq$ 2dep(par).

### 32.5.4 Spanning Forest Expansion

In this section, we give the definition of spanning forest. If we are given a spanning forest of a contracted graph and spanning trees of each contracted component, then we show a procedure which can merge them to get a spanning forest of the original graph. Before go to the details, let us formally define the spanning forest.

Definition 32.5.7 (Rooted Spanning Forest). Let $G=(V, E)$ be an undirected graph. Let par : $V \rightarrow V$ be a set of parent pointers which is compatible (Definition 32.4.4) with $G$. If $\forall u, v \in V, \operatorname{dist}_{G}(u, v)<\infty \Rightarrow \operatorname{par}^{(\infty)}(u)=\operatorname{par}^{(\infty)}(v)$, and $\forall v \in V, \operatorname{part}(v) \neq v \Rightarrow$ $(v, \operatorname{part}(v)) \in E$, then we call par a rooted spanning forest of $G$.

The Algorithm 32.10 shows how to combine the spanning forest in the contracted graph with local spanning trees to get a spanning forest in the graph before contraction. Figure 32.1 shows an example.

Lemma 32.5.12. Let $G_{2}=\left(V_{2}, E_{2}\right)$ be an undirected graph. Let $\widetilde{\operatorname{par}}: V_{2} \rightarrow V_{2}$ be a set of parent pointers (See Definition 32.4.2) which satisfies that $\forall v \in V_{2}$ with $\widetilde{\operatorname{par}}(v) \neq v$, $(v, \widetilde{\operatorname{par}}(v))$ must be in $E_{2}$. Let $G_{1}=\left(V_{1}, E_{1}\right)$ be an undirected graph satisfies $V_{1}=\left\{v \in V_{2} \mid\right.$ $\widetilde{\operatorname{par}}(v)=v\}, E_{1}=\left\{(u, v) \in V_{1} \times V_{1} \mid u \neq v, \exists(x, y) \in E_{2}, \widetilde{\operatorname{par}^{(\infty)}}(x)=u, \widetilde{\operatorname{par}^{(\infty)}}(y)=v\right\}$. Let par : $V_{1} \rightarrow V_{1}$ be a rooted spanning forest (See Definition 32.5.7) of $G_{1}$. Let $f: V_{1} \times V_{1} \rightarrow$ $\{$ null $\} \cup\left(V_{2} \times V_{2}\right)$ satisfy the following property: for $u \neq v \in V_{1}$, if $\operatorname{part}(u)=v$, then $f(u, v) \in\left\{(x, y) \in E_{2} \mid \widetilde{\operatorname{par}}^{(\infty)}(x)=u, \widetilde{\operatorname{par}}^{(\infty)}(y)=v\right\}$, and $f(v, u) \in\left\{(x, y) \in E_{2} \mid\right.$ $\left.\widetilde{\operatorname{par}}{ }^{(\infty)}(x)=v, \widetilde{\operatorname{par}}{ }^{(\infty)}(y)=u\right\}$. Let $\widehat{\operatorname{par}}=$ FORESTEXPANSION $($ par, $\widetilde{\operatorname{par}}, f)$. Then $\widehat{\operatorname{par}}: V_{2} \rightarrow$ $V_{2}$ is a rooted spanning forest of $G_{2}$. In addition, $\operatorname{dep}(\widehat{\operatorname{par}}) \leq(2 \cdot \operatorname{dep}(\widetilde{\operatorname{par}})+1)(\operatorname{dep}(\operatorname{par})+1)$.

Proof. Let $x, y \in V_{2}, u=\widetilde{\operatorname{par}^{(\infty)}}(x), v=\widetilde{\operatorname{par}^{(\infty)}}(y) \in V_{1}$ if $\operatorname{dist}_{G_{2}}(x, y)<\infty$, then since $E_{1}=\left\{\left(u^{\prime}, v^{\prime}\right) \in V_{1} \times V_{1} \mid u^{\prime} \neq v^{\prime}, \exists\left(x^{\prime}, y^{\prime}\right) \in E_{2}, \widetilde{\operatorname{par}}{ }^{(\infty)}\left(x^{\prime}\right)=u^{\prime}, \widetilde{\operatorname{par}}^{(\infty)}\left(y^{\prime}\right)=v^{\prime}\right\}$, it must be true that $\operatorname{dist}_{G_{1}}(u, v)<\infty$. Since par is a spanning forest of $G_{1}$, we have $\operatorname{par}^{(\infty)}(u)=$ $\operatorname{par}^{(\infty)}(v)$. It suffices to say $\forall x \in V_{2}, \widehat{\operatorname{par}}^{(\infty)}(x)=\operatorname{par}^{(\infty)}\left(\widetilde{\operatorname{par}}^{(\infty)}(x)\right)$. We can prove it by
 this case, we have $\widehat{\operatorname{par}}^{(\infty)}(x)=\widetilde{\operatorname{par}}^{(\infty)}(x)=u=\operatorname{par}^{(\infty)}(u)=\operatorname{par}^{(\infty)}\left(\widetilde{\operatorname{par}}^{(\infty)}(x)\right)$, and also we have $\operatorname{dep}_{\widetilde{\mathrm{par}}}(x)=\operatorname{dep}_{\widetilde{\mathrm{par}}}(x)$. Now suppose for all $x \in V_{2}$ with $\operatorname{dep}_{\mathrm{par}}(\widetilde{\operatorname{par}}(\infty)(x)) \leq i-1$, it has $\widehat{\operatorname{par}}^{(\infty)}(x)=\operatorname{par}^{(\infty)}\left(\widetilde{\operatorname{par}}{ }^{(\infty)}(x)\right)$ and $\operatorname{dep}_{\widehat{\operatorname{par}}}(x) \leq i \cdot(2 \operatorname{dep}(\widetilde{\operatorname{par}})+1)$. Let $y \in V_{2}$ satisfy $\operatorname{dep}_{\mathrm{par}}(\widetilde{\operatorname{par}}(\infty)(y))=i$. Let $v=\widetilde{\operatorname{par}^{(\infty)}}(y)$. By line 8 and the properties of $f$, we know $\widetilde{\operatorname{par}}^{(\infty)}\left(x_{v}\right)=v$, and $\widetilde{\operatorname{par}}{ }^{(\infty)}\left(y_{v}\right)=\operatorname{part}(v)$. By line 9, line 10 and Lemma 32.5.11, we have $\widehat{\operatorname{par}}_{v}^{(\infty)}(y)=x_{v}, \widehat{\operatorname{par}}\left(x_{v}\right)=y_{v}$. Thus, there must be $k \leq 2 \operatorname{dep}_{\widehat{\mathrm{par}}}(y)$ such that $\widehat{\operatorname{par}}^{(k)}(y)=x_{v}$. Since $\widehat{\operatorname{par}}^{(\infty)}\left(y_{v}\right)=\operatorname{par}^{(\infty)}(v)$ and $\operatorname{dep}_{\widehat{\operatorname{par}}}\left(y_{v}\right) \leq i \cdot(2 \operatorname{dep}(\widetilde{\operatorname{par}})+1)$, we have $\widehat{\operatorname{par}}^{(\infty)}(y)=$ $\operatorname{par}^{(\infty)}(v)=\operatorname{par}^{(\infty)}\left(\widetilde{\operatorname{par}}{ }^{(\infty)}(y)\right)$ and $\operatorname{dep}_{\widehat{\operatorname{par}}}(y) \leq(i+1) \cdot(2 \operatorname{dep}(\widetilde{\operatorname{par}})+1)$.

In addition, by the properties of $f$ and Lemma 32.5.11, $\forall v \in V_{2}$ with $\widehat{\operatorname{par}}(v) \neq v$, we have $(v, \widehat{\operatorname{par}}(v)) \in E_{2}$. To conclude, $\widehat{\operatorname{par}}: V_{2} \rightarrow V_{2}$ is a spanning forest of $G_{2}$, and $\operatorname{dep}(\widehat{\operatorname{par}}) \leq(\operatorname{dep}(\operatorname{par})+1)(2 \operatorname{dep}(\widetilde{\operatorname{par}})+1)$.

### 32.5.5 Spanning Forest Algorithm

In this section, we show how to apply the ideas shown in connectivity algorithm to get an spanning forest algorithm. Algorithm 32.11 can output a spanning forest of a graph $G$, but the edges are not orientated. Then in the Algorithm 32.12, we assign each forest edge an direction thus it is a rooted spanning forest.

Before we prove the correctness of the algorithms, let us briefly introduce the meaning of each variables appeared in the algorithms.

In Algorithm 32.11, $G_{0}$ is the original input graph, for $i \in\{0\} \cup[r-1], G_{i}^{\prime}$ is obtained by deleting all the small size connected components in $G_{i}$, and $G_{i+1}$ is obtained by contracting some vertices of $G_{i}^{\prime}$. For a vertex $v$ in graph $G_{i}$, if $h_{i}(v)=$ null, then it means that the connected component which contains $v$ is deleted when obtaining $G_{i}^{\prime}$. If $h_{i}(v) \neq$ null, it means that the vertex $v$ is contracted to the vertex $h_{i}(v)$ when obtaining $G_{i+1}$. $\operatorname{par}_{i}$ is a rooted forest (may not be spanning) in graph $G_{i}$, if a tree from the forest is spanning in $G_{i}$, then all the vertex in that tree will be deleted when obtaining $G_{i}^{\prime}$. Otherwise all the vertices in that tree will be contracted to the root, and the root will be one of the vertex in $G_{i+1}$. Since each connected component in $G_{i+1}$ is obtained by contraction of some vertices in a connected component in $G_{i}$, each edge in $G_{i+1}$ must correspond to an edge in $G_{i}$ where the end vertices of the edge are contracted to different vertices. Thus, each edge in $G_{i}$ should correspond to an edge in $G$, and $g_{i}: E_{i} \rightarrow E$ records the such correspondence. $D_{i}$ records the edges added to the spanning forest $F$ in the $i^{\text {th }}$ round. For each vertex $v$ in graph $G_{i}$, $\widetilde{T}_{i}(v)$ is a local shortest path tree (See definition 32.5.1) which is either with a large size or is a spanning tree in the component of $v . L_{i}$ is a set of random leaders in $G_{i}^{\prime}$ such that in each local shortest path tree $\widetilde{T}_{i}(v)$, there is at least one leader shown in the tree. The following
lemmas formally state the properties of the algorithm.

Lemma 32.5.13. Let $G=(V, E)$ be an undirected graph, $m$ be a parameter which is at least $16|V|$, and $r$ be a rounds parameter. If SpanningForest $(G, m, r)$ (Algorithm 32.11) does not return FAIL, then $\operatorname{diam}(G)=\operatorname{diam}\left(G_{0}\right) \geq \operatorname{diam}\left(G_{0}^{\prime}\right) \geq \operatorname{diam}\left(G_{1}\right) \geq \operatorname{diam}\left(G_{1}^{\prime}\right) \geq \cdots \geq$ $\operatorname{diam}\left(G_{r}\right)$.

Proof. By property 3 of Lemma 32.5.3, $\forall[i] \in\{0\} \cup[r-1]$, there is no edge between $V_{i} \backslash V_{i}^{\prime}$ and $V_{i}^{\prime}$. Thus, $\operatorname{diam}\left(G_{i}^{\prime}\right) \leq \operatorname{diam}\left(G_{i}\right)$. Then due to property 1 of Corollary 32.4.8, we have $\operatorname{diam}\left(G_{i+1}\right) \leq \operatorname{diam}\left(G_{i}^{\prime}\right)$.

Lemma 32.5.14. Let $G=(V, E)$ be an undirected graph, $m$ be a parameter which is at least $16|V|$, and $r$ be a rounds parameter. If SpanningForest ( $G, m, r$ ) (Algorithm 32.11) does not return $F A I L$, then $\forall i \in\{0\} \cup[r-1], \operatorname{dep}\left(\operatorname{par}_{i}\right) \leq \min \left(\operatorname{diam}(G),\left\lfloor\left(m / n_{i}\right)^{1 / 2}\right\rfloor\right)$.

Proof. Let $v \in V_{i}$. If $v \in V_{i} \backslash V_{i}^{\prime}$, then due to property 3 of Lemma 32.5.3, we have $V_{\widetilde{T}_{i}(v)}=$ $V_{\widetilde{T}_{i}\left(u_{v}\right)}$. Due to Lemma 32.5.13 and Lemma 32.5.3, we have $\operatorname{dep}_{\operatorname{par}_{i}}(v) \leq \operatorname{dep}\left(\widetilde{T}_{i}\left(u_{v}\right)\right) \leq$ $\min \left(\operatorname{diam}(G),\left\lfloor\left(m / n_{i}\right)^{1 / 2}\right\rfloor\right)$. For $v \in V_{i}$, we define $\operatorname{dist}_{G_{i}}\left(v, L_{i}\right)=\min _{u \in L_{i}} \operatorname{dist}_{G_{i}}(v, u)$. By Lemma 32.5.3, we know $\operatorname{dist}_{G_{i}}\left(v, L_{i}\right)=\operatorname{dist}_{G_{i}}\left(v, z_{i}(v)\right)$. Since $\widetilde{T}_{i}(v)$ is a LSPT (See Definition 32.5.1), by applying Lemma 32.5.9, we know $\operatorname{dist}_{G_{i}}\left(v, L_{i}\right)=\operatorname{dist}_{G_{i}}\left(w_{i}(v), L_{i}\right)+1$, and $\left(v, w_{i}(v)\right) \in E_{i}$. Thus, by induction on $\operatorname{dist}_{G_{i}}\left(v, L_{i}\right)$, we can get $\operatorname{dep}_{\operatorname{par}_{i}}(v) \leq \operatorname{dist}_{G_{i}}\left(v, L_{i}\right)$. By Lemma 32.5.13 and Lemma 32.5.3, we can conclude $\operatorname{dep}\left(\operatorname{par}_{i}\right)(v) \leq \min \left(\operatorname{diam}(G),\left\lfloor\left(m / n_{i}\right)^{1 / 2}\right\rfloor\right)$.

Lemma 32.5.15. Let $G=(V, E)$ be an undirected graph, $m$ be a parameter which is at least $16|V|$, and $r$ be a rounds parameter. If SpanningForest $(G, m, r)$ (Algorithm 32.11) does
not return FAIL, then $\forall i \in\{-1,0\} \cup[r-1]$, we can define

$$
\forall v \in V, h^{(i)}(v)= \begin{cases}v & i=-1 \\ h_{i}\left(h^{(i-1)}(v)\right) & h^{(i-1)}(v) \neq \text { null } \\ \text { null } & \text { otherwise }\end{cases}
$$

Then we have following properties:

1. If $h^{(i)}(v) \neq$ null, then $h^{(i)}(v) \in V_{i+1}$.
2. $\forall u, v \in V, h^{(i)}(u) \neq h^{(i)}(v),(u, v) \in E$, we have $\left(h^{(i)}(u), h^{(i)}(v)\right) \in E_{i+1}$.
3. $\forall(x, y) \in E_{i+1},(u, v)=g_{i+1}(x, y)$, we have $(u, v) \in E, h^{(i)}(u)=x, h^{(i)}(v)=y$.

Proof. For property 1, we can prove it by induction. It is true for $i=-1$. If $h_{0}(v) \neq$ null, we know $h_{0}(v)$ must be assigned at line 19. Due to property 2 of Lemma 32.4.6, $h_{0}(v) \in V_{1}$. Suppose $\forall v \in V, h^{(i-1)}(v) \neq$ null, we have $h^{(i-1)}(v) \in V_{i}$. For a vertex $v$ with $h^{(i)}(v) \neq$ null, according to the definition of $h^{(i)}(v)$, we know $h^{(i-1)}(v) \neq$ null. Let $u=h^{(i-1)}(v)$. u must be a vertex in $G_{i}$ by the induction hypothesis. Since $h^{(i)}(v) \neq$ null, we know $h_{i}(u) \neq$ null. Thus, $h_{i}(u)$ must be assigned at line 19. Due to property 2 of Lemma 32.4.6, $h_{i}(u)$ must be in $G_{i+1}$, which implies $h^{(i)}(v) \in V_{i+1}$.

For property 2 , we can also prove it by induction. It is true for $i=-1$. If $(u, v) \in E$, then due to property 3 of Lemma 32.5.3, either both $u, v$ are in $V_{0}^{\prime}$ or both $u, v$ are in $V_{0} \backslash V_{0}^{\prime}$. If both $u, v$ are in $V_{0} \backslash V_{0}^{\prime}$, then $h_{0}(u)=h_{0}(v)=$ null. Otherwise, if $h_{0}(u) \neq h_{0}(v)$, then due to property 3 of Lemma 32.4.6, $\left(h_{0}(u), h_{0}(v)\right) \in E_{1}$. Now suppose we have $\forall u, v \in V$, if $h^{(i-1)}(u) \neq h^{(i-1)}(v),(u, v) \in E$, then $\left(h^{(i-1)}(u), h^{(i-1)}(v)\right) \in E_{i}$. Let $(u, v) \in E, h^{(i)}(u) \neq$ $h^{(i)}(v)$. Let $x=h^{(i-1)}(u), y=h^{(i-1)}(v)$. Due to property 3 of Lemma 32.5.3, either both $x, y$
are in $V_{i}^{\prime}$ or both are in $V_{i} \backslash V_{i}^{\prime}$. If $x, y \in V_{i} \backslash V_{i}^{\prime}$, then $h_{i}(x)=h_{i}(y)=$ null which contradicts to $h^{(i)}(u) \neq h^{(i)}(v)$. Thus, both of $x, y \in V_{i}^{\prime}$. Then due to property 3 of Lemma 32.4.6, $\left(h_{i}(x), h_{i}(v)\right) \in E_{i+1}$. Thus, $\left(h^{(i)}(u), h^{(i)}(v)\right) \in E_{i+1}$.

For property 3, we can prove it by induction. It is true for $i=-1$. Let us consider the case when $i=0$. Due to property 3 of Lemma 32.4.6 and the definition of $g_{0}, g_{1}$, we have $\forall(x, y) \in E_{1},(u, v)=g_{1}(x, y), h_{0}(u)=x, h_{0}(v)=y,(u, v) \in E$. Now suppose the property holds for $i-1$. Let $(x, y) \in E_{i+1}$. Then $g_{i+1}(x, y)=g_{i}\left(x^{\prime}, y^{\prime}\right)$ for some $\left(x^{\prime}, y^{\prime}\right) \in E_{i}, h_{i}\left(x^{\prime}\right)=$ $x, h_{i}\left(y^{\prime}\right)=y$. Let $(u, v)=g_{i}\left(x^{\prime}, y^{\prime}\right)$. By the induction hypothesis $(u, v) \in E, h^{(i-1)}(u)=$ $x^{\prime}, h^{(i-1)}(v)=y^{\prime}$. Thus, $h^{(i)}(u)=x, h^{(i)}(v)=y$.

Lemma 32.5.16. Let $G=(V, E)$ be an undirected graph, $m$ be a parameter which is at least $16|V|$, and $r$ be a rounds parameter. If SpanningForest $(G, m, r)$ (Algorithm 32.11) does not return $F A I L$, then $\forall i \in\{-1,0\} \cup[r-1]$, we can define

$$
\forall v \in V, h^{(i)}(v)= \begin{cases}v & i=-1 \\ h_{i}\left(h^{(i-1)}(v)\right) & h^{(i-1)}(v) \neq \text { null } \\ \text { null } & \text { otherwise }\end{cases}
$$

Let $\forall i \in\{0\} \cup[r], \widehat{G}_{i}=\left(V_{i}, \widehat{E}_{i}=\left\{(x, y) \mid(u, v) \in \bigcup_{j=i}^{r-1} D_{j}, h^{(j-1)}(u)=x, h^{(j-1)}(v)=y\right\}\right)$. Then $\widehat{G}_{i}$ is a spanning forest of $G_{i}$.

Proof. The proof is by induction. When $i=r$, since $V_{r}=\emptyset, \widehat{G}_{r}=(\emptyset, \emptyset)$ is a spanning forest of $G_{r}$. Now suppose $\widehat{G}_{i+1}$ is a spanning forest of $G_{i+1}$. Let $u, v \in V_{i}$. By property 2,3 of Lemma 32.5.15, we have $\widehat{E}_{i} \subseteq E_{i}$. Thus, if $\operatorname{dist}_{G_{i}}(u, v)=\infty$, then $\operatorname{dist}_{\widehat{G}_{i}}(u, v)=\infty$. If $\operatorname{dist}_{G_{i}}(u, v)<\infty$, there are several cases:

1. If $h_{i}(u)=h_{i}(v)=$ null, then due to line 10 , we know $u_{u}=u_{v}$, and $\widetilde{T}_{i}\left(u_{v}\right)$ is a spanning tree of the component which contains $u, v$. Thus, $\widehat{G}_{i}$ has a spanning tree of the component which contains $u, v$.
2. If $h_{i}(u)=h_{i}(v) \neq$ null, then $\operatorname{par}_{i}:\left\{x \in V_{i} \mid h_{i}(x)=h_{i}(v)\right\} \rightarrow\left\{x \in V_{i} \mid h_{i}(x)=h_{i}(v)\right\}$ is a tree, and $\forall y \in\left\{x \in V_{i} \mid h_{i}(x)=h_{i}(v)\right\}$, if $\operatorname{par}_{i}(y) \neq y$, then $\left(y, \operatorname{par}_{i}(y)\right) \in \widehat{E}_{i}$. Since $\widehat{G}_{i+1}$ does not have any cycle, there is a unique path from $u$ to $v$ in $\widehat{G}_{i}$.
3. If $h_{i}(u) \neq h_{i}(v)$, then neither of them can be null. Since $\widehat{G}_{i+1}$ is a spanning forest on $G_{i+1}$, there must be a unique path from $h_{i}(u)$ to $h_{i}(v)$ in $\widehat{G}_{i+1}$. Suppose the path in $\widehat{G}_{i+1}$ is $h_{i}(u)=p_{1}-p_{2}-\cdots-p_{t}=h_{i}(v)$. Then there must be a sequence of vertices in $G_{i}, u=p_{1,1}, p_{1,2}, p_{2,1}, p_{2,2}, \cdots, p_{t, 1}, p_{t, 2}=v$ such that $h_{i}\left(p_{j, 1}\right)=h_{i}\left(p_{j, 2}\right)=p_{j}$ and $\left(p_{j-1,2}, p_{j, 1}\right) \in \widehat{E}_{i}$. Thus, there is a unique path from $u$ to $v$.

Thus, $\widehat{G}_{i}$ is a spanning forest of $G_{i}$.

Corollary 32.5.17 (Correctness of Algorithm 32.11). Let $G=(V, E)$ be an undirected graph, $m$ be a parameter which is at least $16|V|$, and $r$ be a rounds parameter. If $\operatorname{SpAnningForest}(G, m, r)$ (Algorithm 32.11) does not return FAIL, then $\widehat{G}_{0}=(V, F)$ is a spanning forest of $G$.

Proof. Just apply Lemma 32.5 .16 for $i=0$ case.

Lemma 32.5.18. Let $G=(V, E)$ be an undirected graph, $m$ be a parameter which is at least $16|V|$, and $r$ be a rounds parameter. If SpanningForest $(G, m, r)$ (Algorithm 32.11) does
not return FAIL, then $\forall i \in\{-1,0\} \cup[r-1]$, we can define

$$
\forall v \in V, h^{(i)}(v)= \begin{cases}v & i=-1 \\ h_{i}\left(h^{(i-1)}(v)\right) & h^{(i-1)}(v) \neq \text { null } \\ \text { null } & \text { otherwise }\end{cases}
$$

$\forall i \in\{0\} \cup[r-1], v \in V_{i}$ with $\operatorname{par}_{i}(v) \neq v$, there exists $(x, y) \in F$ such that $h^{(i-1)}(x)=$ $v, h^{(i-1)}(y)=\operatorname{par}_{i}(v)$.

Proof. By line 11, line 19, $\forall i \in\{0\} \cup[r-1], v \in V_{i}, \operatorname{par}_{i}(v) \neq v$, we have $g_{i}\left(v, \operatorname{par}_{i}(v)\right), g_{i}\left(\operatorname{par}_{i}(v), v\right) \in$ $D_{i} \subseteq F$. Since $\left(\operatorname{par}_{i}(v), v\right) \in E_{i}$, by property 3 of Lemma 32.5.15, $(x, y)=g_{i}\left(v, \operatorname{par}_{i}(v)\right)$ satisfies $h^{(i-1)}(x)=v, h^{(i-1)}(y)=\operatorname{par}_{i}(v)$.

Theorem 32.5.19 (Correctness of Algorithm 32.12). Let $G=(V, E)$ be an undirected graph, $m$ be a parameter which is at least $16|V|$, and $r$ be a rounds parameter. If $\operatorname{SpAnNingForest}(G, m, r)$ (Algorithm 32.11) does not return FAIL, then let the output be the input of Orientate (•), (Algorithm 32.12) and the output par : $V \rightarrow V$ of $\operatorname{Orientate}(\cdot)$ will be a rooted spanning forest (See Definition 32.5.7) of $G$. Furthermore, $\operatorname{dep}(\operatorname{par}) \leq O(\operatorname{diam}(G))^{r}$.

Proof. The proof is by induction. We want to show $\widehat{\operatorname{par}}_{i}$ is a rooted spanning forest of $G_{i}$. When $i=r$, since $V_{r}=\emptyset$, the claim is true. Now suppose we have $\widehat{\operatorname{par}}_{i+1}$ is a spanning forest of $G_{i+1}$. Let $\widetilde{G}_{i+1}=\left(\widetilde{V}_{i+1}, E_{i+1}\right)$. It is easy to see $\widetilde{\operatorname{par}}_{i+1}: \widetilde{V}_{i+1} \rightarrow \widetilde{V}_{i+1}$ is a spanning forest of $\widetilde{G}_{i+1}$. An observation is $\widetilde{V}_{i+1}=\left\{v \in V_{i} \mid \operatorname{par}_{i}(v)=v\right\}$. Thus, $\widetilde{\operatorname{par}}_{i+1}, \operatorname{par}_{i}$ satisfies the condition in Lemma 32.5.12 when invoking ForestExpansion $\left(\widetilde{\operatorname{par}}_{i+1}, \operatorname{par}_{i}, f_{i+1}\right)$. By Lemma 32.5.18, we know $f_{i+1}$ also satisfies the condition in Lemma 32.5.12 when we invoke $\operatorname{FORESTEXPANSION}\left(\widetilde{\operatorname{par}}_{i+1}, \operatorname{par}_{i}, f_{i+1}\right)$. Thus, $\widehat{\operatorname{par}}_{i}$ is a rooted spanning forest of $G_{i}$ due to Lemma 32.5.12.

By Lemma 32.5.12, we have $\operatorname{dep}\left(\widehat{\operatorname{par}}_{i}\right) \leq 16 \operatorname{dep}\left(\widehat{\operatorname{par}}_{i+1}\right) \operatorname{diam}(G)$. By induction, we have $\operatorname{dep}(\operatorname{par}) \leq O(\operatorname{diam}(G))^{r}$.

Lemma 32.5.20. Let $G=(V, E)$ be an undirected graph, $m$ be a parameter which is at least $16|V|$, and $r \leq n$ be a round parameter. If $\operatorname{SpAnNingForest}(G, m, r)$ (Algorithm 32.11) does not return FAIL, then with probability at least 0.89, $\sum_{i=0}^{r} n_{i} \leq 40 n$.

Proof. Since $V_{r} \subseteq V_{r-1} \subseteq \cdots \subseteq V_{0}=V$, we have $n_{r} \leq n_{r-1} \leq n_{r-2} \leq \cdots \leq n$. Due to line 14 , line 15 and line 17 , we know $\forall i \in\{0\} \cup[r-1]$, $V_{i+1}=L_{i}$. If $p_{i}<1 / 2$, we know $p_{i}=(30 \log (n)+100) / \gamma_{i}$. Since $\left|V_{\widetilde{T}_{i}}(v)\right| \geq \gamma_{i}$, we can apply Lemma 32.4.3 to get $\operatorname{Pr}\left(\left|L_{i}\right| \leq 1.5 p_{i} n_{i}\right) \geq \operatorname{Pr}\left(\left|L_{i}\right| \leq 0.75 n_{i}\right) \geq 1-1 /(100 n)$. By taking union bound over all $i \in\{0\} \cup[r-1]$, with probability at least 0.99 , if $p_{i}<0.5$, then $n_{i+1} \leq 0.75 n_{i}$. By applying Lemma 32.4.4, condition on $n_{i}$ and $p_{i}=\frac{1}{2}$, we have $\mathbb{E}\left(n_{i+1}\right) \leq 0.75 n_{i}$. By Markov's inequality, with probability at 0.89 , we have $\sum_{i=0}^{r} n_{i} \leq 40 n$.

Now let us define the total iterations of Algorithm 32.11 as the following:

Definition 32.5.8 (Total iterations). Let graph $G=(V, E), m \leq \operatorname{poly}(|V|)$ be a parameter which is at least $16|V|$, and $r$ be a rounds parameter. The total number of iterations of $\operatorname{SpanningForest}(G, m, r)$ (Algorithm 32.11) is defined as $\sum_{i=0}^{r-1}\left(k_{i}+k_{i}^{\prime}\right)$, where $\forall i \in\{0\} \cup[r-1], k_{i}$ denotes the number of iterations (See Definition 32.5.5) of $\operatorname{MultipleLargeTrees}\left(G_{i}, m\right)$ (see line 8), and $k_{i}^{\prime}$ denotes the number of iterations (See Definition 32.4.5) of TreeContraction $\left(G_{i}^{\prime}, \operatorname{par}_{i}\right)$ (see line 17).

Theorem 32.5.21 (Success probability of Algorithm 32.11). Let $G=(V, E)$ be an undirected graph. Let $m \leq \operatorname{poly}(n)$ and $m \geq 16|V|$. Let $r$ be a rounds parameter. Let $c>0$
be a sufficiently large constant. If $r \geq c \log \log _{m / n} n$, then with probability at least 0.79 , SpanningForest $(G, m, r)$ (Algorithm 32.11) does not return FAIL. Furthermore, let $\forall i \in$ $\{0\} \cup[r-1], k_{i}$ be the number of iterations (See Definition 32.5.5) of MultipleLargeTrees $\left(G_{i}, m\right)$ and $k_{i}^{\prime}$ be the number of iterations (See Definition 32.4.5) of TreeContraction $\left(G_{i}^{\prime}, \operatorname{par}_{i}\right.$ : $\left.V_{i}^{\prime} \rightarrow V_{i}^{\prime}\right)$. Let $c_{1}>0$ be a sufficiently large constant. If $m \geq c_{1} n \log ^{8} n$, then with probability at least $0.99, \sum_{i=0}^{r-1} k_{i}+k_{i}^{\prime} \leq O\left(\min \left(\log (\operatorname{diam}(G)) \log _{\log }^{\operatorname{diam}(G)}(n), r \log (\operatorname{diam}(G))\right)\right)$. If $m<c_{1} n \log ^{8} n$, then with probability at least $0.98, \sum_{i=0}^{r-1} k_{i}^{\prime}+k_{i} \leq O(\min (\log (\operatorname{diam}(G))$ $\left.\left.\log \log _{\operatorname{diam}(G)}(n)+(\log \log n)^{2}, r \log (\operatorname{diam}(G))\right)\right)$.

Proof. Due to Lemma 32.5.20, with probability at last 0.89 , we have $\forall i \in[r], n_{i}^{\prime} \leq 40 n$. Thus, we can condition on that $\operatorname{SpanningForest}(G, m, r)$ will not fail on line 21.

Due to Lemma 32.5.6, $k_{i} \leq O\left(\log \left(\operatorname{diam}\left(G_{i}\right)\right)\right) \leq O(\log (\operatorname{diam}(G)))$. Due to Corollary 32.4.8 and Lemma 32.5.14, $k_{i}^{\prime} \leq O(\log (\operatorname{diam}(G)))$. Thus, $\sum_{i \in\{0\} \cup[r-1]} k_{i}^{\prime}+k_{i} \leq O(r \log (\operatorname{diam}(G)))$.

Since $V_{r} \subseteq V_{r-1} \subseteq V_{r-2} \subseteq \cdots \subseteq V_{0}=V$, we have $n_{r} \leq n_{r-1} \leq n_{r-2} \leq \cdots \leq n$. Due to line 14 , line 15 and line 17 , we know $\forall i \in\{0\} \cup[r-1]$, $V_{i+1}=L_{i}$. If $p_{i}<1 / 2$, we know $p_{i}=(30 \log (n)+100) / \gamma_{i}$. Since $\left|V_{\widetilde{T}_{i}}(v)\right| \geq \gamma_{i}$, we can apply Lemma 32.4.3 to get $\operatorname{Pr}\left(\left|L_{i}\right| \leq 1.5 p_{i} n_{i}\right) \geq 1-1 /(100 n)$. By taking union bound over all $i \in\{0\} \cup[r-1]$, with probability at least 0.99 , if $p_{i}<0.5$, then $n_{i+1} \leq 1.5 p_{i} n_{i} \leq 0.75 n_{i}$. Let $\mathcal{E}$ be the event that $\forall i \in\{0\} \cup[r-1]$, if $p_{i}<0.5$, then $n_{i+1} \leq 1.5 p_{i} n_{i}$. Now, we suppose $\mathcal{E}$ happens.

If $p_{0}=0.5$, then $m \leq n \cdot(600 \log n)^{8}$. By applying Lemma 32.4.4, $\mathbb{E}\left(n_{i+1}\right)=\mathbb{E}\left(\left|L_{i}\right|\right) \leq$ $0.75 \mathbb{E}\left(n_{i}\right) \leq \cdots \leq 0.75^{i+1} n$. By Markov's inequality, when $i^{*} \geq 8 \log (6000 \log n) / \log (4 / 3)$, with probability at least $0.99, n_{i^{*}} \leq n /(600 \log n)^{8}$ and thus $p_{i^{*}}<0.5$. Condition on this
event and $\mathcal{E}$, we have

$$
\begin{aligned}
n_{r} & \leq \frac{\left(\frac{\left(\frac{n_{i *}^{1.25}}{m_{0.25}(45 \log n+150)}\right)^{1.25}}{m^{0.25}}(45 \log n+150)\right)^{\cdots}}{\cdots} \quad \quad \quad \text { (Apply } r^{\prime}=r-i^{*} \text { times) } \\
& =n_{i^{*}} /\left(m / n_{i^{*}}\right)^{1.25^{r^{\prime}}-1} \cdot(45 \log n+150)^{4 \cdot\left(1.25^{r^{\prime}}-1\right)} \\
& \leq n /\left(m /\left(n_{i^{*}}(45 \log n+150)^{4}\right)\right)^{1.25^{r^{\prime}-1}} \\
& \leq n /\left(m /\left(n_{i^{*}}(45 \log n+150)^{4}\right)\right)^{1.25^{r^{\prime} / 2}} \leq n /(m / n)^{1.25^{r^{\prime} / 2}} \leq 1 / 2,
\end{aligned}
$$

where the second inequality follows by $n_{i^{*}} \leq n$, the third inequality follows by $r^{\prime} \geq 5$, the forth inequality follows by $n_{i^{*}} \leq n /(600 \log n)^{8}$, and the last inequality follows by $r^{\prime} \geq \frac{2}{\log 1.25} \log \log _{m / n}(2 n)$. Since $16 n \leq m \leq n \cdot(600 \log n)^{8}, \log \log _{m / n} n=\Theta(\log \log n)$. Let $c>0$ be a sufficiently large constant. Thus, when $r \geq c \log \log _{m / n} n \geq i^{*}+r^{\prime}=$ $8 \log (6000 \log n) / \log (4 / 3)+\frac{2}{\log 1.25} \log \log _{m / n}(2 n)$, with probability at least $0.98, n_{r}=0 \mathrm{im}-$ plies that $\operatorname{SpanningForest}(G, m, r)$ will not fail. Due to Lemma 32.5.6, we have $k_{i} \leq$
$O\left(\min \left(\log \left(m / n_{i}\right), \log (\operatorname{diam}(G))\right)\right)$.Thus,

$$
\begin{aligned}
& \sum_{i=0}^{r-1} k_{i} \\
= & \sum_{i=0}^{i^{*}} k_{i}+\sum_{i=i^{*}+1}^{r-1} k_{i} \\
\leq & O\left((\log \log n)^{2}\right)+\sum_{i=i^{*}+1}^{r-1} k_{i} \\
\leq & O\left((\log \log n)^{2}\right)+\sum_{i: i \geq i^{*}+1, m / n_{i} \leq \operatorname{diam}(G)} k_{i}+\sum_{i: i \leq r, m / n_{i}>\operatorname{diam}(G)} k_{i} \\
\leq & O\left((\log \log n)^{2}\right)+O\left(\sum_{i=0}^{\left\lceil\log _{1.25} \log _{2}(\operatorname{diam}(G)) 7\right.} \log \left(2^{1.25^{i}}\right)\right)+O\left(\sum_{i=0}^{\left\lceil\log _{1.25} \log _{\operatorname{diam}(G)}(m)\right\rceil} \log (\operatorname{diam}(G))\right) \\
\leq & O\left((\log \log n)^{2}\right)+O(\log (\operatorname{diam}(G)))+O\left(\log (\operatorname{diam}(G)) \log \log _{\operatorname{diam}(G)}(n)\right) \\
\leq & O\left(\log (\operatorname{diam}(G)) \log \log _{\operatorname{diam}(G)}(n)+(\log \log (n))^{2}\right),
\end{aligned}
$$

where the first inequality follows by $i^{*}=O(\log \log n)$ and $\forall i \leq\left[i^{*}\right], m / n_{i} \leq \operatorname{poly}(\log n)$, the third inequality follows by $m / n_{i+1} \geq\left(m / n_{i}\right)^{1.25} /(45 \log n+100) \geq\left(m / n_{i}\right)^{1.125}$. Due to Corollary 32.4.8 and Lemma 32.5.14, we also have $k_{i}^{\prime} \leq O\left(\min \left(\log \left(m / n_{i}\right), \log (\operatorname{diam}(G))\right)\right)$. Then, by the same argument, we have $\sum_{i=0}^{r-1} k_{i}^{\prime}=O\left(\log (\operatorname{diam}(G)) \log \log _{\operatorname{diam}(G)}(n)+(\log \log (n))^{2}\right)$.

If $m>n \cdot(600 \log n)^{8}$, then $\forall i \in\{0\} \cup[r-1]$, we have $p_{i}<0.5$. Since $\mathcal{E}$ happens.

We have:

$$
\begin{aligned}
n_{r} & \leq \frac{\left(\frac{\left(\frac{n^{1.25}}{m^{0.25}}(45 \log n+150)\right)^{1.25}}{m^{0.25}}(45 \log n+150)\right)}{\cdots} \\
& =\frac{n^{1.25^{r}}}{m^{1.25^{r}-1}}(45 \log n+150)^{4 \cdot\left(1.25^{r}-1\right)} \\
& =n /(m / n)^{1.25^{r}-1} \cdot(45 \log n+150)^{4 \cdot\left(1.25^{r}-1\right)} \\
& =n /\left(m /\left(n(45 \log n+150)^{4}\right)\right)^{1.25^{r}-1} \\
& \leq n /\left(m /\left(n(45 \log n+150)^{4}\right)\right)^{1.25^{r / 2}} \\
& \leq n /\left(m /\left(n(200 \log n)^{4}\right)\right)^{1.25^{r / 2}} \\
& \leq \frac{1}{2},
\end{aligned}
$$

where the second inequality follows by $r \geq 5$, the third inequality follows by $45 \log n+150 \leq$ $200 \log n$, and the last inequality follows by

$$
r \geq c \log \log _{m / n} n \geq 2 \log _{1.25} \log _{(m / n)^{1 / 2}} 2 n \geq 2 \log _{1.25} \log _{m /\left(n(200 \log n)^{4}\right)} 2 n,
$$

for a sufficiently large constant $c>0$. Since $n_{r}$ is an integer, $n_{r}$ must be 0 when $n_{r} \leq 1 / 2$. SpanningForest $(G, m, r)$ will succeed with probability at least 0.99. Due to Lemma 32.5.6, we have $k_{i} \leq O\left(\min \left(\log \left(m / n_{i}\right), \log (\operatorname{diam}(G))\right)\right)$. Thus,

$$
\begin{aligned}
\sum_{i=0}^{r-1} k_{i} & \leq \sum_{m / n_{i} \leq \operatorname{diam}(G)} k_{i}+\sum_{m / n_{i}>\operatorname{diam}(G)} k_{i} \\
& \leq O\left(\sum_{i=0}^{\left\lceil\log _{1.25} \log _{2}(\operatorname{diam}(G))\right\rceil} \log \left(2^{1.25^{i}}\right)\right)+O\left(\sum_{i=0}^{\left\lceil\log _{1.25} \log _{\operatorname{diam}(G)}(m)\right\rceil} \log (\operatorname{diam}(G))\right) \\
& \leq O(\log (\operatorname{diam}(G)))+O\left(\log (\operatorname{diam}(G)) \log \log _{\operatorname{diam}(G)}(n)\right),
\end{aligned}
$$

where the second inequality follows by $m / n_{i+1} \geq\left(m / n_{i}\right)^{1.25} /(45 \log n+100) \geq\left(m / n_{i}\right)^{1.125}$. Due to Corollary 32.4.8 and Lemma 32.5.14, we also have $k_{i}^{\prime} \leq O\left(\min \left(\log \left(m / n_{i}\right), \log (\operatorname{diam}(G))\right)\right)$. Then, by the same argument, we have $\sum_{i=0}^{r-1} k_{i}^{\prime}=O(\log (\operatorname{diam}(G)))+O\left(\log (\operatorname{diam}(G)) \log \log _{\operatorname{diam}(G)}(n)\right)$.

### 32.6 Depth-First-Search Sequence for Tree and Applications 32.6.1 Lowest Common Ancestor and Multi-Paths Generation

Given a rooted forest induced by par : $V \rightarrow V$ which is a set of parent pointers (See Definition 32.4.2) on $V$, and a set of $q$ queries $Q=\left\{\left(u_{1}, v_{1}\right),\left(u_{2}, v_{2}\right), \cdots,\left(u_{q}, v_{q}\right) \mid u_{i}, v_{i} \in V\right\}$, we show an algorithm which can return a mapping lca : $Q \rightarrow(V \cup\{$ null $\}) \times(V \cup\{$ null $\}) \times$ $(V \cup\{$ null $\})$ such that $\forall\left(u_{i}, v_{i}\right) \in Q,\left(p, p_{u_{i}}, p_{v_{i}}\right)=\operatorname{lca}\left(u_{i}, v_{i}\right)$ satisfies the following properties:

1. If $\operatorname{par}^{(\infty)}\left(u_{i}\right)=\operatorname{par}^{(\infty)}\left(v_{i}\right)$, then $p$ is the lowest ancestor of $u_{i}$ and $v_{i}$. Otherwise $p=p_{u_{i}}=p_{v_{i}}=$ null.
2. Suppose $p \neq$ null. If $p \neq u_{i}$, then $p_{u_{i}}$ is an ancestor of $u_{i}$ and $\operatorname{part}\left(p_{u_{i}}\right)=p$. Otherwise, $p_{u_{i}}=$ null.
3. Suppose $p \neq$ null. If $p \neq v_{i}$, then $p_{v_{i}}$ is an ancestor of $v_{i}$ and $\operatorname{part}\left(p_{v_{i}}\right)=p$. Otherwise, $p_{v_{i}}=$ null.

Before we describe the algorithms, let us formally define ancestor and the lowest common ancestor.

Definition 32.6.1 (Ancestor). Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$. For $u, v \in V$, if $\exists k \in \mathbb{Z}_{\geq 0}$ such that $u=\operatorname{par}^{(k)}(v)$, then $u$ is an ancestor of $v$.

Definition 32.6.2 (Common ancestor and the lowest common ancestor). par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$. For $u, v \in V$, if $w$ is an ancestor of $u$ and is also an ancestor of $v$, then $w$ is a common ancestor of $(u, v)$. If a common
ancestor $w$ of $(u, v)$ satisfies $\operatorname{dep}_{\mathrm{par}}(w) \geq \operatorname{dep}_{\mathrm{par}}(x)$ for any common ancestor $x$ of $(u, v)$, then $w$ is the lowest common ancestor (LCA) of $(u, v)$.

Definition 32.6.3 (Path between two vertices). par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$. For $u, v \in V$, if $\operatorname{par}^{(\infty)}(u)=\operatorname{par}^{(\infty)}(v)$, then the path from $u$ to $v$ is a sequence $\left(x_{1}, x_{2}, \cdots, x_{j}, x_{j+1}, \cdots x_{k}\right)$ such that $\forall i \neq i^{\prime} \in[k], x_{i} \neq$ $x_{i^{\prime}}, x_{1}=u, x_{k}=v, x_{j}$ is the lowest common ancestor of $(u, v), \forall i \in[j-1], \operatorname{part}\left(x_{i}\right)=x_{i+1}$, and $\forall i \in\{j+1, j+2, \cdots, k\}, \operatorname{part}\left(x_{i}\right)=x_{i-1}$.

The algorithm which can compute the lowest common ancestor is described in Algorithm 32.13.

Lemma 32.6.1. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$. Let $Q=\left\{\left(u_{1}, v_{1}\right),\left(u_{2}, v_{2}\right), \cdots,\left(u_{q}, v_{q}\right)\right\}$ be a set of $q$ pairs of vertices, and $\forall i \in[q], u_{i} \neq v_{i}$. Let lca $=\mathrm{LCA}($ par, $Q)$ (Algorithm 32.13). Then for any $(u, v) \in Q$, $\left(p, p_{u}, p_{v}\right)=\operatorname{lca}(u, v)$ satisfies the following properties:

1. If $\operatorname{par}^{(\infty)}(u) \neq \operatorname{par}^{(\infty)}(v)$, then $p=p_{u}=p_{v}=$ null.
2. If $u$ (or $v$ ) is the lowest common ancestor of $(u, v)$, then $p=u, p_{u}=$ null, $p_{v} \neq u$ is an ancestor of $v$ such that $\operatorname{part}\left(p_{v}\right)=u$ (or $p=v, p_{v}=$ null, $p_{u} \neq v$ is an ancestor of $u$ such that $\operatorname{part}\left(p_{u}\right)=v$.)
3. If neither $u$ nor $v$ is the lowest common ancestor of $(u, v)$ and $\operatorname{par}^{(\infty)}(u)=\operatorname{par}^{(\infty)}(v)$, then $p$ is the lowest common ancestor of $(u, v), p_{u} \neq p$ is an ancestor of $u, p_{v} \neq p$ is an ancestor of $v$, and $\operatorname{part}\left(p_{u}\right)=\operatorname{part}\left(p_{v}\right)=p$.

Proof. According to Lemma 32.5.7, $r$ should be at most $\lceil\log (\operatorname{dep}(\operatorname{par})+1)\rceil$, $\operatorname{dep}_{\text {par }}: V \rightarrow \mathbb{Z}_{\geq 0}$ records the depth of every vertex in $V$, and $\forall i \in\{0\} \cup[r], v \in V g_{i}(v)=\operatorname{par}^{\left(2^{i}\right)}(v)$. Then property 1 follows by line 5 directly.

Then for all $(u, v) \in Q$ with $\operatorname{par}^{(\infty)}(u)=\operatorname{par}^{(\infty)}(v)$, either $(u, v) \in Q^{\prime}$ or $(v, u) \in Q^{\prime}$. For each $(u, v) \in Q^{\prime}$, we have $\operatorname{dep}_{\mathrm{par}}(u) \geq \operatorname{dep}_{\mathrm{par}}(v)$. For all $(u, v) \in Q^{\prime}$, with $\operatorname{dep}_{\mathrm{par}}(u)>$ $\operatorname{dep}_{\mathrm{par}}(v)$, by induction we can prove that $\forall i \in\{0\} \cup[r-1],(x, y)=h_{i}(u, v)$ satisfies that $x$ is an ancestor of $u, y=v, \operatorname{dep}_{\mathrm{par}}(x)>\operatorname{dep}_{\mathrm{par}}(v)$ and $\operatorname{par}^{\left(2^{i}\right)}(x)$ is an ancestor of $v$. Thus, for $\left(p, p_{u}, p_{v}\right)=\operatorname{lca}(u, v)$, if $v$ is the lowest common ancestor of $u$, then we have $p=v, p_{u}=h_{0}(u), p_{v}=$ null. In this case, $h_{0}(u)$ is an ancestor of $u$, and $\operatorname{dep}_{\mathrm{par}}(u)=$ $\operatorname{dep}_{\mathrm{par}}(v)+1, \operatorname{part}\left(h_{0}(u)\right)=v$. Thus, property 2 holds.

For all $(u, v) \in Q$ with $\operatorname{par}^{(\infty)}(u)=\operatorname{par}^{(\infty)}(v)$, if neither $u$ nor $v$ is the lowest common ancestor of $(u, v)$, then we know either $(u, v)$ or $(v, u)$ is in $Q^{\prime \prime}$. Now let $(u, v) \in Q^{\prime \prime}$. We have $\operatorname{dep}_{\mathrm{par}}\left(h_{r}^{\prime}(u)\right)=\operatorname{dep}_{\mathrm{par}}\left(h_{r}^{\prime}(v)\right), h_{r}^{\prime}(u) \neq h_{r}^{\prime}(v)$, and $h_{r}^{\prime}(u), h_{r}^{\prime}(v)$ are ancestors of $u, v$ respectively. We can prove by induction to get $\forall i \in\{0\} \cup[r], h_{i}^{\prime}(u) \neq h_{i}^{\prime}(v)$ and $\operatorname{par}^{\left(2^{i}\right)}\left(h_{i}^{\prime}(u)\right)=$ $\operatorname{par}^{\left(2^{i}\right)}\left(h_{i}^{\prime}(v)\right)$ is a common ancestor of $(u, v)$. Thus, $p=\operatorname{part}\left(h_{0}^{\prime}(u)\right)=\operatorname{part}\left(h_{0}^{\prime}(v)\right)$ is the lowest common ancestor of $(u, v)$, and $\operatorname{dep}_{\mathrm{par}}\left(h_{0}^{\prime}(u)\right)=\operatorname{dep}_{\mathrm{par}}\left(h_{0}^{\prime}(v)\right)=\operatorname{dep}_{\mathrm{par}}(p)+1$. Since $p_{u}=h_{0}^{\prime}(u), p_{v}=h_{0}^{\prime}(v)$, property 3 holds.

In Algorithm 32.14, we show a generalization of Algorithm 32.8 such that we can find multiple vertex-to-ancestor paths simultaneously.

The following lemma claims the properties of the outputs of Algorithm 32.14. And the proof is similar to the proof of Lemma 32.5.9.

Lemma 32.6.2. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$. Let $Q=\left\{\left(u_{1}, v_{1}\right),\left(u_{2}, v_{2}\right), \cdots,\left(u_{q}, v_{q}\right)\right\} \subseteq V \times V$ satisfy $\forall j \in[q], v_{j}$ is an ancestor (See Definition 32.6.1) of $u_{j}$ in par. Let $\left(\operatorname{dep}_{\mathrm{par}},\left\{P_{j} \mid j \in[q]\right\}\right)=\operatorname{MultiPath}(\operatorname{par}, Q)$ (Algorithm 32.14). Then $\operatorname{dep}_{\mathrm{par}}: V \rightarrow \mathbb{Z}_{\geq 0}$ records the depth of every vertex in $V$ and $\forall j \in[q], P_{j} \subseteq V$ is the set of all vertices on the path from $u_{j}$ to $v_{j}$, i.e. $P_{j}=\{v \in$ $\left.V \mid \exists k_{1}, k_{2} \in \mathbb{Z}_{\geq 0}, v=\operatorname{par}^{\left(k_{1}\right)}\left(u_{j}\right), v_{j}=\operatorname{par}^{\left(k_{2}\right)}(v)\right\}$. Furthermore, $r$ should be at most $\lceil\log (\operatorname{dep}(\operatorname{par})+1)\rceil$.

Proof. By Lemma 32.5.7, since $\left(r, \operatorname{dep}_{\mathrm{par}},\left\{g_{i} \mid i \in\{0\} \cup[r]\right\}\right)=$ FindAncestors(par), we know $r$ should be at most $\lceil\log (\operatorname{dep}(\operatorname{par})+1)\rceil$, $\operatorname{dep}_{\text {par }}: V \rightarrow \mathbb{Z}_{\geq 0}$ records the depth of every vertex in $V$, and $\forall i \in\{0\} \cup[r], v \in V g_{i}(v)=\operatorname{par}^{\left(2^{i}\right)}(v)$.

For $j \in[q]$, let us prove that $P_{j}$ is the vertex set of all the vertices on the path from $u_{j}$ to its ancestor $v_{j}$. We use divide-and-conquer to get $P_{j}$. The following claim shows that $S_{j}^{(i)}$ is a set of segments which is a partition of the path from $u_{j}$ to $v_{j}$, and each segment has length at most $2^{r-i}$.

Claim 32.6.3. $\forall i \in\{0\} \cup[r], j \in[q] S_{j}^{(i)}$ satisfies the following properties:

1. $\exists(x, y) \in S_{j}^{(i)}$ such that $x=u_{j}$.
2. $\exists(x, y) \in S_{j}^{(i)}$ such that $y=v_{j}$.
3. $\forall(x, y) \in S_{j}^{(i)}$, $\operatorname{dep}_{\mathrm{par}}(y)-\operatorname{dep}_{\mathrm{par}}(x) \leq 2^{r-i}$.
4. $\forall(x, y) \in S_{j}^{(i)}$, if $y \neq v_{j}$, then $\exists\left(x^{\prime}, y^{\prime}\right) \in S_{j}^{(i)}, x^{\prime}=y$.
5. $\forall(x, y) \in S_{j}^{(i)}, \exists k \in \mathbb{Z}_{\geq 0}, \operatorname{par}^{(k)}(x)=y$.

Proof. We fix a $j \in[q]$. Our proof is by induction. According to line 4 , all the properties hold when $i=0$. Suppose all the properties hold for $i-1$. For property 1 , by induction we know there exists $(x, y) \in S_{j}^{(i-1)}$ such that $x=u_{j}$. Then by line 9 and line 10 , there must be an $\left(x, y^{\prime}\right)$ in $S_{j}^{(i)}$. For property 2, by induction we know there exists $(x, y) \in S_{j}^{(i-1)}$ such that $y=$ $v_{j}$. Thus, there must be an $\left(x^{\prime}, y\right)$ in $S_{j}^{(i)}$. For property 3 , if $(x, y)$ is added into $S_{j}^{(i)}$ by line 10, then $\operatorname{dep}_{\mathrm{par}}(x)-\operatorname{dep}_{\mathrm{par}}(y) \leq 2^{r-i}$. Otherwise, in line 9 , we have $\operatorname{dep}_{\mathrm{par}}(x)-\operatorname{dep}_{\mathrm{par}}\left(g_{r-i}(x)\right) \leq$ $2^{r-i}, \operatorname{dep}_{\mathrm{par}}\left(g_{r-i}(x)\right)-\operatorname{dep}_{\mathrm{par}}(y) \leq 2^{r-i+1}-2^{r-i}=2^{r-i}$. For property 4 , if $(x, y)$ is added into $S_{j}^{(i)}$ by line 10 , then by induction there is $\left(y, y^{\prime}\right) \in S_{j}^{(i-1)}$, and thus by line 10 and line 9 , there must be $\left(y, y^{\prime \prime}\right) \in S_{j}^{(i)}$. Otherwise, in line 9 will generate two pairs $\left(x, g_{r-i}(x)\right),\left(g_{r-i}(x), y\right)$. For $\left(x, g_{r-i}(x)\right)$, the property holds. For $\left(g_{r-i}(x), y\right)$, there must be $\left(y, y^{\prime}\right) \in S_{i-1}$ and thus there should be $\left(y, y^{\prime \prime}\right) \in S^{(i)}$. For property 5 , since $g_{r-i}(x)=\operatorname{par}^{(r-i)}(x)$, for all pairs generated by line 9 and line 10, the property holds.

By Claim 32.6.3, we know

$$
\begin{aligned}
& S_{j}^{(r)}=\{ \\
&\left(u_{j}, \operatorname{part}\left(u_{j}\right)\right), \\
&\left(\operatorname{part}^{\left.\left(u_{j}\right), \operatorname{par}^{(2)}\left(u_{j}\right)\right),}\right. \\
&\left(\operatorname{par}^{(2)}\left(u_{j}\right), \operatorname{par}^{(3)}\left(u_{j}\right)\right), \\
& \cdots, \\
&\left(\operatorname{par}^{\left(\operatorname{dep}_{\operatorname{par}}\left(u_{j}\right)-\operatorname{dep}_{\operatorname{par}}\left(v_{j}\right)-1\right)}\left(u_{j}\right), \operatorname{par}^{\left(\operatorname{dep}_{\operatorname{par}}\left(u_{j}\right)-\operatorname{dep}_{\operatorname{par}}\left(v_{j}\right)\right)}\left(u_{j}\right)\right) \\
&\}
\end{aligned}
$$

Thus, $P_{j}$ is the set of all the vertices on the path from $u_{j}$ to an ancestor $v_{j}$.

### 32.6.2 Depth-First-Search Sequence for a Tree

Since we can use our spanning tree algorithm to get a rooted tree, in this section, we only consider how to get a Depth-First-Search (DFS) sequence for a rooted tree. Before we go to the details, let us firstly give formal definitions of some useful concepts.

Definition 32.6.4 (Children in the forest). Given a set of parent pointers (See Definition 32.4.2) par : $V \rightarrow V$ on a vertex set $V . \forall u, v \in V, u \neq v$ if $\operatorname{part}(u)=v$, then we say $u$ is a child of $v . \forall v \in V$, we can define $\operatorname{child}_{\text {par }}(v)$ as the set of all children of $v$, i.e. $\operatorname{child}_{\mathrm{par}}(v)=\{u \in V \mid u \neq v, \operatorname{part}(u)=v\}$. Furthermore, if $u$ is the $k^{\text {th }}$ smallest vertex in the children set child $\mathrm{par}_{\mathrm{par}}(v)$, then we say $\operatorname{rank}_{\mathrm{par}}(u)=k$, or $u$ is the $k^{\text {th }}$ child of $v$. If $\operatorname{part}(v)=v$, then $\operatorname{rank}_{\mathrm{par}}(v)=1$. We use child $\mathrm{par}_{\mathrm{par}}(v, k)$ to denote the $k^{\text {th }}$ child of $v$.

For simplicity of the notation, if par : $V \rightarrow V$ is clear in the context, we just use $\operatorname{child}(v), \operatorname{rank}(v)$ and $\operatorname{child}(v, k)$ to denote $\operatorname{child}_{\mathrm{par}}(v), \operatorname{rank}_{\mathrm{par}}(v)$ and $\operatorname{child} \mathrm{par}_{\mathrm{par}}(v, k)$ respectively.

Definition 32.6.5 (Leaves in the forest). Given a set of parent pointers (See Definition 32.4.2) par : $V \rightarrow V$ on a vertex set $V$. If $\operatorname{child}_{\text {par }}(v)=\emptyset$, then $v$ is called a leaf. The set of all the leaves of par is defined as leaves $(\operatorname{par})=\left\{v \mid \operatorname{child}_{\text {par }}(v)=\emptyset\right\}$.

Definition 32.6.6 (Subtree). Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$. Let $v \in V, V^{\prime}=\{u \in V \mid v$ is an ancestor (Definition 32.6.1) of $u\}$. Let par ${ }^{\prime}: V^{\prime} \rightarrow V^{\prime}$ be a set of parent pointers on $V^{\prime}$. If $\forall u \in V^{\prime} \backslash\{v\}, \operatorname{par}^{\prime}(u)=\operatorname{part}(u)$, and $\operatorname{par}^{\prime}(v)=v$, then we say par' is the subtree of $v$ in par. For $u \in V^{\prime}$, we say $u$ is in the subtree of $v$.

Definition 32.6.7 (Depth-First-Search (DFS) sequence). Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$. Let $v$ be a vertex in $V$. If $v$ is a leaf (See Definition 32.6.5) in par, then the DFS sequence of the subtree (See Definition 32.6.6) of $v$ is $(v)$. Otherwise the DFS sequence of the subtree of $v$ in par is recursively defined as

$$
\left(v, a_{1,1}, a_{1,2}, \cdots, a_{1, n_{1}}, v, a_{2,1}, a_{2,2}, \cdots, a_{2, n_{2}}, v, \cdots, a_{k, 1}, a_{k, 2}, \cdots, a_{k, n_{k}}, v\right)
$$

where $k=|\operatorname{child}(v)|$ is the number of children (See Definition 32.6.4) of $v$, and $\forall i \in$ $[k],\left(a_{i, 1}, \cdots, a_{i, n_{i}}\right)$ is the DFS sequence of the subtree of child $(v, i)$, i.e. the $i^{\text {th }}$ child of $v$.

If $\forall u \in V, \operatorname{par}^{(\infty)}(u)=v$, then the subtree of $v$ is exactly par, and thus the DFS sequence of the subtree of $v$ is also called the DFS sequence of par.

Here are some useful facts of the above defined DFS sequence.
Fact 32.6.4. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$, and par has a unique root. Let $A=\left(a_{1}, a_{2}, \cdots, a_{m}\right)$ be the DFS sequence (See Definition 32.6.7) of par. Then, A satisfies the following properties:

1. $\forall v \in V, v$ appears exactly $|\operatorname{child}(v)|+1$ times in $A$.
2. If $a_{i}$ is the $k^{\text {th }}$ time that $v$ appears, and $a_{j}$ is the $(k+1)^{\text {th }}$ time that $v$ appears. Then $\left(a_{i+1}, a_{i+2}, \cdots, a_{j-1}\right)$ is the DFS sequence of the subtree of $\operatorname{child}(v, k)$ (See Definition 32.6.4), the $k^{\text {th }}$ child of $v$. Furthermore, $a_{i+1}$ is the first time that $\operatorname{child}(v, k)$ appears, and $a_{j-1}$ is the last time of $\operatorname{child}(v, k)$ appears.
3. If $a_{i}$ is the first time that $v$ appears, and $a_{j}$ is the last time that $v$ appears. Then $\left(a_{i}, a_{i+1}, \cdots, a_{j}\right)$ is the DFS sequence of the subtree of $v$.

$$
\text { 4. } m=2|V|-1 \text {. }
$$

Proof. The property 1, 2,3 directly follows by Definition 32.6.7.
For property 4 , notice that $\forall u \in V, \operatorname{part}(u) \neq u, u$ can only be a child of $\operatorname{part}(u)$. Thus, $\sum_{v \in V}(|\operatorname{child}(v)|+1)=|V|-1+|V|=2|V|-1$.

Due to the above fact, if $v$ is a leaf in par, then it will only once in the DFS sequence. Thus, we are able to determine the order of all the leaves.

Definition 32.6.8 (The order of the leaves). Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$, and par has a unique root. Let $A=\left(a_{1}, a_{2}, \cdots, a_{m}\right)$ be the DFS sequence (See Definition 32.6.7) of par . Let $u, v$ be two leaves (See definition 32.6.5) of par. If $u$ appears before $v$ in $A$, then we say $u<_{\text {par }} v$.

### 32.6.2.1 Leaf Sampling

Given a set of rooted trees, our goal is to sample a set of leaves for each tree, and to give an order of those sampled leaves. The algorithm is shown in Algorithm 32.15.

Lemma 32.6.5. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$, and par has a unique root. Let $m>0, \delta \in(0,1)$ be parameters, and let $|V| \leq m^{1 / \delta}$. Let $\left(a_{1}, a_{2}, \cdots, a_{s}\right)=\operatorname{LEAFSAMPLing(par,~} m, \delta$ ) (Algorithm 32.15). Then it has following properties:

1. $a_{1}<_{\text {par }} a_{2}<_{\text {par }} \cdots<_{\text {par }} a_{s}$.
2. If $|V| \leq m$ or $\mid$ leaves(par) $\mid \leq 8\left\lceil m^{1 / 3}\right\rceil$, then $\left\{a_{1}, a_{2}, \cdots, a_{s}\right\}=$ leaves(par). Otherwise, with probability at least $1-1 /\left(100 m^{5 / \delta}\right), \forall v \in \operatorname{leaves}(\mathrm{par}) \backslash\left\{a_{1}\right\}$, there is a vertex
$w \in\left\{a_{1}, a_{2}, \cdots, a_{s}\right\}$ such that $w<_{\text {par }} v$ and the number of leaves between $w$ and $v$ is at most $\mid$ leaves(par) $\mid /\left\lceil m^{1 / 3}\right\rceil$, i.e. $\left|\left\{u \in \operatorname{leaves(par)~} \mid w<_{\text {par }} u<_{\text {par }} v\right\}\right| \leq$ | leaves(par) $\mid /\left\lceil m^{1 / 3}\right\rceil$.
3. If $|V|>m$ and $\mid$ leaves(par) $\mid>8\left\lceil m^{1 / 3}\right\rceil$, then with probability at least $1-1 /\left(100 m^{5 / \delta}\right)$, $s=|S|=\left|\left\{a_{1}, a_{2}, \cdots, a_{s}\right\}\right| \leq 960\left\lceil m^{1 / 3}\right\rceil(1+\log (m) / \delta)$.

Proof. Firstly, let us focus on property 1. According to line 11 to line 13 and Lemma 32.5.7, we know $\forall k \in \mathbb{Z}_{\geq 0} \operatorname{rank}_{\operatorname{par}}\left(\operatorname{par}^{(k)}\left(a_{1}\right)\right)=1$, and $\operatorname{par}^{\prime}\left(a_{1}\right)=a_{1}$ which implies that $a_{1}$ is a leaf. Due to the definition of 32.6.7, we know that $a_{1}$ must be the first leaf appeared in the DFS sequence of par. We can prove the property by induction. Suppose we already have $a_{1}<_{\text {par }} a_{2}<_{\text {par }} \cdots<_{\text {par }} a_{i-1}$. According to line 19 and Lemma 32.6.1, $p_{a_{i-1}, a_{i}}$ is the LCA of $\left(a_{i-1}, a_{i}\right) \cdot p_{a_{i-1} a_{i}, a_{i-1}}$ is a child of $p_{a_{i-1}, a_{i}}$ and is an ancestor of $a_{i-1} \cdot p_{a_{i-1} a_{i}, a_{i}}$ is a child of $p_{a_{i-1}, a_{i}}$ and is an ancestor of $a_{i}$. By Fact 32.6.4, since $\operatorname{rank}\left(p_{a_{i-1} a_{i}, a_{i-1}}\right)<\operatorname{rank}\left(p_{a_{i-1} a_{i}, a_{i}}\right)$, we have $a_{i-1}<_{\text {par }} a_{i}$. To conclude, we have $a_{1}<_{\text {par }} a_{2}<_{\text {par }} \cdots<_{\text {par }} a_{s}$.

For property 2 , if $|V| \leq m$ or $\mid$ leaves(par) $\mid \leq 8\left\lceil m^{1 / 3}\right\rceil$, then by line 6 and line 7 , we know $\left\{a_{1}, a_{2}, \cdots, a_{s}\right\}=$ leaves(par).

Now consider the case when $|V|>m$ and $\mid$ leaves(par) $\mid>8\left\lceil m^{1 / 3}\right\rceil$. Let $t=\left\lceil m^{1 / 3}\right\rceil$. Let leaves(par) $=\left\{u_{1}, u_{2}, \cdots, u_{q}\right\}$, and let $u_{1}<_{\text {par }} u_{2}<_{\text {par }} \cdots<_{\text {par }} u_{q}$. Let us partition $u_{1}, \cdots, u_{q}$ into $4 \cdot t$ groups $G_{1}=\left\{u_{1}, u_{2}, \cdots, u_{\lfloor q /(4 t)\rfloor}\right\}, G_{2}=\left\{u_{\lfloor q /(4 t)\rfloor+1}, u_{\lfloor q /(4 t)\rfloor+2}, \cdots, u_{2 \cdot\lfloor q /(4 t)\rfloor}\right\}, \cdots, G_{4 t}=$ $\left\{u_{(4 t-1)\lfloor q /(4 t)\rfloor+1}, u_{(4 t-1)\lfloor q /(4 t)\rfloor+2}, \cdots, u_{q}\right\}$. Then each group has size at least $q /(8 t)$ and at most
$q /(2 t)$. For a certain $G_{i}$, by Chernoff bound, we have

$$
\begin{aligned}
& \operatorname{Pr}\left(\left|G_{i} \cap S\right| \leq \frac{1}{2} \cdot \frac{q}{8 t} \cdot p\right) \\
\leq & \exp \left(-\frac{1}{8} \cdot \frac{q}{8 t} \cdot p\right) \\
\leq & 1 /\left(100 m^{10 / \delta}\right)
\end{aligned}
$$

where the last inequality follows by $p=\min (1,(10+10 \log (m) / \delta) \cdot 64 t / q)$. Notice that $q \leq|V| \leq m^{1 / \delta}$. We can take union bound over all $G_{i}$. Then with probability at least $1-1 /\left(100 m^{5 / \delta}\right), \forall i \in[4 t], G_{i} \cap S \neq \emptyset$. Thus, $\forall v \in$ leaves(par), there is a vertex $w \in$ $\left\{a_{1}, a_{2}, \cdots, a_{s}\right\}$ such that $w<_{\text {par }} v$ and the number of leaves between $w$ and $v$ is at most $\mid$ leaves(par) $\mid /\left\lceil m^{1 / 3}\right\rceil$, i.e. $\mid\left\{u \in\right.$ leaves(par) $\left.\mid w<_{\text {par }} u<_{\text {par }} v\right\}|\leq|$ leaves(par) $\mid /\left\lceil m^{1 / 3}\right\rceil$.

For property 3 , by applying Chernoff bound, we have

$$
\begin{aligned}
& \operatorname{Pr}\left(\mid S \cap \text { leaves }(\operatorname{par})\left|\geq \frac{3}{2}\right| \text { leaves }(\operatorname{par}) \mid \cdot p\right) \\
\leq & \exp \left(\left.-\frac{1}{12} \cdot \right\rvert\, \text { leaves }(\mathrm{par}) \mid \cdot p\right) \\
\leq & 1 /\left(100 \mathrm{~m}^{10 / \delta}\right)
\end{aligned}
$$

where the last inequality follows by $p=\min (1,(10+10 \log (m) / \delta) \cdot 64 t / \mid$ leaves $($ par $) \mid)$.
Since $\left.\frac{3}{2} \right\rvert\,$ leaves(par) $\mid \cdot p \leq 960\left\lceil m^{1 / 3}\right\rceil(1+\log (m) / \delta)$, we complete the proof.

### 32.6.2.2 DFS Subsequence

Let par : $V \rightarrow V$ be a set of parent pointers on a vertex set $V$, and par has a unique root $v$. Let $\left\{u_{1}, u_{2}, \cdots, u_{q}\right\}=$ leaves(par), and $u_{1}<_{\text {par }} u_{2}<_{\text {par }} \cdots<_{\text {par }} u_{q}$. One observation is that the DFS sequence of par can be generated in the following way:

1. The first part of the DFS sequence is the path from $v$ to $u_{1}$.
2. Then it follows by the path from part $\left(u_{1}\right)$ to the LCA of $\left(u_{1}, u_{2}\right)$, the path from one of the child of the LCA of $\left(u_{1}, u_{2}\right)$ to $u_{2}$, the path from part $\left(u_{2}\right)$ to the LCA of $\left(u_{2}, u_{3}\right)$, the path from one of the child of the LCA of $\left(u_{2}, u_{3}\right)$ to $u_{3}, \cdots$, the path from one of the child of the LCA of $\left(u_{q-1}, u_{q}\right)$ to $u_{q}$.
3. The last part of the DFS sequence is a path from $\operatorname{part}\left(u_{q}\right)$ to $v$.

Fact 32.6.6. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$, and par has a unique root $v$. Let $\left\{u_{1}, u_{2}, \cdots, u_{q}\right\}=$ leaves(par) (See Definition 32.6.5), and $u_{1}<_{\text {par }} u_{2}<_{\text {par }} \cdots<_{\text {par }} u_{q}$. Let $A=\left(a_{1}, a_{2}, \cdots, a_{m}\right)$ be the DFS sequence (See Definition 32.6.7) of par. Then,

1. If $u_{1}$ appears at $a_{i}$, then $\left(a_{1}, a_{2}, \cdots, a_{i}\right)$ is the path from $v$ to $u_{1}$.
2. $\forall i \in[q-1]$, if $u_{i}$ appears at $a_{j}$, and $u_{i+1}$ appears at $a_{k}$, then $\exists j<t<k$ such that $a_{t}$ is the LCA of $\left(u_{i}, u_{i+1}\right)$. In addition, $\left(a_{j}, a_{j+1}, \cdots, a_{t}\right)$ is the path from $a_{j}$ to $a_{t}$, and $\left(a_{t}, a_{t+1}, \cdots, a_{k}\right)$ is the path from $a_{t}$ to $a_{k}$.
3. If $u_{q}$ appears at $a_{i}$, then $\left(a_{i}, a_{i+1}, \cdots, a_{m}\right)$ is the path from $u_{q}$ to $v$.

Proof. Property 1, 3 follows by the definition of DFS sequence (See Definition 32.6.7) and a simple induction.

Now consider property 2 . Since $A$ is a DFS sequence, $\forall l \in\{j, j+1, \cdots, k-1\}$, either $\operatorname{part}\left(a_{l}\right)=a_{l+1}$ or $\operatorname{part}\left(a_{l+1}\right)=a_{l}$. Thus, the path between $u_{i}$ and $u_{i+1}$ is a subsequence of
$\left(a_{j}, a_{j+1}, \cdots, a_{k}\right)$. If $\operatorname{part}\left(a_{l+1}\right)=a_{l}$ but $a_{l+1}$ is not on the path between $u_{i}$ and $u_{i+1}$, then there must be a leaf $x$ in the subtree of $a_{l+1}$ which implies $u_{i}<_{\text {par }} x<_{\text {par }} u_{i+1}$, and thus leads to a contradiction. If $\operatorname{part}\left(a_{l}\right)=a_{l+1}$ but $a_{l+1}$ is not on the path between $u_{i}$ and $u_{i+1}$, then both $u_{i}$ and $u_{i+1}$ should be in the subtree of $a_{l}$, and both of $u_{i}$ and $u_{i+1}$ should be in the DFS sequence of the subtree of $a_{l}$. But we know $a_{l+1}$ cannot be in the DFS sequence of the subtree of $a_{l}$. This leads to a contradiction.

Lemma 32.6.7. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$, with a unique root. Let $v \in V$. Let $V^{\prime}=V \backslash\{u \in V \mid v$ is an ancestor (See Definition 32.6.1) of $\imath$ Let $\operatorname{par}^{\prime}: V^{\prime} \rightarrow V^{\prime}$ satisfy $\forall u \in V^{\prime}, \operatorname{par}^{\prime}(u)=\operatorname{part}(u)$. Then the DFS sequence (See Definition 32.6.7) of $\mathrm{par}^{\prime}$ is a subsequence of the DFS sequence of par.

Proof. The proof follows by the property 3 of Fact 32.6.4 directly.

Corollary 32.6.8. Let par : $V \rightarrow V$ be $a$ set of parent pointers (See Definition 32.4.2) on $a$ vertex set $V$, and par has a unique root. Let $v_{1}, v_{2}, \cdots, v_{t}$ be $t$ vertices in $V$. Let $V^{\prime}=V \backslash\{u \in$ $V \mid \exists v \in\left\{v_{1}, \cdots, v_{t}\right\}, v$ is an ancestor (See Definition 32.6.1) of $\left.u\right\}$. Let $\operatorname{par}^{\prime}: V^{\prime} \rightarrow V^{\prime}$ satisfy $\forall u \in V^{\prime}, \operatorname{par}^{\prime}(u)=\operatorname{part}(u)$. Then the DFS sequence (See Definition 32.6.7) of $\operatorname{par}^{\prime}$ is a subsequence of the DFS sequence of par.

Proof. The proof is by induction on $t$. When $t=1$, then the statement is true by Lemma 32.6.7. Suppose the statement is true for $t-1$. Let $V^{\prime \prime}=V \backslash\left\{u \in V \mid \exists v \in\left\{v_{1}, \cdots, v_{t-1}\right\}, v\right.$ is an ancestor of $\left.u\right\}$, and let $\operatorname{par}^{\prime \prime}: V^{\prime \prime} \rightarrow V^{\prime \prime}$ satisfy $\forall v \in V^{\prime \prime}, \operatorname{par}^{\prime \prime}(v)=\operatorname{part}(v)$. By induction hypothesis, the DFS sequence of par" is a subsequence of the DFS sequence of par. If one of the $v_{1}, \cdots, v_{t-1}$ is an ancestor of $v_{t}$, then $\operatorname{par}^{\prime}=\operatorname{par}^{\prime \prime}$, thus, the DFS sequence of $\mathrm{par}^{\prime}$ is a subsequence of
the DFS sequence of par. Otherwise, we have $V^{\prime}=V^{\prime \prime} \backslash\left\{u \in V^{\prime \prime} \mid v_{t}\right.$ is an ancestor of $\left.u\right\}$. By Lemma 32.6.7, the DFS sequence of par' is a subsequence of the DFS sequence of par" ${ }^{\prime \prime}$. Thus, the DFS sequence of par ${ }^{\prime}$ is a subsequence of the DFS sequence of par.

Lemma 32.6.9 (Removing several subtrees). Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$, and par has a unique root. Let $m>0, \delta \in(0,1)$ be parameters, and let $|V| \leq m^{1 / \delta}$. Let $\left(V^{\prime}, A\right)=\operatorname{SuBDFS}($ par, $m, \delta)$ (Algorithm 32.16). Then $\forall u \in V^{\prime}$, we have $\operatorname{part}(u) \in V^{\prime}$. Furthermore, with probability at least $1-1 /\left(100 m^{5 / \delta}\right)$, $\forall u \in V \backslash V^{\prime}$, the number of leaves (See Definition 32.6.5) in the subtree (See Definition 32.6.6) of $u$ is at most $\left\lfloor\mid\right.$ leaves(par) $\left.\mid /\left\lceil m^{1 / 3}\right\rceil\right\rfloor$.

Proof. By Lemma 32.6.5, we know $L \subseteq$ leaves(par), and $l_{1}<_{\text {par }} l_{2}<_{\text {par }} \cdots<_{\text {par }} l_{t}$.
We first prove $\forall u \in V^{\prime}$, $\operatorname{part}(u) \in V^{\prime}$. Our proof is by induction on the leaf $l_{i}$. By Lemma 32.6.1, we have that $\forall i \in[t-1], p_{l_{i}, l_{i+1}}$ is the LCA of $\left(l_{i}, l_{i+1}\right), p_{i, l_{i+1}}$ is an ancestor of $l_{i+1}$, and $p_{i, l_{i+1}} \neq p_{l_{i}, l_{i+1}}, \operatorname{part}\left(p_{i, l_{i+1}}\right)=p_{l_{i}, l_{i+1}}$. By Lemma 32.14, $P_{1}$ contains all the vertices on the path from $l_{1}$ to the root $v . P_{1}$ is the set of all the ancestors of $l_{1}$. Thus, every ancestor $u$ of $l_{1}$ is in $P_{1}$ and satisfies $\operatorname{part}(u) \in V^{\prime} . P_{2}$ contains all the vertices on the path from $l_{1}$ to an ancestor of $l_{1}$. Thus, $P_{2} \subseteq P_{1}$. Suppose now $P_{1} \cup P_{2} \cup \cdots \cup P_{2 i-2}=\{u \in V \mid$ $\exists j \in[i-1], u$ is an ancestor of $\left.l_{j}\right\}$. Notice that $P_{2 i-1}$ contains all the vertices on the path from $l_{i}$ to the ancestor $p_{i-1, l_{i}}$. Since $\operatorname{part}\left(p_{i-1, l_{i}}\right)=p_{l_{i-1}, l_{i}}$ is also an ancestor of $l_{i-1}$, we have $P_{1} \cup P_{2} \cup \cdots \cup P_{2 i-2} \cup P_{2 i-1}=\left\{u \in V \mid \exists j \in[i], u\right.$ is an ancestor of $\left.l_{j}\right\}$. Since $P_{2 i}$ contains all the vertices on the path from $l_{i}$ to an ancestor of $l_{i}$, we have $P_{2 i} \subseteq P_{1} \cup P_{2} \cup \cdots \cup P_{2 i-2} \cup P_{2 i-1}$. To conclude, we have $V^{\prime}=P_{1} \cup P_{2} \cup \cdots \cup P_{2 t}=\left\{u \in V \mid \exists j \in[t], u\right.$ is an ancestor of $\left.l_{j}\right\}$. Thus, $\forall u \in V^{\prime}$, we have $\operatorname{part}(u) \in V^{\prime}$.

By Lemma 32.6.5, with probability at least $1-1 /\left(100 m^{5 / \delta}\right), \forall u \in$ leaves(par) $\backslash$ $L$, there exists $w \in L, w<_{\text {par }} u$ such that $\mid\left\{x \in\right.$ leaves(par) $\left.\mid w<_{\text {par }} x<_{\text {par }} u\right\} \mid \leq$ $\left\lfloor\mid\right.$ leaves(par) $\left.\mid /\left\lceil m^{1 / 3}\right\rceil\right\rfloor$. In the following, we condition on the above event happens. Let $u \in V \backslash V^{\prime}$. Due to Fact 32.6.4, the DFS sequence of the subtree of $u$ in par must be a consecutive subsequence of the DFS sequence of par. Thus, $\exists x, y \in$ leaves(par), the leaves
 number of leaves in the subtree of $u$ is more than $\left\lfloor\mid\right.$ leaves(par) $\left.\mid /\left\lceil m^{1 / 3}\right\rceil\right\rfloor$, then $\exists l_{i} \in L$, $u$ is an ancestor of leaf $l_{i}$. But $l_{i} \in V^{\prime}$ contradicts to $u \notin V^{\prime}$. Thus, the number of leaves in the subtree of $u$ is at most $\left\lfloor\mid\right.$ leaves(par) $\left.\mid /\left\lceil m^{1 / 3}\right\rceil\right\rfloor$.

Lemma 32.6.10 ( $A$ is a subsequence). Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$, and par has a unique root. Let $m>0, \delta \in(0,1)$ be parameters, and let $|V| \leq m^{1 / \delta} . \operatorname{Let}\left(V^{\prime}, A\right)=\operatorname{SuBDFS}(\mathrm{par}, m, \delta)$ (Algorithm 32.16). Then $A$ is a subsequence of the DFS sequence of par. Furthermore, $\forall u \in V^{\prime}$, u appears in $A$ exactly $\left|\operatorname{child}_{\mathrm{par}}(u)+1\right|$ times, and $\forall u \notin V^{\prime}, u$ does not appear in $A$.

Proof. We first show that $A^{\prime}$ is the DFS sequence of par ${ }^{\prime}$.
Claim 32.6.11. $A^{\prime}$ is the DFS sequence of $\mathrm{par}^{\prime}: V^{\prime} \rightarrow V^{\prime}$.

Proof. By Lemma 32.6.5, we know $\left\{l_{1}, l_{2}, \cdots, l_{t}\right\}=L \subseteq$ leaves(par), and $l_{1}<_{\text {par }} l_{2}<_{\text {par }}$ $\cdots<_{\text {par }} l_{t}$. By Lemma 32.6.1, we have that $\forall i \in[t-1], p_{l_{i}, l_{i+1}}$ is the LCA of $\left(l_{i}, l_{i+1}\right), p_{i, l_{i+1}}$ is an ancestor of $l_{i+1}$, and $p_{i, l_{i+1}} \neq p_{l_{i}, l_{i+1}}, \operatorname{part}\left(p_{i, l_{i+1}}\right)=p_{l_{i}, l_{i+1}}$. By Lemma 32.14, $\forall i \in[t]$, $P_{2 i-1}$ and $P_{2 i}$ only contains some ancestors of $l_{i}$. Thus, leaves $\left(\operatorname{par}^{\prime}\right)=L$.

According to Lemma 32.6.9 and Corollary 32.6.8, the DFS sequence of par ${ }^{\prime}$ is a subsequence of the DFS sequence of par. Thus, we still have $l_{1}<\operatorname{par}^{\prime} l_{2}<\operatorname{par}^{\prime}<\operatorname{par}^{\prime} \cdots<\operatorname{par}^{\prime} l_{t}$.

Due to Lemma 32.14, $P_{1}$ contains all the vertices on the path from $l_{1}$ to the root $v, P_{2 t}$ contains all the vertices on the path from $l_{t}$ to the root $v, \forall i \in[t-1], P_{2 i}$ contains all the vertices on the path from $\operatorname{par}^{\prime}\left(l_{i}\right)$ to the LCA of $\left(l_{i}, l_{i+1}\right)$, and $P_{2 i+1}$ contains all the vertices on the path from $l_{i+1}$ to $p_{i, l_{i+1}}$. Thus, $A_{1}^{\prime}$ is the path from the root $v$ to leaf $l_{1}, A_{2 t}^{\prime}$ is the path from $l_{t}$ to the root $v, \forall i \in[t-1], A_{2 i}^{\prime} A_{2 i+1}^{\prime}$ is the path from $\operatorname{par}^{\prime}\left(l_{i}\right)$ to $l_{i+1}$. Due to Fact 32.6.6, $A^{\prime}=A_{1}^{\prime} A_{2}^{\prime} \cdots A_{2 t}^{\prime}$ is the DFS sequence of $\mathrm{par}^{\prime}$.

Let us define some notations. Let $\widetilde{A}=\left\{\widetilde{a}_{1}, \widetilde{a}_{2}, \cdots, \widetilde{a}_{\widetilde{s}}\right\}$ be the DFS sequence of par. $\forall u \in V$, let $\operatorname{st}_{\widetilde{A}}(u)=j$ such that $\widetilde{a}_{j}$ is the first time that $u$ appears in $\widetilde{A}$. We define $\operatorname{ed}_{\widetilde{A}}(u)$ be the position such that $\widetilde{a}_{\text {ed }_{\tilde{A}}(u)}$ is the last time that $u$ appears in $\widetilde{A}$. Similarly, $\forall u \in V^{\prime}$, we can define $\operatorname{st}_{A^{\prime}}(u), \mathrm{st}_{A}(u), \operatorname{ed}_{A^{\prime}}(u), \mathrm{ed}_{A}(u)$ to be the positions of the first time $u$ appears in $A^{\prime}$, the first time $u$ appears in $A$, the last time $u$ appears in $A^{\prime}$, and the last time $u$ appears in $A$ respectively.

Since $v$ is the root (in both par and par'), it suffices to prove that $\left(a_{\mathrm{st}_{A}(v)}, a_{\text {st }_{A}(v)+1}, \cdots, a_{\operatorname{ed}_{A}(v)}\right)$ is a subsequence of $\left(\widetilde{a}_{\text {st }_{\tilde{A}}(v)}, \widetilde{a}_{\text {st }_{\tilde{A}}(v)+1}, \cdots, \widetilde{a}_{\text {ed }_{\tilde{A}}(v)}\right)$. Our proof is by induction on $\operatorname{dep}_{\text {par }}(u)$ for $u \in V^{\prime}$. If $\operatorname{dep}_{\mathrm{par}}(u)=\operatorname{dep}($ par $)$, then $u$ must be a leaf in par (or par, since par' and par are the same on $V^{\prime}$. In this case, $\operatorname{st}_{A}(u)=\operatorname{ed}_{A}(u), \operatorname{st}_{\widetilde{A}}(u)=\operatorname{ed}_{\widetilde{A}}(u)$, and $\left(a_{\mathrm{st}_{A}(u)}\right)=\left(\widetilde{a}_{\mathrm{st}_{\tilde{A}}(u)}\right)=$ $(u)$. Suppose for all $u \in V^{\prime}$ with $\operatorname{dep}_{\text {par }}(u)>d$, we have that $\left(a_{\text {st }_{A}(u)}, \cdots, a_{\operatorname{ed}_{A}(u)}\right)$ is a subsequence of $\left(\widetilde{a}_{\text {st }_{\tilde{A}}(u)}, \cdots, \widetilde{a}_{\text {ed }_{\tilde{A}}(u)}\right)$. Let $u$ be a vertex in $V^{\prime}$ with $\operatorname{dep}_{\text {par }}(u)=d$. If $u$ is a leaf, then it is the same as the previous argument. Now let us consider the case when $u$ is not a leaf. According to Claim 32.6.11, $A^{\prime}$ is the DFS sequence of par' . Due to line $30, A$ is obtained by duplicating each element of $A^{\prime}$ several times. Let $w_{1}, w_{2}, \cdots, w_{k}$ be the children of $u$ in par', and $\operatorname{rank}_{\operatorname{par}^{\prime}}\left(w_{1}\right)=1, \operatorname{rank}_{\operatorname{par}^{\prime}}\left(w_{2}\right)=2, \cdots, \operatorname{rank}_{\operatorname{par}^{\prime}}\left(w_{k}\right)=\left|\operatorname{child}_{\operatorname{par}^{\prime}}(u)\right|$. Then, according to

Fact 32.6.4, $\left(a_{\text {st }_{A}(u)}, \cdots, a_{\operatorname{ed}_{A}(u)}\right)$ should look like:
$\left(u, \cdots, u, a_{\operatorname{st}_{A}\left(w_{1}\right)}, \cdots, a_{\operatorname{ed}_{A}\left(w_{1}\right)}, u, \cdots, u, a_{\operatorname{st}_{A}\left(w_{2}\right)}, \cdots, a_{\operatorname{ed}_{A}\left(w_{2}\right)}, \cdots, a_{\operatorname{st}_{A}\left(w_{k}\right)}, \cdots, a_{\operatorname{ed}_{A}\left(w_{k}\right)}, u, \cdots, u\right)$
where the number of $u$ before $a_{\mathrm{st}_{A}\left(w_{1}\right)}$ is $\operatorname{rank}_{\mathrm{par}}\left(w_{1}\right)$ (see line 26), the number of $u$ before $a_{\text {st }_{A}\left(w_{i}\right)}$ for $i \in[k] \backslash\{1\}$ is $\operatorname{rank}_{\operatorname{par}}\left(w_{i}\right)-\operatorname{rank}_{\mathrm{par}}\left(w_{i-1}\right)$ (see line 28), and the number of $u$ after $a_{\operatorname{ed}_{A}\left(w_{k}\right)}$ is $\left|\operatorname{child}_{\mathrm{par}}(u)\right|-\operatorname{rank}_{\mathrm{par}}\left(w_{k}\right)+1$ (see line 27). Since $\widetilde{A}$ is the DFS sequence of par, according to Fact 32.6.4, the number of $u$ in $\widetilde{A}$ before $\widetilde{a}_{\text {st }_{\tilde{A}}\left(w_{1}\right)}$ is $\operatorname{rank}_{\mathrm{par}}\left(w_{1}\right)$. By our induction hypothesis, $\left(a_{\text {st }_{A}\left(w_{1}\right)}, \cdots, a_{\operatorname{ed}_{A}\left(w_{1}\right)}\right)$ is a subsequence of $\left(\widetilde{a}_{\mathrm{st}_{\tilde{A}}\left(w_{1}\right)}, \cdots, \widetilde{a}_{\text {ed }_{\tilde{A}}\left(w_{1}\right)}\right)$. Thus, $\left(a_{\mathrm{st}_{A}(u)}, \cdots, a_{\operatorname{ed}_{A}\left(w_{1}\right)}\right)$ is a subsequence of $\left(\widetilde{a}_{\mathrm{st}_{\tilde{\mathcal{A}}}(u)}, \cdots, \widetilde{a}_{\mathrm{ed}_{\tilde{A}\left(w_{1}\right)}}\right)$. According to Fact 32.6.4, $\forall i \in$ $[k] \backslash\{1\}$, the number of $u$ in $\widetilde{A}$ between $\widetilde{a}_{\text {ed }_{\tilde{A}}\left(w_{i-1}\right)}$ and $\widetilde{a}_{\mathrm{st}_{\tilde{A}}\left(w_{i}\right)}$ is $\operatorname{rank}_{\mathrm{par}}\left(w_{i}\right)-\operatorname{rank}_{\mathrm{par}}\left(w_{i-1}\right)$. By our induction hypothesis, for all $i \in[k] \backslash\{1\},\left(a_{\operatorname{st}_{A}\left(w_{i}\right)}, \cdots, a_{\operatorname{ed}_{A}\left(w_{i}\right)}\right)$ is a subsequence of $\left(\widetilde{a}_{\text {st }_{\tilde{A}}\left(w_{i}\right)}, \cdots, \widetilde{a}_{\operatorname{ed}_{\tilde{A}}\left(w_{i}\right)}\right)$. Thus, $\left(a_{\operatorname{st}_{A}(u)}, \cdots, a_{\operatorname{ed}_{A}\left(w_{k}\right)}\right)$ is a subsequence of $\left(\widetilde{a}_{\mathrm{st}_{\tilde{A}}(u)}, \cdots, \widetilde{a}_{\left.\mathrm{ed}_{\tilde{A}\left(w_{k}\right)}\right)}\right)$. According to Fact 32.6.4, the number of $u$ in $\widetilde{A}$ after $\widetilde{a}_{\text {ed }_{\tilde{A}}\left(w_{k}\right)}$ is $\left|\operatorname{child}_{\text {par }}(u)\right|-\operatorname{rank}_{\mathrm{par}}\left(w_{k}\right)+1$. Thus, $\left(a_{\mathrm{st}_{A}(u)}, \cdots, a_{\operatorname{ed}_{A}(u)}\right)$ is a subsequence of $\left(\widetilde{a}_{\mathrm{st}_{\tilde{A}}(u)}, \cdots, \widetilde{a}_{\mathrm{ed}_{\tilde{A}(u)}}\right)$. Furthermore, the number of $u$ appears in $A$ is $\left|\operatorname{child}_{\mathrm{par}}(u)\right|-\operatorname{rank}_{\mathrm{par}}\left(w_{k}\right)+1+\operatorname{rank}_{\mathrm{par}}\left(w_{1}\right)+\sum_{i=2}^{k} \operatorname{rank}_{\mathrm{par}}\left(w_{i}\right)-$ $\operatorname{rank}_{\mathrm{par}}\left(w_{i-1}\right)=\left|\operatorname{child}_{\mathrm{par}}(u)\right|+1$.

Since $A^{\prime}$ is the DFS sequence of $\operatorname{par}^{\prime}, \forall u \notin V^{\prime}, u$ does not appear in $A^{\prime}$. Thus, $\forall u \notin V^{\prime}$, $u$ does not appear in $A$.

### 32.6.2.3 DFS Sequence

In this section, we show how to use Algorithm 32.16 as a subroutine to output a DFS sequence. The high level idea is that we use Algorithm 32.16 to generate subsequences of
the DFS sequence in each iteration, and we ensure that the miss part of the DFS sequence must be the DFS sequences of many subtrees. After the $i^{\text {th }}$ iteration, we should ensure that the number of leaves of each subtree which has missing DFS sequence is at most $n / m^{i}$, where $m$ is some parameter depends on some computational resources (e.g. memory size of a machine). The description of the algorithm is shown in Algorithm 32.17. Figure 32.2 shows one step in our algorithm.

Theorem 32.6.12 (Correctness of DFS sequence). Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$, and par has a unique root. Let $n=$ $|V|, m=n^{\delta}$ for some constant $\delta \in(0,1)$. If $A=\operatorname{DFS}($ par, $m)$ (Algorithm 32.17) does not output FAIL, then $A$ is the DFS sequence of par.

Proof. It suffices to prove the following claim.
Claim 32.6.13. Let $i \in\{0\} \cup[r] . A_{i}$ is a subsequence of the DFS sequence of par. $\forall v \in V_{i}$, $\operatorname{part}(v) \in V_{i}$. Furthermore, $\forall v \in V_{i}$, $v$ appears in $A_{i}$ exactly $\left|\operatorname{child}_{\mathrm{par}}(v)\right|+1$ times, and $\forall v \notin V_{i}, v$ does not appear in $A_{i}$.

Proof. Our proof is by induction on $i$. If $i=0$, then by Lemma 32.6.10, $A_{0}$ is a subsequence of the DFS sequence of par, $\forall v \in V_{0}, v$ appears in $A_{0}$ exactly $\left|\operatorname{child}_{\text {par }}(v)\right|+1$ times, and $\forall v \notin V_{0}, v$ does not appear in $A_{0}$. By Lemma 32.6.9, we have $\forall v \in V_{0}, \operatorname{part}(v) \in V_{0}$.

Suppose the claim is true for $i-1$. Let $u \in V_{i}$.
If $u \in V_{i-1}$, then since $V_{i-1} \subseteq V_{i}$, part $(u) \in V_{i}$. Otherwise $u \in V_{i, v}$ for some $v$ with $\operatorname{par}_{i}(v)=v$. If $u=v$, then $\operatorname{part}(v) \in V_{i-1} \subseteq V_{i}$. Otherwise, by Lemma 32.6.9, $\operatorname{part}(u) \in$ $V_{i, v} \subseteq V_{i}$.

Now consider the property of $A_{i}$. If $u \in V_{i-1}$, then since $A_{i-1}$ is a subsequence of $A_{i}$, and by Lemma $32.6 .10 u$ cannot appear in any $A_{i, v}, u$ must appear in $A_{i}$ exactly $|\operatorname{child} \mathrm{par}(u)|+1$ times. Otherwise $u \in V_{i, v}$ for some $v$ with $\operatorname{par}_{i}(v)=v$. By Lemma 32.6.10, $u$ must appear in $A_{i, v}\left|\operatorname{child}_{\operatorname{par}_{i, v}}(u)\right|+1=\left|\operatorname{child}_{\text {par }}(u)\right|+1$ times. Since $u$ cannot appear in $A_{i-1}, u$ must appear in $A_{i}$ exactly $\left|\operatorname{child}_{\text {par }}(u)\right|+1$ times. For $v \in V_{i}^{\prime}$ with $\operatorname{par}_{i}(v)=v$, according to Fact 32.6.4 and $\forall w \in\{x \in V \mid v$ is an ancestor of $x\}$, $w$ cannot be in $V_{i-1}$, the $\operatorname{rank}_{\mathrm{par}}(v)^{\text {th }}$ time appearance of $v$ and the $\left(\operatorname{rank}_{\mathrm{par}}(v)+1\right)^{\text {th }}$ time appearance of $v$ should be adjacent in $A_{i-1}$. Due to Lemma 32.6.10, $A_{i, v}$ is a subsequence of the DFS sequence of the subtree of $v$ in par. Due to Fact 32.6.4, $A_{i}$ is still a subsequence of the DFS sequence of par after insertion of the sequence $A_{i, v}$.

For any $x \notin V_{i}$, by Lemma $32.6 .10, x$ cannot be in any $A_{i, v}$. By our induction hypothesis, $x$ cannot be in $A_{i-1}$. Thus, $x$ cannot be in $A_{i}$.

If the procedure does not output FAIL, then according to the above Claim 32.6.13, $\forall v \in V_{r}=V, v$ appears in $A_{r}=A$ exactly $\left|\operatorname{child}_{\mathrm{par}}(v)\right|+1$ times, and $A_{r}=A$ is a subsequence of the DFS sequence of par. Due to Fact 32.6.4, $A=A_{r}$ is the DFS sequence of par.

The following lemma claims the success probability of Algorithm 32.17.

Theorem 32.6.14 (Success probability). Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$, and par has a unique root. Let $n=|V|, m=n^{\delta}$ for some constant $\delta \in(0,1)$. With probability at least $1-1 /\left(100 n^{4}\right), A=\operatorname{DFS}($ par, $m)$ (Algorithm 32.17) does not output FAIL.

Proof. $\forall i \in[r], v \in V_{i}^{\prime}$ with $\operatorname{par}_{i}(v)=v$, let $\mathcal{E}_{i, v}$ be the event that $\forall u \in V_{i}^{\prime}(v) \backslash V_{i, v}$, the number of leaves in the subtree of $u$ in par is at most $\mid$ leaves $\left(\operatorname{par}_{i, v}\right) \mid / n^{\delta / 3}$. Notice that due to Lemma 32.6.9, if $\operatorname{par}_{i}(v)=v$, then $v$ will be in $V_{i}$. Thus, we use $\mathcal{E}_{v}$ to denote the event $\mathcal{E}_{i, v}$. By Lemma 32.6.9, $\mathcal{E}_{v}$ happens with probability at least $1-1 /\left(100 n^{5}\right)$. By taking union bound over all $v$, with probability at least $1-1 /\left(100 n^{4}\right)$, all the events $\mathcal{E}_{v}$ will happen.

Claim 32.6.15. Condition on all the events $\mathcal{E}_{v}$ happen. $\forall i \in[r], v \in V_{i}^{\prime}$ with $\operatorname{par}_{i}(v)=v$, we have $\mid$ leaves $\left(\operatorname{par}_{i, v}\right) \mid \leq n / n^{(i-1) \delta / 3}$.

Proof. When $i=1$, the claim is obviously true, since $\mid$ leaves $\left(\operatorname{par}_{i, v}\right) \mid \leq n$. Suppose the claim holds for $i-1$. Let $v \in V_{i}^{\prime}$ with $\operatorname{par}_{i}(v)=v$. There must be $v^{\prime} \in V_{i-1}^{\prime}$ with $\operatorname{par}_{i-1}\left(v^{\prime}\right)=v^{\prime}$, and $v \in V_{i-1}^{\prime}\left(v^{\prime}\right) \backslash V_{i-1, v^{\prime}}$. Since $\mathcal{E}_{v^{\prime}}$ happens, the number of leaves in the subtree of $v$ in par is at most | leaves $\left(\operatorname{par}_{i-1, v^{\prime}}\right) \mid / n^{\delta / 3} \leq n / n^{(i-1) \delta / 3}$.

If $V_{r}^{\prime} \neq \emptyset$, then $\exists v \in V_{r}^{\prime}$ with $\operatorname{par}_{i}(v)=v$ and $\left|\operatorname{leaves}\left(\operatorname{par}_{i, v}\right)\right| \geq 1$. If all the events $\mathcal{E}_{v}$ happens, it will contradict to Claim 32.6.15. Thus, if all the events $\mathcal{E}_{v}$ happens, $V_{r}^{\prime}$ must be $\emptyset$, and thus $V_{r}=V$ which implies that the procedure will not fail.

### 32.6.3 Range Minimum Query

Range Minimum Query (RMQ) problem is defined as following: given a sequence of $n$ numbers $a_{1}, a_{2}, \cdots, a_{n}$, the goal is to preprocess the sequence $a$ to get a data structure such that for any query $(p, q),(p<q)$ we can efficiently find the element which is the minimum in $a_{p}, a_{p+1}, \cdots, a_{q}$. A classic method is to preprocess a sparse table $f$ in $\log (n)$ number of iterations such that $\forall i \in[n], j \in[\lceil\log n\rceil] \cup\{0\}, f_{i, j}=\arg \min _{i \leq i^{\prime} \leq \min \left(n, i+2^{j}-1\right)} a_{i}$. To answer
 1) $\rfloor$. In this section, we firstly show a modified data structure. We will compute $\widehat{f}_{i, j}=$ $\arg \min _{i \leq i^{\prime} \leq \min \left(n, i+\left\lceil n^{\delta}\right\rceil^{j}-1\right)} a_{i^{\prime}}$ The Algorithm is shown in Algorithm 32.18. Then we show how to use $\widehat{f}$ to compute $f$ in Algorithm 32.19.

Lemma 32.6.16. Let $a_{1}, a_{2}, \cdots, a_{n}$ be a sequence of numbers, and $\delta \in(0,1)$. Let $\left\{\widehat{f}_{p, q}\right\}$ be the output of $\operatorname{SpARSETABLE}^{+}\left(a_{1}, a_{2}, \cdots, a_{n}, \delta\right)$ (Algorithm 32.18). Then $\forall p \in[n], q \in$ $\{0\} \cup[\lceil 1 / \delta\rceil], \widehat{f}_{p, q}=\arg \min _{p \leq p^{\prime} \leq \min \left(n, i+\left\lceil n^{\delta}\right\rceil^{q}-1\right)} a_{p^{\prime}}$.

Proof. The proof is by induction on $q$. When $q=0$, the statement obviously holds for all $\widehat{f}_{p, 0}$. Suppose all $p \in[n], \widehat{f}_{p, 0}, \widehat{f}_{p, 1}, \cdots, \widehat{f}_{p, q-1}$ satisfy the property. The first observation is that the value of $\widehat{f}_{p, q}$ will be assigned in the procedure when $l=q, j=\lfloor(p-$ $1) / m\rfloor, i^{\prime}=(p-1) \bmod m$. Then by line $8, z_{j, l}^{*}$ will be the position of the minimum value in $a_{j \cdot m+m}, a_{j \cdot m+m+2}, \cdots, a_{j \cdot m+m^{l}-1}$ by our induction hypothesis. Then by line $11, \widehat{f}_{i+i^{\prime}, l}$ will be the position of the minimum value in $a_{j \cdot m+i^{\prime}+1}, a_{j \cdot m+i^{\prime}+2}, \cdots, a_{j \cdot m+i^{\prime}+m^{l}}$. Thus, Since $j \cdot m+i^{\prime}+1=p, \widehat{f}_{p, q}$ satisfies the property.

Lemma 32.6.17. Let $a_{1}, a_{2}, \cdots, a_{n}$ be a sequence of numbers, and $\delta \in(0,1)$. Let $\left\{f_{p, q}\right\}$ be the output of $\operatorname{SparseTable}\left(a_{1}, a_{2}, \cdots, a_{n}, \delta\right)$ (Algorithm 32.19). Then $\forall p \in[n], q \in$ $\{0\} \cup[\lceil\log n\rceil], f_{p, q}=\arg \min _{p \leq p^{\prime} \leq \min \left(n, i+2^{q}-1\right)} a_{p^{\prime}}$.

Proof. Let $m=\left\lceil n^{\delta}\right\rceil$. By Lemma 32.6.16, $\forall x \in[n], y \in\{0\} \cup[[1 / \delta\rceil], \widehat{f}_{x, y}=\arg \min _{x \leq x^{\prime} \leq \min \left(n, i+m^{y}-1\right)} a_{x^{\prime}}$. Thus, by the definition of $S_{t}$, we know $z_{j, t}^{*}=\arg \min _{j \cdot m+m+1 \leq z \leq j \cdot m+2^{t}} a_{z}$. An observation is that the value of $f_{p, q}$ will be assigned in the procedure when $t=q, j=\lfloor(p-1) / m\rfloor, i^{\prime}=(p-1)$ mod $m$. By line 23, we know

$$
f_{p, q}=f_{i+i^{\prime}, t}=\arg \min _{z: i+i^{\prime}+1 \leq z \leq i+i^{\prime}+2^{t}} a_{z}=\arg \min _{p \leq p^{\prime} \leq \min \left(n, i+2^{q}-1\right)} a_{p^{\prime}} .
$$

### 32.6.4 Applications of DFS Sequence

In this section, we briefly discuss some applications of the DFS sequence of a tree.
Since the DFS sequence of a subtree should be a continuous subsequence of the DFS sequence of the tree, one direct application of the DFS sequence is to compute the size of each subtree, i.e. for each subtree with root $v$, we can find the first place $v$ appeared and the last place $v$ appeared, and then calculate the vertices between those two appearances.

Another application of the DFS sequence and the range minimum query is to output a data structure which can answer any LCA query in $O(1)$ time (for both sequential and parallel). This is better than the data structure provided by Section 32.6 .1 which needs $O(\log D)$ time (for both sequential and parallel) to answer the query.

Since it is easy to output a data structure which can answer the depth of each vertex in $O(1)$ time (in both sequential and parallel), together with the lowest common ancestor data structure, we can answer the query of the tree distance between any two vertices in $O(1)$ time (for both sequential and parallel).

### 32.7 The MPC Model

In this section, let us introduce the computational model studied in this paper. Suppose we have $p$ machines indexed from 1 to $p$ each with memory size $s$ words, where $n$ is the number of words of the input and $p \cdot s=O\left(n^{1+\gamma}\right), s=\Theta\left(n^{\delta}\right)$. Here $\delta \in(0,1)$ is a constant, $\gamma \in \mathbb{R}_{\geq 0}$, and a word has $\Theta(\log (s \cdot p))$ bits. Thus, the total space in the system is only $O\left(n^{\gamma}\right)$ factor more than the input size $n$, and each machine has local memory size sublinear in $n$. When $0 \leq \gamma \leq O(1 / \log n)$, the total space is just linear in the input size. The computation proceeds in rounds. At the beginning of the computation, the input is distributed on the local memory of $\Theta(n / s)$ input machines. Input machines and other machines are identical except that input machine can hold a part of the input in its local memory at the beginning of the computation while each of other machines initially holds nothing. In each round, each machine performs computation on the data in its local memory, and sends messages to other machines (including the sender itself when it wants to keep the data) at the end of the round. Although any two machines can communicate directly in any round, the total size of messages (including the self-sent messages) sent or received of a machine in a round is bounded by $s$, its local memory size. In the next round, each machine only holds the received messages in its local memory. At the end of the computation, the output is distributed on the output machines. Output machines and other machines are identical except that output machine can hold a part of the output in its local memory at the end of the computation while each of other machines holds nothing. We call the above model $(\gamma, \delta)$ - MPC model. The model is exactly the same as the model $\operatorname{MPC}(\epsilon)$ defined by [BKS13] with $\epsilon=\gamma /(1+\gamma-\delta)$ and the number of machines $p=O\left(n^{1+\gamma-\delta}\right)$. Since we care more about the total space used by the algorithm, we use $(\gamma, \delta)$ to characterize the model, while in [BKS13] they use parameter
$\epsilon$ to characterize the repetition of the data. The main complexity measure is the number of rounds $R$ required to solve the problem.

### 32.7.1 Basic MPC Algorithms

Sorting One of the most important algorithms in MPC model is sorting. The following theorem shows that there is an efficient sorting algorithm.

Theorem 32.7.1 ([GSZ11, Goo99]). Sorting can be solved in $c / \delta$ rounds in $(0, \delta)$ - MPC model for any constant $\delta \in(0,1)$, where $c \geq 0$ is a universal constant. Precisely, there is an algorithm $\mathcal{A}$ in $(0, \delta)$ - MPC model such that for any set $S$ of $n$ comparable items stored $O\left(n^{\delta}\right)$ per machine on input machines, $\mathcal{A}$ can run in $c / \delta$ rounds and leave the $n$ items sorted on the output machines, i.e. the ouput machine with smaller index holds a smaller part of $O\left(n^{\delta}\right)$ items.

Notice that for any $\delta^{\prime} \geq \delta, O(1)$ number of machines with $\Theta\left(n^{\delta^{\prime}}\right)$ memory can always simulate the computation of $O\left(n^{\delta^{\prime}-\delta}\right)$ number of machines with $\Theta\left(n^{\delta}\right)$ memory. Thus, if an algorithm $\mathcal{A}$ can solve a problem in $(\gamma, \delta)-$ MPC model in $R(n)$ rounds, then $\mathcal{A}$ can be simulated in $\left(\gamma^{\prime}, \delta^{\prime}\right)$ - MPC model still in $R(n)$ rounds with all $\gamma^{\prime} \geq \gamma, \delta^{\prime} \geq \delta$.

Indexing In the indexing problem, a set $S=\left\{x_{1}, x_{2}, \cdots, x_{n}\right\}$ of $n$ items are stored $O\left(n^{\delta}\right)$ per machine on input machines. The output is

$$
S^{\prime}=\{(x, y) \mid x \in S, y-1 \text { is the number of items before } x\}
$$

of $n$ pairs stored $O\left(n^{\delta}\right)$ per machine on output machines. Here, "an item $x^{\prime} \in S$ is before $x \in S^{\prime \prime}$ means that $x^{\prime}$ is held by a input machine with a smaller index, or $x^{\prime}, x$ are stored in the same input machine but $x^{\prime}$ has a smaller local memory address.

Prefix sum In the prefix sum problem, a set $S=\left\{\left(x_{1}, y_{1}\right),\left(x_{2}, y_{2}\right), \cdots,\left(x_{n}, y_{n}\right)\right\}$ of $n$ (item, number) pairs are stored $O\left(n^{\delta}\right)$ per machine on input machines. The output is

$$
S^{\prime}=\left\{\left(x, y^{\prime}\right) \mid(x, y) \in S, y^{\prime}-y=\sum_{(\widetilde{x}, \tilde{y}) \text { is before }(x, y)} \widetilde{y}\right\}
$$

of $n$ pairs stored $O\left(n^{\delta}\right)$ per machine on output machines. Here, "an pair $(\widetilde{x}, \widetilde{y}) \in S$ is before $(x, y) \in S^{\prime \prime}$ means that $(\widetilde{x}, \widetilde{y})$ is held by a input machine with a smaller index, or $(\widetilde{x}, \widetilde{y}),(x, y)$ are stored in the same input machine but $(\widetilde{x}, \widetilde{y})$ has a smaller local memory address. Notice that indexing problem is a special case of prefix sum problem.

Theorem 32.7.2 ([GSZ11]). Indexing/prefix sum problem can be solved in $c / \delta$ rounds in $(0, \delta)$ - MPC model for any constant $\delta \in(0,1)$, where $c \geq 0$ is a universal constant.

Once each item has an index, it is able to reallocate them onto the machines.

Load balance Sometimes, local computations of a machine may generate new data. When some machines are not able to keep the new data generated, we need to do loading balance. Fortunately, this operation can be done in constant number of rounds of computations.

For arbitrary constant $\delta \in(0,1)$, we are able to spend constant number of rounds to reallocate the data in $(0, \delta)$ - MPC model such that if a machine is not empty, the size of its local data is at least $n^{\delta} / k$ and is at most $2 n^{\delta} / k$ where $k>1$ is an arbitrary constant. The method is very simple, we can use the algorithm mentioned in Theorem 32.7.2 to index each data item, and then send them to the corresponding machine.

Predecessor In the predecessor problem, a set $S=\left\{\left(x_{1}, y_{1}\right),\left(x_{2}, y_{2}\right), \cdots,\left(x_{n}, y_{n}\right)\right\}$ of $n$ (item, $0 / 1$ ) pairs are stored $O\left(n^{\delta}\right)$ per machine on input machines. The output machines are all input machines. If an input (also output) machine holds a tuple $\left(x_{i}, y_{i}\right) \in S$ at the beginning of the computation, then at the end of the computation, that machine should still hold the tuple $\left(x_{i}, y_{i}\right)$. In addition, if an input (also output) machine holds a tuple $(x, 0) \in S$ at the beginning of the computation, then at the end of the computation, that machine should hold a tuple $\left(x, x^{\prime}\right)$ such that $\left(x^{\prime}, 1\right) \in S$, and $\left(x^{\prime}, 1\right)$ is the last tuple occurred before $(x, 0)$. Here, " $\left(x^{\prime}, 1\right)$ is before $(x, 0)$ " means that $\left(x^{\prime}, 1\right)$ is held by a input machine with a smaller index, or $\left(x^{\prime}, 1\right),(x, 0)$ are stored in the same input machine but $\left(x^{\prime}, 1\right)$ has a smaller local memory address.

Theorem 32.7.3 ([GSZ11]). Predecessor problem can be solved in c/ $\delta$ rounds in $(0, \delta)-\mathrm{MPC}$ model for any constant $\delta \in(0,1)$, where $c \geq 0$ is a universal constant.

Roughly speaking the algorithm is as the following: firstly, build a $\Theta\left(n^{\delta}\right)$ branching tree on the machines, then follows by bottom-up stages to collect the last $\left(x_{l}, 1\right)$ tuple in each large interval and then follows by top-down stages to compute the predecessors of every prefix. For completeness, we describe the algorithm for predecessor problem in the following:

## Predecessor Algorithm:

- Setups:
- There are $2 p=\Theta\left(n^{\delta}\right)$ machines indexed from 1 to $2 p$ each with local memory size $s=\Theta\left(n^{\delta}\right)$. The machine with index from $p+1$ to $2 p$ are input/output machines.
$-\left(x_{1}, y_{1}\right), \cdots\left(x_{n}, y_{n}\right)$ are stored on input/output machine $p+1$ to $2 p$, where $\forall i \in$ $[n], y_{i} \in\{0,1\}$.
- The goal: If an input machine holds a tuple $(x, y)$ with $y=0$, then it will create a tuple $\left(x, x^{\prime}\right)$ at the end of the computation, where $\left(x^{\prime}, y^{\prime}\right)$ is the last tuple with $y^{\prime}=1$ stored before $(x, y)$.
- Bottom-up stage ( $O(1 / \delta)$ constant rounds):
- Let $d=s / 10$ be the branching factor.
- In the $i^{\text {th }}$ round, each machine $j$ with $j$ in the range $\left\lfloor p / d^{i-1}\right\rfloor+1$ to $\lfloor(2 p-$ 1) $\left./ d^{i-1}\right\rfloor+1$ sends the last $\left(x_{l}, y_{l}\right)$ tuple with $y_{l}=1$ in its local memory to machine $\lfloor(j-1) / d\rfloor+1$. If machine $j$ does not have any tuple with $y_{l}=1$, it just sends an arbitrary tuple to machine $\lfloor(j-1) / d\rfloor+1$.
- Until the end of the computation, machine $j$ sends itself messages to keep the data. The stage ends when machine 1 receives messages.
- Top-down stage $(O(1 / \delta)$ constant rounds):
- Let $d=s / 10$ be the branching factor.
- In the $i^{\text {th }}$ round, each machine $j$ with $j$ in the range $\left\lfloor d^{i-2}\right\rfloor+1$ to $\min \left(d^{i-1}, p\right)$ sends to each machine $h$ in the range $(j-1) d+1$ to $\min (j \cdot d, 2 p)$ a tuple $\left(x_{l}, y_{l}\right)$ which is the last tuple with $y_{l}=1$ appeared before machine $h$.
- The stage ends when machine $2 p$ receives messages.
- The last round:
- Machine $i \in\{p+1, \cdots, 2 p\}$ scans its local memory, for each tuple $(x, y)$ with $y=0$, create a tuple $\left(x, x^{\prime}\right)$ where $\left(x^{\prime}, y^{\prime}\right)$ is the last tuple stored before $(x, y)$ with $y^{\prime}=1$.


### 32.7.2 Data Organization

In this section, we introduce the method to organize the data in the system.

Set Let $S=\left\{x_{1}, x_{2}, \cdots, x_{m}\right\}$ be a set of $m$ items, and each item $x_{i}$ can be described by $O(1)$ number of words. If $x \in S$ is equivalent to that there is a unique machine which holds a pair (" $S$ ", $x$ ) in its local memory, then we say that $S$ is stored in the system. Here " $S$ " is the name of the set $S$ and can be described by $O(1)$ number of words.

Let $\mathcal{S}=\left\{S_{1}, S_{2}, \cdots, S_{m}\right\}$ be a set of $m$ sets, where $\forall i \in[m], S_{i}$ is stored in the system, and the name " $S_{i}$ " of each set $S_{i}$ can be described by $O(1)$ number of words. If $S \in \mathcal{S}$ is equivalent to that there is a unique machine which holds a pair (" $\mathcal{S}$ ", " $S$ ") in its local memory, then we say $\mathcal{S}$ is stored in the system. Here " $\mathcal{S}$ " is the name of $\mathcal{S}$ and can be described by $O(1)$ number of words.

Let $S$ be a set stored in the system. If machine $i$ has a pair ( $" S$ ",$x)$, then we say that the element $x$ of $S$ is held by the machine $i$. If every element of $S$ is held by a machine with index in $\{i, i+1, \cdots, j\}$, then we say $S$ is stored on the machine $i$ to the machine $j$.

The total space needed to store $S$ is $\Theta(m)$.

Mapping Let $f: U \rightarrow H$ be a mapping from a finite set $U$ to a set $H$. In the following, we show how to use a set to represent a mapping.

Definition 32.7.1 (Set representation of a mapping). Let $f: U \rightarrow H$ be a mapping from a finite set $U$ to a set $H$. Let $S=\{(x, y) \mid x \in U, y=f(x)\}$. then the set $S$ is a set representation of the mapping $f$.

Let $U$ be a finite set where each element of $U$ can be described by $O(1)$ number of words. Let $S$ be a set representation of the mapping $f: U \rightarrow H$. If $S$ is stored in the system, then we say $f$ is stored in the system. If $S$ is stored on the machine $i$ to the machine $j$, then $f$ is stored on the machine $i$ to the machine $j$. At any time of the system, there can be at most one set representation $S$ of $f$ stored in the system. Furthermore, the name of $S$ is " $f$ " which is the same as the name of mapping $f$, and can be described by $O(1)$ number of words.

The total space needed to store $f$ is the total space needed to store $S$, and thus is $\Theta(|U|)$.

Sequence Let $A=\left(a_{1}, a_{2}, \cdots, a_{m}\right)$ be a sequence of $m$ elements. In the following, we show how to use a set to represent a sequence.

Definition 32.7.2 (Set representation of a sequence). Let $A=\left(a_{1}, a_{2}, \cdots, a_{m}\right)$ be a sequence of $n$ elements. If a set $S=\left\{\left(x_{1}, y_{1}\right),\left(x_{2}, y_{2}\right), \cdots,\left(x_{m}, y_{m}\right)\right\} \subseteq \mathbb{R} \times\left\{a_{1}, a_{2}, \cdots, a_{m}\right\}$ satisfies $x_{1}<x_{2}<\cdots<x_{m}, y_{1}=a_{1}, y_{2}=a_{2}, \cdots, y_{m}=a_{m}$, then the set $S$ is a set representation of the sequence $A$. Furthermore, if $x_{1}=1, x_{2}=2, \cdots, x_{m}=m$, then $S$ is a standard set representation of $A$.

Let $A$ be a sequence of elements where each element can be described by $O(1)$ number of words. Let $S$ be a set representation of the sequence $A$. If $S$ is stored in the system, then we say $A$ is stored in the system. If $S$ is stored on the machine $i$ to the machine $j$, then $A$ is stored on the machine $i$ to the machine $j$. At any time of the system, there can be at most one set representation $S$ of $A$ stored in the system. Furthermore, the name of $S$ is
" $A$ " which is the same as the name of sequence $A$, and can be described by $O(1)$ number of words.

The total space needed to store $A$ is the total space needed to store $S$, and thus is $\Theta(m)$.

### 32.7.3 Set Operations

In this section, we introduce some MPC model operations for sets.

Duplicates removing There are $n$ tuples stored in the machines. But there are some duplicates of them. The goal is to remove all the duplicates. To achieve this, we can just sort all the tuples. After sorting, if a tuple is different from its previous tuple, then we keep it. Otherwise, we remove the tuple.

Sizes of sets Suppose we have $k$ sets $S_{1}, S_{2}, \cdots, S_{k}$ stored in the system. Our goal is to get the sizes of all the sets. We can firstly sort all the tuples such that the tuples from the same set are consecutive. Then we can calculate the index of each tuple. Every machine can scan all the tuples in its local memory, if $x$ is an element of set $S_{i}$ and has the smallest/largest index $y$, then create a pair ("boundary of ' $S_{i}$ '", $y$ ). Then we sort all the created pairs, then for each set $S_{i}$, there are two pairs ("boundary of ' $S_{i}$ ' ", $y_{1}$ ), ("boundary of ' $S_{i}$ ' ", $y_{2}$ ) stored on the same machine. Each machine can store its local memory. For each pair of tuples ("boundary of ' $S_{i}$ '", $y_{1}$ ), ("boundary of ' $S_{i}$ '", $y_{2}$ ) with $y_{1}<y_{2}$, the machine can generate a new tuple (" $f$ ", (" $S_{i}$ ", $\left.y_{2}-y_{1}+1\right)$ ). Finally, there will be a mapping $f$ stored in the system, where $f\left(S_{i}\right)=\left|S_{i}\right|$. Thus, the total number of rounds is a constant.

Copies of sets $\quad$ Suppose we have $k$ sets $S_{1}, S_{2}, \cdots, S_{k}$ stored in the system. Let $s_{1}, s_{2}, \cdots, s_{k} \in$ $\mathbb{Z}_{\geq 1}$. If a machine holds an element $x \in S_{i}$, then the machine knows the value of $s_{i}$. Our goal is to create sets $S_{1,1}, S_{1,2}, \cdots, S_{1, s_{1}}, S_{2,1}, S_{2,2}, \cdots, S_{2, s_{2}}, \cdots, S_{k, s_{k}}$ and make them stored in the system, where $S_{i, j}$ is a copy of $S_{i}$.

The idea is very simple: for an element $x \in S_{i}$, we need to make $s_{i}$ copies (" $S_{i, 1}$ ", $x$ ), (" $S_{i, 2}$ ", $x$ ), $\cdots,\left(\right.$ " $S_{i, s_{i}}$ ", $x$ ) of tuple (" $S_{i}$ ", $x$ ). But the issue is that $s_{i}$ may be very large such that it is not able to generate all the copies of a tuple on a single machine. For the above reason, we implement it in three steps: firstly we compute the new "position" of each original tuple among all the copies, then send the original tuples to their new "positions", and finally filling the gap by generating copies between any two adjacent original tuples. Precisely, each machine can scan its local memory, and assign each tuple ( " $S_{i}$ ", $x$ ) a weight $s_{i}$. Then we can use prefix sum algorithm (See Theorem 32.7.2) to compute the prefix sum of each tuple (" $S_{i}$ ", $x$ ). The prefix sum $\operatorname{pos}\left(\right.$ " $S_{i}$ ", $x$ ) of a tuple (" $S_{i}$ ", $x$ ) denotes the new "position" of the last copy of this tuple when all the copies are generated. Let $n=\sum_{i=1}^{k} s_{i} \cdot\left|S_{i}\right|$. Let machine 1 to $t$ be $t$ empty machines each maintains $s / 10$ "positions", i.e. machine 1 has "positions" 1 to $s / 10$, machine 2 has "positions" $s / 10+1$ to $2 s / 10$, and so on. Let $t \cdot s / 10=\Theta(n)$. The machine which holds tuple (" $S_{i}$ ", $x$ ) sends the tuple (" $S_{i}$ " $x$ ) to the "position" $\operatorname{pos}\left(\right.$ " $S_{i}$ ", $\left.x\right)-s_{i}+1$, and sends the tuple (" $S_{i, s_{i}}$ ", x) to the "position" $\operatorname{pos}\left(" S_{i}\right.$ ", $x$ ). Then each machine $i \in[t]$ scans its "positions". If a "position" received a tuple, the machine marks that "position" as " 1 ". Otherwise, the machine marks that position as " 0 ". Now we can apply the predecessor algorithm (See Theorem 32.7.3) such that each empty "position" learns its predecessor tuple. If the predecessor tuple of an empty "position" $l$ is (" $S_{i}$ ", $x$ ), and the predecessor tuple is at "position" $l^{\prime}$, then create a tuple ( " $S_{i, l-l}$ ",$x$ ) at this empty position. Thus, at the end of all the computations, $S_{1,1}, S_{1,2}, \cdots, S_{1, s_{1}}, S_{2,1}, S_{2,2}, \cdots, S_{2, s_{2}}, \cdots, S_{k, s_{k}}$ are stored on the system.

Indexing elements in sets Suppose we have $k$ sets $S_{1}, S_{2}, \cdots, S_{k}$ stored in the system. The goal is to compute a mapping $f$ such that $\forall i \in[k], x \in S_{i}, x$ is the $f\left(S_{i}, x\right)^{\text {th }}$ element of
$S_{i}$.
To achieve this goal, we can sort (See Theorem 32.7.1) all the tuples such that the elements from the same set are stored consecutively on several machines. Then we can run indexing algorithm (See Theorem 32.7.2) to compute the global index of each tuple. In the next, each machine scans its local data. If ( " $S_{i}$ ", $x$ ) is in the local memory, and $x$ is the first element of $S_{i}$, then the machine marks this tuple as " 1 ". For other tuples in the local memory, the machine marks them as " 0 ". Then we can invoke predecessor algorithm (See Theorem 32.7.3) on all the tuples. At the end of the computation, each machine scans its all tuples. For a tuple (" $S_{i}$ ", $x$ ) with global index $l$, the machine determine the index of $x$ in $S_{i}$ based on the global index $l^{\prime}$ of its predecessor (" $S_{i}$ ", $x$ ). Precisely, the machine creates a tuple (" $f$ ", $\left(\left(\right.\right.$ " $S_{i}$ ", $\left.\left.\left.x\right), l-l^{\prime}+1\right)\right)$ stored in the memory. Thus at the end of the computation, the desired mapping $f$ is stored in the system.

Set merging Suppose we need to merge several sets $S_{1}, S_{2}, \cdots, S_{k}$ stored on the system, i.e. create a new set $S=\bigcup_{i=1}^{k} S_{i}$. To implement this operation, each machine scans its local memory. If there is a tuple ( " $S_{i}$ ", x) in its memory, then it creates a tuple (" $S$ ", $x$ ). Finally, we just need to remove all the duplicates.

Set membership Suppose we have $k$ sets $S_{1}, S_{2}, \cdots, S_{k}$ stored in the system. There is an another set $Q=\left\{\left(x_{1}, y_{1}\right), \cdots,\left(x_{q}, y_{q}\right)\right\}$ also stored in the system where $x_{i}$ is the name of a set $S$, and $y_{i}$ is an item. The goal is to answer whether $y_{i}$ is in $S$.

To achieve this, we can firstly sort all the tuples. For tuple with form (" $S_{i}$ ", $x$ ), the first key is $S_{i}$, the second key is $x$, and the third key is $-\infty$ which has the highest priority.

For tuple with form (" $Q$ ", $(x, y)$ ), the first key is $x$, the second key is $y$, and the third key is $\infty$ which has the lowest priority. The comparison in the sorting procedure firstly compare the first key, then the second key, and finally the third key. After sorting, for each tuple with form (" $S_{i}$ ", $x$ ), we mark it as " 1 ". For each tuple with form (" $Q$ ", $(x, y)$ ), we mark it as " 0 ". Now we can apply the predecessor algorithm (See Theorem 32.7.3). For each tuple $(" Q$ ", $(x, y))$, if its predecessor is (" $S$ ", $y$ ) where $x$ is the name of " $S$ ", then we create a tuple (" $f$ ", $((x, y), 1))$; Otherwise, we create a tuple (" $f$ ", $((x, y), 0))$. Thus, at the end of the computation, there is a mapping $f$ stored on the system such that for each $(x, y) \in Q$, if $x$ is the name of some set $S_{i}$, and $y \in S_{i}$, then $f(x, y)=1$; Otherwise $f(x, y)=0$.

### 32.7.4 Mapping Operations

In this section, we introduce some MPC model operations for mapping. The most important operation is called Multiple queries.

Multiple queries We have $k$ sets $S_{1}, S_{2}, \cdots, S_{k}$ stored in the system. Without loss of generality, $S_{1}, S_{2}, \cdots, S_{t}(t \leq k)$ are sets representations of mappings (See Definition 32.7.1) $f_{1}: U_{1} \rightarrow H_{1}, f_{2}: U_{2} \rightarrow H_{2}, \cdots, f_{t}: U_{t} \rightarrow H_{t}$ respectively. When a machine does local computation, it may need to query some values which are in the form $f_{i}(u)$ for some $u \in U_{i}$. The following lemma shows that we can answer all the such queries simultaneously in constant number of rounds in $(0, \delta)$ - MPC model for all constant $\delta \in(0,1)$. It means that we can use constant number of rounds to simulate concurrent read operations on a shared memory where $S_{1}, \cdots, S_{k}$ are stored in the shared memory.

Lemma 32.7.4 (Multiple queries). Let $\delta \in(0,1)$ be an arbitrary constant. There is a constant number of rounds algorithm $\mathcal{A}$ in $(0, \delta)$ - MPC model which satisfies the following properties. The input of $\mathcal{A}$ contains two parts. The first part are $k$ sets $S_{1}, S_{2}, \cdots, S_{k}$ stored (See Section 32.7.2 for data organization of sets) on the input machines, where $S_{1}, S_{2}, \cdots, S_{t}$ $(t \leq k)$ are sets representations of mappings (See Definition 32.7.1) $f_{1}: U_{1} \rightarrow H_{1}, f_{2}: U_{2} \rightarrow$ $H_{2}, \cdots, f_{t}: U_{t} \rightarrow H_{t}$ respectively. The second part is a set

$$
Q=\left\{\left(x_{1}, y_{1}, z_{1}\right),\left(x_{2}, y_{2}, z_{2}\right), \cdots,\left(x_{q}, y_{q}, z_{q}\right)\right\}
$$

stored on the input machines, where $\forall(x, y, z) \in Q, x$ is the name " $f_{i}$ " of the mapping $f_{i}$ for some $i \in[t], y$ is an element in $U_{i}$, and $z$ is the index of the input machine which holds the element $(x, y, z)$ of $Q$. The total input size $n=|Q|+\sum_{i=1}^{k}\left|S_{i}\right|$. The output machines
are all the input machines. $\forall i \in[k], x \in S_{i}$, if the element $x$ of $S_{i}$ is held by the input (also output) machine $j$, then at the end of the computation, the element $x$ of $S_{i}$ should still be held by the output (also input) machine $j$. Let $Q^{\prime}$ be the set $\{(x, y, z, w) \mid \exists(x, y, z) \in Q, w=$ $f_{i}(y)$, where $x$ is the name of $\left.f_{i}\right\}$. At the end of the computation, $Q^{\prime}$ is stored on the output (also input) machines such that $\forall(x, y, z, w) \in Q^{\prime}$, the element $(x, y, z, w)$ of $Q^{\prime}$ is held by the machine $z$.

Proof. The idea is that we can firstly use sorting (See Theorem 32.7.1) to make queries and the corresponding values be stored consecutively in several machines. The issue remaining is that there may be many queries queried the same position such that some queries may not be stored in the machine which holds the corresponding value. In this case, we need to find the predecessor by invoking the algorithm shown in Theorem 32.7.3.

The Multiple queries algorithm is shown as the following:

Multiple Queries Algorithm:

- Setups:
- There are $3 p=\Theta\left(n^{\delta}\right)$ machines indexed from 1 to $3 p$ each with local memory size $s=\Theta\left(n^{\delta}\right)$.
- The machine with index from $2 p+1$ to $3 p$ are input/output machines.
- Sets $S_{1}, S_{2}, \cdots, S_{k}, Q$ are stored on machine $2 p+1$ to $3 p . \quad$ Corresponding to Lemma 32.7.4
- The first round:
- Machine $i \in\{2 p+1, \cdots, 3 p\}$ scans its local memory, and send all the tuples with form (" $f_{j}$ ", $(x, y)$ ) or (" $Q$ ", $(x, y, z)$ ) to machine $i-p$, where " $f_{j}$ " is the name of $f_{j}$ (also $S_{j}$ ) for $j \in[t]$. Until the end of the computation, machine $i$ sends itself messages to keep its local data.
- Using constant number $(O(1 / \delta))$ of rounds to sort:
- Use machine 1 to $2 p$ to sort all the tuples stored on machine $p+1$ to $2 p$, and thus at the end of this stage, machine $p+1$ to $2 p$ holds sorted tuples. For tuple with the form (" $f_{j}$ ", $(x, y)$ ), the first key value is " $f_{j}$ ", the second key value is $x$ and the third key value is $-\infty$ which is the highest priority. For tuple with form (" $Q$ ", $(x, y, z)$ ), the first key value is $x$, the second key value is $y$, and the third key value is $\infty$ which is the lowest priority. The comparison in the sorting is: Firstly compare the first key. If they are the same, then compare the second key. If they are still the same, compare the third key.
- Using constant number $(O(1 / \delta))$ of rounds to find predecessors:
- Machine $p+1$ to $2 p$ scans its local memory. For a tuple in the form (" $f_{j}$ ", $(x, y)$ ), the machine marked it as " 1 ". For a tuple in the form (" $Q$ ", $(x, y, z)$ ), the machine marked it as " 0 ".
- Machine 1 to $2 p$ together invoke the Predecessor algorithm (Theorem 32.7.3), where the input is on machine $p+1$ to machine $2 p$.
- The last round:
- Machine $p+1$ to $2 p$ scans its local memory. For each tuple with form (" $Q$ ", $(x, y, z)$ ), it sends machine $z$ a tuple (" $Q^{\prime \prime},(x, y, z, w)$ ), where $x$ is the name of $f_{j}$, and $w=f_{j}(y)$.


### 32.7.5 Sequence Operations

In this section, we introduce some MPC model operations for sequence.

Sequence standardizing Suppose there is a sequence $A$, and one of its set representation (see Definition 32.7 .2 ) $S$ is stored in the system. The goal is to modify the set $S$ such that $S$ is a standard set representation of $A$.

We can compute the index (see Indexing elements in sets in Section 32.7.3) of elements in $S$. Then for each element $(x, y) \in S$, we can query (see Multiple queries in Section 32.7.4) the index of $(x, y)$ in $S$. Suppose the index is $i$, we modify the tuple ("S", $(x, y))$ to (" $\left.S^{\prime \prime},(i, y)\right)$.

Sequence duplicating Suppose there is a sequence $A=\left(a_{1}, a_{2}, \cdots, a_{s}\right)$, and one of its set representation (see Definition 32.7.2) $S$ is stored in the system. Furthermore, there is a mapping $f:[s] \rightarrow \mathbb{Z}_{\geq 0}$ which is also stored in the system. The goal is to get a set $S^{\prime}$ stored in the system such that $S^{\prime}$ is a set representation of the sequence:

$$
(\underbrace{a_{1}, a_{1}, \cdots, a_{1}}_{f(1) \text { times }}, \underbrace{a_{2}, a_{2}, \cdots, a_{2}}_{f(2) \text { times }}, \cdots, \underbrace{a_{s}, a_{s}, \cdots, a_{s}}_{f(s) \text { times }}) .
$$

Firstly, we can standardize (see the above paragraph Sequence standardizing) the set $S$. Then for each tuple (" $S$ ", $\left(i, a_{i}\right)$ ), we create a tuple ( " $S_{i}$ ", $a_{i}$ ), and we can query (see Multiple queries in Section 32.7.4) the value of $f(i)$. Then we can copy (see Copies of sets in Section 32.7.3) set $S_{i} f(i)$ times. For each tuple ( $S_{i, j}$ ", $a_{i}$ ), we create a tuple (" $S^{\prime \prime}$ ", ( $\left.(i, j), a_{i}\right)$ ). Then we can compute the index (see Indexing elements in sets in Section 32.7.3) of each element in $S^{\prime}$. For each tuple (" $S^{\prime \prime}$ ", $\left((i, j), a_{i}\right)$ ), we can query
(see Multiple queries in Section 32.7.4) the index $i^{\prime}$ of it, and then modify the tuple as (" $\left.S^{\prime \prime \prime},\left(i^{\prime}, a_{i}\right)\right)$.

Sequence insertion Suppose there are $k+1$ sequences $A=\left(a_{1}, a_{2}, \cdots, a_{s}\right), A_{1}, \cdots, A_{k}$ which have sets representations (see Definition 32.7.2) $S, S_{1}, \cdots, S_{k}$ respectively and stored on the system. There is also a mapping $f:[k] \rightarrow\{0\} \cup[s]$ stored on the system where $\forall i \neq j \in[k], f(i) \neq f(j)$. The goal is to insert each sequence $A_{i}$ into the sequence $A$, and $A_{i}$ should be between the element $a_{f(i)}$ and $a_{f(i)+1}$.

Firstly, we can standardize (see Sequence standardizing in Section 32.7.3) S. Then we can compute the total size (see Sizes of sets in Section 32.7.3) $N=|S|+\left|S_{1}\right|+\cdots+\left|S_{k}\right|+1$. For each tuple (" $S$ ", $\left(i, a_{i}\right)$ ), we can modify it as (" $S$ ", $\left(i \cdot N, a_{i}\right)$ ). For each tuple (" $S_{i}$ ", $\left(j, a_{i j}\right)$ ), we query (see Multiple queries in Section 32.7.4) the value of $f(i)$, then create a tuple (" $S$ ", $\left.\left(f(i) \cdot N+j, a_{i j}\right)\right)$.

### 32.7.6 Multiple Tasks

In this section, we show that if the entire computational tasks consist of some independent small computational tasks, then we are able to schedule the machines such that the small computational tasks can be computed simultaneously.

Task and multiple tasks problem A computational task here is running a specific algorithm on specific input data.

There are $k$ sets $S_{1}, S_{2}, \cdots, S_{k}$ stored in the system. Let $n=\sum_{i=1}^{k}\left|S_{i}\right|$ be the total input size. There are $h$ independent computational tasks $T_{1}, T_{2}, \cdots, T_{h}$. Each task $T_{i}$ needs to take some sets $\mathcal{S}_{i} \subseteq\left\{S_{1}, S_{2}, \cdots, S_{k}\right\}$ as its input, and is running a $\left(\gamma_{i}, \delta_{i}\right)$ - MPC algorithm in $r_{i}$ rounds where $\gamma_{i} \in \mathbb{R}_{\geq 0}$, constant $\delta_{i} \in(0,1) . \forall i \in[h]$, let $n_{i}=\sum_{S \in s_{i}}|S|$ be the input size of task $T_{i}$. Without loss of generality, we can assume that the input of different tasks are disjoint. Otherwise we can use sets copying technique (See Section 32.7.3) to generate different copies of input sets for the tasks shared the same input set. The goal here is to use the small number of rounds to finish all the tasks. Since we can always use sorting and indexing to extract the desired input data. The most naive way is to compute the tasks one-by-one. This can be trivially done in $r=O\left(\sum_{i=1}^{h} r_{i}\right)$ rounds in $(\gamma, \delta)-$ MPC model for $\gamma=\log _{n}(h)+\max _{i \in[h]} \gamma_{i}, \delta=\max _{i \in[h]} \delta_{i}$. Here we show how to compute all the tasks simultaneously in $r=O\left(\max _{i \in[h]} r_{i}\right)$ rounds in $(\gamma, \delta)$ - MPC model for $\gamma=\log _{n}(m)-1, \delta=\max _{i \in[h]} \delta_{i}$, where $m=\Theta\left(n+\sum_{i=1}^{h} n_{i}^{1+\gamma_{i}}\right)$.

Each machine scans its local memory. If the machine holds a tuple (" $S_{i}$ ", $x$ ), and $S_{i}$ is a part of input of task $T_{j}$, then it creates a tuple (" $W_{j}$ ", ("S $\left.S_{i}, x\right)$ ). Thus, at the end of this step, there are additional $h$ sets $W_{1}, W_{2}, \cdots, W_{h}$ stored in the system. Here $W_{i}, i \in[h]$
contains all the information of input data of task $T_{i}$. Then we can compute a mapping $f$ such that $\forall i \in[h], f\left(W_{i}\right)=\left|W_{i}\right|$ (see Section 32.7.3). Thus, we know the input size of each task. Then each machine scans its local memory. If the machine holds a tuple ("f", (" $\left.\left.W_{i} ",\left|W_{i}\right|\right)\right)$, then it creates a tuple (" $H_{i}$ ", $\left.\left|W_{i}\right|\right)$, i.e. a set $H_{i}=\left\{\left|W_{i}\right|\right\}$. Then for each set $H_{i}=\left\{\left|W_{i}\right|\right\}, i \in[h]$, we can copy (see Section 32.7.3) it $s_{i}=c \cdot\left|W_{i}\right|^{1+\gamma_{i}}$ times for a sufficiently large $c$ to get sets $H_{i, 1}=H_{i, 2}=\cdots=H_{i, s_{i}}=\left|W_{i}\right|$. Each set $H_{i, j}$ is just a placeholder of one unit working space of the task $T_{i}$. Thus, the number of copies of the set $H_{i}$ is the total space needed for the task $T_{i}$. We can sort all the tuples (" $H_{i, j}{ }^{\prime},\left|W_{i}\right|$ ) on machines with index in $I=\{2,5,8,11, \cdots, 3 p-1\}$, where local memory $s=\Theta\left(n^{\delta}\right)$, total required memory $m=\Theta\left(n+\sum_{i=1}^{h} n_{i}^{1+\gamma_{i}}\right)$, and $p=\Theta(m / s)$ For each machine with index $q \in I$, the tuples on that machine must be in the following form
(" $H_{i, j}$ ", $\left.\left|W_{i}\right|\right),\left(" H_{i, j+1} ",\left|W_{i}\right|\right), \cdots,\left(" H_{i, s_{i}}\right.$ ", $\left.\left|W_{i}\right|\right),\left(" H_{i+1,1} ",\left|W_{i+1}\right|\right), \cdots,\left(" H_{i+1, s_{i+1}}\right.$ ", $\left.\left|W_{i+1}\right|\right)$, $\left(" H_{i+2,1}\right.$ ", $\left.\left|W_{i+2}\right|\right), \cdots,\left(" H_{i+2, s_{i+2}}\right.$ ", $\left.\left|W_{i+2}\right|\right), \cdots,\left(" H_{i^{\prime}, 1^{\prime}},\left|W_{i^{\prime}}\right|\right), \cdots\left(\right.$ " $H_{i^{\prime}, j^{\prime}}$ ", $\left.\left|W_{i^{\prime}}\right|\right)$.

Then machine $q$ just sends all the tuples (" $H_{i, j}$ ", $\left.\left|W_{i}\right|\right),\left(" H_{i, j+1}\right.$ ", $\left.\left|W_{i}\right|\right), \cdots,\left(" H_{i, s_{i}}\right.$ ", $\left.\left|W_{i}\right|\right)$ to machine $q-1$, and sends all the tuples (" $H_{i^{\prime}, 1}$ ", $\left.\left|W_{i^{\prime}}\right|\right)$, (" $H_{i^{\prime}, 2}$ ", $\left.\left|W_{i^{\prime}}\right|\right), \cdots\left(\right.$ " $H_{i^{\prime}, j^{\prime}}$ ", $\left.\left|W_{i^{\prime}}\right|\right)$ to machine $q+1$. Thus, $\forall i \in[h]$,

1. either all the $H_{i, 1}, H_{i, 2}, \cdots, H_{i, s_{i}}$ are stored on consecutive machines, machine $q$ to machine $q^{\prime}$, and any of machine $q$ to machine $q^{\prime}$ does not hold other tuples,
2. or there is a unique machine $q$ which holds all the sets $H_{i, 1}, H_{i, 2}, \cdots, H_{i, s_{i}}$.

For each machine $q \in[3 p]$, if $H_{i, 1}$ is held by machine $q$, then it creates a tuple ("st", (" $T_{i}$ ", q)). If $H_{i, s_{i}}$ is held by machine $q$, then it creates a tuple ("ed", (" $T_{i}$ ", q)). The mapping st, ed
then are stored in the system, where $\operatorname{st}\left(T_{i}\right)$ is the index of the first machine assigned to task $T_{i}$, and $\operatorname{ed}\left(T_{i}\right)$ is the index of the last machine assigned to task $T_{i}$. Recall that $W_{i}$ contains all the information of the input data to task $T_{i}$. The remaining task is to move the input data of task $T_{i}$ to the machines with index from $\operatorname{st}\left(T_{i}\right)$ to $\operatorname{ed}\left(T_{i}\right)$. According to Section 32.7.3, we can compute a mapping $f^{\prime}$, such that $f^{\prime}\left(W_{i}, x\right)$ records the index of $x \in W_{i}$ in set $W_{i}$. Now, each machine scans its local memory. For each tuple (" $W_{j}$ ", (" $S_{i}$ ", $x$ ) ), the machine needs to query the value of $f^{\prime}\left(W_{j},\left(" S_{i}\right.\right.$ ", $\left.x\right)$ ), the value of $\operatorname{st}\left(T_{j}\right)$ and the value of $\operatorname{ed}\left(T_{j}\right)$. By Lemma 32.7.4, these queries can be handled simultaneously in constant number of rounds. Then the machine can send the tuple (" $S_{i}$ ", $x$ ) to the corresponding machine based on the value of $f^{\prime}\left(W_{j},\left(" S_{i}\right.\right.$ ", x) ), st $\left(T_{j}\right)$, and $\operatorname{ed}\left(T_{j}\right)$. Finally, $\forall i \in[h]$, since $\delta \geq \delta_{i}$ and $\left(\operatorname{ed}\left(T_{i}\right)-\operatorname{st}\left(T_{i}\right)+1\right) \cdot s=\Theta\left(n_{i}^{1+\gamma_{i}}\right)$, the machines with index from $\operatorname{st}\left(T_{i}\right)$ to ed $\left(T_{i}\right)$ can simulate task $T_{i}$ in $r_{i}$ number of rounds.

### 32.8 Implementations in MPC Model

In this section, we show how to implement all the previous batch algorithms in MPC model.

### 32.8.1 Neighbor Increment Operation

Lemma 32.8.1. Let graph $G=(V, E), n=|V|, N=|V|+|E|$ and $m=\Theta\left(N^{\gamma}\right)$ for some arbitrary $\gamma \in[0,2]$. NeighborIncrement $(m, G)$ (Algorithm 32.1) can be implemented in $(\gamma, \delta)$ - MPC model for any constant $\delta \in(0,1)$. Furthermore, the parallel running time is $O(r)$, where $r$ is the number of iterations (see Definition 32.4.1) of NeighborIncrement $(m, G)$.
 Then for each (" $S_{v}^{(0) ", ~ u) ~ w e ~ c a n ~ c o m p u t e ~ t h e ~ i n d e x ~(s e e ~ I n d e x i n g ~ e l e m e n t s ~ i n ~ s e t s ~ a n d ~}$ Multiple queries) of $u$ in set $S_{v}^{(0)}$. If the index of $u$ in set $S_{v}^{(0)}$ is at least $\left\lceil(m / n)^{1 / 2}\right\rceil$, then


Now let us discuss how to implement line 14 and line 17 in the $i^{\text {th }}$ iteration. Firstly, we can compute the size of every set stored in the system (see Sizes of sets). Then for each
 $S_{u}^{(i-1)}$. If $\left|S_{u}^{(i-1)}\right| \geq\left\lceil(m / n)^{1 / 2}\right\rceil$, then create a tuple ("temp $v_{v}^{i "}, u$ ). We can index (see Indexing elements in sets) all the elements in set temp ${ }_{v}^{i}$, and only keep the element with index 1 . Thus, temp ${ }_{v}^{i}$ has a only element $u$, and we need to create a set $S_{v}^{(i)}=S_{u}^{(i-1)}$. Notice that there may be many $v \in V$ which needs need to implement $S_{v}^{(i)}=S_{u}^{(i-1)}$. Thus, for each tuple ("temp $v_{v}^{i "}, u$ ), we create a tuple (" $\left.\operatorname{target}_{u}^{i} ", v\right) . v \in \operatorname{target}{ }_{u}^{i}$ means that $S_{v}^{(i)}$ needs a copy of $S_{u}^{(i-1)}$. Thus, $\left|\operatorname{target}_{u}^{i}\right|$ means that $S_{u}^{(i-1)}$ needs to copy $\left|\operatorname{target}_{u}^{i}\right|$ times. For each tuple
 set $S_{u}^{(i-1)}$ can be copied (see Copies of sets) $\left|\operatorname{target}_{u}^{i}\right|$ times. For each tuple ("target ${ }_{u}^{i}$ ", $v$ ), we query (see Multiple queries) the index (see Indexing elements in sets) of $v$ in set $\operatorname{target}_{u}^{i}$, and then create a tuple (" $f^{i}$ ", (("target $\left.\left.{ }_{u}^{i "}, x\right), v\right)$ ), where $x$ is the index of $v$ in $\operatorname{target}{ }_{u}^{i}$. Thus $f^{i}$ is a mapping such that $f^{i}\left(\operatorname{target}_{u}^{i}, x\right)$ is the $x^{\text {th }}$ element in $\operatorname{target}_{u}^{i}$. For

 Duplicates removing) of elements of for every set $S_{v}^{(i)}$. For each tuple ("temp ${ }_{v}^{i}$ ", u), query (see Multiple queries) the size (see Sizes of sets) of $S_{v}^{(i)}$ and $S_{u}^{(i-1)}$. If $\left|S_{v}^{(i)}\right|>\left|S_{u}^{(i-1)}\right|$, then we create a tuple (" $g^{i "},(v,(u$, "delete" $\left.))\right)$; Otherwise, create a tuple (" $g^{i "},(v,(u$, "keep" $))$ ). Finally, for each tuple $\left(S_{v}^{(i)}, x\right)$, we query (see Multiple queries) $(u, o)=g^{i}(v)$, if $u=x$ and $o=$ "delete", the machine deletes the tuple $\left(S_{v}^{(i)}, x\right)$.

Next, let us discuss how to implement line 20. Similar as before, we can compute the
 the corresponding machine queries (see Multiple queries) the size of $S_{u}^{(i-1)}$. If $\left|S_{u}^{(i-1)}\right| \geq$ $\left\lceil(m / n)^{1 / 2}\right\rceil$, then create a tuple ("temp ${ }_{v}^{i "}, u$ ). For each tuple (" $V$ ", $v$ ), we can create a tuple ("temp ${ }_{v}^{i "}$, null). Then for each tuple (" $V$ ", $v$ ) we can query (see Multiple queries) the size (see Sizes of sets) of temp ${ }_{v}^{i}$. If $\left|\operatorname{temp}_{v}^{i}\right|=1$, then we create a tuple (" $f^{\prime i}$ ", 1 ); Otherwise, we create a tuple (" $f^{\prime \prime} ", 0$ ). Thus, mapping $f^{\prime i}$ is stored in the system, and $f^{\prime \prime}(v)=1$ if and only if $\forall u \in S_{v}^{(i-1)},\left|S_{u}^{(i-1)}\right|<\left\lceil(m / n)^{1 / 2}\right\rceil$. For each tuple (" $S_{v}^{(i-1) \text { ", }} u$ ), we query (see Multiple queries) the value $f^{\prime i}(v)$. If $f^{\prime i}(v)=1$, we create a tuple ("target ${ }_{u}^{i}$ ", $v$ ). Thus, $v \in \operatorname{target}^{i}{ }_{u}$ means that $S_{u}^{(i-1)}$ should be a part of $S_{v}^{(i)}$. $\left|\operatorname{target}{ }_{u}^{i}\right|$ means that $S_{u}^{(i-1)}$ needs to copy $\left|\operatorname{target}{ }_{u}^{i}\right|$

sets) of $\operatorname{target}{ }_{u}^{i}$. Then we can copy (see Copies of sets) each set $S_{u}^{(i-1)}\left|\operatorname{target}_{u}^{i}\right|$ times. Then for each tuple (" $\operatorname{target}_{u}^{i}$ ", $v$ ), we can query (see Multiple queries) the index $x$ (see Indexing elements in sets) of $v$ in set $\operatorname{target}_{u}^{i}$, and then create a tuple (" $f^{i}$ ", (( " $\left.\left.\operatorname{target}_{u}^{i} ", x\right), v\right)$ ) which means that the $x^{\text {th }}$ element of $\operatorname{target}_{u}^{i}$ is $f^{i}\left(\operatorname{target}_{u}^{i}, x\right)=v$. For each tuple ( " $S_{u, j}^{(i-1) \text { ", }, x) \text {, we }}$ query (see Multiple queries) the value $v=f^{i}\left(\operatorname{target}_{u}^{i}, j\right.$ ), and then create a tuple (" $S_{v}^{(i)}$ ", $x$ ). We then remove the duplicates (see Duplicates removing) of elements of for every set $S_{v}^{(i)}$.

Finally, let us consider how to implement line 24 . It is very simple, we only need to query the sizes of sets. For each tuple (" $V$ ", v), query (see Multiple queries) the size (see Sizes of sets) of $S_{v}^{(i)}$ and $S_{v}^{(i-1)}$, if $v$ satisfies the condition, create a tuple ("Done", v). Every machine queries (see Multiple queries) the size (see Sizes of sets) of Done. If it is $|V|$, then all the machines know that they finish the loop. In the end, for each tuple $\left(S_{v}^{(r)}, u\right)$ we create tuples (" $\left.E^{\prime \prime},(u, v)\right),\left(" E^{\prime "},(v, u)\right)$, and for each tuple (" $E$ ", $(u, v)$ ) we create tuple (" $E^{\prime \prime},(u, v)$ ). Then we then remove the duplicates (see Duplicates removing) of elements of $E$.

In the $i^{\text {th }}$ iteration, we only need to maintain sets $V, E, S_{v}^{(i-1)}$. Since all the copy operation will create at most $n \cdot(m / n)^{1 / 2} \cdot(m / n)^{1 / 2}=m$ tuples, the total space needed is $\Theta(m)$ plus the space needed to maintain $V, E, S_{v}^{(i)}$. By Property 4 of Lemma 32.4.1, $\left|S_{v}^{(i)}\right| \leq m / n$. Thus, the total space is $\Theta(m)+|V|+|E|+\sum_{v \in V}\left|S_{v}^{(i)}\right|=\Theta(m)+N=\Theta(m)$.

The above implementation shows that the parallel time is $O(r)$, where $r$ is the number of iterations (see Definition 32.4.1).

### 32.8.2 Tree Contraction Operation

In this section, we show how to implement Algorithm 32.2 in MPC model.

Lemma 32.8.2. Let graph $G=(V, E)$ and par : $V \rightarrow V$ be a set of parent points (see Definition 32.4.2) on the vertex set $V$. TreeContraction(G, par) (Algorithm 32.2) can be implemented in $(0, \delta)$ - MPC model for any constant $\delta \in(0,1)$. Furthermore, the parallel running time is $O(r)$, where $r$ is the number of iterations (see Definition 32.4.5) of TreeContraction( $G$, par).

Proof. Let $N=|V|+|E|$. Then the total space is $\Theta(N)$.
Initially, each machine scans its local memory. If there is a tuple (" $V$ ", $v$ ), then it queries the value of $\operatorname{part}(v)$. It needs $O(1)$ parallel time to answer all the queries (see Multiple
 the initialization stage, mapping $g^{(0)}$, par, set $V, E$ are stored in the system.

In the $l^{\text {th }}$ iteration, Each machine scans its local memory. If there is a tuple (" $V$ ", $v$ ), then it queries the value of $g^{(l-1)}(v)$. This can be done by Multiple queries. Then it queries the value of $\operatorname{part}\left(g^{(l-1)}(v)\right)$. This can also be done by Multiple queries. If $\operatorname{part}\left(g^{(l-1)}(v)\right)=$ $g^{(l-1)}(v)$, it creates a tuple ("Done", $v$ ). Then the machines can compute the sizes (see Section 32.7.3) of $V$ and Done. Each machine queries the size of $V$ and Done. This can be done by Multiple queries. Then if $|V|=\mid$ Done $\mid$, every machine knows that the iterations are finished. Otherwise, the machine which holds (" $V$ ", $v$ ) queries the value of $g^{(l-1)}\left(g^{(l-1)}\right)(v)$. This can be done by Multiple queries. And then it creates a tuple (" $\left.g^{(l) "},\left(v, g^{(l-1)}\left(g^{(l-1)}\right)(v)\right)\right)$.

At the end, if a machine holds a tuple (" $V$ ", $v$ ), then the queries part $(v)$. If $v=$
$\operatorname{part}(v)$, it creates a tuple (" $\left.V^{\prime \prime}, v\right)$. If a machine holds a tuple (" $E$ ", $(u, v)$ ), then it queries $g^{(r)}(u), g^{(r)}(v)$, and creates a tuple (" $\left.E^{\prime "},\left(g^{(r)}(u), g^{(r)}(v)\right)\right)$.

Since at the end of each iteration $l$, the system only stores mappings par : $V \rightarrow V, g^{(r)}$ : $V \rightarrow V$, and sets $V, E$, the total space used is at most $O(N)$. Thus, we can implement the algorithm in $(0, \delta)$ - MPC model.

The total parallel time is $O(r)$. By Corollary 32.4.8, $r=O(\operatorname{dep}($ par $))$. Thus, the total parallel time is $O(\operatorname{dep}(\mathrm{par}))$.

### 32.8.3 Graph Connectivity

Theorem 32.8.3. Let graph $G=(V, E), n=|V|, N=|V|+|E|$ and $m=\Theta\left(N^{\gamma}\right)$ for some arbitrary $\gamma \in[0,2]$. Let $r>0$ be a round parameter. Connectivity ( $G, m, r$ ) (Algorithm 32.3) can be implemented in $(\gamma, \delta)$ - MPC model for any constant $\delta \in(0,1)$. Furthermore, the parallel running time is $O(R)$, where $R$ is the total number of iterations (see Definition 32.4.6) of $\operatorname{Connectivity~}(G, m, r)$.

Proof. Initially, we store sets $V_{0}, E_{0}, V, E$ and mapping $h_{0}$ in the system. Now consider the $i^{\text {th }}$ round. Due to Lemma 32.8.1, line 9 can be implemented in total space $\Theta(m)$ and with $O\left(k_{i}\right)$ parallel time, where $k_{i}$ is the number of iterations (See Definition 32.4.1) of NeighborIncrement $\left(m, G_{i-1}\right)$. To store $V_{i}^{\prime}$ and $E_{i}^{\prime}$, we need total space $\Theta(m)$. Line 10 can be implemented by operations described in Sizes of sets and Multiple queries (see Section 32.7). Line 11 can be implemented by the operations described in Set membership and Multiple queries. To implement line 14 , for each tuple (" $V_{i}^{\prime \prime "}, v$ ), we can create a tuple ("l", $(v, x)$ ) where $x=1$ with probability $p_{i}, x=0$ with probability $1-p_{i}$. To calculate $p_{i}$, the machine only needs to know $n_{i-1}$. This can be done by the operations described in Sizes of sets and Multiple queries. Line 15 and line 16 can be implemented by operations described in Set membership and Multiple queries. For line 17, set $L_{i} \cap\left(\Gamma_{G_{i}^{\prime}}(v) \cup\{v\}\right)$ can be computed by operations described in Set membership and Multiple queries. Then, by operations in Indexing elements in sets and Multiple queries, we can get $\min _{u \in L_{i} \cap\left(\Gamma_{G_{i}^{\prime}}(v) \cup\{v\}\right)} u$. Finally, by operation described in Multiple queries, $\forall v \in V_{i}^{\prime \prime}$ with $v \notin L_{i}$, the tuple (" $\operatorname{par}_{i} ",(v, x)$ ) can be created, where $x=\min _{u \in L_{i} \cap\left(\Gamma_{G_{i}^{\prime}}(v) \cup\{v\}\right)} u$. Due to Lemma 32.8.2, line 18 can be implemented in total $\Theta(m)$ space and $O\left(r_{i}^{\prime}\right)$ parallel running
time, where $r_{i}^{\prime}$ is the number of iterations (see Definition 32.4.5) of TreeContraction $\left(G_{i}^{\prime \prime}, \operatorname{par}_{i}\right)$. Line 21 can be implemented by operations in Set membership, Indexing elements in sets and Multiple queries. Line 22 can be implemented by operations in Set membership and Multiple queries. Line 23 can be implemented by Multiple queries. For other $v \in V$ with $h_{i}(v)=$ null assigned by line 8 , we can use the operations in Set membership and Multiple queries to find those $v$, and create a tuple (" $h_{i}$ ", $v$, null).

Thus, in the $i^{\text {th }}$ round, the parallel time needed is $O\left(k_{i}+r_{i}^{\prime}\right)$. At the end of the $i^{\text {th }}$ round, we only need to keep sets $V_{i}, E_{i}, V, E$ and mapping $h_{i}$ in the system. It will take total space at most $O(m)$.

Due to Lemma 32.8.2, line 26 can be implemented in at most $O(m)$ total space and $O(\log r)$ parallel time.

Thus, the total parallel time is $O\left(\log r+\sum_{i=1}^{r}\left(k_{i}+r_{i}^{\prime}\right)\right)=O\left(\sum_{i=1}^{r}\left(k_{i}+r_{i}^{\prime}\right)\right)$. By definition 32.4.6, the total parallel time is $O(R)$, where $R$ is the total number of iterations of $\operatorname{Connectivity}(G, m, r)$. The total space in the computation is always at most $\Theta(m)$.

Here, we are able to conclude the following theorem for graph connectivity problem.

Theorem 32.8.4. For any $\gamma \in[0,2]$ and any constant $\delta \in(0,1)$, there is a randomized $(\gamma, \delta)$ - MPC algorithm (see Algorithm 32.3) which can output the connected components for any graph $G=(V, E)$ in $O\left(\min \left(\log D \cdot \log \left(1 / \gamma^{\prime}\right), \log n\right)\right)$ parallel time, where $D$ is the diameter of $G, n=|V|, N=|V|+|E|$ and $\gamma^{\prime}=(1+\gamma) \log _{n} \frac{2 N}{n^{1 /(1+\gamma)}}$. The success probability is at least 0.98. In addition, if the algorithm fails, then it will return FAIL.

Proof. The implementation of Algorithm 32.3 in MPC model is shown by Theorem 32.8.3.

The correctness of Algorithm 32.3 is proved by Theorem 32.4.9. The total parallel time of Algorithm 32.3 is proved by Theorem 32.4.15.

### 32.8.4 Algorithms for Local Shortest Path Trees

In this section, we mainly explained how to implement local shortest path tree algorithms described in Section 32.5.1 and Section 32.5.2.

Lemma 32.8.5. Let $G=(V, E)$ be an undirected graph, $s_{1}, s_{2} \in \mathbb{Z}_{\geq 0}$, and $v \in V$. Let $\widetilde{T}=\left(V_{\widetilde{T}}, \operatorname{par}_{\widetilde{T}}\right)$ with root $v$ and radius $s_{1}$ be a local complete shortest path tree (see Definition 32.5.3) in $G$, and $\operatorname{dep}_{\widetilde{T}}: V_{\widetilde{T}} \rightarrow \mathbb{Z}_{\geq 0}$ be the depth of every vertex in $\widetilde{T}$. $\forall u \in V_{\widetilde{T}}$, let $T(u)$ with root $u$ and radius $s_{2}$ be a local complete shortest path tree in $G$, and $\operatorname{dep}_{T(u)}$ : $V_{T(u)} \rightarrow \mathbb{Z}_{\geq 0}$ be the depth of every vertex in $T(u)$. Then TreeExpansion $\left(\widetilde{T}, \operatorname{dep}_{\widetilde{T}},\{T(u) \mid\right.$ $\left.u \in V_{\widetilde{T}}\right\}$, $\left\{\operatorname{dep}_{T(u)} \mid u \in V_{\widetilde{T}}\right\}$ ) (Algorithm 32.4) can be implemented in ( $0, \delta$ ) - MPC model for any constant $\delta \in(0,1)$ in $O(1)$ parallel time.

Proof. For line 3, we apply operation shown in Copies of sets to copy each $V_{T(u)}$, then we can merge (see Set merging) all the copies to get $V_{\widehat{T}}$. To implement line 4 and line 5, we only need to apply the operation shown in Multiple queries. To implement line 6, for each tuple (" $V_{T(u)}$ ", $x$ ), we can firstly check whether $x \in V_{\widehat{T}} \backslash V_{\widetilde{T}}$ by operations described in Set membership and Multiple queries. If $x \in V_{\widehat{T}} \backslash V_{\widetilde{T}}$, then we can query the values of $\operatorname{dep}_{\widetilde{T}}(u)$ and $\operatorname{dep}_{T(u)}(x)$ by operations shown in Multiple queries. Then we create a tuple ("temp ${ }_{x}$ ", $\left(\operatorname{dep}_{\widetilde{T}}(u)+\operatorname{dep}_{T(u)}(x), u\right)$ ). By Indexing elements in sets and Multiple queries, we can find the element with the smallest index in set temp ${ }_{x}$, and thus that element is $\left(\operatorname{dep}_{\widetilde{T}}\left(u_{x}\right)+\operatorname{dep}_{T\left(u_{x}\right)}(x), u_{x}\right)$. Finally, the remaining things in line 6 and line 7 can be done by the operations described by Multiple queries.

For all the operations, the total space is always linear. The parallel time needed for the above operations is also a constant.

Lemma 32.8.6. Let graph $G=(V, E), n=|V|, N=|V|+|E|$ and $m=\Theta\left(N^{\gamma}\right)$ for some arbitrary $\gamma \in[0,2]$. MultiRadiusLCSPT $(G, m)$ (Algorithm 32.5) can be implemented in $(\gamma, \delta)$ - MPC model for any constant $\delta \in(0,1)$. Furthermore, the parallel running time is $O(r)$, where $r$ is the number of iterations (see Definition 32.5.4) of MultiRadiusLCSPT $(G, m)$.

Proof. To implement line 4 to line 6 , we can scan all the tuples (" $E$ ", $(u, v)$ ), then query the size of $\{v\} \cup \Gamma_{G}(v)$ and the size of $\{u\} \cup \Gamma_{G}(u)$, where these operations are described in Sizes of sets and Multiple queries. Then based on the sizes, we decide whether we need to create the corresponding tuples for $V_{T_{0}(v)}, V_{T_{0}(u)}, \operatorname{par}_{T_{0}(u)}, \operatorname{part}\left(T_{0}(v)\right)$.

Now consider the main loop. We focus on the $i^{\text {th }}$ round. Line 12 can be implemented by the operation described in Multiple queries. To implement line 13, for each tuple (" $V_{T_{i-1}(v)}$ ", $u$ ), we can query (see Multiple queries) whether $T_{i-1}(u)$ is null. If $T_{i-1}(u)$ is null, then we create a tuple (" $\operatorname{temp}_{v}$ ", $u$ ). Then for each tuple (" $V$ ", $v$ ), we can query the size of temp ${ }_{v}$ by operations described in Sizes of sets and Multiple queries. If the size is not 0 , then $T_{i}(v)$ must be null. Line 15 can be implemented by coping input for different tasks and running tasks in parallel, where it only needs operations shown in Copies of sets, Multiple queries and Multiple Tasks (see Section 32.7.6). According to Lemma 32.8.5, it only needs $O(1)$ parallel time. Line 16 and line 19 only need the operation shown in Multiple queries.

Thus, the total parallel time is $O(r)$ where $r$ is the number of iterations (see Definition 32.5.4) of MultiRadiusLCSPT $(G, m)$. For the total space, we stored the sets $V_{T_{i}(v)}$ for all $i \in[r], v \in V$ and mappings $\operatorname{par}_{T_{i}(v)}, \operatorname{dep}_{T_{i}(v)}$ for all $i \in[r], v \in V$. By Lemma 32.5.2, the total space to store all of them is at most $O\left(r \cdot n \cdot(m / n)^{1 / 4}\right)=O(m)$. In the $i^{\text {th }}$ round of
the main loop, line 15 may make copies of the set. By Lemma 32.5.2, the input size of each task will be at most $O\left((m / n)^{1 / 4} \cdot(m / n)^{1 / 4}\right)$. Since the there are at most $n$ tasks, the total space needed is at most $O(m)$.

Lemma 32.8.7. Let graph $G=(V, E), n=|V|, N=|V|+|E|$ and $m=\Theta\left(N^{\gamma}\right)$ for some arbitrary $\gamma \in[0,2]$. MultipleLargeTrees $(G, m)$ (Algorithm 32.6) can be implemented in $(\gamma, \delta)$ - MPC model for any constant $\delta \in(0,1)$. Furthermore, the parallel running time is $O(r)$, where $r$ is the number of iterations (see Definition 32.5.5) of MultipleLargeTrees( $G, m$ ).

Proof. By Lemma 32.8.6, line 3 can be implemented in total space $m$ and $O(r)$ parallel time where $r$ is the number of iterations (see Definition 32.5.4) of MultiRadiusLCSPT( $G, m$ ). Line 4 to line 6 can be implemented by the operation described by Multiple queries. The implementation of line 7 to line 17 is similar as the implementation of the main loop of Algorithm 32.5 (See Lemma 32.8.6 for details of the implementation). The implementation of line 18 and line 19 only needs the operation described in Indexing elements in sets and Multiple queries. Line 22 can be implemented by copying input sets for different tasks and running multiple tasks in parallel, where the operations needed are described in Copies of sets, Multiple queries and Multiple Tasks (see Section 32.7.6). Line 24 to line 28 can be implemented by the operations described in Copies of sets, Set membership, Indexing elements in sets, and Multiple queries.

The total parallel time of the first loop is $O(r)$ since it has $r$ rounds. The second loop can be done in one round. Thus the parallel time of the second loop is $O(1)$. Then the total parallel time is $O(r)$. Due to Lemma 32.8.6 and Lemma 32.5.6, $r$ is the number of iterations (see Definition 32.5.5) of MultipleLargeTrees $(G, m)$.

We stored all the $V_{T_{i}(v)}, V_{\widetilde{T}_{i}(v)}, \operatorname{par}_{T_{i}(v)}, \operatorname{par}_{\widetilde{T}_{i}(v)}, \operatorname{dep}_{T_{i}(v)}, \operatorname{dep}_{\widetilde{T}_{i}(v)}$ in the system. By Lemma 32.5.3 and Lemma 32.5.2, the total space needed to store them is at most $O(m)$. Furthermore, at any round, the size of all the input copies for multiple tasks is at most $n \cdot(m / n)^{1 / 4} \cdot(m / n)^{1 / 4}=O(m)$. Thus, the total space needed is $O(m)$.

### 32.8.5 Path Generation and Root Changing

Lemma 32.8.8. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$. Let $n=|V|$. FindAncestors(par) (Algorithm 32.7) can be implemented in $(\gamma, \delta)-$ MPC model for any $\gamma \geq \frac{\log n}{\log \log n}$ and any constant $\delta \in(0,1)$. The parallel running time is $O(r)$, where $r$ is the number of iterations (see Definition 32.5.6) of FindAncestors(par).

Proof. The structure of the whole algorithm is the same as the Algorithm 32.2 (see Lemma 32.8.2). All the steps can be done by operation described in Multiple queries.

Since the number of rounds needed is $r$, the parallel time is $O(r)$. For the total space, we need to store all the mappings $g_{1}, \cdots, g_{r}$. At the end of the $i^{\text {th }}$ round, we need to store mapping $h_{i}$. According to Lemma 32.5.7, $r=O(\log n)$ Thus, the total space is $O(r n)=O(n \log n)$.

Lemma 32.8.9. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$. Let $q$ be a vertex in $V$, and $n=|V|$. FindPath(par, q) (Algorithm 32.8) can be implemented in $(\gamma, \delta)-$ MPC model for any $\gamma \geq \frac{\log n}{\log \log n}$ and any constant $\delta \in(0,1)$. The parallel running time is $O(r)$, where $r$ is the number of iterations (see Definition 32.5.6) of FindAncestors(par) (Algorithm 32.7).

Proof. By Lemma 32.8.8, FindAncestors(par) can be implemented in $(\gamma, \delta)$ - MPC model for $\gamma \geq \frac{\log n}{\log \log n}$ and any constant $\delta \in(0,1)$. All the other other steps in the algorithm can be done by operation described in Multiple queries. Notice that, after each round, we need to do load balancing which can be done by operation described in Load balance.

The number of rounds must be smaller than $O(r)$, where $r$ should be the number of iterations of FindAncestors(par) according to Lemma 32.8.8.

We store all the mappings $g_{i}, \operatorname{dep}_{\mathrm{par}}$ in the system. They need $O(n \log n)$ total space. In the $i^{\text {th }}$ round, we only need to additionally store set $S_{i}$ which has size at most $O(n)$. Thus, the total space needed is at most $O(n \log n)$.

Lemma 32.8.10. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$. Let $q$ be a vertex in $V$. RootChange(par, $q$ ) (Algorithm 32.9) can be implemented in $(\gamma, \delta)-$ MPC model for any $\gamma \geq \frac{\log n}{\log \log n}$ and any constant $\delta \in(0,1)$. The parallel running time is $O(r)$, where $r$ is the number of iterations (see Definition 32.5.6) of FindAncestors(par) (Algorithm 32.7).

Proof. By Lemma 32.8.9, FindPath(par, $q$ ) can be implemented in $(\gamma, \delta)$ - MPC model. The remaining steps in the procedure can be implemented by the operation described by Multiple queries, and has $O(1)$ parallel running time.

The total space needed is the total space needed for $\operatorname{FindPath}($ par, $q$ ) plus the space needed to store mapping $h, \widehat{\text { par. Thus the total space needed is } O(n \log n)+O(n)=}$ $O(n \log n)$.

The parallel running time is linear in the parallel running time of $\operatorname{FindPath}(p a r, q)$. Then, by Lemma 32.8.9, the parallel running time is $O(r)$ where $r$ is the number of iterations (see Definition 32.5.6) of FindAncestors(par).

### 32.8.6 Spanning Forest Algorithm

Lemma 32.8.11. Let $G_{2}=\left(V_{2}, E_{2}\right)$ be an undirected graph. Let $\widetilde{\mathrm{par}}: V_{2} \rightarrow V_{2}$ be a set of parent pointers (See Definition 32.4.2) which satisfies that $\forall v \in V_{2}$ with $\widetilde{\operatorname{par}}(v) \neq v$, ( $v, \widetilde{\operatorname{par}}(v)$ ) must be in $E_{2}$. Let $G_{1}=\left(V_{1}, E_{1}\right)$ be an undirected graph satisfies $V_{1}=\left\{v \in V_{2} \mid\right.$ $\widetilde{\operatorname{par}}(v)=v\}, E_{1}=\left\{(u, v) \in V_{1} \times V_{1} \mid u \neq v, \exists(x, y) \in E_{2}, \widetilde{\operatorname{par}}^{(\infty)}(x)=u, \widetilde{\operatorname{par}}^{(\infty)}(y)=v\right\}$. Let par: $V_{1} \rightarrow V_{1}$ be a rooted spanning forest (See Definition 32.5.7) of $G_{1}$. Let $f: V_{1} \times$ $V_{1} \rightarrow\{$ null $\} \cup\left(V_{2} \times V_{2}\right)$ satisfy the following property: for $u \neq v \in V_{1}$, if $\operatorname{part}(u)=v$, then $f(u, v) \in\left\{(x, y) \in E_{2} \mid \widetilde{\operatorname{par}}^{(\infty)}(x)=u, \widetilde{\operatorname{par}}^{(\infty)}(y)=v\right\}$, and $f(v, u) \in\{(x, y) \in$ $\left.E_{2} \mid \widetilde{\operatorname{par}}^{(\infty)}(x)=v, \widetilde{\operatorname{par}}^{(\infty)}(y)=u\right\}$. Let $n=\left|V_{2}\right|$. Then ForestExpansion(par, $\left.\widetilde{\operatorname{par}}, f\right)$ (Algorithm 32.10) can be implemented in ( $\gamma, \delta$ ) - MPC model for any $\gamma \geq \log n / \log \log n$ and any constant $\delta \in(0,1)$ in parallel running time $O(R)$, where $R=\log (\operatorname{dep}(\widetilde{\operatorname{par}}))$.

Proof. Due to Lemma 32.8.2, line 3 can be done in $O(R)$ parallel time for $R=\log (\operatorname{dep}(\widetilde{\mathrm{par}}))$. Line 9 corresponds to multiple tasks, we can implement them parallelly by operations described in Multiple queries, and Multiple Tasks (see Section 32.7.6). By Lemma 32.8.10, the total space needed is at most $O(n \log n)$ and the parallel running time is at most $O(R)$ where $R=\log (\operatorname{dep}(\widetilde{\text { par }}))$.

Theorem 32.8.12. Let graph $G=(V, E), n=|V|, N=|V|+|E|$ and $m=\Theta\left(N^{\gamma}\right)$ for some arbitrary $\gamma \in[0,2]$. Let $r>0$ be a round parameter. $\operatorname{Spanning} \operatorname{Forest}(G, m, r)$ (Algorithm 32.11) can be implemented in $(\gamma, \delta)-$ MPC model for any constant $\delta \in(0,1)$. Furthermore, the parallel running time is $O(R)$, where $R$ is the total number of iterations (see Definition 32.5.8) of SpanningForest( $G, m, r$ ).

Proof. At the beginning of the algorithm, we just store sets $V, E, V_{0}, E_{0}$ and mapping $g_{0}$ in the system.

Consider the $i^{\text {th }}$ round of the loop. By Lemma 32.8.7, line 8 can be implemented in total space $\Theta(m)$ and in parallel running time $O\left(k_{i}\right)$ where $k_{i}$ is the number of iterations (see Definition 32.5.5) of MultipleLargeTrees $\left(G_{i}, m\right)$. Line 9 can be implemented by operations described in Sizes of sets, Set membership, and Multiple queries. Line 10 can be implemented by operations described in Indexing elements in sets, Set membership, and Multiple queries. In line 12 , to calculate $\gamma_{i}$, we need to query $n_{i}$, this can be done by operations described in Sizes of sets and Multiple queries. In line 14, to compute $L_{i}$, we only need operations described in Set membership and Multiple queries. Line 15 can be implemented by operations shown in Set membership, Indexing elements in sets and Multiple queries. By Lemma 32.8.9, for line 16, there are multiple tasks each can be implemented in $O\left(\left|V_{\widetilde{T}_{i}(v)}\right| \log \left|V_{\widetilde{T}_{i}(v)}\right|\right)$ total space, and $O\left(k_{i}\right)$ parallel time. We can schedule these multiple tasks (see Section 32.7.6) such that we can finish them in parallel in $O\left(k_{i}\right)$ parallel time. According to Lemma 32.8.2, for line 17, we can implement it in $O\left(n_{i}\right)=O(n)$ total space, and in $O\left(k_{i}^{\prime}\right)$ parallel time, where $k_{i}^{\prime}$ is the number of iterations (see Definition 32.4.5) of TreeContraction $\left(G_{i}^{\prime}, \operatorname{par}_{i}\right)$. Line 19 can be done by the operation described in Multiple queries. Line 20 can be done by the operation described in Indexing elements in sets and Multiple queries.

Thus, the parallel time is $O(R)$, where $R=\sum_{i=0}^{r-1}\left(k_{i}+k_{i}^{\prime}\right)$. By definition of the total number of iterations (see Definition 32.5.8) of $\operatorname{SpanningForest}(G, m, r) . R$ is the total number of iterations of SpanningForest $(G, m, r)$.

For the space, we store all the sets $V, E, V_{i}, D_{i}$ and mappings $\operatorname{par}_{i}, h_{i}$ in all the rounds.

Notice that $\sum_{i=0}^{r}\left|V_{i}\right| \leq 40|V|$. Thus this part takes only $O(N)$ space. In the $i^{\text {th }}$ round, we additionally store all the sets $V_{\widetilde{T}_{i}(v)}, V_{i}^{\prime}, E_{i}^{\prime}, L_{i}$ and all the mappings $\operatorname{par}_{\widetilde{T}_{i}(v)}, \operatorname{dep}_{\widetilde{T}_{i}(v)}, l_{i}, z_{i}$. The total space for this part is at most $O(m)$. For line 16 , it creates multiple tasks. The input of each task is at most $\left|V_{\widetilde{T}_{i}(v)}\right| \leq\left(m / n_{i}\right)^{1 / 2}$. There are at most $n_{i}$ tasks, and by Lemma 32.8.9, each task will need space at most $O\left(\left|V_{\widetilde{T}_{i}(v)}\right| \log \left|V_{\widetilde{T}_{i}(v)}\right|\right)$. Thus, the space for this part is at most $O(m)$. To conclude, the total space needed is at most $O(m)$.

Theorem 32.8.13. Let graph $G=(V, E), n=|V|, N=|V|+|E|$ and $m=\Theta\left(N^{\gamma}\right)$ for some arbitrary $\gamma \in[0,2]$. Let $r>0$ be a round parameter. If $\operatorname{SpanningForest}(G, m, r)$ (Algorithm 32.11) does not return FAIL, then let the output be the input of Orientate (•) (Algorithm 32.12), and $\operatorname{Orientate}(\cdot)$ can be implemented in $(\gamma, \delta)$ - MPC model for any constant $\delta \in(0,1)$. Furthermore, the parallel running time is $O(R)$, where $R$ is the total number of iterations (see Definition 32.5.8) of SpanningForest( $G, m, r$ ).

Proof. Line 4 to line 7 can be implemented by operations described in Multiple queries. Notice that there is a trick here, if $f_{i}(u, v)=$ null, we do not need to store the tuple (" $f_{i}$ ", $((u, v)$, null $)$ ) in the system. The total space needed to store all the mappings $f_{i}$ and all the sets $F_{i}$ for $i \in\{0\} \cup[r]$ is at most $\sum_{i=0}^{r}\left|V_{i}\right|=O(m)$.

Line 10 and line 11 can be implemented by operations described in Set membership and Multiple queries.

We now look at the second loop, and focus on round $i$. Line 12 can be implemented by Lemma 32.8.11. The total space needed is at most $O\left(\left|V_{i}\right| \cdot\left(m /\left|V_{i}\right|\right)^{1 / 2} \cdot \log \left(m /\left|V_{i}\right|\right)\right)=O(m)$. The parallel running time needed is at most $O\left(k_{i}\right)$, where $k_{i}$ is the number of iterations (see

Definition 32.5.5) of MultipleLargeTrees $\left(G_{i}, m\right), G_{i}$ is the intermediate graph in the procedure $\operatorname{SpanningForest}(G, m, r)$.

Thus, the parallel running time is $O(R)$, where $R$ is the total number of iterations (see Definition 32.5.8) of SpanningForest $(G, m, r)$. The total space needed is $O(m)$.

Now, we are able to conclude the following theorem for spanning forest problem.
Theorem 32.8.14. For any $\gamma \in[0,2]$ and any constant $\delta \in(0,1)$, there is a randomized $(\gamma, \delta)$ - MPC algorithm (see Algorithm 32.11 and Algorithm 32.12) which can output the rooted spanning forest for any graph $G=(V, E)$ in $O\left(\min \left(\log D \cdot \log \frac{1}{\gamma^{\prime}}, \log n\right)\right)$ parallel time, where $D$ is the diameter of $G, n=|V|, N=|V|+|E|$ and $\gamma^{\prime}=(1+\gamma) \log _{n} \frac{2 N}{n^{1 /(1+\gamma)}}$. The success probability is at least 0.98. In addition, if the algorithm fails, then it will return FAIL.

Proof. Algorithm 32.11 outputs all the edges in the spanning forest and all the contraction information. Algorithm 32.12 takes the output of Algorithm 32.11 as its input, and outputs a rooted spanning forest.

The implementation of Algorithm 32.11 and Algorithm 32.12 in MPC model is shown by Theorem 32.8.12 and Theorem 32.8.13 respectively. The correctness of Algorithm 32.11 and Algorithm 32.12 is proved by Corollary 32.5.17 and Theorem 32.5.19 respectively. The parallel time of Algorithm 32.11 and Algorithm 32.12 is proved by Theorem 32.5.21.

A byproduct of our spanning forest algorithm is an estimator of the diameter of the graph.

Theorem 32.8.15. For any $\gamma \in[0,2]$ and any constant $\delta \in(0,1)$, there is a randomized $(\gamma, \delta)-$ MPC algorithm which can output an diameter estimator $D^{\prime}$ for any graph $G=(V, E)$ in $O\left(\min \left(\log D \cdot \log \left(1 / \gamma^{\prime}\right), \log n\right)\right)$ parallel time such that $D \leq D^{\prime} \leq D^{O\left(\log \left(1 / \gamma^{\prime}\right)\right)}$, where $D$ is the diameter of $G, n=|V|, N=|V|+|E|$ and $\gamma^{\prime}=(1+\gamma) \log _{n} \frac{2 N}{n^{1 /(1+\gamma)}}$. The success probability is at least 0.98. In addition, if the algorithm fails, then it will return FAIL.

Proof. By Theorem 32.8.14, we can find a rooted spanning forest. By Theorem 32.5.19, the depth of that rooted spanning forest is at most $D^{O\left(\log \left(1 / \gamma^{\prime}\right)\right)}$. Then we can implement a doubling algorithm (e.g. Modified Lemma 32.8.8, Algorithm 32.7 without maintaining useless $g_{l}$ ) with log in depth parallel time to output the depth of that spanning forest.

### 32.8.7 Lowest Common Ancestor and Multi-Paths Generation

Lemma 32.8.16. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$. Let $Q=\left\{\left(u_{1}, v_{1}\right),\left(u_{2}, v_{2}\right), \cdots,\left(u_{q}, v_{q}\right)\right\}$ be a set of $q$ pairs of vertices, and $\forall i \in[q], u_{i} \neq v_{i}$. Let $n=|V|, N=n+q$. LCA(par, $Q$ ) (Algorithm 32.13) can be implemented in $(\gamma, \delta)-$ MPC model for any $\gamma \geq \log \log N / \log N$ and any constant $\delta \in(0,1)$ in $O(\log (\operatorname{dep}(\mathrm{par})))$ parallel running time.

Proof. By Lemma 32.8.8, line 3 can be implemented in space $O(N \log N)$ and $O(\log (\operatorname{dep}($ par $)))$ parallel running time. It is easy to see that all the other steps in the procedure can be done by the operations shown in Multiple queries.

Thus, the total space needed is $O(N \log N)$ and the parallel running time is $O(\log (\operatorname{dep}(\operatorname{par})))$.

Lemma 32.8.17. Let par: $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$. Let $Q=\left\{\left(u_{1}, v_{1}\right),\left(u_{2}, v_{2}\right), \cdots,\left(u_{q}, v_{q}\right)\right\} \subseteq V \times V$ satisfy $\forall j \in[q]$, $v_{j}$ is an ancestor (See Definition 32.6.1) of $u_{j}$ in par. Let $n=|V|, N=n+q$. MultiPath(par, $Q$ ) (Algorithm 32.14) can be implemented in ( $\gamma, \delta$ ) - MPC model for any $\gamma$ with $N \log N+$ $\sum_{i=1}^{q}\left(\operatorname{dep}_{\mathrm{par}}\left(u_{i}\right)-\operatorname{dep}_{\mathrm{par}}\left(v_{i}\right)+1\right)=O\left(N^{\gamma}\right)$ and any constant $\delta \in(0,1)$ in $O(\operatorname{dep}(\operatorname{par}))$ parallel running time.

Proof. By Lemma 32.8.8, line 3 can be implemented in space $O(N \log N)$ and $O(\log (\operatorname{dep}(\operatorname{par})))$ parallel running time. It is easy to see that all the other steps in the procedure can be done by the operations shown in Multiple queries. Notice that after each round, we need to do load balancing (see Load balance) to make each machine have large enough available local
memory. The total space needed is to store all the pathes and the output of line 3. Notice that in round $i$, we do not need to keep $S_{j}^{\left(i^{\prime}\right)}$ for $i^{\prime}<i-1$, thus, the space to keep $S_{j}^{(i)}$ for all $j \in[q]$ only needs $O\left(\sum_{j=1}^{q}\left(\operatorname{dep}_{\mathrm{par}}\left(u_{j}\right)-\operatorname{dep}_{\mathrm{par}}\left(v_{j}\right)+1\right)\right)$ space.

Thus, the total space needed is at most $O\left(N \log N+\sum_{i=1}^{q}\left(\operatorname{dep}_{\mathrm{par}}\left(u_{i}\right)-\operatorname{dep}_{\mathrm{par}}\left(v_{i}\right)+\right.\right.$ $1))=O\left(N^{\gamma}\right)$. The parallel running time is then $O(\operatorname{dep}(\operatorname{par}))$.

### 32.8.8 Leaf Sampling

Lemma 32.8.18. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$, and par has a unique root. Let $n=|V|$. Let $\delta$ be an arbitrary constant in $(0,1)$, and let $m=\left\lceil n^{\delta}\right\rceil$. Then LeafSAmpling(par, $m, \delta$ ) (Algorithm 32.15) can be implemented in $(\gamma, \delta)$ - MPC model for any $\gamma \geq \log \log n / \log n$. Furthermore, with probability at least $1-1 /\left(100 m^{5 / \delta}\right)$, the parallel running time is at most $O(\log \operatorname{dep}(\mathrm{par}))$.

Proof. To implement line 4 , for each $v \in V$, we can add $\operatorname{part}(v)$ to a temporary set $X$. Then each $v$ can check whether $v$ is a leaf by checking whether $v$ is in $X$, and this can be done by the operations shown in Set membership and Multiple queries.

To implement line 5 , for each $v \in V$, we can add $v$ to the set $\operatorname{child}_{\operatorname{par}}(\operatorname{part}(v))$. Then rank can be computed by the operations shown in Indexing elements in sets and Multiple queries. For line 6 , we can implement it on a single machine, since a single machine has local memory $\Theta(m)$. For line 7 to line 9 , for each $x \in L$, we add $x$ into $S$ with probability $p$, where $p$ can be computed by querying the size of $L$ (see Sizes of sets and Multiple queries). Line 10 can be implemented by operation described in Indexing elements in sets, Set membership, and Multiple queries. By Lemma 32.8.2, line 11 can be implemented in total space $O(N \log N)$ and $O(\log \operatorname{dep}($ par $))$ parallel time. By Property 3 of Lemma 32.6.5, with probability at least $1-1 /\left(100 m^{5 / \delta}\right),|S|^{2}=O(m)$. Thus, $Q$ can be stored on a single machine. By Lemma 32.8.16, line 15 can be implemented in total space $O(n \log n+|Q|)=O(n \log n)$ and in $O(\log \operatorname{dep}($ par $))$ parallel time. By Lemma 32.8.8, line 17 can be implemented in total space $O(n \log n)$ and in $O(\log \operatorname{dep}(\operatorname{par}))$ parallel time. Then line 18 to line 22 can be implemented on a single machine.

Thus, the total space needed is at most $O(n \log n)$. The parallel time is at most $O(\log \operatorname{dep}($ par $))$

### 32.8.9 DFS Sequence

Lemma 32.8.19. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$, and par has a unique root. Let $n=|V|$. Let $\delta$ be an arbitrary constant in $(0,1)$, and let $m=\left\lceil n^{\delta}\right\rceil$. SubDFS (par, $\left.m, \delta\right)($ Algorithm 32.16) can be implemented in $(\gamma, \delta)-\mathrm{MPC}$ model for any $\gamma \geq \log \log n / \log n$. Furthermore, with probability at least $1-1 /\left(100 m^{5 / \delta}\right)$, the parallel running time is at most $O(\log \operatorname{dep}(\mathrm{par}))$.

Proof. By Lemma 32.8.18, line 5 can be implemented in total space $O(n \log n)$ and with probability at least $1-1 /\left(100 m^{5 / \delta}\right)$ has parallel running time $O(\log \operatorname{dep}(p a r))$. By Lemma 32.8.16, line 7 can be implemented in total space $O(n \log n)$ and in parallel running time $O(\log \operatorname{dep}($ par $))$. Line 9 can be implemented by operation shown in Multiple queries. By Lemma 32.8.17, since all the pathes are disjoint (except the first path and the last path intersecting on the root) and $V$ has $n$ vertices, line 10 can be implemented in $O(n \log n)$ total space and in $O(\log \operatorname{dep}($ par $))$ parallel running time. Loop in line 13 and Loop in line 16 can be implemented in parallel, and can be implemented by operations shown in Indexing elements in sets and Multiple queries. Line 20 can be implemented by operations shown in Indexing elements in sets and Multiple queries. Now we describe the implementation of line 21. Firstly, we can standardize (see Sequence standardizing) the sequence $A^{\prime}$. For each tuple (" $A^{\prime \prime},(j, u)$ ), create a tuple ("tempu", $)$ ). Thus, "tempu" is a set which contains all the positions that $u$ appeared. For each tuple ("tempu", $j$ ), we query (see Multiple queries) the index $i$ (see Indexing elements in sets) of $j$ in set ("tempu", $j$ ), and create a tuple ("pos", ((u,i),j)). Thus, the desired mapping pos is stored in the system. The loop in line 24 is implemented in parallel. Line 25 can be implemented by the operations shown
in Set membership and Multiple queries. Line 26 to line 28 can be implemented by the operation shown in Multiple queries. Finally, line 30 can be implemented by Multiple queries and Sequence duplicating.

The total space used in the procedure is at most $O(n \log n)$. The parallel running time is $O(\log \operatorname{dep}($ par $))$.

Theorem 32.8.20. Let par : $V \rightarrow V$ be a set of parent pointers (See Definition 32.4.2) on a vertex set $V$, and par has a unique root. Let $n=|V|, m=n^{\delta}$ for some arbitrary constant $\delta \in(0,1)$. DFS(par, $m$ ) (Algorithm 32.17) can be implemented in $(\gamma, \delta)$ - MPC model for any $\gamma \geq \log \log n / \log n$. With probability at least 0.99 , the parallel running time is $O(\log (\operatorname{dep}(\operatorname{par})))$.

Proof. By Lemma 32.8.19, line 5 can be implemented in total space $O(n \log n)$. With probability at least $1-1 /\left(100 n^{5}\right)$, the parallel running time is $O(\log (\operatorname{dep}(\operatorname{par})))$. Line 8 to line 10 can be implemented by operations shown in Set membership and Multiple queries. By Lemma 32.8.2, line 11 can be implemented in $O(n)$ total space, and $O(\log \operatorname{dep}(p a r))$ parallel running time. The loop in line 14 contains multiple tasks (see Section 32.7.6 Multiple Tasks), thus we can implement those tasks in parallel. By Lemma 32.8.19, line 17 can be implemented in total space $O\left(\left|V_{i}^{\prime}(v)\right| \log \left|V_{i}^{\prime}(v)\right|\right)$. Furthermore, with probability at least $1-1 /\left(100 n^{5}\right)$, the parallel running time is $O(\log (\operatorname{dep}(\operatorname{par})))$. Thus, the total space needed for those tasks is at most $O(n \log n)$. Line 19 can be implemented by operations shown in Indexing elements in sets, Sequence insertion and Multiple queries.

Thus, the total space needed is $O(n \log n)$. By taking union bound over all the task SubDFS, with probability at least 0.99 , the parallel running time is $O(\log \operatorname{dep}(\mathrm{par}))$.

Now we are able to conclude the following theorem.

Theorem 32.8.21. For any $\gamma \in[\beta, 2]$ and any constant $\delta \in(0,1)$, there is a randomized $(\gamma, \delta)$-MPC algorithm (Algorithm 32.17) which can output a Depth-First-Search sequence for any tree graph $G=(V, E)$ in $O\left(\min \left(\log D \cdot \log \left(1 / \gamma^{\prime}\right), \log n\right)\right)$ parallel time, where $n=|V|$, $\beta=\Theta(\log \log n / \log n)$, $D$ is the diameter of $G$, and $\gamma^{\prime}=\gamma+\Theta(1 / \log n)$. The success probability is at least 0.98. In addition, if the algorithm fails, then it will return FAIL.

Proof. Firstly, by Theorem 32.8.14, we can find a rooted tree. Algorithm 32.17 can output the DFS sequence for a rooted tree.

The implementation and parallel time of Algorithm 32.17 is shown by Theorem 32.8.20. The correctness of Algorithm 32.17 is proved by Theorem 32.6.12. The success probability of Algorithm 32.17 is proved by Theorem 32.6.14.

### 32.8.10 Range Minimum Query

Lemma 32.8.22. Let $A=\left(a_{1}, a_{2}, \cdots, a_{n}\right)$ be a sequence of numbers. Let $\delta$ be an arbitrary constant in $(0,1)$. SparseTABLE ${ }^{+}\left(a_{1}, a_{2}, \cdots, a_{n}, \delta\right)$ (Algorithm 32.18) can be implemented in $(0, \delta)$ - MPC model with $O(1)$ parallel running time.

Proof. Let $A$ be the sequence $\left(a_{1}, a_{2}, \cdots, a_{n}\right)$. The algorithm takes $O(1 / \delta)$ rounds. $m$ is the local space of a machine. There are $\Theta(n / m)$ machines each holds a consecutive $\Theta(m)$ elements of sequence $A$. Now consider the round $l$. Machine $j \in\{0\} \cup[\lceil n / m\rceil]$ needs to compute $\widehat{f}_{j \cdot m+1, l}, \widehat{f}_{j \cdot m+2, l}, \cdots, \widehat{f}_{j \cdot m+m-1, l}$. The number of queries machine $j$ made in line 8 and line 11 is at most $\sum_{t=1}^{[1 / \delta\rceil}\left|S_{t}\right|+2 m \leq O(m / \delta)=O(m)$. Thus, there are total $O(n)$ queries. These queries can be answered simultaneously by operation shown in Multiple queries.

Thus, the total space needed is $O(n)$, and the parallel running time is $O(1)$.
Lemma 32.8.23. Let $a_{1}, a_{2}, \cdots, a_{n}$ be a sequence of numbers. Let $\delta$ be an arbitrary constant in $(0,1) . \operatorname{SparseTable}\left(a_{1}, a_{2}, \cdots, a_{n}, \delta\right)$ (Algorithm 32.19) can be implemented in $(\gamma, \delta)-$ MPC model for any $\gamma \geq \log \log n / \log n$ in $O(1)$ parallel time.

Proof. By Lemma 32.8.22, line 4 can be implemented in $O(n)$ total space and $O(1)$ parallel time. The loop in line 15 is similar to Algorithm 32.18. Each machine $j$ needs to compute $f_{j \cdot m+1, t}, \cdots, f_{j \cdot m+m-1, t}$ for all $t \in[\lceil\log n\rceil] \cup\{0\}$. The difference from Algorithm 32.18 is that, it can compute for all $t$ at the same time since it only depends on the value of $\widehat{f}$. The number of queries made by each machine is $O(m \log n)$. Thus, the total number of queries is at most $O(n \log n)$. These queries can be answered simultaneously by operation shown in Multiple queries.

Thus, the total space needed is $O(n \log n)$, and the parallel running time is $O(1)$.

### 32.9 Minimum Spanning Forest

In this section, we discuss how to apply our connectivity/spanning forest algorithm to the Minimum Spanning Forest (MSF) and Bottleneck Spanning Forest (BSF) problem.

The input of MSF/BSF problem is an undirected graph $G=(V, E)$ together with a weight function $w: E \rightarrow \mathbb{Z}$, where $E$ contains $m$ edges $e_{1}, e_{2}, \cdots, e_{m}$ with $w\left(e_{1}\right) \leq w\left(e_{2}\right) \leq$ $\cdots \leq w\left(e_{m}\right)$. The goal of MSF is to output a spanning forest such that the sum of weights of the edges in the forest is minimized. The goal of BSF is to output a spanning forest such that the maximum weight of the edges in the forest is minimized. $D$ is the diameter of the minimum spanning forest. If there are multiple choices of the minimum spanning forest, then let $D$ be the minimum diameter among all the minimum spanning forests.

For simplicity, in all of our proofs, we only discuss the case when all the edges have different weights, i.e. $w\left(e_{1}\right)<w\left(e_{2}\right)<\cdots<w\left(e_{m}\right)$. In this case, the minimum spanning forest is unique. It is easy to extend our algorithms to the case when there are edges with the same weight. We omit the proof for this fact.

Firstly, we show that $D$ is an upper bound of the diameter of $G^{\prime}$ where the vertex set of $G^{\prime}$ is the vertex set of $G$, and the edge set of $G^{\prime}$ is $\left\{e_{1}, e_{2}, \cdots, e_{i}\right\}$ for some arbitrary $i \in[m]$.

Lemma 32.9.1. Given a graph $G=(V, E)$ for $E=\left\{e_{1}, e_{2}, \cdots, e_{m}\right\}$ together with a weight function $w$ which satisfies $w\left(e_{1}\right)<w\left(e_{2}\right)<\cdots<w\left(e_{m}\right)$, then the diameter of $G^{\prime}=\left(V, E^{\prime}\right)$ is at most $D$, where $D$ is the diameter of the minimum spanning forest of $G$, and $E^{\prime}$ only contains the first $i$ edges of $E$, i.e. $e_{1}, e_{2}, \cdots, e_{i}$ for some arbitrary $i \in[m]$.

Proof. The proof follows by Kruskal's algorithm directly.

Our algorithms is based on the following simple but useful Lemma.

Lemma 32.9.2. Given a graph $G=(V, E)$ for $E=\left\{e_{1}, e_{2}, \cdots, e_{m}\right\}$ together with a weight function $w$ which satisfies $w\left(e_{1}\right) \leq w\left(e_{2}\right) \leq \cdots \leq w\left(e_{m}\right), \forall 1 \leq i<j \leq m$ an edge e from $\left\{e_{i}, e_{i+1}, \cdots, e_{j}\right\}$ is in the minimum spanning forest of $G$ if and only if $e^{\prime}$ from $\left\{e_{i}^{\prime}, e_{i+1}^{\prime}, \cdots, e_{j}^{\prime}\right\}$ is in the minimum spanning forest of $G^{\prime}$, where the vertices of $G^{\prime}$ is obtained by contracting all the edges $e_{1}, e_{2}, \cdots, e_{i-1}$ of $G$, and $e^{\prime}, e_{i}^{\prime}, e_{i+1}^{\prime}, \cdots, e_{j}^{\prime}$ are the edges (or vertices) in $G^{\prime}$ which corresponds to the edges e, $e_{i}, e_{i+1}, \cdots, e_{j}$ before contraction.

Proof. The proof follows by Kruskal's algorithm directly.

A natural way to apply Lemma 32.9.2 to parallel minimum spanning forest algorithm is that we can divide the edges into several groups, and recursively solve the minimum spanning forest for each group of edges. More precisely, suppose we have total space $\Theta(\mathrm{km})$, we can divide $E$ into $k$ groups $E_{1}, E_{2}, \cdots, E_{k}$, where $E_{i}=\left\{e_{(i-1) \cdot m / k+1}, e_{(i-1) \cdot m / k+2}, \cdots, e_{i \cdot m / k}\right\}$. We can compute graph $G_{1}, G_{2}, \cdots, G_{k}$ where the vertices of $G_{i}$ is obtained by contracting all the edges from $e_{1}$ to $e_{(i-1) \cdot m / k}$, the edges of $G_{i}$ are corresponding to the edges in $E_{i}$. Then by Lemma 32.9.2, we can obtain the whole minimum spanning forest by solving these $k$ size $O(m / k)$ minimum spanning forest problems. For each sub-problem, we can assign it $\Theta(m)$ working space, thus each sub-problem still has $\Theta(k)$ factor more total space. Therefore, we can recursively apply the above argument.

Theorem 32.9.3. For any $\gamma \in[0,2]$ and any constant $\delta \in(0,1)$, there is a randomized $(\gamma, \delta)$ - MPC algorithm which can output a minimum spanning forest for any weighted graph $G=(V, E)$ with weights $w: E \rightarrow \mathbb{Z}$ in $O\left(\min \left(\log D \cdot \log \left(1 / \gamma^{\prime}\right), \log n\right) \cdot 1 / \gamma^{\prime}\right)$ parallel time,
where $n=|V|, \forall e \in E,|w(e)| \leq \operatorname{poly}(n), D$ is the diameter of a minimum spanning forest of $G$, and $\gamma^{\prime}=\gamma / 2+\Theta(1 / \log n)$. The success probability is at least 0.98 . In addition, if the algorithm fails, then it will return FAIL.

Proof. Let $n=|V|, m=|E|$. Let $E=\left\{e_{1}, \cdots, e_{m}\right\}$ with $w\left(e_{1}\right) \leq w\left(e_{2}\right) \leq \cdots \leq w\left(e_{m}\right)$. The total space in the system is $\Theta\left(m^{1+\gamma}\right)$. Let $k=\Theta\left(m^{\gamma / 2}\right)$. By our previous discussion, we can divide $E$ into $k$ groups $E_{1}, E_{2}, \cdots, E_{k}$, where $E_{i}=\left\{e_{(i-1) \cdot m / k+1}, e_{(i-1) \cdot m / k+2}, \cdots, e_{i \cdot m / k}\right\}$. By Lemma 32.9.1 and Theorem 32.8.4, we can use $O\left(\min \left(\log D \cdot \log \left(1 / \gamma^{\prime}\right), \log n\right)\right)$ parallel time and $\Theta\left(k m^{1+\gamma / 2}\right)$ total space to compute graph $G_{1}, G_{2}, \cdots, G_{k}$ where the vertices of $G_{i}$ is obtained by contracting all the edges from $e_{1}$ to $e_{(i-1) \cdot m / k}$, the edges of $G_{i}$ are corresponding to the edges in $E_{i}$ after contraction.

By Lemma 32.9.2, it suffices to recursively solve the minimum spanning forest problem for each group $G_{i}$. Since each time, we split the edges into $k$ groups, the recursion will have at most $O\left(1 / \gamma^{\prime}\right)$ levels. At the end of the recursion, we are able to determine for every edge $e$ whether $e$ is in the minimum spanning forest.

Now let us consider the success probability. Although Theorem 32.8.4 is a randomized algorithm, the parallel time is always bounded by $\min \left(\log D \cdot \log \left(1 / \gamma^{\prime}\right), \log n\right)$. If we repeat the algorithm until it succeeds, the expectation of number of trials is a constant. Furthermore, for each level of the recursion, we can regard the graphs in all the tasks composed one large graph. Thus, in real implementation, in each level of the recursion, we will only invoke one connectivity procedure. Thus in expectation, the total parallel time is $O(\min (\log D$. $\left.\left.\log \left(1 / \gamma^{\prime}\right), \log n\right) \cdot 1 / \gamma^{\prime}\right)$. By applying Markov's inequality, we complete the proof.

In the following theorem, we show that Lemma 32.9.2 can also be applied in approximate minimum spanning forest problem.

Theorem 32.9.4. For any $\gamma \in[\beta, 2]$ and any constant $\delta \in(0,1)$, there is a randomized $(\gamma, \delta)$ - MPC algorithm which can output a $(1+\epsilon)$ approximate minimum spanning forest for any weighted graph $G=(V, E)$ with weights $w: E \rightarrow \mathbb{Z}_{\geq 0}$ in $O(\min (\log D$. $\left.\left.\log \left(1 / \gamma^{\prime}\right), \log n\right)\right)$ parallel time, where $n=|V|, N=|V|+|E|, \beta=\Theta\left(\log \left(\epsilon^{-1} \log n\right) / \log n\right)$, $\forall e \in E,|w(e)| \leq \operatorname{poly}(n), D$ is the diameter of a minimum spanning forest of $G$, and $\gamma^{\prime}=(1+\gamma-\beta) \log _{n} \frac{2 N}{n^{1 /(1+\gamma-\beta)}}$. The success probability is at least 0.98. In addition, if the algorithm fails, then it will return FAIL.

Proof. For each edge $e \in E$, we can round $w(e)$ to $w^{\prime}(e)$ such that $w^{\prime}(e)=0$ when $w(e)=0$, and $w^{\prime}(e)=(1+\epsilon)^{i}$ when $w(e) \neq 0$, and $i$ is the smallest integer such that $w(e) \leq(1+\epsilon)^{i}$.

Since $|w(e)| \leq \operatorname{poly}(n)$ for all $e \in E$, there are only $k=O(\log (n) / \epsilon)$ different values of $w^{\prime}(e)$. We can divide $E$ into $k$ groups, where the $i^{\text {th }}$ group $E_{i}$ contains all edges with the $i^{\text {th }}$ largest weight in $w^{\prime}$. By Lemma 32.9.1 and Theorem 32.8.4, we can use $O(\min (\log D$. $\left.\left.\log \left(1 / \gamma^{\prime}\right), \log n\right)\right)$ parallel time and $\Theta\left(k N^{1+\gamma-\beta}\right)=\Theta\left(N^{1+\gamma}\right)$ total space to compute graph $G_{1}, G_{2}, \cdots, G_{k}$ where the vertices of $G_{i}$ is obtained by contracting all the edges from $E_{1}$ to $E_{i-1}$, the edges of $G_{i}$ are corresponding to the edges in $E_{i}$ after contraction.

Then, for each $G_{i}$, since all the edges have the same $w^{\prime}$ weight, any spanning forest of $G_{i}$ is a minimum spanning forest of $G_{i}$. By Theorem 32.8.14, we can use $O(\min (\log D$. $\left.\log \left(1 / \gamma^{\prime}\right), \log n\right)$ ) parallel time and $\Theta\left(k N^{1+\gamma-\beta}\right)=\Theta\left(N^{1+\gamma}\right)$ total space to compute the spanning forest for each graph $G_{1}, G_{2}, \cdot, G_{k}$. By Lemma 32.9.2, the union of all the minimum spanning forest with respect to $w^{\prime}$ must be the minimum spanning forest of $G$ with respect
to $w^{\prime}$. Since all the weights $w$ are nonnegative integers, $w^{\prime}$ is a $(1+\epsilon)$ approximation to $w$. Therefore, our output minimum spanning forest with respect to $w^{\prime}$ is a $(1+\epsilon)$ approximation to the minimum spanning forest with respect to $w$.

For the success probability, we can apply the similar argument made in the proof of Theorem 32.9.3 to prove that the success probability is at least 0.98 .

In the following, we show that if we only need to find the largest edge in the minimum spanning tree, then we are able to get a better parallel time. It is an another application of our

Theorem 32.9.5. For any $\gamma \in[0,2]$ and any constant $\delta \in(0,1)$, there is a randomized $(\gamma, \delta)$ - MPC algorithm which can output a bottleneck spanning forest for any weighted graph $G=(V, E)$ with weights $w: E \rightarrow \mathbb{Z}$ in $O\left(\min \left(\log D \cdot \log \left(1 / \gamma^{\prime}\right), \log n\right) \cdot \log \left(1 / \gamma^{\prime}\right)\right)$ parallel time, where $n=|V|, \forall e \in E,|w(e)| \leq \operatorname{poly}(n), D$ is the diameter of a minimum spanning forest of $G$, and $\gamma^{\prime}=\gamma / 2+\Theta(1 / \log n)$. The success probability is at least 0.98 . In addition, if the algorithm fails, then it will return FAIL.

Proof. Let $n=|V|, m=|E|$. Let $E=\left\{e_{1}, \cdots, e_{m}\right\}$ with $w\left(e_{1}\right) \leq w\left(e_{2}\right) \leq \cdots \leq w\left(e_{m}\right)$. The total space in the system is $\Theta\left(m^{1+\gamma}\right)$. Let $k=\Theta\left(m^{\gamma / 2}\right)$. By our previous discussion, we can divide $E$ into $k$ groups $E_{1}, E_{2}, \cdots, E_{k}$, where $E_{i}=\left\{e_{(i-1) \cdot m / k+1}, e_{(i-1) \cdot m / k+2}, \cdots, e_{i \cdot m / k}\right\}$. By Lemma 32.9.1 and Theorem 32.8.4, we can use $O\left(\min \left(\log D \cdot \log \left(1 / \gamma^{\prime}\right), \log n\right)\right)$ parallel time and $\Theta\left(k m^{1+\gamma / 2}\right)$ total space to compute graph $G_{1}, G_{2}, \cdots, G_{k}$ where the vertices of $G_{i}$ is obtained by contracting all the edges from $e_{1}$ to $e_{(i-1) \cdot m / k}$, the edges of $G_{i}$ are corresponding to the edges in $E_{i}$ after contraction.

By Lemma 32.9.2, the edge with largest weight must be in the group $E_{i}$ for some $i$ with $G_{i+1}=G_{i+2}$. Thus, we reduce the problem size to $m / k$. By Remark 32.1.2, we can finish the recursion in $O\left(\log \left(1 / \gamma^{\prime}\right)\right)$ phases.

Suppose the bottleneck is $e_{i}$, then by Theorem 32.8.14, we can find a spanning forest by only using edges from $\left\{e_{1}, \cdots, e_{i}\right\}$ in $O\left(\min \left(\log D \cdot \log \left(1 / \gamma^{\prime}\right), \log n\right)\right)$ parallel time and in $\Theta\left(m^{1+\gamma / 2}\right)$ total space. Thus, the resulting spanning forest is a bottleneck spanning forest.

For the success probability, we can apply the similar argument made in the proof of Theorem 32.9.3 to prove that the success probability is at least 0.98 .

### 32.10 Directed Reachability vs. Boolean Matrix Multiplication

In this section, we discuss the directed graph reachability problem which is a directed graph problem highly related to the undirected graph connectivity. In the all-pair directed graph reachability problem, we are given a directed graph $G=(V, E)$, the goal is to answer for every pair $(u, v) \in V \times V$ whether there is a directed path from $u$ to $v$. There is a simple standard way to reduce Boolean Matrix Multiplication to all-pair directed graph reachability problem. In the Boolean Matrix Multiplication problem, we are given two boolean matrices $A, B \in\{0,1\}^{n \times n}$, the goal is to compute $C=A \cdot B$, where $\forall i, j \in[n], C_{i, j}=\bigvee_{k \in[n]} A_{i, k} \wedge B_{k, j}$. The reduction is as the following. We create $3 n$ vertices $u_{1}, u_{2}, \cdots, u_{n}, v_{1}, v_{2}, \cdots, v_{n}, w_{1}, w_{2}, \cdots, w_{n}$. For every $i, j \in[n]$, if $A_{i, j}=1$, then we add an edge from $u_{i}$ to $v_{j}$, and if $B_{i, j}=1$, then we add an edge from $v_{i}$ to $w_{j}$. Thus, $C_{i, j}=1$ is equivalent to there is a path from $u_{i}$ to $w_{j}$. Thus, if we can solve all-pair directed graph reachability problem in $O(T)$ sequential time, then we can solve Boolean Matrix Multiplication in $O(T)$ time. For the current status of sequential running time of Boolean Matrix Multiplication problem, we refer readers to [LG14] and the references therein.

Now, consider the multi-query directed graph reachability problem. In this problem, we are given a directed graph $G=(V, E)$ together with $|V|+|E|$ queries where each query queries the reachability from vertex $u$ to vertex $v$. The goal is to answer all these queries. A similar problem in the undirected graph is called multi-query undirected graph connectivity problem. In this problem, we are given an undirected graph $G=(V, E)$ together with $|V|+|E|$ queries where each query queries the connectivity between vertex $u$ and vertex $v$.

According to Theorem 32.8.4 and Lemma 32.7.4, there is a polynomial local running time fully scalable $\sim \log D$ parallel time $(0, \delta)$ - MPC algorithm for multi-query undirected
graph connectivity problem. Here polynomial local running time means that there is a constant $c>0$ (independent from $\delta$ ) such that every machine in one round can only have $O\left(\left(n^{\delta}\right)^{c}\right)$ local computation.

For multi-query directed graph reachability problem, we show that if there is a polynomial local running time fully scalable $(\gamma, \delta)$ - MPC algorithm which can solve multi-query reachability problem in $O\left(n^{\alpha}\right)$ parallel time, then we can solve all-pair directed graph reachability problem in $O\left(n^{2} \cdot n^{2 \gamma+\alpha+\epsilon}\right)$ sequential running time for any arbitrarily small constant $\epsilon>0$. Especially, if the algorithm is in $(0, \delta)-$ MPC model, and the parallel time is $n^{o(1)}$, then we will have an $O\left(n^{2+\epsilon+o(1)}\right)$ sequential running time algorithm for Boolean Matrix Multiplication which implies a break through in this field.

Suppose we have a such MPC algorithm. Let the input size be $\Theta(m)$, i.e. the number of edges is $\Theta(m)$, and the number of queries is also $\Theta(m)$. Then the total space is $\Theta\left(m^{1+\gamma}\right)$. Let $\delta=\epsilon /(c-2)$. Then the number of machines is $\Theta\left(m^{1+\gamma-\delta}\right)$. Now we just simulate this $(\gamma, \delta)$ - MPC algorithm sequentially, the total running time is $O\left(m^{1+\gamma-\delta} \cdot m^{c \delta} \cdot n^{\alpha}\right)=O(m$. $\left.n^{2 \gamma+\epsilon+\alpha}\right)$. To answer reachability for all pairs, we need total $O\left(n^{2} \cdot m \cdot n^{2 \gamma+\epsilon+\alpha} / m\right)=O\left(n^{2}\right.$. $\left.n^{2 \gamma+\alpha+\epsilon}\right)$ time. Therefore, we can use this algorithm to solve Boolean Matrix Multiplication in $O\left(n^{2} \cdot n^{2 \gamma+\alpha+\epsilon}\right)$ time.

Theorem 32.10.1. If there is a polynomial local running time fully scalable $(\gamma, \delta)$ - MPC algorithm which can answer $|V|+|E|$ pairs of reachability queries simultaneously for any directed graph $G=(V, E)$ in $O\left(|V|^{\alpha}\right)$ parallel time, then there is a sequential algorithm which can compute the multiplication of two $n \times n$ boolean matrices in $O\left(n^{2} \cdot n^{2 \gamma+\alpha+\epsilon}\right)$ time, where $\epsilon>0$ is a constant which can be arbitrarily small.

Proof. See above discussions.

### 32.11 Discussion on a Previous Conjectured Fast Algorithm

In this section, we discuss the hard example for the algorithm described by [RMCS13]. In [RMCS13], they conjectured that their Hash-to-Min connectivity algorithm can finish in $O(\log D)$ rounds. The description of their algorithm is as the following:

1. The input graph is $G=(V, E)$.
2. For each vertex $v \in V$, initialize a set $S_{v}^{(0)}=v$.
3. in round $i$ :
(a) Each vertex $v$ find $u \in S_{v}^{(i-1)}$ which has the minimum label, i.e. $u=\min _{x \in S_{v}^{(i-1)}} x$.
(b) $v$ sends $u$ the all the vertices in $S_{v}^{(i-1)}$.
(c) $v$ sends every $x \in S_{v}^{((i-1))} \backslash\{u\}$ the vertex $u$.
(d) Let $S_{v}^{(i)}$ be $\{v\}$ union the set of all the vertices received.
(e) If for all $v, S_{v}^{(i)}$ is the same as $S_{v}^{(i-1)}$, then finish the procedure.

The above procedure can be seen as the modification of the graph: in each round, all the vertices together create a new graph. For each vertex $v$, let $u$ be the neighbor of $v$ with the minimum label, and if $x$ is a neighbor of $v$, then add an edge between $x$ and $u$ in the new graph. So in each round, each vertex just communicates with its neighbors to update the new minimum neighbor it learned. At the end of the algorithm, it is obvious that the minimum vertex in each component will have all the other vertices in that component, and for each non minimum vertex, it will have the minimum vertex in the same component.

A hard example for this algorithm is shown by Figure 32.3. The example is a thin and tall grid graph with a vertex connected to all the vertices in the first column. The total number of vertices is $n$. The grid graph has $D=\frac{1}{2} \log n$ columns and $n / D$ rows. We index each column from left to right by 1 to $D$. We index each row from top to down by 1 to $n / D$. The single large degree vertex has label 0 . The $i^{\text {th }}$ row has the vertices with label $(i-1) \cdot D+1$ to $i \cdot D$ from the first column to the $D^{\text {th }}$ column. We claim that if vertex $v$ is the $i^{\text {th }}$ row and $j^{\text {th }}$ column, then before round $k$ for $2^{k}<i, k<j$, the neighbors of $v$ will only in column $j-1$, column $j$ and column $j+1$. Furthermore, the minimum neighbor of $v$ in column $j-1$ will be $v-\left(2^{k-1}-1\right) \cdot D-1$. The minimum neighbor of $v$ in column $j$ will be $v-2^{k-1} \cdot D$. The minimum neighbor of $v$ in column $j+1$ will be $v-D \cdot\left(2^{k-1}-1\right)+1$. This claim is true when $k=1$. Then by induction, we can prove the claim. Thus, it will take at least $\Theta(D)$ rounds to finish the procedure where $D=\Theta(\log n)$.

If we randomly label the vertices at the beginning, then consider the case we copy that hard structure at least $n^{n+2}$ times, then with high probability, there is a component which has the labels with the order as the same as described above. In this case, the procedure needs $\Omega(\log \log N)$ rounds, where $N=n^{n+3}$ is the total number of the vertices.

Also notice that, even we give more total space to this algorithm, this algorithm will not preform better. In our connectivity algorithm, if we have $\Omega\left(n^{1+\epsilon}\right)$ total space for some arbitrary constant $\epsilon>0$, then our parallel running time is $O(\log D)$.

### 32.12 Alternative Approach for Leader Selection

In this section, we show that there is a different way to select leaders (see Section 32.4.2). The number of leaders selected by this approach will depend on the sum of inverse degrees of all the vertices. Let us first introduce the concept of Min Parent Forest.

### 32.12.1 Min Parent Forest

Let $G=(V, E)$ be an undirected graph where $V$ denotes the vertex set of $G$, and $E$ denotes the edge set of $G$. Each vertex $v \in V$ has a weight $w(v) \in \mathbb{R}$, and it also has a unique label from $\mathbb{Z}$. For convenience, for each vertex $v \in V$, we also use $v$ to denote its label. Let $\Gamma_{G}(v)$ denote the set of neighbors of $v$, i.e. $\Gamma_{G}(v)=\{u \in V \mid(u, v) \in E\}$. If $G$ is clear in the context, we just use $\Gamma(v)$ to denote $\Gamma_{G}(v)$. The size of $\Gamma(v),|\Gamma(v)|$, is called the degree of $v$. Let $f_{G, w}: V \rightarrow V$ be the "min-weight-parent" function defined as the following:

1. If $w(v)=\min _{u \in \Gamma(v) \cup\{v\}} w(u)$, then $f_{G, w}(v)=v$.
2. Otherwise, let $u^{*} \in \Gamma(v)$ be the vertex which has the smallest weight, i.e. $w\left(u^{*}\right)=$ $\min _{u \in \Gamma(v)} w(u)$. If there is more than one choice of $u^{*}$, let $u^{*}$ be the one with the smallest label. And $f_{G, w}(v)$ is defined to be $u^{*}$.

We call ( $V, f_{G, w}$ ) the min-parent-forest of graph $G$ with vertex weights $w$. We can then define $i$-step "min-weight-parent" function. For $v \in V$, we define $f_{G, w}^{(0)}(v)=v$. For $i \in \mathbb{Z}_{>0}$, we can define $f_{G, w}^{(i)}$ as the following:

$$
\forall v \in V, f_{G, w}^{(i)}(v)=f_{G, w}\left(f_{G, w}^{(i-1)}(v)\right)
$$

In the following, we define the concept of roots in the min-parent-forest.

Definition 32.12.1 (Roots in the forest). Let $v \in V$, and let $\left(V, f_{G, w}\right)$ be the min-parentforest of graph $G=(V, E)$ with vertex weights $w$. If $f_{G, w}(v)=v$, then $v$ is a root in the forest $\left(V, f_{G, w}\right)$.

The depth of a vertex $v$ is defined as the distance on the tree between $v$ and the corresponding root in the forest.

Definition 32.12.2 (The depth of $v$ ). Let $v \in V$, and let $\left(V, f_{G, w}\right)$ be the min-parent-forest of graph $G=(V, E)$ with vertex weights $w$. The depth of $v$ in the forest $\left(V, f_{G, w}\right)$ is the smallest $i \in \mathbb{Z}_{\geq 0}$ such that $f_{G, w}^{(i)}(v)=f_{G, w}^{(i+1)}(v)$. We use $\operatorname{dep}_{G, w}(v)$ to denote the depth of $v$ in $\left(V, f_{G, w}\right)$. We call $f_{G, w}^{\left(\operatorname{dep}_{G, w}(v)\right)}(v)$ the root of $v$. For the simplicity of the notation, we also use $f_{G, w}^{(\infty)}(v)$ to denote the root of $v$.

The above definition is well defined since if $f_{G, w}^{(i+1)}(v) \neq f_{G, w}^{(i)}(v)$ then $w\left(f_{G, w}^{(i+1)}(v)\right)$ should be strictly smaller than $w\left(f_{G, w}^{(i)}(v)\right)$ by the definition of $f_{G, w}$ and $f_{G, w}^{(j)}$ for all $j \in \mathbb{Z}_{\geq 0}$. Therefore, there must exist $i$ such that $f_{G, w}^{(i)}(v)=f_{G, w}^{(i+1)}(v)$.

The depth of the forest is the largest depth among all the vertices.

Definition 32.12.3 (The depth of the min-parent-forest). The depth $\operatorname{dep}(G, w)$ of the forest ( $V, f_{G, w}$ ) is defined as:

$$
\operatorname{dep}(G, w)=\max _{v \in V} \operatorname{dep}_{G, w}(v)
$$

If the weights $w$ of vertices of $G$ are some i.i.d. random variables, then with high probability, the depth of $\left(V, f_{G, w}\right)$ is only $O(\log |V|)$. Precisely, we have the following Lemma.

Lemma 32.12.1 (The depth of the random min-parent-forest). Let $G=(V, E)$ be an undirected graph with $n$ vertices where $V=\left\{v_{1}, v_{2}, \cdots, v_{n}\right\}$, and the labels satisfies $v_{1}<$ $v_{2}<\cdots<v_{n}$. Let $w\left(v_{1}\right), w\left(v_{2}\right), \cdots, w\left(v_{n}\right)$ be $n$ i.i.d. random variables drawn uniformly from $[N]$. If $N>n^{2} / \delta$ for some $\delta \in(0,1)$, then for any $t \geq 60 \log n$,

$$
\operatorname{Pr}_{w \sim[N]^{n}}(\operatorname{dep}(G, w) \leq t) \geq 1-\delta-e^{-\frac{1}{2} t} .
$$

Proof. Let $w\left(v_{1}\right), w\left(v_{2}\right), \cdots, w\left(v_{n}\right)$ be $n$ i.i.d. random variables drawn uniformly from $[N]$. Let $\left(V, f_{G, w}\right)$ be the min-parent-forest of $(G, w)$. For a fixed $s \in V$, we create a set of random variables $z_{1}, z_{2}, \cdots, z_{n}$ by the following deterministic procedure:

1. Let $z_{1}=w(s), k=0, S_{k}=\{s\}, u_{k}=s, i=2, \operatorname{pos}(s) \leftarrow 1$.
2. Let $S_{k+1}=S_{k}$.
3. For $j=1 \rightarrow n$,
if $v_{j} \in \Gamma\left(u_{k}\right)$ and $v_{j} \notin S_{k}$ then let $\operatorname{pos}\left(v_{j}\right) \leftarrow i, S_{k+1} \leftarrow S_{k+1} \cup\left\{v_{j}\right\}, z_{i}=w\left(v_{j}\right)$, $i \leftarrow i+1$.
4. If $f_{G, w}\left(u_{k}\right) \neq u_{k}$, then let $u_{k+1}=f_{G, w}\left(u_{k}\right), k \leftarrow k+1$ and go to step 2 .
5. Otherwise, for $j=1 \rightarrow n$,
if $v_{j} \notin S_{k+1}$ then let $\operatorname{pos}\left(v_{j}\right) \leftarrow i, z_{i}=w\left(v_{j}\right), i \leftarrow i+1$.

It is easy to observe that $k$ is exactly $\operatorname{dep}_{G, w}(s)$ at the end of the above procedure. The reason is that $u_{0}=s=f_{G, w}^{(0)}(s), \forall j \in[k], u_{j}=f_{G, w}\left(u_{j-1}\right)=f_{G, w}^{(j)}(s)$ and $f_{G, w}\left(u_{k}\right)=u_{k}$.

Fact 32.12.2. $\forall v \in V, w(v)=z_{\operatorname{pos}(v)}$, where pos: $[V] \rightarrow[n]$ and $\operatorname{pos}^{-1}:[n] \rightarrow[V]$.

Claim 32.12.3. $\forall j \in\{0,1, \cdots, k+1\}, S_{j}=\left\{u_{0}\right\} \cup \bigcup_{p=0}^{j-1} \Gamma\left(u_{p}\right)$.

Proof. We can prove this by induction. The statement is obviously true for $S_{0}$ since $S_{0}=$ $\left\{u_{0}\right\}$. Now suppose the claim is true for $S_{j-1}$. Then according to the step 3 of the procedure $S_{j}=S_{j-1} \cup\left(\Gamma\left(u_{j-1}\right) \backslash S_{j-1}\right)=S_{j-1} \cup \Gamma\left(u_{j-1}\right)=\left\{u_{0}\right\} \cup \bigcup_{p=0}^{j-1} \Gamma\left(u_{p}\right)$.

Claim 32.12.4. $\forall j \in\{0,1, \cdots, k\}, w\left(u_{j}\right)=\min _{v \in S_{j}} w(v)$.

Proof. Since $\forall j \in[k], u_{j}=f_{G, w}\left(u_{j-1}\right)$, we have $w\left(u_{j}\right)=\min _{v \in \Gamma\left(u_{j-1}\right) \cup\left\{u_{j-1}\right\}} w(v)$. Then we have $w\left(u_{j}\right)=\min _{v \in\left\{u_{0}\right\} \cup \bigcup_{p=0}^{j-1} \Gamma\left(u_{p}\right)} w(v)=\min _{v \in S_{j}} w(v)$, where the last equality follows by Claim 32.12.3.

We use $\operatorname{pos}^{-1}(i)$ to denote vertex $v$ which satisfies $\operatorname{pos}(v)=i$. According to the step 3 , it is easy to see $\forall j \in\{0,1, \cdots, k+1\}$, we have $\left\{\operatorname{pos}^{-1}(i) \mid i \in\left[\left|S_{j}\right|\right]\right\}=S_{j}$.

Claim 32.12.5. $\forall j \in\{0,1, \cdots, k\}, z_{\operatorname{pos}\left(u_{j}\right)}=\min _{p \in\left[\operatorname{pos}\left(u_{j}\right)\right]} z_{p}$.

Proof. $z_{\operatorname{pos}\left(u_{j}\right)}=w\left(u_{j}\right)=\min _{v \in S_{j}} w(v)=\min _{v \in S_{j}} z_{\operatorname{pos}(v)}=\min _{p \in\left[\left|S_{j}\right|\right]} z_{p}=\min _{p \in\left[\operatorname{pos}\left(u_{j}\right)\right]} z_{p}$, where the second equality follows by Claim 32.12.4, and the last equality follows by $u_{j} \in S_{j}$, so $\operatorname{pos}\left(u_{j}\right) \leq\left|S_{j}\right|$.

Now we define an another set of random variables $y_{1}, y_{2}, \cdots, y_{n}$, where $\forall i \in[n], y_{i} \in$ $\{0,1\}$ and $y_{i}=1$ if and only if $z_{i}=\min _{j \in[i]} z_{j}$. According to Claim 32.12.5, we have that $\forall i \in\{0,1, \cdots, k\}, y_{\operatorname{pos}\left(u_{i}\right)}=1$. Thus, $\operatorname{dep}_{G, w}(s)=k \leq \sum_{i=1}^{n} y_{i}$. To upper bound $\operatorname{dep}_{G, w}(s)$, it suffices to upper bound $\sum_{i=1}^{n} y_{i}$.

Before we look at $y_{1}, \cdots, y_{n}$, we firstly focus on the properties of $z_{1}, \cdots, z_{n}$ :

Claim 32.12.6. $z_{1}, z_{2}, \cdots, z_{n}$ are $n$ i.i.d random variables drawn uniformly from $[N]$.

Proof. A key observation is that if $z_{1}, z_{2}, \cdots, z_{n}$ are given, then we can recover $w\left(v_{1}\right), w\left(v_{2}\right), \cdots, w\left(v_{n}\right)$ exactly by the following deterministic procedure:

1. Let $w(s)=z_{1}, k=0, S_{k}=\{s\}, u_{k}=s, i=2$.
2. Let $S_{k+1}=S_{k}$.
3. For $j=1 \rightarrow n$,
if $v_{j} \in \Gamma\left(u_{k}\right)$ and $v_{j} \notin S_{k}$ then let $S_{k+1} \leftarrow S_{k+1} \cup\left\{v_{j}\right\}, w\left(v_{j}\right)=z_{i}, i \leftarrow i+1$.
4. If $f_{G, w}\left(u_{k}\right) \neq u_{k}$, then let $u_{k+1}=f_{G, w}\left(u_{k}\right), k \leftarrow k+1$ and go to step 2 .
5. Otherwise, for $j=1 \rightarrow n$,
if $v_{j} \notin S_{k+1}$ then let $\operatorname{pos}\left(v_{j}\right) \leftarrow i, w\left(v_{j}\right)=z_{i}, i \leftarrow i+1$.

Notice that after step $3, \forall v \in \Gamma\left(u_{k}\right) \cup\left\{u_{k}\right\}, w(v)$ is already recovered, thus we can implement step 4 . Thus, the above procedure is a valid procedure. Since $z_{1}, \cdots, z_{n}$ are generated by $w\left(v_{1}\right), \cdots, w\left(v_{n}\right)$, we can also know $z_{1}, \cdots, z_{n}$ by given $w\left(v_{1}\right), \cdots, w\left(v_{n}\right)$. This means that $H\left(z_{1}, z_{2}, \cdots, z_{n} \mid w\left(v_{1}\right), w\left(v_{2}\right), \cdots, w\left(v_{n}\right)\right)=H\left(w\left(v_{1}\right), w\left(v_{2}\right), \cdots, w\left(v_{n}\right) \mid z_{1}, z_{2}, \cdots, z_{n}\right)=0$, where $H(\cdot)$ is the information entropy. Notice that

$$
\begin{aligned}
& I\left(z_{1}, z_{2}, \cdots, z_{n} ; w\left(v_{1}\right), w\left(v_{2}\right), \cdots, w\left(v_{n}\right)\right) \\
& =H\left(z_{1}, z_{2}, \cdots, z_{n}\right)-H\left(z_{1}, z_{2}, \cdots, z_{n} \mid w\left(v_{1}\right), w\left(v_{2}\right), \cdots, w\left(v_{n}\right)\right) \\
& =H\left(w\left(v_{1}\right), w\left(v_{2}\right), \cdots, w\left(v_{n}\right)\right)-H\left(w\left(v_{1}\right), w\left(v_{2}\right), \cdots, w\left(v_{n}\right) \mid z_{1}, z_{2}, \cdots, z_{n}\right)
\end{aligned}
$$

where $I(\cdot)$ is the mutual information. Thus, $H\left(z_{1}, z_{2}, \cdots, z_{n}\right)=H\left(w\left(v_{1}\right), w\left(v_{2}\right), \cdots, w\left(v_{n}\right)\right)=$ $n \log N$. For $i \in[n]$, since the size of the support of $z_{i}$ is at $\operatorname{most} N, H\left(z_{i}\right) \leq \log N$ where the equality holds if and only if $z_{i}$ is uniformly distributed on $[N]$. Also notice that $H\left(z_{1}, z_{2}, \cdots, z_{n}\right) \leq \sum_{i=1}^{n} H\left(z_{i}\right)$, where the equality holds if and only if $z_{i}$ are independent. Since $\sum_{i=1}^{n} H\left(z_{i}\right) \leq n \log N$, we have $H\left(z_{1}, z_{2}, \cdots, z_{n}\right)=\sum_{i=1}^{n} H\left(z_{i}\right)$, and for each $i \in[n]$, $H\left(z_{i}\right)=\log N$. Thus, $z_{1}, z_{2}, \cdots, z_{n}$ are i.i.d. random variables drawn uniformly from $[N]$.

Claim 32.12.7. If $N>n^{2} / \delta$ for some $\delta \in(0,1)$, then with probability at least $1-\delta$, $\forall i \neq j \in[n]$, we have $w\left(v_{i}\right) \neq w\left(v_{j}\right)$.

Proof. Recall that $w\left(v_{1}\right), w\left(v_{2}\right), \cdots, w\left(v_{n}\right)$ are $n$ i.i.d. random variables drawn uniformly from $N$. For any $i \neq j \in[n]$, the $\operatorname{Pr}\left(w\left(v_{i}\right) \neq w\left(v_{j}\right)\right)=1 / N$, thus $\mathbb{E}(\mid\{(i, j) \in[n] \times[n] \mid i \neq$ $\left.\left.j, w\left(v_{i}\right) \neq w\left(v_{j}\right)\right\} \mid\right) \leq n^{2} / N$. By Markov's inequality,

$$
\operatorname{Pr}\left(\left|\left\{(i, j) \in[n] \times[n] \mid i \neq j, w\left(v_{i}\right) \neq w\left(v_{j}\right)\right\}\right| \geq 1\right) \leq n^{2} / N \leq \delta
$$

Thus,

$$
\operatorname{Pr}\left(\forall i \neq j \in[n], z_{i} \neq z_{j}\right) \geq 1-\delta
$$

Claim 32.12.8. Let $\mathcal{E}$ be the event that $\forall i \neq j \in[n], w\left(v_{i}\right) \neq w\left(v_{j}\right)$. Then, for any $t \geq$ $3 \sum_{i=1}^{n} \frac{1}{i}$, we have

$$
\operatorname{Pr}_{w \sim[N]^{n}}\left(\left.\sum_{i=1}^{n} y_{i} \geq t+\sum_{i=1}^{n} \frac{1}{i} \right\rvert\, \mathcal{E}\right) \leq e^{-\frac{3}{4} t} .
$$

Proof. Note that $\mathcal{E}$ happened if and only if we have $\forall i \neq j \in[n], z_{i} \neq z_{j}$. Due to Claim 32.12.6, $z_{1}, z_{2}, \cdots, z_{n}$ are i.i.d. random variables drawn uniformly from $[N]$, then conditioned on $\mathcal{E}$, $y_{1}, y_{2}, \cdots, y_{n}$ are independent, and the probability that $y_{i}=1$ is $1 / i$. Thus, we have:

$$
\begin{aligned}
& \operatorname{Pr}\left(\left.\sum_{i=1}^{n} y_{i} \geq \sum_{i=1}^{n} \frac{1}{i}+t \right\rvert\, \mathcal{E}\right) \\
= & \operatorname{Pr}\left(\sum_{i=1}^{n}\left(y_{i}-\mathbb{E}\left(y_{i} \mid \mathcal{E}\right)\right) \geq t \mid \mathcal{E}\right) \\
\leq & \exp \left(-\frac{\frac{1}{2} t^{2}}{\sum_{i=1}^{n} \mathbb{V}\left(y_{i} \mid \mathcal{E}\right)+\frac{1}{3} t}\right) \\
\leq & \exp \left(-\frac{\frac{1}{2} t^{2}}{\sum_{i=1}^{n} \frac{1}{i}+\frac{1}{3} t}\right) \\
\leq & \exp \left(-\frac{\frac{1}{2} t^{2}}{\frac{2}{3} t}\right) \\
= & \exp \left(-\frac{3}{4} t\right)
\end{aligned}
$$

where the first equality follows by $\mathbb{E}\left(y_{i} \mid \mathcal{E}\right)=1 / i$. The first inequality follows by Berinstein inequality. The second inequality follows by

$$
\sum_{i=1}^{n} \mathbb{V}\left(y_{i} \mid \mathcal{E}\right) \leq \sum_{i=1}^{n} \mathbb{E}\left(y_{i}^{2} \mid \mathcal{E}\right)=\sum_{i=1}^{n} \mathbb{E}\left(y_{i} \mid \mathcal{E}\right)=\sum_{i=1}^{n} \frac{1}{i}
$$

The third inequality follows by $\sum_{i=1}^{n} \frac{1}{i} \leq \frac{1}{3} t$.

For a fixed vertex $s \in V$, due to Claim 32.12.8, for any $t \geq 3 \sum_{i=1}^{n} 1 / i$, we have

$$
\begin{equation*}
\operatorname{Pr}\left(\operatorname{dep}_{G, w}(s) \geq \sum_{i=1}^{n} 1 / i+t \mid \varepsilon\right) \leq e^{-\frac{3}{4} t} \tag{32.1}
\end{equation*}
$$

Thus, for any $t \geq 60 \log n$,

$$
\begin{aligned}
& \operatorname{Pr}_{w \sim[N]^{n}}\left(\exists s \in V \text { s.t. } \operatorname{dep}_{G, w}(s) \geq t\right) \\
\leq & \operatorname{Pr}\left(\exists s \in V \text { s.t. } \operatorname{dep}_{G, w}(s) \geq 5 t / 6+\sum_{i=1}^{n} 1 / i\right) \\
= & \operatorname{Pr}\left(\exists s \in V \text { s.t. } \operatorname{dep}_{G, w}(s) \geq 5 t / 6+\sum_{i=1}^{n} 1 / i \mid \mathcal{E}\right) \operatorname{Pr}(\mathcal{E}) \\
& +\operatorname{Pr}\left(\exists s \in V \text { s.t. } \operatorname{dep}_{G, w}(s) \geq 5 t / 6+\sum_{i=1}^{n} 1 / i \mid \neg \mathcal{E}\right) \operatorname{Pr}(\neg \mathcal{E}) \\
\leq & \operatorname{Pr}\left(\exists s \in V \text { s.t. } \operatorname{dep}_{G, w}(s) \geq 5 t / 6+\sum_{i=1}^{n} 1 / i \mid \mathcal{E}\right)+\operatorname{Pr}(\neg \mathcal{E}) \\
\leq & \operatorname{Pr}\left(\exists s \in V \text { s.t. } \operatorname{dep}_{G, w}(s) \geq 5 t / 6+\sum_{i=1}^{n} 1 / i \mid \mathcal{E}\right)+\delta \\
\leq & \sum_{s \in V} \operatorname{Pr}\left(\operatorname{dep}_{G, w}(s) \geq 5 t / 6+\sum_{i=1}^{n} 1 / i \mid \mathcal{E}\right)+\delta \\
\leq & n e^{-\frac{5}{8} t}+\delta \\
\leq & e^{-\frac{1}{2} t}+\delta
\end{aligned}
$$

where the first inequality follows by $\frac{1}{6} t \geq 10 \log n \geq \sum_{i=1}^{n} 1 / i$. The third inequality follows by Claim 32.12.7. The forth inequality follows by union bound. The fifth inequality follows by Equation (32.1). The sixth inequality follows by $e^{-\frac{1}{8} t} \leq \frac{1}{n}$.

Thus, we can conclude that for any $t \geq 60 \log n$, we have $\operatorname{Pr}(\operatorname{dep}(G, w) \leq t) \geq$ $1-\delta-e^{-\frac{1}{2} t}$.

Lemma 32.12.9 (The number of roots of the random min-parent-forest). Let $G=(V, E)$ be an undirected graph with $n$ vertices where $V=\left\{v_{1}, v_{2}, \cdots, v_{n}\right\}$, and the labels satisfies
$v_{1}<v_{2}<\cdots<v_{n}$. Let $w\left(v_{1}\right), w\left(v_{2}\right), \cdots, w\left(v_{n}\right)$ be $n$ i.i.d. random variables drawn uniformly from $[N]$. Let $\delta \in(0,1)$. If $N>n^{3}$, then

$$
\operatorname{Pr}_{w \sim[N]^{n}}\left(\left|\left\{v \in V \mid f_{G, w}(v)=v\right\}\right| \geq \frac{2}{\delta} \sum_{v \in V} \frac{1}{|\Gamma(v)|+1}\right) \leq \delta
$$

Proof. Let $w\left(v_{1}\right), w\left(v_{2}\right), \cdots, w\left(v_{n}\right)$ be $n$ i.i.d. random variables drawn uniformly from $[N]$. Let $\mathcal{E}$ be the event that $\forall i \neq j \in[n], w\left(v_{i}\right) \neq w\left(v_{j}\right)$. Notice that for $i \neq j$, the probability that $w\left(v_{i}\right)=w\left(v_{j}\right)$ is $1 / N$. Thus, $\mathbb{E}\left(\left|\left\{(i, j) \in[n] \times[n] \mid i \neq j, w\left(v_{i}\right)=w\left(v_{j}\right)\right\}\right|\right) \leq n^{2} / N$. Thus, if $N>n^{2}$, then $\operatorname{Pr}(\neg \mathcal{E})=\operatorname{Pr}\left(\left|\left\{(i, j) \in[n] \times[n] \mid i \neq j, w\left(v_{i}\right)=w\left(v_{j}\right)\right\}\right| \geq 1\right) \leq n^{2} / N \leq \frac{1}{n}$. Now, we fix a vertex $v \in V$,

$$
\begin{aligned}
& \operatorname{Pr}\left(f_{G, w}(v)=v\right) \\
= & \operatorname{Pr}\left(f_{G, w}(v)=v \mid \mathcal{E}\right) \operatorname{Pr}(\mathcal{E})+\operatorname{Pr}\left(f_{G, w}(v)=v \mid \neg \mathcal{E}\right) \operatorname{Pr}(\neg \mathcal{E}) \\
\leq & \operatorname{Pr}\left(f_{G, w}(v)=v \mid \mathcal{E}\right)+\operatorname{Pr}(\neg \mathcal{E}) \\
\leq & \operatorname{Pr}\left(w(v)=\min _{u \in\{v\} \cup \Gamma(v)} w(u) \mid \mathcal{E}\right)+\frac{1}{n} \\
\leq & \frac{1}{|\Gamma(v)|+1}+\frac{1}{n} \\
\leq & \frac{2}{|\Gamma(v)|+1}
\end{aligned}
$$

where the third inequality follows by the symmetry of all the variables $w(u)$ for $u \in\{v\} \cup \Gamma(v)$ so condition on all the $w$ are different, with probability $\frac{1}{1+|\Gamma(v)|}, w(v)$ is the smallest one. The last inequality follows by $|\Gamma(v)|+1 \leq|V|=n$.

Thus, $\mathbb{E}\left(\left|\left\{v \in V \mid f_{G, w}(v)=v\right\}\right|\right) \leq \sum_{v \in V} \frac{2}{|\Gamma(v)|+1}$. Let $\delta \in(0,1)$, then by Markov's inequality,

$$
\operatorname{Pr}\left(\left|\left\{v \in V \mid f_{G, w}(v)=v\right\}\right| \geq \frac{2}{\delta} \sum_{v \in V} \frac{1}{|\Gamma(v)|+1}\right) \leq \delta
$$

### 32.12.2 Leader Selection via Min Parent Forest

Given a graph, we can randomly assign each vertex a weight, thus we have a min-parent-forest, then we select those roots in the min-parent-forest as leaders, and try to contract all the vertices to the leaders. If we replace line 13 to line 17 of Algorithm 32.3 by Algorithm 32.20. We can get a new algorithm with the following guarantees.

Theorem 32.12.10. Suppose we replace line 13 to line 17 of Algorithm 32.3 by Algorithm 32.20.

Let $G=(V, E)$ be an undirected graph, $m=\Omega(n)$, and $r \leq n$ be the rounds parameter where $n$ is the number of vertices in $G$. Let $c>0$ be a sufficiently large constant. If $r \geq c \log \log _{m / n}(n)$, then with probability at least $2 / 3$, the modified $\operatorname{Connectivity~}(G, m, r)$ (Algorithm 32.3) will not return FAIL, and the total number of iterations (see Definition 32.4.6) of the modified Connectivity $(G, m, r)$ is at most $O(r \cdot(\log D+\log \log n))$, where $D=\operatorname{diam}(G)$.

Proof. Let $k_{i}$ denote the number of iterations (see Definition 32.4.1) of NeighborIncrement $\left(m, G_{i-1}\right)$. By Lemma 32.4.2, we have $k_{i} \leq O(\log D)$. Thus, $\sum_{i=1}^{r} k_{i}=O(r \cdot \log D)$.

According to Lemma 32.12.1, with probability at least $1-\frac{2}{100 r}, \operatorname{dep}\left(G_{i}^{\prime \prime}, w_{i}\right) \leq O(\log n)$. By Lemma 32.4.6, with probability at least $1-\frac{2}{100 r}$, the number of iteration of Treecontraction $\left(G_{i}^{\prime \prime}, \operatorname{par}_{i}\right.$ (see Definition 32.4.5) $r_{i}^{\prime} \leq O(\log \log n)$ By taking union bound over all $i \in[r]$, then with probability at least $1-\frac{1}{50}, \sum_{i=1}^{r} r_{i}^{\prime} \leq O(r \cdot \log \log n)$.

Due to the Property 3 of Lemma 32.4.2, $\forall i \in[r], \forall v \in V_{i}^{\prime \prime}, u \in \Gamma_{G_{i}^{\prime}}(v)$, we have $u \in V_{i}^{\prime \prime}$ which means that $u \in \Gamma_{G_{i}^{\prime \prime}}(v)$. Thus, $\left|\Gamma_{G_{i}^{\prime \prime}}(v)\right| \geq\left\lceil\left(m / n_{i-1}\right)^{1 / 2}\right\rceil-1$. Then due
to Lemma 32.12.9, we have that with probability at most $\frac{1}{8}, n_{i} \geq 16 n_{i-1}^{3 / 2} / m^{1 / 2}$. Since $m / n \geq m / n_{i} \geq 1024$, we have that with probability at most $\frac{1}{8}, n_{i} \geq n_{i-1}^{11 / 10} / m^{1 / 10}$. Let $y_{1}, y_{2}, \cdots, y_{r}$ be $r$ random variables. If $n_{i} \geq n_{i-1}^{11 / 10} / m^{1 / 10}$, then $y_{i}=1$, otherwise $y_{i}=0$. We have $\mathbb{E}\left(\sum_{i=1}^{r} y_{i}\right) \leq \frac{r}{8}$. By Markov's inequality, we have $\operatorname{Pr}\left(\sum_{i=1}^{r} y_{i} \geq \frac{r}{2}\right) \leq \frac{1}{4}$. Thus, with probability at least $\frac{3}{4}, \sum_{i=1}^{r} y_{i} \leq \frac{r}{2}$. Notice that when $y_{i}=0$, then $n_{i} \leq n_{i-1}^{11 / 10} / m^{1 / 10}$, and when $y_{i}=1$, we have $n_{i} \leq n_{i-1}$. So if there are at least $\frac{r}{2}$ number of $y_{i} \mathrm{~S}$ which are 0 , then

$$
\begin{aligned}
n_{r} & \leq \frac{\left(\frac{\left(\frac{n^{1.1} m^{0.1}}{m^{0.1}}\right.}{m^{1.1}}\right)^{\cdots}}{\cdots} \quad \text { (Apply } r / 2 \text { times) } \\
& =\frac{n^{1.1^{r / 2}}}{m^{1.1^{r / 2-1}}} \\
& =n /(m / n)^{1.1^{r / 2-1}} \\
& \leq n /(m / n)^{1.1^{r / 4}} \\
& \leq \frac{1}{2}
\end{aligned}
$$

where the last inequality follows by $r \geq \frac{4}{\log 1.1}\left(\log \log _{m / n}(2 n)\right)$. Since $n_{r}$ is an integer, when $n_{r} \leq \frac{1}{2}, n_{r}=0$. Thus, we can conclude that if $r \geq c \cdot \log \log _{m / n} n$ for a sufficiently large constant $c>0$, then with probability at least $\frac{3}{4}-\frac{1}{50} \geq \frac{2}{3}$, the modified $\operatorname{Connectivity~}(G, m, r)$ will not output FAIL.

Notice that though the theoretical guarantees of the min-parent-forest leader selection method is worse than the random leader sampling, the merit of min-parent-forest leader selection method is that it can have an "early start".

Consider the case when the total space size $m$ is $\Theta(n)$. In this case, random leader sampling will always sample a half of the vertices as the leaders until the total space $m$ is
poly $(\log n)$ larger than the number of vertices. However, min-parent-forest leader selection method can make a large progress at the beginning, it will choose the number of leaders to be about the sum of inverse degrees. Furthermore, the depth of the min-parent-forest may not always have $\log n$ depth. Thus, it is an interesting question which leader selection approach has better performance in practice.

```
Algorithm 32.3 Graph Connectivity
    procedure Connectivity \((G=(V, E), m, r) \quad \triangleright\) Theorem 32.4.9, Theorem 32.4.15
        Output: FAIL or col : \(V \rightarrow V\).
        \(n \leftarrow|V|\)
        \(\forall v \in V, h_{0}(v) \leftarrow\) null.
        \(G_{0}=\left(V_{0}, E_{0}\right)=G\), i.e. \(V_{0}=V, E_{0}=E\).
        \(n_{0}=n\).
        for \(i=1 \rightarrow r\) do
                \(\forall v \in V, h_{i}(v) \leftarrow\) null. \(\quad \triangleright h_{i}(v)\) is the vertex that \(v\) contracts to
                \(G_{i}^{\prime}=\left(V_{i}^{\prime}, E_{i}^{\prime}\right)=\operatorname{NeIGHBORInCREMENT}\left(m, G_{i-1}\right) . \quad \triangleright\) Algorithm 32.1
                Compute \(V_{i}^{\prime \prime}=\left\{v \in V_{i}^{\prime}| | \Gamma_{G_{i}^{\prime}}(v) \mid \geq\left\lceil\left(m / n_{i-1}\right)^{1 / 2}\right\rceil-1\right\}\).
                Compute \(E_{i}^{\prime \prime}=\left\{(u, v) \in E_{i-1} \mid u \in V_{i}^{\prime \prime}, v \in V_{i}^{\prime \prime}\right\}\).
                \(G_{i}^{\prime \prime}=\left(V_{i}^{\prime \prime}, E_{i}^{\prime \prime}\right) . \quad G_{i}^{\prime \prime}\) is obtained by removing all the small components of \(G_{i}\)
                Let \(\gamma_{i}=\left\lceil\left(m / n_{i-1}\right)^{1 / 2}\right\rceil, p_{i}=\min \left((30 \log (n)+100) / \gamma_{i}, 1 / 2\right)\).
                Let \(l_{i}: V_{i}^{\prime \prime} \rightarrow\{0,1\}\) be a random hash function such that \(\forall v \in V_{i}^{\prime \prime}, l_{i}(v)\) are i.i.d.
    Bernoulli random variables, and \(\operatorname{Pr}\left(l_{i}(v)=1\right)=p_{i}\).
        Let \(L_{i}=\left\{v \in V_{i}^{\prime \prime} \mid l_{i}(v)=1\right\} \cup\left\{v \in V_{i}^{\prime \prime} \mid \forall u \in \Gamma_{G_{i}^{\prime}}(v) \cup\{v\}, l_{i}(u)=0\right\} . \triangleright L_{i}\) are
    leaders
        \(\forall v \in V_{i}^{\prime \prime}\) with \(v \in L_{i}\), let \(\operatorname{par}_{i}(v)=v\).
        \(\forall v \in V_{i}^{\prime \prime}\) with \(v \notin L_{i}\), let \(\operatorname{par}_{i}(v)=\min _{u \in L_{i} \cap\left(\Gamma_{G_{i}^{\prime}}(v) \cup\{v\}\right)} u . \quad \triangleright\) Non-leader finds a
    leader.
        Let \(\left(\left(V_{i}, E_{i}\right), g_{i}^{\left(r_{i}^{\prime}\right)}\right)=\operatorname{TrEEConTRACTION}\left(G_{i}^{\prime \prime}, \operatorname{par}_{i}\right) . \quad \triangleright\) Algorithm 32.2
            \(G_{i}=\left(V_{i}, E_{i}\right)\).
            \(n_{i}=\left|V_{i}\right|\).
            For each \(v \in V_{i}^{\prime} \backslash V_{i}^{\prime \prime}\), let \(h_{i}(v) \leftarrow \min _{u \in \Gamma_{G_{i}^{\prime}}(v) \cup\{v\}} u\). \(\triangleright\) Contract small component
    to one vertex
            For each \(v \in V_{i}^{\prime \prime} \backslash V_{i}\), let \(h_{i}(v) \leftarrow g_{i}^{\left(r_{i}^{\prime}\right)}(v) . \quad \triangleright\) Contract non-leader to leader
            For each \(v \in V\), if \(h_{i-1}(v) \neq\) null, then let \(h_{i}(v)=h_{i-1}(v)\).
        end for
        If \(n_{r} \neq 0\), return FAIL.
        \(((\widehat{V}, \widehat{E}), \operatorname{col})=\operatorname{TreeContraction}\left(G, h_{r}\right) . \quad \triangleright\) Algorithm 32.2
        return col.
    end procedure
```

```
Algorithm 32.4 Local Complete Shortest Path Tree Expansion
    procedure TreeExpansion \(\left(\widetilde{T}, \operatorname{dep}_{\tilde{T}},\left\{T(u) \mid u \in V_{\widetilde{T}}\right\},\left\{\operatorname{dep}_{T(u)} \mid u \in V_{\widetilde{T}}\right\}\right)\)
    Lemma 32.5.1
        \(\triangleright \widetilde{T}=\left(V_{\widetilde{T}}, \operatorname{par}_{\widetilde{T}}\right)\) with root \(v\) and radius \(s_{1}\) is a LCSPT in graph \(G=(V, E)\).
                \(\triangleright \operatorname{dep}_{\widetilde{T}}: V_{\widetilde{T}} \rightarrow \mathbb{Z}_{\geq 0}\) records the depth of every vertex in \(\widetilde{T}\).
            \(\triangleright \forall u \in V_{\widetilde{T}}, T(u)=\left(V_{T(u)}, \operatorname{par}_{T(u)}\right)\) with root \(u\) and radius \(s_{2}\) is a LCSPT in \(G\).
                    \(\triangleright \forall u \in V_{\widetilde{T}}, \operatorname{dep}_{T(u)}: V_{T(u)} \rightarrow \mathbb{Z}_{\geq 0}\) records the depth of every vertex in \(T(u)\).
        Output: \(\widehat{T}=\left(V_{\widehat{T}}, \operatorname{par}_{\widehat{T}}\right)\), \(\operatorname{dep}_{\widehat{T}}\).
        Let \(V_{\widehat{T}}=\bigcup_{u \in \widetilde{T}} V_{T(u)}\).
        \(\forall x \in V_{\widetilde{T}}, \operatorname{par}_{\widehat{T}}(x) \leftarrow \operatorname{par}_{\widetilde{T}}(x)\).
        \(\forall x \in V_{\widetilde{T}}, h(x) \leftarrow \operatorname{dep}_{\widetilde{T}}(x)\).
        \(\forall x \in V_{\widehat{T}} \backslash V_{\widetilde{T}}, u_{x} \leftarrow \underset{u: u \in V_{\widetilde{T}}, x \in V_{T(u)}}{\arg \min } \operatorname{dep}_{\widetilde{T}}(u)+\operatorname{dep}_{T(u)}(x), \operatorname{par}_{\widehat{T}}(x) \leftarrow \operatorname{par}_{T\left(u_{x}\right)}(x)\).
                                    \(\triangleright u_{x}\) is on the shortest path from \(x\) to \(v\).
        \(\forall x \in V_{\widehat{T}} \backslash V_{\widetilde{T}}, h(x) \leftarrow \operatorname{dep}_{\widetilde{T}}\left(u_{x}\right)+\operatorname{dep}_{T\left(u_{x}\right)}(x)\).
        return \(\widehat{T}=\left(V_{\widehat{T}}, \operatorname{par}_{\widehat{T}}\right)\), and return \(h: V_{\widehat{T}}: \rightarrow \mathbb{Z}_{\geq 0}\) as \(\operatorname{dep}_{\widehat{T}}\).
    end procedure
```

```
Algorithm 32.5 Doubling Algorithm for Local Complete Shortest Path Trees
    procedure MultiRadiusLCSPT \((G=(V, E), m) \quad \triangleright\) Lemma 32.5.2
            Output: \(r,\left\{T_{i}(v) \mid i \in\{0\} \cup[r], v \in V\right\},\left\{\operatorname{dep}_{T_{i}(v)} \mid i \in\{0\} \cup[r], v \in V, T_{i}(v) \neq \operatorname{null}\right\}\)
            Initialization:
            \(\forall v \in V\), if \(\left|\{v\} \cup \Gamma_{G}(v)\right|<\left\lceil(m / n)^{1 / 4}\right\rceil\), then let \(T_{0}(v) \leftarrow\left(\{v\} \cup \Gamma_{G}(v), \operatorname{par}_{T_{0}(v)}\right)\),
            where \(\operatorname{par}_{T_{0}(v)}:\{v\} \cup \Gamma_{G}(v) \rightarrow\{v\} \cup \Gamma_{G}(v)\), and \(\forall u \in\{v\} \cup \Gamma_{G}(v), \operatorname{par}_{T_{0}(v)}(u)=v\).
            \(\forall v \in V\), if \(\left|\{v\} \cup \Gamma_{G}(v)\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil\), then let \(T_{0}(v) \leftarrow\) null.
            \(\forall v \in V\), if \(T_{0}(v) \neq\) null, let \(\operatorname{dep}_{T_{0}(v)}: V_{T_{0}(v)} \rightarrow \mathbb{Z}_{\geq 0}\) s.t. \(\operatorname{dep}_{T_{0}(v)}(v)=0, \forall u \in\)
    \(\Gamma_{G}(v), \operatorname{dep}_{T_{0}(v)}(u)=1\).
            \(r=1\).
```


## Main Loop:

```
for true do
for \(v \in V\) do \(\triangleright\) If \(T_{r}(v) \neq\) null, \(T_{r}(v)\) is a local complete shortest path tree with radius \(2^{r}\).
if \(T_{r-1}(v)\) is null then \(T_{r}(v) \leftarrow\) null.
else if \(\exists u \in V_{T_{r-1}(v)}, T_{r-1}(u)\) is null then \(T_{r}(v) \leftarrow\) null.
else
\[
\left(T_{r}(v), \operatorname{dep}_{T_{r}(v)}\right)=\operatorname{TREEEXPANSION}\left(T_{r-1}(v), \operatorname{dep}_{T_{r-1}(v)}{ }_{u \in V_{T_{r-1}(v)}}^{\cup}\left\{T_{r-1}(u)\right\}, \underset{u \in V_{T_{r-1}(v)}}{\cup}\left\{\operatorname{dep}_{T_{r-1}(u)}\right\}\right) .
\]
\(\triangleright\) Algorithm 32.4
\[
\text { If }\left|V_{T_{r}(v)}\right| \geq\left\lceil(m / n)^{1 / 4}\right\rceil \text {, let } T_{r}(v) \leftarrow \text { null. }
\]
end if
end for
if \(\forall v \in V\) either \(T_{r}(v)=\) null or \(\left|V_{T_{r}(v)}\right|=\left|V_{T_{r-1}(v)}\right|\) then
return \(r,\left\{T_{i}(v) \mid i \in\{0\} \cup[r], v \in V\right\},\left\{\operatorname{dep}_{T_{i}(v)} \mid i \in\{0\} \cup[r], v \in V, T_{i}(v) \neq\right.\) null\}
end if
\(r \leftarrow r+1\).
end for
end procedure
```

```
Algorithm 32.6 Large Local Shortest Path Trees
    procedure MultipleLargeTrees \((G=(V, E), m) \quad \triangleright\) Lemma 32.5.3, Lemma 32.5.6
        Output: \(\{\widetilde{T}(v) \mid v \in V\},\left\{\operatorname{dep}_{\widetilde{T}(v)} \mid v \in V\right\}\).
        \(\left(r,\left\{T_{i}(v) \mid i \in\{0\} \cup[r], v \in V\right\},\left\{\operatorname{dep}_{T_{i}(v)} \mid i \in\{0\} \cup[r], v \in V, T_{i}(v) \neq\right.\right.\) null \(\left.\}\right)=\operatorname{MultiRADusLCSPT}(G, m) . \quad \triangleright\)
    Algorithm 32.5
        \(\forall v \in V\) with \(T_{r}(v) \neq\) null let \(\widetilde{T}(v)=T_{r}(v)\) and \(\operatorname{dep}_{\widetilde{T}(v)} \leftarrow \operatorname{dep}_{T_{r}(v)}\).
        \(\forall v \in V\) with \(T_{r}(v)=\) null, let \(\widetilde{T}_{0}(v)=\left(\{v\}, \operatorname{par}_{\widetilde{T}_{0}(v)}\right), s_{0}(v)=0\), and \(\operatorname{dep}_{\widetilde{T}_{0}(v)}:\{v\} \rightarrow\)
    \(\mathbb{Z}_{\geq 0}\),
        where \(\operatorname{par}_{\widetilde{T}_{0}(v)}:\{v\} \rightarrow\{v\}\) satisfies \(\operatorname{par}_{\widetilde{T}_{0}(v)}(v)=v\), and \(\operatorname{dep}_{\widetilde{T}_{0}(v)}(v)=0\).
        for \(i=1 \rightarrow r\) do
            for \(v \in\left\{u \in V \mid T_{r}(u)=\right.\) null \(\}\) do
                if \(\forall u \in V_{\widetilde{T}_{i-1}(v)}, T_{r-i}(u) \neq\) null then
                    \(\left(\widetilde{T}_{i}(v), \operatorname{dep}_{\tilde{T}_{i}(v)}\right)=\operatorname{TrREEXPANSION}\left(\tilde{T}_{i-1}(v), \operatorname{dep} \tilde{T}_{i-1}(v),{ }_{u \in V_{\tilde{T}_{i-1}(v)}}^{\cup}\left\{T_{r-i}(u)\right\},{ }_{u \in V_{\tilde{T}_{i-1}(v)}}^{\cup}\left\{\operatorname{dep}_{T_{r-i}(u)}\right\}\right)\).
                                    \(\triangleright\) Algorithm 32.4
                If \(\left|V_{\widetilde{T}_{i}(v)}\right|<\left\lceil(m / n)^{1 / 4}\right\rceil\), then let \(s_{i}(v)=s_{i-1}(v)+2^{r-i}\).
                Otherwise, let \(s_{i}(v)=s_{i-1}(v), \widetilde{T}_{i}(v) \leftarrow \widetilde{T}_{i-1}(v), \operatorname{dep}_{\widetilde{T}_{i}(v)} \leftarrow \operatorname{dep}_{\widetilde{T}_{i-1}(v)}\).
                else
                        Let \(s_{i}(v)=s_{i-1}(v), \widetilde{T}_{i}(v)=\widetilde{T}_{i-1}(v), \operatorname{dep}_{\widetilde{T}_{i}(v)} \leftarrow \operatorname{dep}_{\widetilde{T}_{i-1}(v)}\).
                end if
            end for
        end for \(\triangleright \widetilde{T}_{r}(v)\) is a LCSPT with root \(v\) and the largest radius s.t.
    \(\left|V_{\widetilde{T}_{r}(v)}\right|<\left\lceil(m / n)^{1 / 4}\right\rceil\).
        \(\forall v \in V\), if \(\left|\Gamma_{G}(v) \cup\{v\}\right| \leq\left\lceil(m / n)^{1 / 4}\right\rceil\), then let \(N(v)=\Gamma_{G}(v) \cup\{v\}\).
        Otherwise arbitrarily choose \(N(v) \subseteq \Gamma_{G}(v) \cup\{v\}\) with \(|N(v)|=\left\lceil(\mathrm{m} / n)^{1 / 4}\right\rceil\).
        for \(v \in\left\{u \in V \mid T_{r}(u)=\right.\) null \(\}\) do \(\triangleright \operatorname{Expand} \widetilde{T}_{r}(v)\) a little bit to get large enough \(\widetilde{T}\).
            if \(\forall u \in V_{\widetilde{T}_{r}(v)}, T_{0}(u) \neq\) null then
```



```
    Algorithm 32.4
        else
            Select an arbitrary \(u_{v} \in V_{\widetilde{T}_{r}(v)}\) with \(T_{0}\left(u_{v}\right)=\) null.
            Let \(V_{\widetilde{T}(v)}=N\left(u_{v}\right) \cup V_{\widetilde{T}_{r}(v)}\).
            \(\forall x \in V_{\widetilde{T}_{r}(v)}\), let \(\operatorname{par}_{\widetilde{T}(v)}(x)=\operatorname{par}_{\widetilde{T}_{r}(v)}(x), \operatorname{dep}_{\widetilde{T}(v)}(x)=\operatorname{dep}_{\widetilde{T}_{r}(v)}(x)\).
            \(\forall x \in N\left(u_{v}\right), x \notin V_{\widetilde{T}_{r}(v)}\), let \(\operatorname{par}_{\widetilde{T}(v)}(x)=u_{v}, \operatorname{dep}_{\widetilde{T}(v)}(x)=\operatorname{dep}_{\widetilde{T}_{r}(v)}\left(u_{v}\right)+1\).
            Let \(\widetilde{T}(v)=\left(V_{\widetilde{T}(v)}, \operatorname{par}_{\widetilde{T}(v)}\right)\).
            end if
        end for
        return \(\{\widetilde{T}(v) \mid v \in V\},\left\{\operatorname{dep}_{\widetilde{T}(v)} \mid v e^{24} \mathcal{F}^{2}\right\}\).
    end procedure
```

```
Algorithm 32.7 Depth and Ancestors of Every Vertex
    procedure FindAncestors (par : \(V \rightarrow V\) ) \(\triangleright\) Lemma 32.5.7
    For \(v \in V\) let \(g_{0}(v)=\operatorname{part}(v)\). If \(\operatorname{part}(v)=v\), let \(h_{0}(v)=0\). Otherwise, let \(h_{0}(v)=\)
    null.
        Let \(l=0\).
        for \(\exists v \in V, h_{l}(v)=\) null do
            \(l \leftarrow l+1\).
            for \(v \in V\) do
                    Let \(g_{l}(v)=g_{l-1}\left(g_{l-1}(v)\right) . \quad \triangleright g_{l}\) is \(\operatorname{par}^{\left(2^{l}\right)}\).
                    if \(h_{l-1}(v) \neq\) null then \(h_{l}(v)=h_{l-1}(v)\).
                    else if \(h_{l-1}\left(g_{l-1}(v)\right) \neq\) null then \(h_{l}(v)=h_{l-1}\left(g_{l-1}(v)\right)+2^{l-1}\).
                else \(h_{l}(v)=\) null.
                end if
            end for
        end for
        Let \(r=l\), \(\operatorname{dep}_{\mathrm{par}} \leftarrow h_{r}\).
        return \(r, \operatorname{dep}_{\text {par }},\left\{g_{i}: V \rightarrow V \mid i \in\{0\} \cup[r]\right\} . \quad \triangleright \operatorname{dep}_{\text {par }}: V \rightarrow \mathbb{Z}_{\geq 0}\)
    end procedure
```

```
Algorithm 32.8 Path in a Tree
    procedure FindPath (par : \(V \rightarrow V, q \in V\) ) \(\triangleright\) Lemma 32.5.9
    Output: \(\operatorname{dep}_{\mathrm{par}}: V \rightarrow \mathbb{Z}_{\geq 0}, P \subseteq V, w \in V \cup\{\) null \(\}\).
    \(\left(r, \operatorname{dep}_{\mathrm{par}},\left\{g_{i} \mid i \in\{0\} \cup[r]\right\}\right)=\) FindAncestors(par) \(\quad \triangleright\) Algorithm 32.7
    Let \(S_{0}=\left\{\left(q, g_{r}(q)\right)\right\}, k=\left\lceil\log \left(\operatorname{dep}_{\mathrm{par}}(q)\right)\right\rceil . \quad \triangleright S_{0}\) contains \((q\), the root of \(q)\)
    for \(i=1 \rightarrow k\) do \(\triangleright S_{i}\) is a set of segments partitioned the path from \(q\) to the root of
    \(q\)
            Let \(S_{i} \leftarrow \emptyset\).
            for \((x, y) \in S_{i-1}\) do
                if \(\operatorname{dep}_{\mathrm{par}}(x)-\operatorname{dep}_{\mathrm{par}}(y)>2^{k-i}\) then \(S_{i} \leftarrow S_{i} \cup\left\{\left(x, g_{k-i}(x)\right),\left(g_{k-i}(x), y\right)\right\}\).
                else \(S_{i} \leftarrow S_{i} \cup\{(x, y)\}\).
                end if
            end for
        end for \(\quad \triangleright S_{k}\) only contains segments with length at most 1
        Let \(P \leftarrow\{q\}\)
        for \((x, y) \in S_{k}\) do
            Let \(P \leftarrow P \cup\{y\}\)
        end for
        Find \(w \in P\) with \(\operatorname{dep}_{\text {par }}(w)=1\). If \(w\) does not exist, let \(w \leftarrow\) null.
        return \(\left(\operatorname{dep}_{\mathrm{par}}, P, w\right)\)
    end procedure
```

```
Algorithm 32.9 Root Changing
    procedure RootChange(par : \(V \rightarrow V, q \in V\) ) \(\triangleright\) Lemma 32.5.11
        Output: \(\widehat{\text { par }: ~} V \rightarrow V\).
        \(\left(\operatorname{dep}_{\mathrm{par}}, P, w\right)=\operatorname{FindPATh}(\) par, \(q) . \quad \triangleright\) Algorithm 32.8
        \(\forall v \in V \backslash P\), let \(\widehat{\operatorname{par}}(v)=\operatorname{part}(v)\).
        Let \(\widehat{\operatorname{par}}(q)=q\).
        Let \(h:\{0\} \cup\left[\operatorname{dep}_{\mathrm{par}}(q)\right] \rightarrow P\) such that \(\forall i \in\{0\} \cup\left[\operatorname{dep}_{\mathrm{par}}(q)\right], h(i)=x\) where
    \(\operatorname{dep}_{\mathrm{par}}(x)=i\).
        for \(v \in P \backslash\{q\}\) do \(\triangleright\) Reverse par of all the vertices on the path from \(q\) to the root
    of \(q\).
            Let \(\widehat{\operatorname{par}}(v)=h\left(\operatorname{dep}_{\mathrm{par}}(v)+1\right)\).
        end for
        return \(\widehat{p a r}\).
    end procedure
```



Figure 32.1: Each tree with green edges on the top-left is a rooted tree of each contracted component. For example, there are five components $\{1,2,3\},\{4,5,6,7\},\{8,9$, $10,11,12\},\{13,14,15\},\{16,17\}$. The dashed edges in the bottom-left figure is a root spanning tree of five components. The red edges in the top-right figure correspond to the dashed edges in the bottom-left figure before contraction. In bottom-right figure, by changing (see blue edges) the root of each contracted tree, we get a rooted spanning tree in the original graph

```
Algorithm 32.10 Spanning Forest Expansion
    procedure ForestExpansion(par : \(V_{1} \rightarrow V_{1}, \widetilde{\text { par }: ~} V_{2} \rightarrow V_{2}, f: V_{1} \times V_{1} \rightarrow\{\) null \(\} \cup\)
    \(\left.\left(V_{2} \times V_{2}\right)\right)\)
                    \(\triangleright\) Lemma 32.5.12
    2: \(\quad\) Output: \(\widehat{\text { par }: ~} V_{2} \rightarrow V_{2}\).
        \(\left(\left(V_{2}^{\prime}, \emptyset\right), \widetilde{\operatorname{par}^{(\infty)}}\right)=\) TREECONTRACTION \(\left(\left(V_{2}, \emptyset\right), \widetilde{\text { par }}\right) . \quad \triangleright\) Algorithm 32.2
        for \(v \in V_{1}\) do
            Let \(V_{2}(v)=\left\{u \in V_{2} \mid \widetilde{\operatorname{par}}^{(\infty)}(u)=v\right\}\).
            Let \(\widetilde{\operatorname{par}}_{v}: V_{2}(v) \rightarrow V_{2}(v)\) such that \(\forall u \in V_{2}(v), \widetilde{\operatorname{par}}_{v}(u)=\widetilde{\operatorname{par}}(u)\).
            if \(\operatorname{part}(v) \neq v\) then
                Let \(\left(x_{v}, y_{v}\right)=f(v, \operatorname{part}(v))\).
                \(\widehat{\operatorname{par}}_{v}=\) RootChange \(\left(\widetilde{\operatorname{par}}_{v}, x_{v}\right) . \quad \triangleright\) Algorithm 32.9
                Let \(\widehat{\operatorname{par}}\left(x_{v}\right)=y_{v}\), and \(\forall u \in V_{2}(v) \backslash\left\{x_{v}\right\}, \widehat{\operatorname{par}}(u)=\widehat{\operatorname{par}}_{v}(u)\).
            else \(\forall u \in V_{2}(v)\) let \(\widehat{\operatorname{par}}(u)=\widetilde{\operatorname{par}}_{v}(u)\).
            end if
        end for
        return \(\widehat{\text { par. }}\)
    end procedure
```

```
Algorithm 32.11 Undirected Graph Spanning Forest
    procedure \(\operatorname{SpanningForest}(G=(V, E), m, r) \triangleright\) Corollary 32.5.17, Theorem 32.5.21
        Output: FAIL or \(\left\{V_{i} \subseteq V \mid i \in\{0\} \cup[r]\right\},\left\{\operatorname{par}_{i}: V_{i} \rightarrow V_{i} \mid i \in\{0\} \cup[r-1]\right\},\left\{h_{i}:\right.\)
    \(V_{i} \rightarrow V_{i+1} \cup\{\) null \(\left.\} \mid i \in\{0\} \cup[r-1]\right\}, F \subseteq E\).
        \(n_{0}=n=|V|, G_{0}=\left(V_{0}, E_{0}\right)=(V, E)\).
        Let \(g_{0}: E_{0} \rightarrow E\) be an identity map.
        Let \(n_{0}^{\prime}=n_{0}\).
        for \(i=0 \rightarrow r-1\) do
            \(D_{i} \leftarrow \emptyset\).
            \(\left(\left\{\widetilde{T}_{i}(v) \mid v \in V_{i}\right\},\left\{\operatorname{dep}_{\widetilde{T}_{i}(v)} \mid v \in V_{i}\right\}\right)=\operatorname{MultipleLargeTrees}\left(G_{i}, m\right) . \quad \triangleright\)
    Algorithm 32.6
            Let \(V_{i}^{\prime}=\left\{v \in V_{i}| | V_{\widetilde{T}_{i}(v)} \mid \geq\left\lceil\left(m / n_{i}\right)^{1 / 4}\right\rceil\right\}, E_{i}^{\prime}=\left\{(u, v) \in E_{i} \mid u, v \in V_{i}^{\prime}\right\}, G_{i}^{\prime}=\)
    \(\left(V_{i}^{\prime}, E_{i}^{\prime}\right)\).
            \(\forall v \in V_{i} \backslash V_{i}^{\prime}\), let \(h_{i}(v)=\) null, \(u_{v}=\min _{u \in V_{\widetilde{T}_{i}(v)}} u\). Let \(\operatorname{par}_{i}(v)=\operatorname{par}_{\widetilde{T}_{i}\left(u_{v}\right)}(v)\).
            \(\forall v \in V_{i} \backslash V_{i}^{\prime}\), if \(\operatorname{par}_{i}(v) \neq v\), then \(D_{i} \leftarrow D_{i} \cup\left\{g_{i}\left(\operatorname{par}_{i}(v), v\right), g_{i}\left(v, \operatorname{par}_{i}(v)\right)\right\}\).
            Let \(\gamma_{i}=\left\lceil\left(m / n_{i}\right)^{1 / 4}\right\rceil, p_{i}=\min \left((30 \log (n)+100) / \gamma_{i}, 1 / 2\right)\).
            Let \(l_{i}: V_{i}^{\prime} \rightarrow\{0,1\}\) be chosen randomly s.t. \(\forall v \in V_{i}^{\prime}, l_{i}(v)\) are i.i.d. Bernoulli
    random variables with \(\operatorname{Pr}\left(l_{i}(v)=1\right)=p_{i}\).
            Let \(L_{i}=\left\{v \in V_{i}^{\prime} \mid l_{i}(v)=1\right\} \cup\left\{v \in V_{i}^{\prime} \mid \forall u \in V_{\widetilde{T}_{i}(v)}, l_{i}(u)=0\right\}\).
            For \(v \in V_{i}^{\prime}\), let \(z_{i}(v)=\arg \min _{u \in L_{i} \cap V_{\widetilde{T}_{i}(v)}} \operatorname{dep}_{\widetilde{T}_{i}(v)}(u)\). If \(z_{i}(v)=v\), let \(\operatorname{par}_{i}(v)=v\).
            Otherwise, \(\left(\operatorname{dep}_{\widetilde{T}_{i}(v)}, P_{i}(v), w_{i}(v)\right)=\operatorname{FinDPATH}\left(\operatorname{par}_{\widetilde{T}_{i}(v)}, z_{i}(v)\right)\), and let \(\operatorname{par}_{i}(v)=\)
    \(w_{i}(v)\).
                                    \(\triangleright\) Algorithm 32.8
17: \(\quad \operatorname{Let}\left(\left(V_{i+1}, E_{i+1}\right), \operatorname{par}_{i}^{(\infty)}\right)=\operatorname{TrEeContraction}\left(G_{i}^{\prime}, \operatorname{par}_{i}: V_{i}^{\prime} \rightarrow V_{i}^{\prime}\right) . \quad \triangleright\)
    Algorithm 32.2
            \(G_{i+1}=\left(V_{i+1}, E_{i+1}\right), n_{i+1}=\left|V_{i+1}\right|\).
            \(\forall v \in V_{i}^{\prime}, h_{i}(v)=\operatorname{par}_{i}^{(\infty)}(v)\). If \(\operatorname{par}_{i}(v) \neq v\), then \(D_{i} \leftarrow D_{i} \cup\)
    \(\left\{g_{i}\left(\operatorname{par}_{i}(v), v\right), g_{i}\left(v, \operatorname{par}_{i}(v)\right)\right\}\).
            Let \(g_{i+1}: E_{i+1} \rightarrow E\) satisfy \(g_{i+1}(u, v)=\min _{(x, y) \in E_{i}, h_{i}(x)=u, h_{i}(y)=v} g_{i}(x, y)\).
            Let \(n_{i+1}^{\prime}=n_{i}^{\prime}+n_{i+1}\). If \(n_{i+1}^{\prime}>40 n\), then return FAIL.
        end for
        If \(n_{r} \neq 0\), return FAIL.
        Let \(F=\bigcup_{i \in\{0\} \cup[r-1]} D_{i}\).
        return \(\left\{V_{i} \mid i \in\{0\} \cup[r]\right\},\left\{\operatorname{par}_{i} \mid i \in\{0\} \cup[r-1]\right\},\left\{h_{i} \mid i \in\{0\} \cup[r-1]\right\}, F\).
    end procedure
```

```
Algorithm 32.12 Rooted Spanning Forest
    procedure ORIEntate \(\left(\left\{V_{i} \mid i \in\{0\} \cup[r]\right\},\left\{\operatorname{par}_{i} \mid i \in\{0\} \cup[r-1]\right\},\left\{h_{i} \mid i \in\{0\} \cup[r-\right.\right.\)
    1] \(\}, F)\)
                                    \(\triangleright\) Takes the output of Algorithm 32.11 as input.
                                    \(\triangleright\) Theorem 32.5.19
    2: \(\quad\) Output: par : \(V_{0} \rightarrow V_{0}\).
    3: \(\quad\) Let \(F_{0}=F\).
    4: \(\quad\) for \(i=0 \rightarrow r-1\) do
    5: \(\quad\) Initialize \(F_{i+1} \leftarrow \emptyset, f_{i+1}: V_{i+1} \times V_{i+1} \rightarrow\{\) null \(\}\).
    6: \(\quad \forall(u, v) \in F_{i}, h_{i}(u) \neq h_{i}(v)\), let \(F_{i+1} \leftarrow F_{i+1} \cup\left\{\left(h_{i}(u), h_{i}(v)\right)\right\}, f_{i+1}\left(h_{i}(u), h_{i}(v)\right) \leftarrow\)
    \((u, v)\).
        end for
        \(\widehat{\operatorname{par}}_{r}: \emptyset \rightarrow \emptyset\).
        for \(i=r \rightarrow 1\) do \(\quad \triangleright \widehat{\operatorname{par}}_{i}\) is the spanning forest of \(G_{i}\).
            Let \(\widetilde{V}_{i}=V_{i} \cup\left\{v \in V_{i-1} \mid h_{i-1}(v)=\right.\) null, \(\left.\operatorname{par}_{i-1}(v)=v\right\}\).
            Let \(\widetilde{\operatorname{par}}_{i}: \widetilde{V}_{i} \rightarrow \widetilde{V}_{i}\) satisfy \(\forall v \in V_{i}, \widetilde{\operatorname{par}}_{i}(v)=\widehat{\operatorname{par}}_{i}(v)\), and \(\forall v \in \widetilde{V}_{i} \backslash V_{i}, \widetilde{\operatorname{par}}_{i}(v)=v\).
            Let \(\widehat{\operatorname{par}}_{i-1}=\) FORESTEXPANSION \(\left(\widetilde{\operatorname{par}}_{i}, \operatorname{par}_{i-1}, f_{i}\right) . \quad \triangleright\) Algorithm 32.10
        end for
        Return \(\widehat{p a r}_{0}\) as par.
    end procedure
```

```
Algorithm 32.13 Lowest Common Ancestor
    procedure LCA \(\left(\right.\) par \(\left.: V \rightarrow V, Q=\left\{\left(u_{1}, v_{1}\right),\left(u_{2}, v_{2}\right), \cdots,\left(u_{q}, v_{q}\right)\right\}\right) \quad \triangleright\) Lemma 32.6.1
        Output: lca : \(Q \rightarrow(V \cup\{\) null \(\}) \times(V \cup\{\) null \(\}) \times(V \cup\{\) null \(\})\)
        \(\left(r, \operatorname{dep}_{\mathrm{par}},\left\{g_{i} \mid i \in\{0\} \cup[r]\right\}\right)=\) FindAncestors(par). \(\quad \triangleright\) Algorithm 32.7
        \(\forall(u, v) \in Q\), if \(u=v\) then let lca \((u, v)=(u\), null, null), \(Q \leftarrow Q \backslash\{(u, v)\}\).
        \(\forall(u, v) \in Q, g_{r}(u) \neq g_{r}(v)\), let lca \((u, v)=\) (null, null, null).
        Let \(Q^{\prime}=\emptyset\).
        \(\forall(u, v) \in Q, g_{r}(u)=g_{r}(v)\), if \(\operatorname{dep}_{\mathrm{par}}(u) \geq \operatorname{dep}_{\mathrm{par}}(v)\), then let \(Q^{\prime} \leftarrow Q^{\prime} \cup\{(u, v)\} ;\)
    Otherwise let \(Q^{\prime} \leftarrow Q^{\prime} \cup\{(v, u)\}\).
        Let \(h_{r}: Q^{\prime} \rightarrow Q^{\prime}\) be an identity mapping.
        for \(i=r-1 \rightarrow 0\) do \(\quad \triangleright\) Move \(u\) to the almost same depth as \(v\).
            For each \((u, v) \in Q^{\prime}\), let \((x, v)=h_{i+1}(u, v)\). If \(\operatorname{dep}_{\mathrm{par}}(x)-2^{i}>\operatorname{dep}_{\mathrm{par}}(y)\), then let
    \(h_{i}(u, v)=\left(g_{i}(x), v\right)\); Otherwise let \(h_{i}(u, v)=(x, v)\).
        end for
        For each \((u, v) \in Q^{\prime}\) with \(\operatorname{part}\left(h_{0}(u)\right)=v\), if \((u, v) \in Q\), then let lca \((u, v)=\)
    \(\left(v, h_{0}(u)\right.\), null \()\); Otherwise lca \((v, u)=\left(v\right.\), null, \(\left.h_{0}(u)\right)\).
            Let \(Q^{\prime \prime}=\emptyset\).
            For each \((u, v) \in Q^{\prime}\) with \(\operatorname{part}\left(h_{0}(u)\right) \neq v, \operatorname{dep}_{\text {par }}\left(h_{0}(u)\right)>\operatorname{dep}_{\text {par }}(v)\) let \(Q^{\prime \prime} \leftarrow\)
    \(Q^{\prime \prime} \cup\{(u, v)\}, h_{r}^{\prime}(u, v) \leftarrow\left(\operatorname{part}\left(h_{0}(u)\right), v\right)\).
    5: \(\quad\) For each \((u, v) \in Q^{\prime}\) with \(\operatorname{dep}_{\mathrm{par}}(u)=\operatorname{dep}_{\mathrm{par}}(v)\) let \(Q^{\prime \prime} \leftarrow Q^{\prime \prime} \cup\{(u, v)\}, h_{r}^{\prime}(u, v) \leftarrow\)
    \((u, v)\).
        for \(i=r-1 \rightarrow 0\) do \(\quad \triangleright\) Move \(u, v\) to the lowest common ancestor.
            For each \((u, v) \in Q^{\prime \prime}\), let \((x, y)=h_{i+1}^{\prime}(u, v)\). If \(g_{i}(x) \neq g_{i}(y)\), then let \(h_{i}^{\prime}(u, v)=\)
    \(\left(g_{i}(x), g_{i}(y)\right)\); Otherwise let \(h_{i}^{\prime}(u, v)=(x, y)\).
        end for
        For each \((u, v) \in Q^{\prime \prime}\), if \((u, v) \in Q\), then let lca \((u, v)=\left(\operatorname{part}\left(h_{0}^{\prime}(u)\right), h_{0}^{\prime}(u), h_{0}^{\prime}(v)\right)\);
    Otherwise lca \((v, u)=\left(\operatorname{part}\left(h_{0}^{\prime}(v)\right), h_{0}^{\prime}(v), h_{0}^{\prime}(u)\right)\).
        return lca.
    end procedure
```

```
Algorithm 32.14 Multiple Paths
    procedure MultiPath(par : \(\left.V \rightarrow V, Q=\left\{\left(u_{1}, v_{1}\right),\left(u_{2}, v_{2}\right), \cdots,\left(u_{q}, v_{q}\right)\right\}\right) \quad \triangleright\)
    Lemma 32.6.2
        Output: \(\operatorname{dep}_{\text {par }}: V \rightarrow \mathbb{Z}_{\geq 0},\left\{P_{i} \subseteq V \mid i \in[q]\right\}\).
        \(\left(r, \operatorname{dep}_{\mathrm{par}},\left\{g_{i} \mid i \in\{0\} \cup[r]\right\}\right)=\) FindAncestors(par). \(\quad \triangleright\) Algorithm 32.7
        \(\forall j \in[q]\), let \(S_{j}^{(0)}=\left\{\left(u_{j}, v_{j}\right) \mid\left(u_{j}, v_{j}\right) \in Q\right\}\).
        for \(i=1 \rightarrow r\) do
            for \(j=1 \rightarrow q\) do \(\triangleright S_{j}^{(i)}\) is a set of segments partitioned the path from \(u_{j}\) to \(v_{j}\).
                Let \(S_{j}^{(i)} \leftarrow \emptyset\).
                for \((x, y) \in S_{j}^{(i-1)}\) do
                if \(\operatorname{dep}_{\mathrm{par}}(x)-\operatorname{dep}_{\mathrm{par}}(y)>2^{r-i}\) then \(S_{j}^{(i)} \leftarrow S_{j}^{(i)} \cup\left\{\left(x, g_{r-i}(x)\right),\left(g_{r-i}(x), y\right)\right\}\).
                    else \(S_{j}^{(i)} \leftarrow S_{j}^{(i)} \cup\{(x, y)\}\).
                    end if
            end for
            end for
        end for
                            \(\triangleright S_{j}^{(r)}\) only contains segments with length 1
            Let \(\forall j \in[q], P_{j} \leftarrow\left\{u_{j}\right\}\).
            for \(j=1 \rightarrow q\) do
                for \((x, y) \in S_{j}^{(r)}\) do
                    Let \(P_{j} \leftarrow P_{j} \cup\{y\}\).
            end for
        end for
    end procedure
```

```
Algorithm 32.15 Leaf Sampling
    procedure LeafSAmpling (par : \(V \rightarrow V, m, \delta)\)
        \(\triangleright\) Lemma 32.6.5
        Output: \(A=\left(a_{1}, a_{2}, \cdots, a_{s}\right)\).
        Let \(t=\left\lceil m^{1 / 3}\right\rceil\).
        Compute \(L=\) leaves(par).
        Compute rank : \(V \rightarrow \mathbb{Z}_{\geq 0}\) such that \(\forall v \in V, \operatorname{rank}(v)=\operatorname{rank}_{\operatorname{par}}(v) . \triangleright\) Definition 32.6.4
        If \(|V| \leq m\), let \(\left\{a_{1}, a_{2}, \cdots, a_{s}\right\}=L\), and return \(A=\left(a_{1}, a_{2}, \cdots, a_{s}\right)\) which satisfies
    \(a_{1}<_{\text {par }} a_{2}<_{\text {par }} \cdots<_{\text {par }} a_{s} . \quad \triangleright<_{\text {par }}\) follows Definition 32.6.8
        If \(|L| \leq 8 t\), let \(S=L\).
        Let \(p=\min (1,640(1+\log (m) / \delta) t /|L|)\).
        If \(|L|>t\), sample each \(v \in L\) with probability \(p\) independently. let \(S\) be the set of
    samples.
        Compute \(\operatorname{par}^{\prime}: V \rightarrow V\) such that \(\forall v \in V\), if \(\operatorname{child}_{\text {par }}(v) \neq \emptyset\), then \(\operatorname{par}^{\prime}(v)=\)
    \(\operatorname{child}_{\text {par }}(v, 1)\); Otherwise let \(\operatorname{par}^{\prime}(v)=v . \quad \triangleright \operatorname{par}^{\prime}(v)\) points to \(v\) 's first child in par.
        \(\left(r^{\prime}, \operatorname{dep}_{\mathrm{par}^{\prime}}: V \rightarrow \mathbb{Z}_{\geq 0},\left\{g_{i}^{\prime}: V \rightarrow V \mid i \in\{0\} \cup\left[r^{\prime}\right]\right\}\right)=\) FindAncestors \(\left(\operatorname{par}^{\prime}\right)\).
        Find \(w \in V\) with \(\operatorname{part}(w)=w\). \(\triangleright\) Find the root.
        Let \(a_{1}=g_{r^{\prime}}^{\prime}(w), S \leftarrow S \cup\left\{a_{1}\right\}\). \(\quad\) Find the first leaf.
        Let \(Q=\{(u, v) \mid(u, v) \in S \times S, u \neq v\}\).
        Let lca \(=\mathrm{LCA}(\) par,\(Q) . \quad \triangleright\) Algorithm 32.13
        Let \(s=|S|\).
        \(\left(r, \operatorname{dep}_{\mathrm{par}}: V \rightarrow \mathbb{Z}_{\geq 0},\left\{g_{i}: V \rightarrow V \mid i \in\{0\} \cup[r]\right\}\right)=\) FindAncestors \((\) par \()\).
        for \(i=2 \rightarrow s\) do \(\quad \triangleright\) Determine the order of sampled leaves.
            For all \(x, y \in S \backslash\left\{a_{1}, a_{2}, \cdots, a_{i-1}\right\}\), let \(\left(p_{x, y}, p_{x y, x}, p_{x y, y}\right)=\operatorname{lca}(x, y)\).
            Find \(x^{*} \in S \backslash\left\{a_{1}, a_{2}, \cdots, a_{i-1}\right\}\) s.t. \(\forall y \in S \backslash\left\{a_{1}, a_{2}, \cdots, a_{i-1}, x^{*}\right\}, \operatorname{rank}\left(p_{x^{*} y, x^{*}}\right)<\)
    \(\operatorname{rank}\left(p_{x^{*} y, y}\right)\).
            Let \(a_{i}=x^{*}\).
        end for
        return \(A=\left(a_{1}, a_{2}, \cdots, a_{s}\right)\).
    end procedure
```

```
Algorithm 32.16 DFS subsequence
    procedure SubDFS(par : \(V \rightarrow V, m, \delta) \quad \triangleright\) Lemma 32.6.9, Lemma 32.6.10
        Output: \(V^{\prime} \subseteq V, A=\left(a_{1}, a_{2}, \cdots, a_{s}\right)\).
        If \(V=\{v\}\), return \(V^{\prime}=V, A=(v)\).
        Let \(v\) be the root in par, i.e. \(\operatorname{part}(v)=v\).
        \(L=\left(l_{1}, l_{2}, \cdots, l_{t}\right)=\operatorname{LEAFSAMPLING}(\) par \(, m, \delta) . \quad \triangleright\) Algorithm 32.15
        \(Q=\left\{\left(l_{i}, l_{i+1}\right) \mid i \in[t-1]\right\}\).
        lca \(=\operatorname{LCA}(\) par,\(Q) . \quad \triangleright\) Algorithm 32.13
        \(\forall i \in[t-1],\left(p_{l_{i}, l_{i+1}}, p_{i, l_{i}}, p_{i, l_{i+1}}\right)=\operatorname{lca}\left(l_{i}, l_{i+1}\right)\).
        \(Q^{\prime}=\left\{\left(l_{1}, v\right),\left(\operatorname{part}\left(l_{1}\right), p_{l_{1}, l_{2}}\right),\left(l_{2}, p_{1, l_{2}}\right),\left(\operatorname{part}\left(l_{2}\right), p_{l_{2}, l_{3}}\right),\left(l_{3}, p_{2, l_{3}}\right), \cdots,\left(l_{t}, p_{t-1, l_{t}}\right),\left(\operatorname{part}\left(l_{t}\right), v\right)\right\}\).
        \(\left(\operatorname{dep}_{\mathrm{par}},\left\{P_{i} \mid i \in[2 t]\right\}\right)=\operatorname{MultiPath}\left(\right.\) par, \(\left.Q^{\prime}\right) . \quad \triangleright\) Algorithm 32.14
        \(V^{\prime}=\bigcup_{i=1}^{2 t} P_{i}\).
        Let \(\operatorname{par}^{\prime}: V^{\prime} \rightarrow V^{\prime}\) satisfy \(\forall v \in V^{\prime}, \operatorname{par}^{\prime}(v)=\operatorname{part}(v)\).
        for \(i \in\{1,3,5, \cdots, 2 t-1\}\) do
            Compute \(A_{i}^{\prime}=\left(u_{1}, u_{2}, \cdots, u_{\left|P_{i}\right|}\right)\) such that \(\left\{u_{1}, u_{2}, \cdots, u_{\left|P_{i}\right|}\right\}=P_{i}\) and
    \(\operatorname{dep}_{\text {par }}\left(u_{1}\right)<\operatorname{dep}_{\mathrm{par}}\left(u_{2}\right)<\cdots<\operatorname{dep}_{\mathrm{par}}\left(u_{\left|P_{i}\right|}\right)\)
            end for
            for \(i \in\{2,4,6, \cdots, 2 t\}\) do
                Compute \(A_{i}^{\prime}=\left(u_{1}, u_{2}, \cdots, u_{\left|P_{i}\right|}\right)\) such that \(\left\{u_{1}, u_{2}, \cdots, u_{\left|P_{i}\right|}\right\}=P_{i}\) and
    \(\operatorname{dep}_{\text {par }}\left(u_{1}\right)>\operatorname{dep}_{\text {par }}\left(u_{2}\right)>\cdots>\operatorname{dep}_{\text {par }}\left(u_{\left|P_{i}\right|}\right)\)
            end for
            Let \(A^{\prime}=A_{1}^{\prime} A_{2}^{\prime} \cdots A_{2 t}^{\prime}\). \(\quad A^{\prime}\) is the concatenation of \(A_{1}^{\prime}, A_{2}^{\prime}, \cdots, A_{2 t}^{\prime}\).
            \(\forall u \in V^{\prime}\), compute \(\operatorname{rank}_{\mathrm{par}}(u)\) and \(\operatorname{rank}_{\mathrm{par}^{\prime}}(u)\).
            \(\forall u \in V^{\prime}, i \in\left[\left|\operatorname{child}_{\mathrm{par}^{\prime}}\right|+1\right]\) compute \(\operatorname{pos}(u, i)=j\) such that the \(j^{\text {th }}\) element in \(A^{\prime}\) is
    the \(i^{\text {th }}\) time that \(u\) appears.
            Let \(b\) be the length of \(A^{\prime}\).
            Initialize \(c:[b] \rightarrow \mathbb{Z}_{\geq 0}\). \(\triangleright c\) determine the number of copies needed for each element
    in \(A^{\prime}\)
            for \(u \in V^{\prime} \backslash\{v\}\) do
            If \(u \in\) leaves \(\left(\operatorname{par}^{\prime}\right)\), let \(c(\operatorname{pos}(u, 1))=1 . \quad \triangleright\) A leaf should only have one copy.
            If \(\operatorname{rank}_{\operatorname{par}^{\prime}}(u)=1\), let \(c\left(\operatorname{pos}\left(\operatorname{par}^{\prime}(u), 1\right)\right)=\operatorname{rank}_{\operatorname{par}}(u)\).
            If \(\operatorname{rank}_{\operatorname{par}^{\prime}}(u)=\left|\operatorname{child}_{\operatorname{par}^{\prime}}\left(\operatorname{par}^{\prime}(u)\right)\right|\), let \(c\left(\operatorname{pos}\left(\operatorname{par}^{\prime}(u), \operatorname{rank}_{\operatorname{par}^{\prime}}(u)+1\right)\right)=\)
    \(\left|\operatorname{child}_{\mathrm{par}}(\operatorname{part}(u))\right|+1-\operatorname{rank}_{\mathrm{par}}(u)\).
            If \(1 \leq \operatorname{rank}_{\operatorname{par}^{\prime}}(u)<\left|\operatorname{child}_{\mathrm{par}^{\prime}}\left(\operatorname{par}^{\prime}(u)\right)\right|\), let \(c\left(\operatorname{pos}\left(\operatorname{par}^{\prime}(u), \operatorname{rank}_{\operatorname{par}^{\prime}}(u)+1\right)\right)=\)
    \(\operatorname{rank}_{\mathrm{par}}\left(\operatorname{child}_{\mathrm{par}^{\prime}}\left(\operatorname{par}^{\prime}(u), \operatorname{rank}_{\operatorname{par}^{\prime}}(u)+1\right)\right)-\operatorname{rank}_{\mathrm{par}}(u)\).
            end for
            For each \(j \in[b]\), duplicate the \(j^{\text {th }}\) element of \(A^{\prime} c(j)\) times. Let \(A\) be the obtained
    sequence.
        return \(V^{\prime}, A\).
    end procedure


Figure 32.2: Given a tree that has 42 vertices (top-left), we label all the vertices from 1 to 42 . Firstly, we sample some leaves (red vertices, i.e. \(\{5,13,24,30,32,34,36,37,40,42\}\) ) in the tree (top-right tree). Then we find a DFS sequence of the tree (the tree formed by all the blue and red vertices in the bottom-left tree) which only contains all the sampled leaves and their ancestors. Finally, we recursively find the DFS sequences of remaining subtrees(bottom-right).
```

Algorithm 32.17 DFS sequence
procedure DFS(par : $V \rightarrow V, m$ )
$\triangleright$ Theorem 32.6.12, Theorem 32.6.14
Output: FAIL or $A=\left(a_{1}, a_{2}, \cdots, a_{2|V|-1}\right)$.
$n=|V|, \delta=1 / \log _{m} n$.
Let $\operatorname{par}_{0}=$ par.
$\left(V_{0}, A_{0}\right)=\operatorname{SuBDFS}\left(\operatorname{par}_{0}, m, \delta\right) . \quad \triangleright$ Algorithm 32.16
Let $r=\lceil 3 / \delta\rceil+2$.
for $i=1 \rightarrow r$ do $\quad \triangleright v \in V_{i} \Leftrightarrow v$ appears in $A_{i}$
$\triangleright$ If $v \in V_{i}$, then $v$ appears $\left|\operatorname{child}_{\text {par }}(v)\right|+1$ times in $A_{i}$
Let $V_{i}^{\prime}=V \backslash V_{i-1}$.
Initialize $\operatorname{par}_{i}: V_{i}^{\prime} \rightarrow V_{i}^{\prime}$.
For $v \in V_{i}^{\prime}$, if $\operatorname{part}(v) \in V_{i-1}$, let $\operatorname{par}_{i}(v)=v$; Otherwise, let $\operatorname{par}_{i}(v)=\operatorname{part}(v)$.
$\left(\left(V_{i}^{\prime \prime}, \emptyset\right), \operatorname{par}_{i}^{(\infty)}\right)=\operatorname{TrEECONTRACTION}\left(\left(V_{i}^{\prime}, \emptyset\right), \operatorname{par}_{i}\right) . \quad \triangleright$ Algorithm 32.2
$V_{i} \leftarrow V_{i-1}$.
$A_{i} \leftarrow A_{i-1}$.
for $v \in V_{i}^{\prime}, \operatorname{par}_{i}(v)=v$ do $\quad \triangleright$ The DFS sequence of the subtree of $v$ in par is
missing.
Let $V_{i}^{\prime}(v)=\left\{u \in V_{i}^{\prime} \mid \operatorname{par}_{i}^{(\infty)}(u)=v\right\}$.
Let $\operatorname{par}_{i, v}: V_{i}^{\prime}(v) \rightarrow V_{i}^{\prime}(v)$ satisfy $\forall u \in V_{i}^{\prime}(v), \operatorname{par}_{i, v}(u)=\operatorname{par}_{i}(u)$.
Let $\left(V_{i, v}, A_{i, v}\right)=\operatorname{SuBDFS}\left(\operatorname{par}_{i, v}, m, \delta\right) . \quad \triangleright$ Algorithm 32.16
$V_{i} \leftarrow V_{i} \cup V_{i, v}$.
Insert $A_{i, v}$ after the $\operatorname{rank}_{\mathrm{par}}(v)^{\mathrm{th}}$ time appearance of $v$ in $A_{i}$.
end for
end for
If $V_{r}=V$, return $A_{r}$ as $A$. Otherwise, return FAIL.
end procedure

```
```

Algorithm 32.18 A Sparser Table for RMQ
procedure SparseTable ${ }^{+}\left(a_{1}, a_{2}, \cdots, a_{n}, \delta\right)$
$\triangleright$ Lemma 32.6.16
2 :
Initially, for all $i \in[n]$ let $\widehat{f}_{i, 0}=i . \forall i>n, j \in \mathbb{Z}$, let $\widehat{f}_{i, j}=0$, and let $a_{0}=\infty$. Let
$m=\left\lceil n^{\delta}\right\rceil$.
For $t \in[\lceil 1 / \delta\rceil]$, let $S_{t}=\left\{x \mid \exists y \in[m-1], x=y \cdot m^{t}\right\}$.
for $l=1 \rightarrow\lceil 1 / \delta\rceil$ do
for $j=0 \rightarrow\lceil n / m\rceil$ do
$i \leftarrow j \cdot m+1$.
$z_{j, l}^{*} \leftarrow \arg \min _{z: t \in[l-1], x \in S_{t}, z=\widehat{f_{j} \cdot m+1+x, t}} a_{z}$.
for $i^{\prime}=0 \rightarrow \min (m-1, n-i)$ do
$T \leftarrow\left\{x \in \mathbb{Z} \mid i+i^{\prime} \leq x \leq i+m-1\right\} \cup\left\{x \in \mathbb{Z} \mid i+m^{l} \leq x \leq\right.$
$\left.i+m^{l}+i^{\prime}-1\right\} \cup\left\{z_{j, l}^{*}\right\}$
$\widehat{f}_{i+i^{\prime}, l}=\arg \min _{z \in T} a_{z}$.
end for
end for
$l \leftarrow l+1$.
end for
return $\widehat{f}_{i, j}$ for $i \in[n], j \in\{0\} \cup[[1 / \delta\rceil]$.
end procedure

```
```

Algorithm 32.19 A Sparse Table for RMQ
procedure $\operatorname{SPARSETABLE}\left(a_{1}, a_{2}, \cdots, a_{n}, \delta\right) \quad \triangleright$ Lemma 32.6.17
$\triangleright$ Output: $f_{i, j}$ for $i \in[n], j \in\{0\} \cup[\lceil\log n\rceil]$.
Initially, for all $i \in[n]$ let $f_{i, 0}=i . \forall i>n, j \in \mathbb{Z}$, let $f_{i, j}=0$, and let $a_{0}=\infty$. Let
$m=\left\lceil n^{\delta}\right\rceil$.
Let $\left\{\widehat{f}_{p, q} \mid p \in[n], q \in\{0\} \cup\lceil 1 / \delta\rceil\right\}=\operatorname{SpARSETABLE}^{+}\left(a_{1}, a_{2}, \cdots, a_{n}, \delta\right) . \quad \triangleright$
Algorithm 32.18
Let all undefined $\widehat{f}_{p, q}$ be 0 .
for $t \in[\lceil\log n\rceil]$ do
if $2^{t} \leq m$ then
$k_{t} \leftarrow-1$
$S_{t} \leftarrow \emptyset$
else
$k_{t} \leftarrow\left\lfloor\log _{m}\left(2^{t}-m\right)\right\rfloor$
$S_{t} \leftarrow\left\{x \mid x \in\left[2^{t}-m-m^{k_{t}}+1\right]\right.$ s.t. $x \equiv 1\left(\bmod m^{k_{t}}\right)$ or $\left(2^{t}-m-x\right) \equiv-1$
$\left.\left(\bmod m^{k_{t}}\right)\right\}$
end if
end for
for $j=0 \rightarrow\lceil n / m\rceil$ do
for $t=0 \rightarrow\lceil\log n\rceil$ do
$i \leftarrow j \cdot m+1$.
$z_{j, t}^{*} \leftarrow \arg \min _{z: x \in S_{t}, z=\widehat{f_{j} \cdot m+m+x, k_{t}}} a_{z}$.
for $i^{\prime}=0 \rightarrow \min (m-1, n-i)$ do
$T_{1} \leftarrow\left\{x \in \mathbb{Z} \mid i+i^{\prime} \leq x \leq \min \left(i+m-1, i+i^{\prime}+2^{t}-1\right)\right\}$
$T_{2} \leftarrow\left\{x \in \mathbb{Z} \mid \max \left(i+2^{t}, i+i^{\prime}\right) \leq x \leq i+2^{t}+i^{\prime}-1\right\}$
$T \leftarrow T_{1} \cup T_{2} \cup\left\{z_{j, t}^{*}\right\}$
$f_{i+i^{\prime}, t}=\arg \min _{z \in T} a_{z}$.
end for
end for
end for
return $f_{i, j}$ for $i \in[n], j \in\{0\} \cup[[\log n\rceil]$.
end procedure

```


Figure 32.3: A hard example for [RMCS13]. For each \(i \in\{2,3, \cdots, n / D-1\}\) and \(j \in\) \(\{1,2, \cdots, D-1\}\), node \((i-1) \cdot D+j\) has degree 4 . For node \(D\) and \(n\), they have degree 2 . Node 0 has degree \(D\). All the other nodes have degree 3 .
```

Algorithm 32.20 Leader Selection via Min Parent Forest
1: Let $N=100 \mathrm{rn}^{10}$.
2: $\forall v \in V_{i}^{\prime}$, let $w_{i}(v)$ be i.i.d. random variables drawn uniformly from [ $N$ ].
3: $\forall v \in V_{i}^{\prime \prime}$, let $\operatorname{par}_{i}(v)=f_{G_{i}^{\prime}, w_{i}}(v) . \quad \triangleright\left(V_{i}^{\prime}, f_{G_{i}^{\prime}, w_{i}}\right)$ is a min-parent-forest of $G_{i}^{\prime}$ with $w_{i}$.

```


Figure 32.4: An example where \(\#\) roots \(\approx \sum_{i=1}^{20} 1 /\left(d\left(v_{i}\right)+1\right)\). For each node, it has two numbers, the first number is the ID, and the second number is weight. \(\sum_{i=1}^{20} 1 /\left(d\left(v_{i}\right)+1\right)=\) \(1 / 4+1 / 3+1 / 3+1 / 6+1 / 5+1 / 5+1 / 5+1 / 5+1 / 5+1 / 6+1 / 4+1 / 6+1 / 8+1 / 7+1 / 6\) \(+1 / 9+1 / 8+1 / 7+1 / 6+1 / 4 \approx 3.89\) and \#roots \(=4\).

\begin{abstract}
Appendices
\end{abstract}

\section*{Appendix A}

\section*{Probability and Inequality Tools}

\section*{A. 1 Scalar version probability tools}

Lemma A.1.1 (Chernoff bounds [Che52]). Let \(x_{1}, \cdots, x_{n}\) be independent random variables. Assume that \(x_{i} \in[0,1]\) always, for each \(i \in[n]\). Let \(x=\sum_{i=1}^{n}\) and \(\mu=\mathbb{E}[x]=\sum_{i=1}^{n} \mathbb{E}\left[x_{i}\right]\). Then for any \(\epsilon>0\)
\[
\operatorname{Pr}[x \geq(1+\epsilon) \mu] \leq \exp \left(-\frac{\epsilon^{2}}{2+\epsilon} \mu\right), \quad \text { and } \quad \operatorname{Pr}[x \leq(1-\epsilon) \mu] \leq \exp \left(-\frac{\epsilon^{2}}{2} \mu\right)
\]

Lemma A.1.2 (Bernstein inequality [Ber24]). Let \(x_{1}, \cdots, x_{n}\) be independent zero-mean variables. Suppose \(\left|x_{i}\right| \leq M\) almost surely, for all \(i \in[n]\). Then for all positive \(t\),
\[
\operatorname{Pr}\left[\sum_{i=1}^{n} x_{i}>t\right] \leq \exp \left(-\frac{t^{2} / 2}{\sum_{i=1}^{n} \mathbb{E}\left[x_{i}^{2}\right]+M t / 3}\right) .
\]

Lemma A.1.3 (Hanson-Wright inequality [HW71, Wri73]). Let \(x=\left(x_{1}, \cdots, x_{n}\right) \in \mathbb{R}^{n}\) be a random vector with independent components \(x_{i}\) which satisfy \(\mathbb{E}\left[x_{i}\right]=0\) and \(\left\|x_{i}\right\|_{\psi_{2}} \leq K\) \(\left(\right.\) where \(\left.\left\|x_{i}\right\|_{\psi_{2}}:=\sup _{p \geq 1} p^{-1 / 2}\left(\mathbb{E}\left[\left|x_{i}\right|^{p}\right]\right)^{1 / p}\right)\). Let \(A\) be an \(n \times n\) matrix. Then, for every \(t \geq 0\),
\[
\operatorname{Pr}\left[\left|x^{\top} A x-\mathbb{E}\left[x^{\top} A x\right]\right|>t\right] \leq 2 \exp \left(-c \cdot \min \left(\frac{t^{2}}{K^{4}\|A\|_{F}^{2}}, \frac{t}{K^{2}\|A\|}\right)\right) .
\]

Lemma A.1.4 (Khintchine's inequality [Haa81]). Let \(\sigma_{1}, \cdots, \sigma_{n}\) be i.i.d. sign random variables, and let \(z_{1}, \cdots, z_{n}\) be real numbers. Then there are constants \(C, C^{\prime}>0\) so that
\[
\operatorname{Pr}\left[\left|\sum_{i=1}^{n} z_{i} \sigma_{i}\right| \geq C t\|z\|_{2}\right] \leq \exp \left(-C^{\prime} t^{2}\right)
\]

Lemma A.1.5 (Hoeffding inequality [Hoe63]). Let \(x_{i} \in\left[a_{i}, b_{i}\right]\) denote independent random variables, let \(x=\frac{1}{n} \sum_{i=1}^{n} x_{i}\) then
\[
\operatorname{Pr}[|x-\mathbb{E}[x]| \geq t] \leq 2 \exp \left(-\frac{2 n^{2} t^{2}}{\sum_{i=1}^{n}\left(b_{i}-a_{i}\right)^{2}}\right)
\]

Lemma A.1.6 (Lemma 1 on page 1325 of [LM00]). Let \(x \sim X_{k}^{2}\) be a chi-squared distributed random variable with \(k\) degrees of freedom. Each one has zero mean and \(\sigma^{2}\) variance. Then
\[
\begin{aligned}
& \operatorname{Pr}\left[x-k \sigma^{2} \geq(2 \sqrt{k t}+2 t) \sigma^{2}\right] \leq \exp (-t) \\
& \operatorname{Pr}\left[k \sigma^{2}-x \geq 2 \sqrt{k t} \sigma^{2}\right] \leq \exp (-t)
\end{aligned}
\]

\section*{A.1.1 Reverse Chernoff bound}

In this Section, we prove that the classical Chernoff bound is tight in some regimes. There are several different proofs, e.g. [Mou10, Slu77, AB09, YZD12]. For completeness, we provide a proof from [YZD12].

Fact A.1.7. If \(1 \leq l \leq k-1\), then
\[
\binom{k}{l} \geq \frac{1}{e \sqrt{2 \pi l}}\left(\frac{k}{l}\right)^{l}\left(\frac{k}{k-l}\right)^{k-l} .
\]

Proof. By Stirling's approximation, \(i!=\sqrt{2 \pi i}(i / e)^{i} e^{\lambda}\) for some \(\lambda \in[1 /(12 i+1), 1 /(12 i)]\).

Thus,
\[
\begin{aligned}
\binom{k}{l} & =\frac{k!}{l!(k-l)!} \\
& \geq \frac{\sqrt{2 \pi k}\left(\frac{k}{e}\right)^{k}}{\sqrt{2 \pi l}\left(\frac{l}{e}\right)^{l} \sqrt{2 \pi(k-l)}\left(\frac{k-l}{e}\right)^{k-l}} \exp \left(\frac{1}{12 k+1}-\frac{1}{12 l}-\frac{1}{12(k-l)}\right) \\
& \geq \frac{\sqrt{2 \pi k}\left(\frac{k}{e}\right)^{k}}{\sqrt{2 \pi l}\left(\frac{l}{e}\right)^{l} \sqrt{2 \pi(k-l)}\left(\frac{k-l}{e}\right)^{k-l}} e^{-1} \\
& \geq \frac{1}{\sqrt{2 \pi l}}\left(\frac{k}{l}\right)^{l}\left(\frac{k}{k-l}\right)^{k-l} e^{-1} .
\end{aligned}
\]
where the first step follows by definition, the second step follows by Stirling's approximation, the third step follows by \(\frac{1}{1+a+b}+1 \geq \frac{1}{a}+\frac{1}{b}\) for \(a \geq 12, b \geq 12\).

Now, we are ready to prove the following result

Lemma A.1.8 ([Mou10, Slu77, AB09, YZD12]). Let \(X\) be the average of \(k\) independent, Bernoulli random variables with mean \(p\). For any \(\epsilon \in(0,1 / 2]\) and \(p \in(0,1 / 2]\), assuming \(\epsilon^{2} p k \geq 3\), we have
\[
\operatorname{Pr}[X \leq(1-\epsilon) p] \geq \exp \left(-9 \epsilon^{2} p k\right)
\]
and
\[
\operatorname{Pr}[X \geq(1+\epsilon) p] \geq \exp \left(-9 \epsilon^{2} p k\right)
\]

Proof. Note that \(\operatorname{Pr}[X \leq(1-\epsilon) p]\) equals the sum \(\sum_{i=0}^{\lfloor(1-\epsilon) p k\rfloor} \operatorname{Pr}[X=i / k]\), and \(\operatorname{Pr}[X=i / k]=\) \(\binom{k}{i} p^{i}(1-p)^{k-i}\).

Fix \(l=\lfloor(1-2 \epsilon) p k\rfloor+1\). The terms in the sum are increasing, so the terms with index \(i \geq \ell\) each have value at least \(\operatorname{Pr}[X=l / k]\), so their sum has total value at least \((\epsilon p k-2) \operatorname{Pr}[X=l / k]\). To complete the proof, we show that
\[
(\epsilon p k-2) \operatorname{Pr}[X=l / k] \geq \exp \left(-9 \epsilon^{2} p k\right)
\]

The assumptions \(\epsilon^{2} p k \geq 3\) and \(\epsilon \leq 1 / 2\) give \(\epsilon p k \geq 6\), so we have
\[
\begin{aligned}
(\epsilon p k-2) \operatorname{Pr}[X=l / k] & \geq \frac{2}{3} \epsilon p k\binom{k}{l} p^{l}(1-p)^{k-l} \\
& \geq \underbrace{\frac{2}{3} \epsilon p k \frac{1}{\sqrt{2 \pi l}}}_{A} \cdot \underbrace{\left(\frac{k}{l}\right)^{l}\left(\frac{k}{k-l}\right)^{k-l} p^{l}(1-p)^{k-l}}_{B}
\end{aligned}
\]

Below, we will show that \(A \geq \exp \left(-\epsilon^{2} p k\right)\) and \(B \geq \exp \left(-8 \epsilon^{2} p k\right)\).
Claim A.1.9. \(A \geq \exp \left(-\epsilon^{2} p k\right)\).

Proof. The assumptions \(\epsilon^{2} p k \geq 3\) and \(\epsilon \leq 1 / 2\) imply that \(p k \geq 12\).
By \(l \leq p k+1\) (from definition), and \(p k \geq 12\), thus \(l \leq 1.1 p k\).
Therefore, we have
\[
\begin{array}{rlr}
A & \geq \frac{2}{3 e} \epsilon \sqrt{p k /(2.2 \pi)} & \\
& \geq \frac{2}{3 e} \sqrt{3 /(2.2 \pi)} & \text { by } \epsilon \sqrt{p k} \geq \sqrt{3} \\
& \geq 0.1 & \\
& \geq \exp (-3) & \\
& \geq \exp \left(-\epsilon^{2} p k\right) &
\end{array}
\]

This completes the proof of the claim.

Claim A.1.10. \(B \geq \exp \left(-8 \epsilon^{2} p k\right)\).

Proof. Fix \(\delta\) such that \(l=(1-\delta) p k\). The choice of \(l\) implies \(\delta \leq 2 \epsilon\), so the claim will hold as long as \(B \geq \exp \left(-2 \delta^{2} p k\right)\). Manipulting this latter inequality, we get
\[
\begin{aligned}
& B^{-1 / l} \leq \exp \left(\frac{2 \delta^{2} p k}{l}\right) \\
\Longleftrightarrow & \frac{l}{p k}\left(\frac{k-l}{(1-p) k}\right)^{k / l-1} \leq \exp \left(\frac{2 \delta^{2} p k}{l}\right) .
\end{aligned}
\]

Substituting \(l=(1-\delta) p k\) and simplifying, it is equivalent to
\[
(1-\delta)\left(1+\frac{\delta p}{1-p}\right)^{\frac{1}{(1-\delta) p}-1} \leq \exp \left(\frac{2 \delta^{2}}{1-\delta}\right)
\]

Taking the logarithm of both sides, we have
\[
\ln (1-\delta)+\left(\frac{1}{(1-\delta) p}-1\right)\left(\ln \left(1+\frac{\delta p}{1-p}\right)\right) \leq \frac{2 \delta^{2}}{1-\delta}
\]

Since \(\ln (1+z) \leq z\), it suffices to prove
\[
\begin{aligned}
& -\delta+\left(\frac{1}{(1-\delta) p}-1\right)\left(\frac{\delta p}{1-p}\right) \leq \frac{2 \delta^{2}}{1-\delta} \\
\Longleftrightarrow & \frac{\delta^{2}}{(1-p)(1-\delta)} \leq \frac{2 \delta^{2}}{1-\delta}
\end{aligned}
\]

Since \(p \leq 1 / 2\), this finishes the proof of the claim.

Combining the above two claims, we obtain the desired probability lower bound for the \((1-\epsilon)\) side. We can prove it for the \((1+\epsilon)\) side similarly.

Lemma A.1.11 (Anti-concentration of Gaussian distribution). Let \(X \sim N\left(0, \sigma^{2}\right)\), that is, the probability density function of \(X\) is given by \(\phi(x)=\frac{1}{\sqrt{2 \pi \sigma^{2}}} e^{-\frac{x^{2}}{2 \sigma^{2}}}\). Then
\[
\operatorname{Pr}[|X| \leq t] \in\left(\frac{2}{3} \frac{t}{\sigma}, \frac{4}{5} \frac{t}{\sigma}\right)
\]

\section*{A. 2 Matrix version probability tools}

We state a version from [Tro11c], and there is also another version can be found in [OT09].

Lemma A.2.1 (Matrix Freedman). Consider a matrix martingale \(\left\{Y_{i}: i=0,1,2, \cdots\right\}\) whose values are self-adjoint matrices with dimension \(n\), and let \(\left\{X_{i}: i=1,2,3, \cdots\right\}\) be the difference sequence. Assume that the difference sequence is uniformly bounded in the sense that
\[
\lambda_{\max }\left(X_{i}\right) \leq R, \text { almost surely for } i=1,2,3, \cdots .
\]

Define the predictable quadratic variation process of the martingale :
\[
W_{i}=\sum_{j=1}^{i} \underset{j-1}{\mathbb{E}}\left[X_{j}^{2}\right], \text { for } i=1,2,3, \cdots
\]

Then, for all \(t \geq 0\) and \(\sigma^{2}>0\),
\[
\operatorname{Pr}\left[\exists i \geq 0: \lambda_{\max }\left(Y_{i}\right) \geq t \text { and }\left\|W_{i}\right\| \leq \sigma^{2}\right] \leq n \cdot \exp \left(-\frac{t^{2} / 2}{\sigma^{2}+R t / 3}\right)
\]

Lemma A.2.2 (Matrix Bernstein, Theorem 6.1.1 in [Tro15]). Consider a finite sequence \(\left\{X_{1}, \cdots, X_{m}\right\} \subset \mathbb{R}^{n_{1} \times n_{2}}\) of independent, random matrices with common dimension \(n_{1} \times n_{2}\). Assume that
\[
\mathbb{E}\left[X_{i}\right]=0, \forall i \in[m] \quad \text { and } \quad\left\|X_{i}\right\| \leq M, \forall i \in[m]
\]

Let \(Z=\sum_{i=1}^{m} X_{i}\). Let \(\operatorname{Var}[Z]\) be the matrix variance statistic of sum:
\[
\operatorname{Var}[Z]=\max \left\{\left\|\sum_{i=1}^{m} \mathbb{E}\left[X_{i} X_{i}^{\top}\right]\right\|,\left\|\sum_{i=1}^{m} \mathbb{E}\left[X_{i}^{\top} X_{i}\right]\right\|\right\}
\]

Then
\[
\mathbb{E}[\|Z\|] \leq\left(2 \operatorname{Var}[Z] \cdot \log \left(n_{1}+n_{2}\right)\right)^{1 / 2}+M \cdot \log \left(n_{1}+n_{2}\right) / 3 .
\]

Furthermore, for all \(t \geq 0\),
\[
\operatorname{Pr}[\|Z\| \geq t] \leq\left(n_{1}+n_{2}\right) \cdot \exp \left(-\frac{t^{2} / 2}{\operatorname{Var}[Z]+M t / 3}\right)
\]

\section*{A. 3 Inequalities}

We state the Hölder's inequality for complex numbers. We will use the corresponding version \(p=q=2\) of Cauchy-Schwarz inequality for complex numbers.

Lemma A.3.1 (Hölder's inequality). If \(S\) is a measurable subset of \(\mathbb{R}^{n}\) with the Lebesgue measure, and \(f\) and \(g\) are measurable complex-valued functions on \(S\), then
\[
\int_{S}|f(x) g(x)| \mathrm{d} x \leq\left(\int_{S}|f(x)|^{p} \mathrm{~d} x\right)^{\frac{1}{p}}\left(\int_{S}|g(x)|^{q} \mathrm{~d} x\right)^{\frac{1}{q}}
\]

Fact A.3.2. For any vector \(x \in \mathbb{R}^{n}\), we have
\[
\|x\|_{2} \leq\|x\|_{1} \leq \sqrt{n}\|x\|_{2}
\]

For any matrix \(A \in \mathbb{R}^{n \times d}\), we have
\[
\|A\|_{F} \leq\|A\|_{1} \leq \sqrt{n d}\|A\|_{F}
\]
and
\[
\|A\| \leq\|A\|_{F} \leq \sqrt{d} \cdot\|A\|
\]

\section*{Appendix B}

\section*{Coauthors Index}

\section*{B. 1 Coauthors Index}

We provide a summary of all the coauthors related to this thesis. I would like to thank all of my coauthors \(x\) times, where \(x\) is the number of pages of this thesis. The purpose of writing my Ph.D. thesis is to thank and memorize helps from all the coauthors.
- Part I
- Zeyuan Allen-Zhu
* [AZLS18, AZLS19]
- Michael B. Cohen
* [CLS19]
- Yin Tat Lee
* [CLS19, LSZ19, LSV18]
- Yuanzhi Li
* [AZLS18, AZLS19]
- Santosh S. Vempala
* [LSV18]
- Qiuyi Zhang
* [LSZ19]
- Part II
- Ankit Garg
* [GLSS18]
- Rasmus Kyng
* [KS18, KLS19]
- Yin Tat Lee
* [GLSS18]
- Kyle Luh
* [KLS19]
- Nikhil Srivastava
* [GLSS18]
- Part III
- Xue Chen
* [CKPS16]
- Danial Kane
* [CKPS16]
- Eric Price
* [PS15, CKPS16]
- Vasileios Nakos
* [NS19, NSW18, NSW19a, NSW19b]
- Zhengyu Wang
* [NSW18, NSW19a, NSW19b]
- Part IV
- Ilya Razenshteyn
* [RSW16]
- David P. Woodruff
* [RSW16, SWZ16, SWZ17, SWZ18, SWZ19b, SWZ19a]
- Huan Zhang
* [SWZ16]
- Peilin Zhong
* [SWZ17, SWZ18, SWZ19b, SWZ19a]
- Part VI
- Alexandr Andoni
* [ASS \(\left.{ }^{+} 18\right]\)
- Huaian Diao
* DJS \(^{+}\)19]
- Rajesh Jayaram
* DJJS \(\left.^{+} 19\right]\)
- Jerry Li
* [LSS19]
- Eric Price
* [PSW17]
- Aviad Rubinstein
* [RS19b, RSSS19]
- Saeed Seddighin
* [RSSS19]
- Ruoqi Shen
* [LSS19]
- Clifford Stein
* [ASS \(\left.{ }^{+} 18\right]\)
- Wen Sun
* [DJS \(\left.{ }^{+} 19\right]\)
- Xiaorui Sun
* [RSSS19]
- Ruosong Wang
* [SWY \({ }^{+}\)19]
- Zhengyu Wang
* \(\left[\mathrm{ASS}^{+} 18\right]\)
- David P. Woodruff
* [PSW17, DJS \(\left.{ }^{+} 19\right]\)
- Lin F. Yang
* \(\left[\mathrm{SWY}^{+} 19\right]\)
- Xin Yang
* [SY19]
- Hongyang Zhang
* SWY \(^{+}\)19]
- Peilin Zhong
* \(\left.\mathrm{ASS}^{+} 18, \mathrm{SWY}^{+} 19\right]\)

\section*{Bibliography}
[AA91] Charles A. Akemann and Joel Anderson. Lyapunov theorems for operator algebras. Mem. Amer. Math. Soc., 94(458):iv+88, 1991. 593
[AAA \(\left.{ }^{+} 16\right]\) Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, et al. Deep speech 2: End-to-end speech recognition in English and Mandarin. In International Conference on Machine Learning (ICML), pages 173-182, 2016. 20, 250, 351
\(\left[\mathrm{AAB}^{+} 07\right]\) Evrim Acar, Canan Aykut-Bingöl, Haluk Bingol, Rasmus Bro, and Bülent Yener. Multiway analysis of epilepsy tensors. In Proceedings 15th International Conference on Intelligent Systems for Molecular Biology (ISMB) \& 6th European Conference on Computational Biology (ECCB), Vienna, Austria, July 21-25, 2007, pages 10-18, 2007. 57, 1469
[AAMM18] Atiye Alaeddini, Siavash Alemzadeh, Afshin Mesbahi, and Mehran Mesbahi. Linear model regression on time-series data: Non-asymptotic error bounds and applications. arXiv preprint arXiv:1807.06611, 2018. 257
[AB09] Sanjeev Arora and Boaz Barak. Computational complexity: a modern approach. Cambridge University Press, 2009. 1273, 2506, 2507
[AB17] Amir Abboud and Arturs Backurs. Towards hardness of approximation for polynomial time problems. In LIPIcs-Leibniz International Proceedings in Informatics, volume 67. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017. 2191, 2214, 2216, 2218
\(\left[\mathrm{ABB}^{+} 17\right]\) Sepehr Assadi, MohammadHossein Bateni, Aaron Bernstein, Vahab S. Mirrokni, and Cliff Stein. Coresets meet EDCS: algorithms for matching and vertex cover on massive graphs. In arXiv preprint. http://arxiv.org/pdf/ 1711.03076, 2017. 2316
[ABEZ16] Amir Adler, David Boublil, Michael Elad, and Michael Zibulevsky. A deep learning approach to block-based compressed sensing of images. arXiv preprint arXiv:1606.01519, 2016. 988
[ \(\left.\mathrm{ABF}^{+} 16\right]\) Jason Altschuler, Aditya Bhaskara, Gang Fu, Vahab Mirrokni, Afshin Rostamizadeh, and Morteza Zadimoghaddam. Greedy column subset selection: New bounds and distributed algorithms. In International Conference on Machine Learning (ICML). https://arxiv.org/pdf/1605.08795, 2016. 1476
[ABGM14] Sanjeev Arora, Aditya Bhaskara, Rong Ge, and Tengyu Ma. Provable bounds for learning some deep representations. In International Conference on Machine Learning (ICML), pages 584-592, 2014. 358
[ABIW09] Alexandr Andoni, Khanh Do Ba, Piotr Indyk, and David P Woodruff. Efficient sketches for earth-mover distance, with applications. In 50th Annual IEEE Symposium on Foundations of Computer Science, pages 324-330, 2009. 1262
[ABSV14] Pranjal Awasthi, Avrim Blum, Or Sheffet, and Aravindan Vijayaraghavan. Learning mixtures of ranking models. In Advances in Neural Information Processing Systems (NIPS). https://arxiv.org/pdf/1410.8750, 2014. 57, 1469
[ABW15] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. Tight hardness results for LCS and other sequence similarity measures. In IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015, pages 59-78, 2015. 1914, 2191, 2214
[ABW17] Pranjal Awasthi, Maria-Florina Balcan, and Colin White. General and robust communication-efficient algorithms for distributed clustering. In arXiv preprint. https://arxiv.org/pdf/1703.00830, 2017. 1784, 1788
[AC06] Nir Ailon and Bernard Chazelle. Approximate nearest neighbors and the fast johnson-lindenstrauss transform. In Proceedings of the thirty-eighth annual ACM symposium on Theory of computing (STOC), pages 557-563. ACM, 2006. 1907, 1922, 1966
[ACCL07] Nir Ailon, Bernard Chazelle, Seshadhri Comandur, and Ding Liu. Estimating the distance to a monotone function. Random Structures \(8 \mathcal{\delta}\) Algorithms, 31(3):371-383, 2007. 2218
\(\left[A C D{ }^{+} 16\right]\) Ittai Abraham, Shiri Chechik, Daniel Delling, Andrew V Goldberg, and Renato F Werneck. On dynamic approximate shortest paths for planar graphs
with worst-case costs. In Proceedings of the Twenty-Seventh Annual ACMSIAM Symposium on Discrete Algorithms (SODA), pages 740-753. Society for Industrial and Applied Mathematics, 2016. 2094
[ACGH18] Sanjeev Arora, Nadav Cohen, Noah Golowich, and Wei Hu. A convergence analysis of gradient descent for deep linear neural networks. arXiv preprint arXiv:1810.02281, 2018. 21, 251, 257, 351
[AÇKY05] Evrim Acar, Seyit A Çamtepe, Mukkai S Krishnamoorthy, and Bülent Yener. Modeling and multiway analysis of chatroom tensors. In International Conference on Intelligence and Security Informatics, pages 256-268. Springer, 2005. 57, 1469
[ACW17] Haim Avron, Kenneth L Clarkson, and David P Woodruff. Sharper bounds for regression and low-rank approximation with regularization. ., 2017. manuscript. \(1335,1338,1342,1343,1344\)
[ACY06] Evrim Acar, Seyit A Camtepe, and Bülent Yener. Collective sampling and analysis of high order tensors for chatroom communications. In International Conference on Intelligence and Security Informatics, pages 213-224. Springer, 2006. 57, 1469
[ADGM16] Anima Anandkumar, Yuan Deng, Rong Ge, and Hossein Mobahi. Homotopy analysis for tensor pca. In arXiv preprint. https://arxiv.org/pdf/1610. 09322, 2016. 58, 1470
[ADH \(\left.{ }^{+} 19\right]\) Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Ruslan Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. arXiv preprint arXiv:1904.11955, 2019. 1908
[ADHP09] Daniel Aloise, Amit Deshpande, Pierre Hansen, and Preyas Popat. Nphardness of euclidean sum-of-squares clustering. Machine learning, 75(2):245248, 2009. 2093, 2097
[ADK \(\left.{ }^{+} 16\right]\) Ittai Abraham, David Durfee, Ioannis Koutis, Sebastian Krinninger, and Richard Peng. On fully dynamic graph sparsifiers. In 2016 IEEE 57 th Annual Symposium on Foundations of Computer Science (FOCS), pages 335-344. IEEE, 2016. 2094
[ADL \(\left.{ }^{+} 12\right]\) Patrick Amestoy, Iain S. Duff, Jean-Yves L'Excellent, Yves Robert, FrançoisHenry Rouet, and Bora Uçar. On computing inverse entries of a sparse matrix in an out-of-core environment. SIAM J. Scientific Computing, 34(4), 2012. 1925
[AFdLGTL09] Santiago Aja-Fernández, Rodrigo de Luis Garcia, Dacheng Tao, and Xuelong Li. Tensors in image processing and computer vision. Springer Science \& Business Media, 2009. 57, 1469
[AFH \({ }^{+}\)12] Anima Anandkumar, Dean P Foster, Daniel J Hsu, Sham M Kakade, and Yi-Kai Liu. A spectral algorithm for latent dirichlet allocation. In Advances in Neural Information Processing Systems(NIPS), pages 917-925. https:// arxiv.org/pdf/1204.6703, 2012. 57, 1469, 1803, 1814
[AFK \(\left.{ }^{+} 01\right]\) Yossi Azar, Amos Fiat, Anna R. Karlin, Frank McSherry, and Jared Saia. Spectral analysis of data. In STOC, 2001. 48, 1036
[AFKM01] Dimitris Achlioptas, Amos Fiat, Anna R. Karlin, and Frank McSherry. Web search via hub synthesis. In FOCS, 2001. 48, 1036
[AFLG15] Andris Ambainis, Yuval Filmus, and François Le Gall. Fast matrix multiplication: limitations of the coppersmith-winograd method. In Proceedings of the forty-seventh annual ACM symposium on Theory of computing(STOC), pages 585-593. ACM, 2015. 12, 4
[AFS12] Andreas Argyriou, Rina Foygel, and Nathan Srebro. Sparse prediction with the \(k\)-support norm. In Advances in Neural Information Processing Systems (NIPS), pages 1457-1465, 2012. 1968, 1971, 1984
[AG09] Kook Jin Ahn and Sudipto Guha. Graph sparsification in the semi-streaming model. In International Colloquium on Automata, Languages, and Programming (ICALP), pages 328-338. Springer, 2009. 26, 549
[AG14] Nima Anari and Shayan Oveis Gharan. The kadison-singer problem for strongly rayleigh measures and applications to asymmetric tsp. arXiv preprint arXiv:1412.1143, 2014. 596, 597, 607, 612, 613
[AG15] Nima Anari and Shayan Oveis Gharan. Effective-resistance-reducing flows, spectrally thin trees, and asymmetric tsp. In Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on, pages 20-39. IEEE, 2015. 554, 562, 563
[AG18] Kook Jin Ahn and Sudipto Guha. Access to data and number of iterations: Dual primal algorithms for maximum matching under resource constraints. ACM Transactions on Parallel Computing (TOPC), 4(4):17, 2018. 2315, 2316
[AGH \(\left.{ }^{+} 14\right]\) Animashree Anandkumar, Rong Ge, Daniel J. Hsu, Sham M. Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. In Journal of Machine Learning Research, volume 15(1), pages 2773-2832. https://arxiv.org/pdf/1210.7559, 2014. 57, 58, 1469, 1800, 1803, 1810, 1814, 1850
[AGHK14] Animashree Anandkumar, Rong Ge, Daniel J Hsu, and Sham M Kakade. A tensor approach to learning mixed membership community models. In Journal of Machine Learning Research, volume 15(1), pages 2239-2312. https://arxiv.org/pdf/1302.2684, 2014. 57, 1469
[AGKM12] Sanjeev Arora, Rong Ge, Ravindran Kannan, and Ankur Moitra. Computing a nonnegative matrix factorization - provably. In Proceedings of the 44 th Symposium on Theory of Computing Conference (STOC), New York, NY, USA, May 19 - 22, 2012, pages 145-162. https://arxiv.org/pdf/1111. 0952, 2012. 1044, 1046, 1050, 1088, 1131, 1506, 1899
\(\left[\mathrm{AGM}^{+} 10\right]\) Arash Asadpour, Michel X. Goemans, Aleksander Mądry, Shayan Oveis Gharan, and Amin Saberi. An \(O(\log n / \log \log n)\)-approximation algorithm for the asymmetric traveling salesman problem. In Proceedings of the Twentyfirst Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages

379-389, Philadelphia, PA, USA, 2010. Society for Industrial and Applied Mathematics. 550, 555
[AGM12a] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Analyzing graph structure via linear measurements. In Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms, pages 459-467. Society for Industrial and Applied Mathematics, 2012. 2094
[AGM12b] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Graph sketches: sparsification, spanners, and subgraphs. In Proceedings of the 31st ACM SIGMOD-SIGACT-SIGAI symposium on Principles of Database Systems, pages 5-14. ACM, 2012. 2094
[AGM13] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Spectral sparsification in dynamic graph streams. In Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, pages 1-10. Springer, 2013. 2094
[AGMR16] Sanjeev Arora, Rong Ge, Tengyu Ma, and Andrej Risteski. Provable learning of noisy-or networks. In Proceedings of the 49th Annual Symposium on the Theory of Computing (STOC). ACM, https://arxiv.org/pdf/1612. 08795, 2016. 57, 1469
[AGR16] Nima Anari, Shayan Oveis Gharan, and Alireza Rezaei. Monte carlo markov chain algorithms for sampling strongly rayleigh distributions and determinantal point processes. In Conference on Learning Theory (COLT), pages 103-115, 2016. 555, 622
[AGS03] Adi Akavia, Shafi Goldwasser, and Shmuel Safra. Proving hard-core predicates using list decoding. In FOCS, volume 44, pages 146-159, 2003. 38, 919
[AH16] Zeyuan Allen-Zhu and Elad Hazan. Variance Reduction for Faster NonConvex Optimization. In Proceedings of the 33rd International Conference on Machine Learning, ICML '16, 2016. Full version available at http:// arxiv.org/abs/1603.05643. 15, 77
\(\left[\right.\) AHL \(\left.^{+} 18\right]\) Sanjeev Arora, Elad Hazan, Holden Lee, Karan Singh, Cyril Zhang, and Yi Zhang. Towards provable control for unknown linear dynamical systems. In ICLR workshop, 2018. 257
[AHPV04] Pankaj K Agarwal, Sariel Har-Peled, and Kasturi R Varadarajan. Approximating extent measures of points. Journal of the ACM (JACM), 51(4):606635, 2004. 2094
[AHWW16] Amir Abboud, Thomas Dueholm Hansen, Virginia Vassilevska Williams, and Ryan Williams. Simulating branching programs with edit distance and friends: or: a polylog shaved is a lower bound made. In Proceedings of the 48 th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016, pages 375-388, 2016. 1914, 2191, 2214
[AIK08] Alexandr Andoni, Piotr Indyk, and Robert Krauthgamer. Earth mover distance over high-dimensional spaces. In Proceedings of the nineteenth annual

ACM-SIAM symposium on Discrete algorithms, pages 343-352. Society for Industrial and Applied Mathematics, 2008. 1262
[AK17] Sepehr Assadi and Sanjeev Khanna. Randomized composable coresets for matching and vertex cover. In SPAA. https://arxiv.org/pdf/1705.08242, 2017. 2315, 2316
[AKDM10] E. Acar, T. G. Kolda, D. M. Dunlavy, and M. Morup. Scalable Tensor Factorizations for Incomplete Data. In arXiv preprint. https://arxiv.org/ pdf/1005.2197, 2010. 58, 1469
[AKK \({ }^{+}\)17] Naman Agarwal, Sham Kakade, Rahul Kidambi, Yin Tat Lee, Praneeth Netrapalli, and Aaron Sidford. Leverage score sampling for faster accelerated regression and erm. arXiv preprint arXiv:1711.08426, 2017. 13, 74
\(\left[\mathrm{AKM}^{+}\right.\)19] Haim Avron, Michael Kapralov, Cameron Musco, Christopher Musco, Ameya Velingker, and Amir Zandieh. A universal sampling method for reconstructing signals with simple fourier transforms. In STOC, 2019. 920
[AKO10] Alexandr Andoni, Robert Krauthgamer, and Krzysztof Onak. Polylogarithmic approximation for edit distance and the asymmetric query complexity. In Proceedings of the Fifty-First IEEE Annual Symposium on Foundations of Computer Science, 2010. 1914, 2029, 2191, 2214, 2215
[AKO11] Alexandr Andoni, Robert Krauthgamer, and Krzysztof Onak. Streaming algorithms via precision sampling. In Foundations of Computer Science
(FOCS), 2011 IEEE 52nd Annual Symposium on, pages 363-372. IEEE, https://arxiv.org/pdf/1011.1263, 2011. 1568, 1593
[AKPS19] Deeksha Adil, Rasmus Kyng, Richard Peng, and Sushant Sachdeva. Iterative refinement for \(\ell_{p}\)-norm regression. In Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 1405-1424. SIAM, 2019. 14, 16, 74, 78
[AL13] Nir Ailon and Edo Liberty. An almost optimal unrestricted fast johnsonlindenstrauss transform. ACM Transactions on Algorithms (TALG), 9(3):21, 2013. 1922
[AL17] Zeyuan Allen-Zhu and Yuanzhi Li. Follow the Compressed Leader: Faster Online Learning of Eigenvectors and Faster MMWU. In ICML, 2017. Full version available at http://arxiv.org/abs/1701.01722. lxxi, 265
[AL18] Zeyuan Allen-Zhu and Yuanzhi Li. Neon2: Finding Local Minima via FirstOrder Oracles. In NeurIPS, 2018. Full version available at http://arxiv. org/abs/1711.06673. lxxi, 265
[AL19a] Zeyuan Allen-Zhu and Yuanzhi Li. Can SGD Learn Recurrent Neural Networks with Provable Generalization? ArXiv e-prints, abs/1902.01028, February 2019. Full version available at http://arxiv.org/abs/1902.01028. 357, 366
[AL19b] Zeyuan Allen-Zhu and Yuanzhi Li. What Can ResNet Learn Efficiently, Going Beyond Kernels? ArXiv e-prints, abs/1905.10337, May 2019. 256
[ALA16] Kamyar Azizzadenesheli, Alessandro Lazaric, and Animashree Anandkumar. Reinforcement learning of POMDPs using spectral methods. In 29th Annual Conference on Learning Theory (COLT), pages 193-256. https://arxiv. org/pdf/1602.07764, 2016. 57, 1469
[ALB13] Mohammad Gheshlaghi Azar, Alessandro Lazaric, and Emma Brunskill. Sequential transfer in multi-armed bandit with finite set of models. In \(A d\) vances in Neural Information Processing Systems(NIPS), pages 2220-2228. https://arxiv.org/pdf/1307.6887, 2013. 57, 1469
[Ald90] David Aldous. The random walk construction of uniform spanning trees and uniform labelled trees. In SIAM Journal on Discrete Mathematics, pages 450-465, 1990. 29, 550
[All12a] Genevera Allen. Sparse higher-order principal components analysis. In AISTATS, volume 15, 2012. 58, 1470
[All12b] Genevera I Allen. Regularized tensor factorizations and higher-order principal components analysis. In arXiv preprint. https://arxiv.org/pdf/1202. 2476, 2012. 58, 1470
[All17a] Zeyuan Allen-Zhu. Katyusha: The first direct acceleration of stochastic gradient methods. In Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, pages 1200-1205. ACM, 2017. 15, 77
[All17b] Zeyuan Allen-Zhu. Natasha: Faster Non-Convex Stochastic Optimization via Strongly Non-Convex Parameter. In Proceedings of the 34th International

Conference on Machine Learning, ICML '17, 2017. Full version available at http://arxiv.org/abs/1702.00763. 15, 77
[All18a] Zeyuan Allen-Zhu. Katyusha X: Practical Momentum Method for Stochastic Sum-of-Nonconvex Optimization. In Proceedings of the 35th International Conference on Machine Learning, ICML '18, 2018. Full version available at http://arxiv.org/abs/1802.03866. 15, 77
[All18b] Zeyuan Allen-Zhu. Natasha 2: Faster Non-Convex Optimization Than SGD. In Proceedings of the 32nd Conference on Neural Information Processing Systems, NIPS '18, 2018. Full version available at http://arxiv.org/abs/ 1708.08694. 15, 77
[ALL18c] Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and Generalization in Overparameterized Neural Networks, Going Beyond Two Layers. arXiv preprint arXiv:1811.04918, November 2018. 256, 270, 357, 366
[Alm19] Josh Alman. Limits on the universal method for matrix multiplication. In CCC (Best Student Paper). arXiv preprint arXiv:1812.08731, 2019. 16, 78
[ALS \(\left.{ }^{+} 18\right]\) A. Andoni, C. Lin, Y. Sheng, P. Zhong, and R. Zhong. Subspace embedding and linear regression with Orlicz norm. In ICML, pages 224-233, 2018. 1967, 1968, 1970, 1971, 1972, 1976
[AM05] Dimitris Achlioptas and Frank McSherry. On spectral learning of mixtures of distributions. In COLT, 2005. 48, 1036
[AM07] Dimitris Achlioptas and Frank McSherry. Fast computation of low-rank matrix approximations. J. ACM, 54(2):9, 2007. 1338, 1468
[AMR \({ }^{+}\)12] Marcel R Ackermann, Marcus Märtens, Christoph Raupach, Kamil Swierkot, Christiane Lammersen, and Christian Sohler. Streamkm++: A clustering algorithm for data streams. Journal of Experimental Algorithmics (JEA), 17:2-4, 2012. 2094
[AMS11] Christoph Ambühl, Monaldo Mastrolilli, and Ola Svensson. Inapproximability results for maximum edge biclique, minimum linear arrangement, and sparsest cut. SIAM Journal on Computing, 40(2):567-596, 2011. 1081
[AMT10] Haim Avron, Petar Maymounkov, and Sivan Toledo. Blendenpik: Supercharging lapack's least-squares solver. SIAM J. Scientific Computing, 32(3):1217-1236, 2010. 1927
\(\left[A N N^{+} 17\right]\) A. Andoni, H. L. Nguyen, A. Nikolov, I. Razenshteyn, and E. Waingarten. Approximate near neighbors for general symmetric norms. In STOC, pages 902-913, 2017. 1967
[ANN \({ }^{+}\)18] A. Andoni, A. Naor, A. Nikolov, I. Razenshteyn, and E. Waingarten. Hölder homeomorphisms and approximate nearest neighbors. In FOCS, pages 159169, 2018. 1967
[ANOY14] Alexandr Andoni, Aleksandar Nikolov, Krzysztof Onak, and Grigory Yaroslavtsev. Parallel algorithms for geometric graph problems. In STOC. http://arxiv.org/pdf/1401.0042, 2014. 2315, 2316
[ANW14] Haim Avron, Huy Nguyen, and David Woodruff. Subspace embeddings for the polynomial kernel. In Advances in Neural Information Processing Systems(NIPS), pages 2258-2266, 2014. 1482, 1486, 1518, 1519
[AO09] Alexandr Andoni and Krzysztof Onak. Approximating edit distance in nearlinear time. In Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing, 2009. 1914, 2214, 2215
[AO12] Alexandr Andoni and Krzysztof Onak. Approximating edit distance in nearlinear time. SIAM Journal on Computing, 41(6):1635-1648, 2012. 2191
[AOGSS18] Nima Anari, Shayan Oveis Gharan, Amin Saberi, and Nikhil Srivastava. Approximating the largest root and applications to interlacing families. In Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1015-1028. SIAM, Philadelphia, PA, 2018. 598
[APY09] Noga Alon, Rina Panigrahy, and Sergey Yekhanin. Deterministic approximation algorithms for the nearest codeword problem. In Algebraic Methods in Computational Complexity, 2009. lxxxi, 1431, 1457
[AR18] Amir Abboud and Aviad Rubinstein. Fast and deterministic constant factor approximation algorithms for les imply new circuit lower bounds. In LIPIcs-Leibniz International Proceedings in Informatics, volume 94. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018. 2191, 2214, 2216, 2218
[AS16] Noga Alon and Joel H. Spencer. The probabilistic method. Wiley Series in

Discrete Mathematics and Optimization. John Wiley \& Sons, Inc., Hoboken, NJ, fourth edition, 2016. 30, 590
[ASS \({ }^{+}\)18] Alexandr Andoni, Zhao Song, Clifford Stein, Zhengyu Wang, and Peilin Zhong. Parallel graph connectivity in log diameter rounds. In 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS), pages 674-685. IEEE, 2018. xxviii, 1916, 2314, 2514, 2515, 2516
[ASS19] Josh Alman, Aaron Schild, and Zhao Song. Fast sparsification and linear algebra for geometric graphs. Manuscript, 2019. xxix
[ASSN08] Abiodun M Aibinu, Momoh JE Salami, Amir A Shafie, and Athaur Rahman Najeeb. MRI reconstruction using discrete Fourier transform: a tutorial. World Academy of Science, Engineering and Technology, 42:179, 2008. 36, 917
[ASW18] Sepehr Assadi, Xiaorui Sun, and Omri Weinstein. Massively parallel algorithms for finding well-connected components in sparse graphs. In ArXiv preprint. https://arxiv.org/pdf/1805.02974, 2018. 2320
[Aue30] H. Auerbach. On the area of convex curves with conjugate diameters. PhD thesis, University of Lwów, 1930. 2019
[AW02] Rudolf Ahlswede and Andreas Winter. Strong converse for identification via quantum channels. IEEE Transactions on Information Theory, 48(3):569579, 2002. 25, 26, 30, 503, 506, 548, 553, 590, 2141
[AW14] Charles Akemann and Nik Weaver. A Lyapunov-type theorem from KadisonSinger. Bull. Lond. Math. Soc., 46(3):517-524, 2014. 593
[AW18a] Josh Alman and Virginia Vassilevska Williams. Further limitations of the known approaches for matrix multiplication. In ITCS. arXiv preprint arXiv:1712.07246, 2018. 16, 78
[AW18b] Josh Alman and Virginia Vassilevska Williams. Limits on all known (and some unknown) approaches to matrix multiplication. In 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS). IEEE, 2018. \(12,16,4,78\)
[AWW14] Amir Abboud, Virginia Vassilevska Williams, and Oren Weimann. Consequences of faster alignment of sequences. In Automata, Languages, and Programming - 41 st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I, pages 39-51, 2014. 1914, 2191
[AY16] Zeyuan Allen-Zhu and Yang Yuan. Improved SVRG for Non-Strongly-Convex or Sum-of-Non-Convex Objectives. In Proceedings of the 33rd International Conference on Machine Learning, ICML '16, 2016. Full version available at http://arxiv.org/abs/1506.01972. 15, 77
[AZLS18] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. On the convergence rate of training recurrent neural networks. In arXiv preprint. https://arxiv.org/ pdf/1810.12065, 2018. xxv, 8, 24, 255, 272, 294, 311, 1908, 1909, 1910, 2140, 2143, 2512
[AZLS19] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In \(I C M L\). https://arxiv.org/ pdf/1811.03962, 2019. xxv, 7, 24, 352, 355, 357, 363, 366, 395, 495, 1908, 1909, 2140, 2143, 2512
[Ban38] Stefan Banach. Über homogene polynome in \(\left(l^{2}\right)\). Studia Mathematica, 7(1):36-44, 1938. 1676, 1677
[Bas08] Surender Baswana. Streaming algorithm for graph spanners-single pass and constant processing time per edge. Information Processing Letters, 106(3):110-114, 2008. 2094
[Bas14] Saugata Basu. Algorithms in real algebraic geometry: a survey. arXiv preprint arXiv:1409.1534, 2014. 1048, 1052, 1131
[Baw75] Vijay S Bawa. Optimal rules for ordering uncertain prospects. Journal of Financial Economics, 2(1):95-121, 1975. 654
[BB08] Léon Bottou and Olivier Bousquet. The tradeoffs of large scale learning. In Advances in neural information processing systems, pages 161-168, 2008. 13, 74, 598
[ \(\left.\mathrm{BBB}^{+} 19 \mathrm{a}\right]\) Frank Ban, Vijay Bhattiprolu, Karl Bringmann, Pavel Kolev, Euiwoong Lee, and David P. Woodruff. A PTAS for \(\ell_{p}\)-low rank approximation. In \(S O D A\), 2019. 56, 1371, 1428
[ \(\mathrm{BBB}^{+}\)19b] Frank Ban, Vijay Bhattiprolu, Karl Bringmann, Pavel Kolev, Euiwoong Lee, and David P Woodruff. A ptas for \(\ell_{p}\)-low rank approximation. In Proceed-
ings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 747-766. SIAM, 2019. 2027
[ \(\left.\mathrm{BBC}^{+} 17\right]\) Jaroslaw Blasiok, Vladimir Braverman, Stephen R Chestnut, Robert Krauthgamer, and Lin F Yang. Streaming symmetric norms via measure concentration. In Proceedings of the 49th Annual Symposium on the Theory of Computing(STOC). ACM, https://arxiv.org/pdf/1511.01111, 2017. 1717, 1971, 1976, 1983, 1984, 2000, 2001, 2018
\(\left[\mathrm{BBD}^{+} 16\right]\) Aritra Banik, Binay K. Bhattacharya, Sandip Das, Tsunehiko Kameda, and Zhao Song. The p-center problem in tree networks revisited. In SWAT, pages 6:1-6:15, 2016. xxix
\(\left[\mathrm{BBD}^{+} 18\right]\) Sergey Bereg, Binay Bhattacharya, Sandip Das, Tsunehiko Kameda, Priya Ranjan Sinha Mahapatra, and Zhao Song. Optimizing squares covering a set of points. Theoretical Computer Science, 729:68-83, 2018. xxix
[BBK13] Pavle VM Blagojević, Boris Bukh, and Roman Karasev. Turán numbers for \(k_{s, t}\)-free graphs: Topological obstructions and algebraic constructions. Israel Journal of Mathematics, 197(1):199-214, 2013. 2284
[BBL09] Julius Borcea, Petter Brändén, and Thomas Liggett. Negative dependence and the geometry of polynomials. Journal of the American Mathematical Society, 22(2):521-567, 2009. 28, 549, 557, 558, 566, 568, 579, 588
[BBLM14] MohammadHossein Bateni, Aditya Bhaskara, Silvio Lattanzi, and Vahab Mirrokni. Distributed balanced clustering via mapping coresets. In Advances
in Neural Information Processing Systems (NIPS), pages 2591-2599, 2014. 1784, 1788
[BBM05] Peter L Bartlett, Olivier Bousquet, and Shahar Mendelson. Local rademacher complexities. The Annals of Statistics, 33(4):1497-1537, 2005. 13, 74
[BCDH10] Richard G Baraniuk, Volkan Cevher, Marco F Duarte, and Chinmay Hegde. Model-based compressive sensing. IEEE Transactions on Information Theory, 56(4):1982-2001, 2010. 635
[ \(\left.\mathrm{BCG}^{+} 12\right]\) Petros Boufounos, Volkan Cevher, Anna C Gilbert, Yi Li, and Martin J Strauss. What's the frequency, Kenneth?: Sublinear Fourier sampling off the grid. In Algorithmica(A preliminary version of this paper appeared in the Proceedings of RANDOM/APPROX 2012, LNCS 7408, pp.61-72), pages 1-28. Springer, 2012. 40, 41, 622, 676, 686, 690, 691, 763, 920, 1334
[ \(\left.\mathrm{BCI}^{+} 16\right]\) Vladimir Braverman, Stephen R Chestnut, Nikita Ivkin, Jelani Nelson, Zhengyu Wang, and David P Woodruff. Bptree: an \(\ell_{2}\) heavy hitters algorithm using constant memory. In Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS). https://arxiv.org/pdf/1603.00759, 2016. 1779
[BCIW16] Vladimir Braverman, Stephen R Chestnut, Nikita Ivkin, and David P Woodruff. Beating countsketch for heavy hitters in insertion streams. In Proceedings of the 48 th Annual Symposium on the Theory of Computing (STOC). https://arxiv.org/pdf/1511.00661, 2016. 1779
[BCKY16] Vladimir Braverman, Stephen R Chestnut, Robert Krauthgamer, and Lin F Yang. Sketches for matrix norms: Faster, smaller and more general. In arXiv preprint. https://arxiv.org/pdf/1609.05885, 2016. 1717
[BCL05] Zheng-Jian Bai, Raymond H Chan, and Franklin T Luk. Principal component analysis for distributed data sets with updating. In Advanced Parallel Processing Technologies, pages 471-483. Springer, 2005. 1318, 1784, 1788
[BCLL18] Sébastien Bubeck, Michael B Cohen, Yin Tat Lee, and Yuanzhi Li. An homotopy method for \(\ell_{p}\) regression provably beyond self-concordance and in input-sparsity time. In Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC), pages 1130-1137. ACM, 2018. 11, 14, 16, 3, 74, 78, 2033, 2046, 2051, 2067
[BCMN14] Sayan Bhattacharya, Parinya Chalermsook, Kurt Mehlhorn, and Adrian Neumann. New approximability results for the robust k-median problem. In Scandinavian Workshop on Algorithm Theory, pages 50-61. Springer, 2014. 2094
[BCMV14] Aditya Bhaskara, Moses Charikar, Ankur Moitra, and Aravindan Vijayaraghavan. Smoothed analysis of tensor decompositions. In Proceedings of the 46th Annual ACM Symposium on Theory of Computing, pages 594-603. ACM, https://arxiv.org/pdf/1311.3651, 2014. 58, 1470
[BCR87] Jacek Bochnak, Michel Coste, and Marie-Françoise Roy. Géométrie algébrique réelle, volume 12. Springer Science \& Business Media, 1987. 1051
[BCS97] Peter Bürgisser, Michael Clausen, and Amin Shokrollahi. Algebraic complexity theory, volume 315. Springer Science \& Business Media, 1997. 59, 1471
[BCS18] Mitali Bafna, Chi-Ning Chou, and Zhao Song. An exposition of dinur-khot-kindler-minzer-safra's proof for the 2-to-2 games conjecture. In Notes from Boaz Barak's Feb 23rd and March 2nd 2018 talks at CMSA, 2018. xxix
[BCV14] Aditya Bhaskara, Moses Charikar, and Aravindan Vijayaraghavan. Uniqueness of tensor decompositions with applications to polynomial identifiability. In 27th Annual Conference on Learning Theory (COLT), pages 742-778. https://arxiv.org/pdf/1304.8087, 2014. liii, 57, 1469, 1470, 1471, 1474, 1491
[BCWY16] V. Braverman, S. R. Chestnut, David P. Woodruff, and Lin F. Yang. Streaming space complexity of nearly all functions of one variable on frequency vectors. In PODS, pages 261-276, 2016. 1967
[BD08] Thomas Blumensath and Mike E Davies. Iterative thresholding for sparse approximations. Journal of Fourier analysis and Applications, 14(5-6):629654, 2008. lxvii, lxviii, 37, 38, 918, 919, 923
[BD09a] Thomas Blumensath and Mike E Davies. Iterative hard thresholding for compressed sensing. Applied and computational harmonic analysis, 27(3):265274, 2009. 34, 38, 617, 918, 925
[BD09b] Thomas Blumensath and Mike E Davies. A simple, efficient and near optimal algorithm for compressed sensing. 2009 IEEE International Conference on Acoustics, Speech and Signal Processing, 2009. 38, 918
[BD10] Thomas Blumensath and Mike E Davies. Normalized iterative hard thresholding: Guaranteed stability and performance. IEEE Journal of selected topics in signal processing, 4(2):298-309, 2010. 38, 918
[BD13] J. Paul Brooks and José H. Dulá. The \(\ell_{1}\)-norm best-fit hyperplane problem. Appl. Math. Lett., 26(1):51-55, 2013. 52, 53, 1109, 1110, 1370
[BDB13] J. Paul Brooks, José H. Dulá, and Edward L Boone. A pure \(\ell_{1}\)-norm principal component analysis. Computational statistics \(\mathcal{E}\) data analysis, 61:83-98, 2013. l, lxxv, 53, 1110, 1286, 1289, 1290, 1298, 1299, 1334, 1370
\(\left[\mathrm{BDE}^{+} 19\right]\) Soheil Behnezhad, Laxman Dhulipala, Hossein Esfandiari, Jakub Lacki, and Vahab Mirrokni. Near-optimal massively parallel graph connectivity. In 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS). IEEE, 2019. 1916
[BDG16] Nikhil Bansal, Daniel Dadush, and Shashwat Garg. An algorithm for komlós conjecture matching banaszczyk's bound. In 2016 IEEE 57 th Annual Symposium on Foundations of Computer Science (FOCS), pages 788-799. IEEE, 2016. 596
[BDGL18] Nikhil Bansal, Daniel Dadush, Shashwat Garg, and Shachar Lovett. The
gram-schmidt walk: A cure for the banaszczyk blues. In STOC. https: //arxiv.org/pdf/1708.01079, 2018. 596
\(\left[\mathrm{BDK}^{+} 14\right]\) Binay K. Bhattacharya, Minati De, Tsunehiko Kameda, Sasanka Roy, Vladyslav Sokol, and Zhao Song. Back-up 2-center on a path/tree/cycle/unicycle. In COCOON, pages 417-428, 2014. xxix
[BDL16] Amitabh Basu, Michael Dinitz, and Xin Li. Computing approximate PSD factorizations. In arXiv preprint. https://arxiv.org/pdf/1602.07351, 2016. 1046, 1131, 1506
[BDM11] Christos Boutsidis, Petros Drineas, and Malik Magdon-Ismail. Near optimal column-based matrix reconstruction. In IEEE 52nd Annual Symposium on Foundations of Computer Science (FOCS), 2011, Palm Springs, CA, USA, October 22-25, 2011, pages 305-314. https://arxiv.org/pdf/1103.0995, 2011. 1114, 1433, 1476, 1604
[BDMO03] Brain Babcock, Mayur Datar, Rajeev Motwani, and Liadan O'Callaghan. Maintaining variance and k-medians over data stream windows. In Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems, pages 234-243. ACM, 2003. 2094
[BDN15] Jean Bourgain, Sjoerd Dirksen, and Jelani Nelson. Toward a unified theory of sparse dimensionality reduction in euclidean space. In Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015, pages 499-508, 2015. 52, 1109, 1370, 1426, 1922
[ \(\left.\mathrm{BEG}^{+} 18\right]\) Mahdi Boroujeni, Soheil Ehsani, Mohammad Ghodsi, MohammadTaghi HajiAghayi, and Saeed Seddighin. Approximating edit distance in truly subquadratic time: Quantum and mapreduce. In Proceedings of the TwentyNinth Annual ACM-SIAM Symposium on Discrete Algorithms, 2018. 62, 1914, 2191, 2212, 2214, 2215, 2220, 2231, 2286
[Ber24] Sergei Bernstein. On a modification of chebyshev's inequality and of the error formula of laplace. Ann. Sci. Inst. Sav. Ukraine, Sect. Math, 1(4):38-49, 1924. 25, 548, 2141, 2505
[Ber44] Joseph Berkson. Application of the logistic function to bio-assay. Journal of the American Statistical Association, 39(227):357-365, 1944. 16, 171
[BES06] Tuğkan Batu, Funda Ergun, and Cenk Sahinalp. Oblivious string embeddings and edit distance approximations. In Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm, 2006. 1914, 2191, 2214, 2215
[BFFN17] Jack Baker, Paul Fearnhead, Emily B Fox, and Christopher Nemeth. Control variates for stochastic gradient mcmc. arXiv preprint arXiv:1706.05439, 2017. 181
[BFL16] Vladimir Braverman, Dan Feldman, and Harry Lang. New frameworks for offline and streaming coreset constructions. arXiv preprint arXiv:1612.00889, 2016. 246, 2090, 2094, 2098, 2099, 2136, 2137
\(\left[\mathrm{BFL}^{+} 17\right]\) Vladimir Braverman, Gereon Frahling, Harry Lang, Christian Sohler, and Lin F Yang. Clustering high dimensional dynamic data streams. In \(I C M L\).
https://arxiv.org/pdf/1706.03887, 2017. lxxvii, 2089, 2090, 2094, 2095, 2105, 2137, 2138
[BFN \(\left.{ }^{+} 17\right]\) Richard G Baraniuk, Simon Foucart, Deanna Needell, Yaniv Plan, and Mary Wootters. Exponential decay of reconstruction error from binary measurements of sparse signals. IEEE Transactions on Information Theory, 63(6):3368-3385, 2017. 636
[BFR16] Joris Bierkens, Paul Fearnhead, and Gareth Roberts. The zig-zag process and super-efficient sampling for bayesian analysis of big data. arXiv preprint arXiv:1607.03188, 2016. 181
[BG17] Nikhil Bansal and Shashwat Garg. Algorithmic discrepancy beyond partial coloring. In Proceedings of the 49 th Annual ACM SIGACT Symposium on Theory of Computing, pages 914-926. ACM, 2017. 21, 250, 257, 351, 596, 2140
[BGMSS18] Alon Brutzkus, Amir Globerson, Eran Malach, and Shai Shalev-Shwartz. SGD learns over-parameterized networks that provably generalize on linearly separable data. In ICLR, 2018. 351, 358
[BGS98] Mihir Bellare, Oded Goldreich, and Madhu Sudan. Free bits, pcps, and nonapproximability-towards tight results. SIAM Journal on Computing, 27(3):804-915, 1998. 1269
[BGS15] Surender Baswana, Manoj Gupta, and Sandeep Sen. Fully dynamic maximal
matching in \(\mathrm{o}(\log n)\) update time. SIAM Journal on Computing, 44(1):88113, 2015. 2094
[BH89] Paul Beame and Johan Håstad. Optimal bounds for decision problems on the CRCW PRAM. J. ACM, 36(3):643-670, 1989. 2315
[Bha97] Rajendra Bhatia. Matrix analysis, volume 169. Springer Science \& Business Media, 1997. 517
[BHL18] Peter Bartlett, Dave Helmbold, and Phil Long. Gradient descent with identity initialization efficiently learns positive definite linear transformations. In International Conference on Machine Learning, pages 520-529, 2018. 21, 251, 257, 351, 1451
[BHNT15] Sayan Bhattacharya, Monika Henzinger, Danupon Nanongkai, and Charalampos Tsourakakis. Space-and time-efficient algorithm for maintaining dense subgraphs on one-pass dynamic streams. In Proceedings of the fortyseventh annual ACM symposium on Theory of computing, pages 173-182. ACM, 2015. 2094
[BI08] Radu Berinde and Piotr Indyk. Sparse recovery using sparse random matrices. preprint, 2008. 625
[BI15] Arturs Backurs and Piotr Indyk. Edit Distance Cannot Be Computed in Strongly Subquadratic Time (unless SETH is false). In Proc. of the 47 th Annual ACM SIGACT Symposium on Theory of Computing (STOC), pages 51-58, 2015. 1914, 2191, 2214, 2215
[Bin80] Dario Bini. Border rank of a \(\mathrm{p} \times \mathrm{q} \times 2\) tensor and the optimal approximation of a pair of bilinear forms. Automata, languages and programming, pages 98108, 1980. 59, 1471
[Bin86] Dario Bini. Border rank of \(\mathrm{m} \times \mathrm{n} \times(\mathrm{mn}-\mathrm{q})\) tensors. Linear Algebra and Its Applications, 79:45-51, 1986. 59, 1471
[BIR08a] Radu Berinde, Piotr Indyk, and Milan Ruzic. Practical near-optimal sparse recovery in the \(\ell_{1}\) norm. In Communication, Control, and Computing, 2008 46th Annual Allerton Conference on, pages 198-205. IEEE, 2008. 34, 617
[BIR08b] Radu Berinde, Piotr Indyk, and Milan Ruzic. Practical near-optimal sparse recovery in the 11 norm. In Communication, Control, and Computing, 2008 46th Annual Allerton Conference on, pages 198-205. IEEE, 2008. 989
[BIRW16] Arturs Backurs, Piotr Indyk, Ilya Razenshteyn, and David P Woodruff. Nearly-optimal bounds for sparse recovery in generic norms, with applications to k-median sketching. In Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, pages 318-337. SIAM, 2016. 620, 1262, 2094
[BJ12] J. Paul Brooks and Sapan Jot. Pcal1: An implementation in r of three methods for \(\ell_{1}\)-norm principal component analysis. Optimization Online preprint, 2012. 53, 1110, 1370
[BJPD17] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G Dimakis. Compressed
sensing using generative models. In ICML. https://arxiv.org/pdf/1703. 03208, 2017. 988
[BJW19] Ainesh Bakshi, Rajesh Jayaram, and David P Woodruff. Learning two layer rectified neural networks in polynomial time. In COLT. http://arxiv.org/ pdf/:1811.01885, 2019. 2140
[BK96] András A. Benczúr and David R. Karger. Approximating s-t minimum cuts in \(\widetilde{O}\left(n^{2}\right)\) time. In Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing (STOC), pages 47-55, New York, NY, USA, 1996. ACM. 26, 29, 549, 556
[BK02] Piotr Berman and Marek Karpinski. Approximating minimum unsatisfiability of linear equations. In Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 6-8, 2002, San Francisco, CA, USA., pages 514-516, 2002. 1431
[BK15] Karl Bringmann and Marvin Kunnemann. Quadratic conditional lower bounds for string problems and dynamic time warping. In Proc. of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pages 79-97, 2015. 1914, 2191
[BKLW14] Maria-Florina Balcan, Vandana Kanchanapally, Yingyu Liang, and David Woodruff. Improved distributed principal component analysis. In Advances in Neural Information Processing Systems (NIPS). https://arxiv.org/pdf/ 1408.5823, 2014. 1318, 1784, 1788
[BKS12a] Surender Baswana, Sumeet Khurana, and Soumojit Sarkar. Fully dynamic randomized algorithms for graph spanners. ACM Transactions on Algorithms (TALG), 8(4):35, 2012. 2094
[BKS12b] Binay Bhattacharya, Tsunehiko Kameda, and Zhao Song. Computing minmax regret 1-median on a tree network with positive/negative vertex weights. In The 23rd International Symposium on Algorithms and Computation (ISAAC 2012), December 19-21, 2012 National Taiwan University, Taipei, Taiwan, 2012. xxix
[BKS12c] Binay Bhattacharya, Tsunehiko Kameda, and Zhao Song. Minmax regret 1-center on a path/cycle/tree. In The Sixth International Conference on Advanced Engineering Computing and Applications in Sciences(ADVCOMP 2012), September 23-28, 2012 - Barcelona, Spain, 2012. xxix
[BKS13] Paul Beame, Paraschos Koutris, and Dan Suciu. Communication steps for parallel query processing. In Proceedings of the 32nd ACM SIGMOD-SIGACT-SIGAI symposium on Principles of database systems, pages 273284. ACM, 2013. 2315, 2316, 2318, 2411
[BKS14] Binay Bhattacharya, Tsunehiko Kameda, and Zhao Song. Improved minmax regret 1-center algorithms for cactus networks with c cycles. In LATIN, pages 330-341, 2014. xxix
[BKS15a] Boaz Barak, Jonathan A Kelner, and David Steurer. Dictionary learning and tensor decomposition via the sum-of-squares method. In Proceedings
of the Forty-Seventh Annual ACM on Symposium on Theory of Computing (STOC), pages 143-151. ACM, https://arxiv.org/pdf/1407.1543, 2015. 58, 1470
[BKS15b] Binay Bhattacharya, Tsunehiko Kameda, and Zhao Song. Minmax regret 1-center algorithms for path/tree/unicycle/cactus networks. In Discrete Applied Mathematics, pages 195: 18-30, 2015. xxix
[BKW17] Karl Bringmann, Pavel Kolev, and David P. Woodruff. Approximation algorithms for \(\ell_{0}\)-low rank approximation. In Advances in Neural Information Processing Systems (NIPS), pages 6651-6662, 2017. 56, 1428, 1431, 1436
\(\left[\mathrm{BLG}^{+} 15\right]\) Aurélien Bellet, Yingyu Liang, Alireza Bagheri Garakani, Maria-Florina Balcan, and Fei Sha. A distributed frank-wolfe algorithm for communicationefficient sparse learning. In Proceedings of the 2015 SIAM International Conference on Data Mining (ICDM), pages 478-486. SIAM, https://arxiv. org/pdf/1404.2644, 2015. 1784, 1788
[BLLM19] Jaroslaw Blasiok, Patrick Lopatto, Kyle Luh, and Jake Marcinek. Sparse reconstruction from hadamard matrices: A lower bound. In arXiv preprint. https://arxiv.org/pdf/1903.12135.pdf, 2019. 925
[BLO05] James V Burke, Adrian S Lewis, and Michael L Overton. A robust gradient sampling algorithm for nonsmooth, nonconvex optimization. SIAM Journal on Optimization, 15(3):751-779, 2005. 20, 250
[BLS \({ }^{+}\)16a] Maria-Florina Balcan, Yingyu Liang, Le Song, David Woodruff, and Bo Xie. Communication efficient distributed kernel principal component analysis. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD), pages 725-734. ACM, https: //arxiv.org/pdf/1503.06858, 2016. 1784, 1788
[BLS \({ }^{+}\)16b] Maria-Florina Balcan, Yingyu Liang, Le Song, David Woodruff, and Bo Xie. Distributed kernel principal component analysis. KDD, 2016. 1318
[BLS \({ }^{+}\)19] Maria-Florina Balcan, Yingyu Liang, Zhao Song, David P. Woodruff, and Hongyang Zhang. Non-convex matrix completion and related problems via strong duality. JMLR, 2019. xxix
[BM86] Y. Bresler and A. Macovski. Exact maximum likelihood parameter estimation of superimposed exponential signals in noise. IEEE Transactions on Acoustics, Speech, and Signal Processing, 34(5):1081-1089, Oct 1986. 762
[BM16] Boaz Barak and Ankur Moitra. Noisy tensor completion via the sum-ofsquares hierarchy. In Proceedings of the 29th Conference on Learning Theory, COLT 2016, New York, USA, June 23-26, 2016, pages 417-445. https: //arxiv.org/pdf/1501.06521, 2016. 58, 1470
[BMD09] Christos Boutsidis, Michael W Mahoney, and Petros Drineas. An improved approximation algorithm for the column subset selection problem. In Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 968-977. Society for Industrial and Applied Mathematics, https://arxiv.org/pdf/0812.4293, 2009. 1114, 1433, 1476, 1604
[BNJ01] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent dirichlet allocation. In NIPS, 2001. 49, 1036, 1814, 1887
[BNR \({ }^{+}\)15] Guillaume Bouchard, Jason Naradowsky, Sebastian Riedel, Tim Rocktäschel, and Andreas Vlachos. Matrix and tensor factorization methods for natural language processing. In ACL (Tutorial Abstracts), pages 16-18, 2015. 57, 1469
[BO10] V. Braverman and R. Ostrovsky. Zero-one frequency laws. In STOC, pages 281-290, 2010. 1967
[Bou10] Petros T Boufounos. Reconstruction of sparse signals from distorted randomized measurements. In Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on, pages 3998-4001. IEEE, 2010. 596
[Bou11] Christos Boutsidis. Topics in matrix sampling algorithms. In Ph.D. Thesis. arXiv preprint. https://arxiv.org/pdf/1105.0709, 2011. 1501
[Bou14] Jean Bourgain. An improved estimate in the restricted isometry problem. In Geometric Aspects of Functional Analysis, pages 65-70. Springer, 2014. lxvii, lxviii, 16, 37, 38, 171, 623, 762, 918, 919, 923, 925, 1922
[BP12] Karl Bringmann and Konstantinos Panagiotou. Efficient sampling methods for discrete distributions. In International Colloquium on Automata, Languages, and Programming, pages 133-144. Springer, 2012. 1807
[BPD18] Ashish Bora, Eric Price, and Alexandros G. Dimakis. AmbientGAN: Generative models from lossy measurements. In International Conference on

Learning Representations. https://openreview.net/forum?id=Hy7fDog0b, 2018. 988
[BPR96] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. On the combinatorial and algebraic complexity of quantifier elimination. J. ACM, 43(6):10021045, 1996. 1046, 1052, 1053, 1131, 1339, 1483, 1506
[BPR05] Saugata Basu, Richard Pollack, and Marie-Francoise Roy. Algorithms in real algebraic geometry, volume 20033. Springer, 2005. 1051, 1052, 1065, 1066, 1131
[BR65] George R Blakley and Prabir Roy. A hölder type inequality for symmetric matrices with nonnegative entries. Proceedings of the American Mathematical Society, 16(6):1244-1245, 1965. 2284
[BR93] Avrim L Blum and Ronald L Rivest. Training a 3-node neural network is npcomplete. In Machine learning: From theory to applications (A preliminary version of this paper was appeared in NIPS 1989), pages 9-28. Springer, 1993. 258
[BR94] Mihir Bellare and John Rompel. Randomness-efficient oblivious sampling. In Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on, pages 276-287. IEEE, 1994. 2119
[BR19] Joshua Brakensiek and Aviad Rubinstein. Constant-factor approximation of near-linear edit distance in near-linear time. Manuscript, 2019. 2193
[Brä18] Petter Brändén. Hyperbolic polynomials and the kadison-singer problem. arXiv preprint arXiv:1809.03255, 2018. 1912, 1913
[BRB08] Yann-Ael Le Borgne, Sylvain Raybaud, and Gianluca Bontempi. Distributed principal component analysis for wireless sensor networks. Sensors, 2008. 1318, 1784, 1788
[Bro89] Andrei Broder. Generating random spanning trees. In Proceedings of the 30th annual Symposium on Foundations of Computer Science (FOCS), pages 442447, 1989. 29, 550
[BS80a] Jon Louis Bentley and James B Saxe. Decomposable searching problems i. static-to-dynamic transformation. Journal of Algorithms, 1(4):301-358, 1980. 2094
[BS80b] Jon Louis Bentley and James B Saxe. Generating sorted lists of random numbers. ACM Transactions on Mathematical Software (TOMS), 6(3):359364, 1980. 1807, 1818, 1819, 1820
[BS12] Markus Blaser and Chandan Saha. Lecture 6: Multipoint evaluation of a polynomial. Max-Planck-Institut für Informatik Class Notes, Computational Number Theory and Algebra, pages 1-4, 2012. 882
[BS15] Srinadh Bhojanapalli and Sujay Sanghavi. A new sampling technique for tensors. In arXiv preprint. https://arxiv.org/pdf/1502.05023, 2015. 57, 1469, 1797, 1798, 1804
[BS16] Aaron Bernstein and Cliff Stein. Faster fully dynamic matchings with small approximation ratios. In Proceedings of the Twenty-Seventh Annual ACMSIAM Symposium on Discrete Algorithms, pages 692-711. Society for Industrial and Applied Mathematics, 2016. 2094
[BSS12] Joshua Batson, Daniel A Spielman, and Nikhil Srivastava. Twice-ramanujan sparsifiers. In SIAM Journal on Computing, volume 41(6), pages 1704-1721. https://arxiv.org/pdf/0808.0163, 2012. 29, 551, 596, 607, 1560, 1913, 2031
[BSST13] Joshua Batson, Daniel A Spielman, Nikhil Srivastava, and Shang-Hua Teng. Spectral sparsification of graphs: theory and algorithms. Communications of the \(A C M, 56(8): 87-94,2013.5,552,1911\)
[BW14] Christos Boutsidis and David P Woodruff. Optimal cur matrix decompositions. In Proceedings of the 46 th Annual ACM Symposium on Theory of Computing (STOC), pages 353-362. ACM, https://arxiv.org/pdf/1405.7910, 2014. 1114, 1115, 1433, 1476, 1477, 1484, 1489, 1520, 1550, 1552, 1561, 1604, 2027
[BW18] Ainesh Bakshi and David Woodruff. Sublinear time low-rank approximation of distance matrices. In Advances in Neural Information Processing Systems, pages 3782-3792, 2018. 2027
[BWZ16] Christos Boutsidis, David P Woodruff, and Peilin Zhong. Optimal principal component analysis in distributed and streaming models. In Proceedings of the 48 th Annual ACM SIGACT Symposium on Theory of Computing
(STOC), pages 236-249. ACM, https://arxiv.org/pdf/1504.06729, 2016. 1310, 1318, 1319, 1338, 1511, 1779, 1784, 1788, 2027, 2094
[BYJK \({ }^{+}\)02] Ziv Bar-Yossef, TS Jayram, Ravi Kumar, D Sivakumar, and Luca Trevisan. Counting distinct elements in a data stream. In International Workshop on Randomization and Approximation Techniques in Computer Science, pages 1-10. Springer, 2002. 2094
[BYJKK04] Ziv Bar-Yossef, TS Jayram, Robert Krauthgamer, and Ravi Kumar. Approximating edit distance efficiently. In Proceedings of the Forty-Fifth Annual IEEE Symposium on Foundations of Computer Science, 2004. 1914, 2191, 2214, 2215
[BZI17] Sina Bittens, Ruochuan Zhang, and Mark A Iwen. A deterministic sparse FFT for functions with structured Fourier sparsity. Advances in Computational Mathematics, to appear., 2017. 38, 919
[CAKM16] Vincent Cohen-Addad, Philip N Klein, and Claire Mathieu. Local search yields approximation schemes for k -means and k -median in euclidean and minor-free metrics. In Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on, pages 353-364. IEEE, 2016. 2094
[Cau42] Augustin Louis Cauchy. Mémoire sur l'emploi du calcul des limites dans l'intégration des équations aux dérivées partielles. ., 1842. 233
[CBJC14] Sheng Cai, Mayank Bakshi, Sidharth Jaggi, and Minghua Chen. Super:

Sparse signals with unknown phases efficiently recovered. arXiv preprint arXiv:1401.4269, 2014. 635
[CC70] J Douglas Carroll and Jih-Jie Chang. Anaylsis of individual differences in multidimensional scaling via an n-way generalization of eckart-young decomposition. Psychometrika, 35(3):283-319, 1970. 57, 1469
[CC10] Cesar F Caiafa and Andrzej Cichocki. Generalizing the column-row matrix decomposition to multi-way arrays. Linear Algebra and its Applications, 433(3):557-573, 2010. 58, 1470, 1476
[CCAY \(\left.{ }^{+} 18\right]\) Xiang Cheng, Niladri S Chatterji, Yasin Abbasi-Yadkori, Peter L Bartlett, and Michael I Jordan. Sharp convergence rates for langevin dynamics in the nonconvex setting. arXiv preprint arXiv:1805.01648, 2018. 17, 172
[CCBJ18] Xiang Cheng, Niladri S Chatterji, Peter L Bartlett, and Michael I Jordan. Underdamped langevin memc: A non-asymptotic analysis. In COLT. arXiv preprint arXiv:1707.03663, 2018. 17, 18, 172, 174, 176
[CCF02] Moses Charikar, Kevin Chen, and Martin Farach-Colton. Finding frequent items in data streams. In Automata, Languages and Programming, pages 693-703. Springer, 2002. 35, 619, 624, 635, 688, 924, 996, 1009
[CCFC04] Moses Charikar, Kevin Chen, and Martin Farach-Colton. Finding frequent items in data streams. Theoretical Computer Science, 312(1):3-15, 2004. 2052, 2053
[CCGG98] Moses Charikar, Chandra Chekuri, Ashish Goel, and Sudipto Guha. Rounding via trees: deterministic approximation algorithms for group steiner trees and k-median. In Proceedings of the thirtieth annual ACM symposium on Theory of computing, pages 114-123. ACM, 1998. 2094
[CD05] Zizhong Chen and Jack J. Dongarra. Condition numbers of gaussian random matrices. SIAM Journal on Matrix Analysis and Applications, 27(3):603-620, 2005. 1045, 1056
[CDD09] Albert Cohen, Wolfgang Dahmen, and Ronald DeVore. Compressed sensing and best \(k\)-term approximation. Journal of the American mathematical society, 22(1):211-231, 2009. lxvii, lxviii, 37, 619, 923
[CDG \(\left.{ }^{+} 18\right]\) Diptarka Charkraborty, Debarati Das, Elazar Goldenberg, Michal Koucky, and Michael Saks. Approximating edit distance within constant factor in truly sub-quadratic time. In FOCS. https://iuuk.mff.cuni.cz/ ~diptarka/publications/approxEdit.pdf, 2018. 63, 1915, 2191, 2192, 2212, 2214, 2215, 2220, 2221, 2223, 2227, 2231, 2295, 2297, 2298
[CDHB09] Volkan Cevher, Marco F Duarte, Chinmay Hegde, and Richard Baraniuk. Sparse signal recovery using markov random fields. In Advances in Neural Information Processing Systems, pages 257-264, 2009. 988
[CDL13] Albert Cohen, Mark A Davenport, and Dany Leviatan. On the stability and accuracy of least squares approximations. Foundations of computational mathematics, 13(5):819-834, 2013. 761, 764
[CDMI \({ }^{+}\)13] Kenneth L Clarkson, Petros Drineas, Malik Magdon-Ismail, Michael W Mahoney, Xiangrui Meng, and David P Woodruff. The fast cauchy transform and faster robust linear regression. In Proceedings of the TwentyFourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 466-477. Society for Industrial and Applied Mathematics, https: //arxiv.org/pdf/1207.4684, 2013. 52, 1109, 1117, 1170, 1186, 1203, 1336, 1717, 2030
\(\left[\mathrm{CEM}^{+} 15\right]\) Michael B Cohen, Sam Elder, Cameron Musco, Christopher Musco, and Madalina Persu. Dimensionality reduction for k-means clustering and low rank approximation. In Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing (STOC), pages 163-172. ACM, https://arxiv.org/pdf/1410.6801, 2015. 48, 1036, 1561, 2080
[CEPR07] Raphaël Clifford, Klim Efremenko, Ely Porat, and Amir Rothschild. \(k\) mismatch with don't cares. Algorithms-ESA 2007, pages 151-162, 2007. 33, 617
[CF14] Emmanuel J Candès and Carlos Fernandez-Granda. Towards a mathematical theory of super-resolution. Communications on Pure and Applied Mathematics, 67(6):906-956, 2014. 41, 676, 686, 763
[CFHW17] Marek Cygan, Fedor V Fomin, Danny Hermelin, and Magnus Wahlström. Randomization in parameterized complexity (dagstuhl seminar 17041). In Dagstuhl Reports, volume 7:1. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017. 1898
\(\left[\mathrm{CFM}^{+} 18\right]\) Niladri S Chatterji, Nicolas Flammarion, Yi-An Ma, Peter L Bartlett, and Michael I Jordan. On the theory of variance reduction for stochastic gradient monte carlo. In \(I C M L, 2018.17,172,181\)
[CG19] Yuan Cao and Quanquan Gu. A Generalization Theory of Gradient Descent for Learning Over-parameterized Deep ReLU Networks. arXiv preprint arXiv:1902.01384, 2019. 357, 366
[CGCB14] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. arXiv preprint arXiv:1412.3555, 2014. 22, 352
[CGK \({ }^{+}\)17a] Flavio Chierichetti, Sreenivas Gollapudi, Ravi Kumar, Silvio Lattanzi, Rina Panigrahy, and David P Woodruff. Algorithms for \(\ell_{p}\) low rank approximation. In ICML. arXiv preprint arXiv:1705.06730, 2017. 56, 1371, 1374, 1375, 1377, 1428, 1433, 1434, 1435
[CGK \({ }^{+}\)17b] Flavio Chierichetti, Sreenivas Gollapudi, Ravi Kumar, Silvio Lattanzi, Rina Panigrahy, and David P Woodruff. Algorithms for \(\ell_{p}\) low rank approximation. In ICML. https://arxiv.org/pdf/1705.06730, 2017. 1334
[CGK \({ }^{+}\)17c] Flavio Chierichetti, Sreenivas Gollapudi, Ravi Kumar, Silvio Lattanzi, Rina Panigrahy, and David P Woodruff. Algorithms for \(\ell_{p}\) low rank approximation. In ICML. arXiv preprint arXiv:1705.06730, 2017. 2027, 2082
[CGKK18] Moses Charikar, Ofir Geri, Michael P. Kim, and William Kuszmaul. On estimating edit distance: Alignment, dimension reduction, and embeddings. In

45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic, pages 34:1-34:14, 2018. 2192
[CGL+19] Lijie Chen, Shafi Goldwasser, Kaifeng Lyu, Guy N Rothblum, and Aviad Rubinstein. Fine-grained complexity meets \(\mathrm{IP}=\mathrm{PSPACE}\). In \(S O D A\). arXiv preprint arXiv:1805.02351, 2019. 2191, 2214, 2216, 2218
[CGV13] Mahdi Cheraghchi, Venkatesan Guruswami, and Ameya Velingker. Restricted isometry of Fourier matrices and list decodability of random linear codes. SIAM Journal on Computing, 42(5):1888-1914, 2013. lxvii, lxviii, 37, 923, 925
[CH09] Graham Cormode and Marios Hadjieleftheriou. Finding the frequent items in streams of data. Communications of the \(A C M, 52(10): 97-105,2009\). 33, 617, 631
[Cha00] Bernard Chazelle. The discrepancy method. Cambridge University Press, Cambridge, 2000. Randomness and complexity. 30, 590
[Cha12] Sourav Chatterjee. Matrix estimation by universal singular value thresholding. pre-print, 2012. http://arxiv.org/abs/1212.1247. 48, 1036, 1043
[Che52] Herman Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. The Annals of Mathematical Statistics, pages 493-507, 1952. 25, 30, 502, 548, 590, 2141, 2505
[Che09] Ke Chen. On coresets for k-median and k-means clustering in metric and euclidean spaces and their applications. SIAM Journal on Computing, 39(3):923-947, 2009. 2089, 2094
[CHW12] Kenneth L. Clarkson, Elad Hazan, and David P. Woodruff. Sublinear optimization for machine learning. J. ACM, 59(5):23, 2012. 1799
[CI17] Mahdi Cheraghchi and Piotr Indyk. Nearly optimal deterministic algorithm for sparse walsh-hadamard transform. ACM Transactions on Algorithms (TALG), 13(3):34, 2017. 38, 919, 1904
[CIHB09] Volkan Cevher, Piotr Indyk, Chinmay Hegde, and Richard G Baraniuk. Recovery of clustered sparse signals from compressive measurements. Technical report, RICE UNIV HOUSTON TX DEPT OF ELECTRICAL AND COMPUTER ENGINEERING, 2009. 635
[Cip00] Barry A Cipra. The best of the 20th century: Editors name top 10 algorithms. SIAM news, 33(4):1-2, 2000. 1
\(\left[\mathrm{CJK}^{+} 04\right]\) Graham Cormode, Theodore Johnson, Flip Korn, Shan Muthukrishnan, Oliver Spatscheck, and Divesh Srivastava. Holistic udafs at streaming speeds. In Proceedings of the 2004 ACM SIGMOD international conference on Management of data, pages 35-46. ACM, 2004. 33, 617
[CK11] Imre Csiszar and János Körner. Information theory: coding theorems for discrete memoryless systems. Cambridge University Press, 2011. 1219
[CKK \({ }^{+}\)18] Michael B Cohen, Jonathan Kelner, Rasmus Kyng, John Peebles, Richard Peng, Anup B Rao, and Aaron Sidford. Solving directed laplacian systems in nearly-linear time through sparse LU factorizations. In 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS), pages 898909. IEEE, 2018. 6
\(\left[\mathrm{CKM}^{+} 11\right]\) Paul Christiano, Jonathan A Kelner, Aleksander Madry, Daniel A Spielman, and Shang-Hua Teng. Electrical flows, laplacian systems, and faster approximation of maximum flow in undirected graphs. In Proceedings of the forty-third annual ACM symposium on Theory of computing, pages 273-282. ACM, 2011. 6
\(\left[\mathrm{CKM}^{+}{ }^{+} 14\right]\) Michael B Cohen, Rasmus Kyng, Gary L Miller, Jakub W Pachocki, Richard Peng, Anup B Rao, and Shen Chen Xu. Solving sdd linear systems in nearly \(\mathrm{m} \log 1 / 2 \mathrm{n}\) time. In Proceedings of the forty-sixth annual ACM symposium on Theory of computing, pages 343-352. ACM, 2014. 6
[CKP \({ }^{+}\)17] Michael B Cohen, Jonathan Kelner, John Peebles, Richard Peng, Anup B Rao, Aaron Sidford, and Adrian Vladu. Almost-linear-time algorithms for markov chains and new spectral primitives for directed graphs. In Proceedings of the 49 th Annual ACM SIGACT Symposium on Theory of Computing (STOC), pages 410-419. ACM, 2017. 25, 548
[CKPS16] Xue Chen, Daniel M Kane, Eric Price, and Zhao Song. Fourier-sparse interpolation without a frequency gap. In Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on, pages 741-750. IEEE, 2016.
xxvi, 10, 40, 47, 164, 622, 920, 965, 977, 996, 1016, 1474, 1901, 1905, 2308, 2309, 2310, 2312, 2513, 2514
[CKSZ17] Volkan Cevher, Michael Kapralov, Jonathan Scarlett, and Amir Zandieh. An adaptive sublinear-time block sparse Fourier transform. In Proceedings of the 49 th Annual Symposium on the Theory of Computing (STOC). ACM, https://arxiv.org/pdf/1702.01286, 2017. 623, 636, 637
[CKY97] Sunglok Choi, Taemin Kim, and Wonpil Yu. Performance evaluation of ransac family. Journal of Computer Vision, 24(3):271-300, 1997. 990
[CL06] Fan Chung and Linyuan Lu. Concentration inequalities and martingale inequalities: a survey. Internet Mathematics, 3(1):79-127, 2006. 509
[Cla95] Kenneth L Clarkson. Las vegas algorithms for linear and integer programming when the dimension is small. Journal of the ACM (JACM), 42(2):488-499, 1995. 13, 5
[Cla05] Kenneth L Clarkson. Subgradient and sampling algorithms for \(\ell_{1}\) regression. In Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms (SODA), pages 257-266, 2005. 14, 52, 74, 1109, 1117, 1124, 1134, 1516, 2032
\(\left[\mathrm{CLK}^{+} 15\right]\) Fengyu Cong, Qiu-Hua Lin, Li-Dan Kuang, Xiao-Feng Gong, Piia Astikainen, and Tapani Ristaniemi. Tensor decomposition of eeg signals: a brief review. Journal of neuroscience methods, 248:59-69, 2015. 57, 1469
[CLLM12] Kai-min Chung, Henry Lam, Zhenming Liu, and Michael Mitzenmacher. Chernoff-Hoeffding bounds for Markov chains: Generalized and simplified. arXiv preprint arXiv:1201.0559, 2012. 26, 502, 537
\(\left[\mathrm{CLM}^{+} 15\right]\) Michael B Cohen, Yin Tat Lee, Cameron Musco, Christopher Musco, Richard Peng, and Aaron Sidford. Uniform sampling for matrix approximation. In Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science (ITCS), pages 181-190. ACM, https://arxiv.org/pdf/1408.5099, 2015. 1133, 1515
\(\left[\mathrm{CLM}^{+} 16\right]\) Michael B Cohen, Yin Tat Lee, Gary Miller, Jakub Pachocki, and Aaron Sidford. Geometric median in nearly linear time. In Proceedings of the fortyeighth annual ACM symposium on Theory of Computing (STOC), pages 9-21. ACM, 2016. 11, 3
\(\left[\mathrm{CLM}^{+} 18\right]\) Artur Czumaj, Jakub Lacki, Aleksander Madry, Slobodan Mitrovic, Krzysztof Onak, and Piotr Sankowski. Round compression for parallel matching algorithms. In STOC. https://arxiv.org/pdf/1707.03478, 2018. 2315, 2316
[CLMW11] Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? Journal of the ACM (JACM),58(3):11, 2011. 1, 56, 1300, 1334, 1370, 1428
[CLRS09] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. Introduction to Algorithms. MIT press, 2009. 2270
[CLS19] Michael B Cohen, Yin Tat Lee, and Zhao Song. Solving linear programs in the current matrix multiplication time. In Proceedings of the 51st Annual ACM Symposium on Theory of Computing (STOC). https://arxiv.org/ pdf/1810.07896.pdf, 2019. xxiv, lxvi, lxviii, 3, 4, 13, 15, 24, 35, 77, 79, 87, 88, 89, 119, 123, 126, 127, 131, 133, 135, 137, 138, 139, 144, 624, 1895, 1896, 2046, 2050, 2067, 2512
[CLS20] Sitan Chen, Jerry Li, and Zhao Song. Tbd. Manuscript, 2020. xxix
[CLZ17] Longxi Chen, Yipeng Liu, and Ce Zhu. Iterative block tensor singular value thresholding for extraction of low rank component of image data. In ICASSP 2017, 2017. 57, 58, 1469, 1470
[CM04] Graham Cormode and Shan Muthukrishnan. An improved data stream summary: the count-min sketch and its applications. In Proceedings of the Annual Conference on Foundations of Software Technology and Theoretical Computer Science, 2004. 35, 624, 2029, 2053
[CM06] Graham Cormode and S Muthukrishnan. Combinatorial algorithms for compressed sensing. In Information Sciences and Systems, 200640 th Annual Conference on, pages 198-201. IEEE, 2006. 35, 619, 624
\(\left[\mathrm{CMDL}^{+} 15\right]\) Andrzej Cichocki, Danilo Mandic, Lieven De Lathauwer, Guoxu Zhou, Qibin Zhao, Cesar Caiafa, and Huy Anh Phan. Tensor decompositions for signal processing applications: From two-way to multiway component analysis. IEEE Signal Processing Magazine, 32(2):145-163, 2015. 57, 1469
[CMM11] Alexandra Carpentier, Odalric-Ambrym Maillard, and Rémi Munos. Sparse recovery with brownian sensing. In Advances in Neural Information Processing Systems, pages 1782-1790, 2011. 988
[CMM15] Michael B Cohen, Cameron Musco, and Christopher Musco. Ridge leverage scores for low-rank approximation. arXiv preprint arXiv:1511.07263, 2015. 1926
[CMN96] Charles J Colbourn, Wendy J Myrvold, and Eugene Neufeld. Two algorithms for unranking arborescences. Journal of Algorithms, 20(2):268-281, 1996. 29, 550
[CMP16] Michael B Cohen, Cameron Musco, and Jakub Pachocki. Online row sampling. In International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX). arXiv preprint arXiv:1604.05448, 2016. 554
[CMS13] Michael S Crouch, Andrew McGregor, and Daniel Stubbs. Dynamic graphs in the sliding-window model. In European Symposium on Algorithms, pages 337-348. Springer, 2013. 2094
[CMSV17] Michael B Cohen, Aleksander Mądry, Piotr Sankowski, and Adrian Vladu. Negative-weight shortest paths and unit capacity minimum cost flow in \(\widetilde{O}\left(m^{10 / 7} \log W\right)\) time. In Proceedings of the Twenty-Eighth Annual ACMSIAM Symposium on Discrete Algorithms(SODA), pages 752-771. SIAM, 2017. 11, 3
[CMTV17] Michael B Cohen, Aleksander Madry, Dimitris Tsipras, and Adrian Vladu. Matrix scaling and balancing via box constrained newton's method and interior point methods. In Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on, pages 902-913. IEEE, 2017. 11, 3
[CNW15] Michael B Cohen, Jelani Nelson, and David P Woodruff. Optimal approximate matrix product in terms of stable rank. In Proceedings of the 43 rd International Colloquium on Automata, Languages and Programming (ICALP), Rome, Italy, July 12-15, 2016. https://arxiv.org/pdf/1507.02268, 2015. 770, 1562, 1565
[CNW16] Michael B Cohen, Jelani Nelson, and David P. Woodruff. Optimal approximate matrix product in terms of stable rank. In Proceedings of the 43 rd International Colloquium on Automata, Languages and Programming (ICALP 2016), Rome, Italy, July 12-15, arXiv preprint arXiv:1507.02268, 2016. 1922
[Coh16a] Michael B. Cohen. Nearly tight oblivious subspace embeddings by trace inequalities. In Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), Arlington, VA, USA, January 10-12, 2016, pages 278-287, 2016. 25, 52, 548, 1109, 1370, 1426, 1922, 1966
[Coh16b] Michael B Cohen. Ramanujan graphs in polynomial time. In 2016 IEEE 57 th Annual Symposium on Foundations of Computer Science (FOCS), pages 276-281. IEEE, 2016. 595, 596
[Com09] P. Comon. Tensor Decompositions, State of the Art and Applications. ArXiv e-prints, 2009. 57, 1469
[Con78] John B. Conway. Functions of one Complex Variable, volume I of Graduate Texts in Mathematics (Second Edition). Springer New York, Heidelberg, Berlin, 1978. 519
[COP03] Moses Charikar, Liadan O'Callaghan, and Rina Panigrahy. Better streaming algorithms for clustering problems. In Proceedings of the thirty-fifth annual ACM symposium on Theory of computing, pages 30-39. ACM, 2003. 2094
[Cox58] David R Cox. The regression analysis of binary sequences. Journal of the Royal Statistical Society. Series B (Methodological), pages 215-242, 1958. 14, 74
[CP15] Michael B. Cohen and Richard Peng. \(\ell_{p}\) row sampling by lewis weights. In Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing (STOC), STOC '15, pages 183-192, New York, NY, USA, 2015. https://arxiv.org/pdf/1412.0588. 52, 1109, 1116, 1133, 1185, 1200, 1253, 1338, 1341, 1488, 1515, 1717, 1966, 1971
[CPS18] Minmin Chen, Jeffrey Pennington, and Samuel S. Schoenholz. Dynamical isometry and a mean field theory of RNNs: Gating enables signal propagation in recurrent neural networks. arXiv:1806.05394, 2018. 358
[CR09] Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. Foundations of Computational Mathematics, 9(6):717772, 2009. 51, 1038
[CRT06a] Emmanuel J Candès, Justin Romberg, and Terence Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. IEEE Transactions on information theory, 52(2):489-509, 2006. 623, 917
[CRT06b] Emmanuel J Candes, Justin K Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. Communications on pure and applied mathematics, 59(8):1207-1223, 2006. lxvi, lxviii, 9, 33, \(34,35,617,619,624,762,917,989\)
[CS07] Maria Chudnovsky and Paul Seymour. The roots of the independence polynomial of a clawfree graph. J. Combin. Theory Ser. B, 97(3):350-357, 2007. 600
[Csi18] Dominik Csiba. Data sampling strategies in stochastic algorithms for empirical risk minimization. arXiv preprint arXiv:1804.00437, 2018. 13, 74
[CT65] James W Cooley and John W Tukey. An algorithm for the machine calculation of complex Fourier series. Mathematics of computation, 19(90):297-301, 1965. 37, 164, 676, 918
[CT05] Emmanuel Candes and Terence Tao. Decoding by linear programming. arXiv preprint math/0502327, 2005. 925
[CT06] Emmanuel J Candes and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? IEEE transactions on
information theory, 52(12):5406-5425, 2006. lxvii, lxviii, 37, 38, 917, 918, 919, 923, 925
[CV95] Corinna Cortes and Vladimir Vapnik. Support-vector networks. Machine learning, 20(3):273-297, 1995. 14, 74
[CV14] Joon Hee Choi and S. Vishwanathan. Dfacto: Distributed factorization of tensors. In NIPS, pages 1296-1304, 2014. 1797
[CV15] Nicoló Colombo and Nikos Vlassis. Fastmotif: spectral sequence motif discovery. Bioinformatics, pages 2623-2631, 2015. 57, 1469
[CVMBB14] Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. arXiv preprint arXiv:1409.1259, 2014. 22, 352
\(\left[\mathrm{CVMG}^{+} 14\right]\) Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In \(E M N L P\), pages \(1724-1734,2014.353\)
[CW87] Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. In Proceedings of the nineteenth annual ACM symposium on Theory of computing(STOC), pages 1-6. ACM, 1987. 12, 16, 4, 78, 761, \(775,792,1133,1426,1468,1515\)
[CW90] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. In Journal of Symbolic Computation, volume 9, pages 251-280, 1990. 1921, 1927
[CW09] Kenneth L. Clarkson and David P. Woodruff. Numerical linear algebra in the streaming model. In Proceedings of the 41 st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31-June 2, 2009, pages 205-214, 2009. 52, 1045, 1109, 1125, 1163, 1310, 1779, 2044
[CW13] Kenneth L. Clarkson and David P. Woodruff. Low rank approximation and regression in input sparsity time. In Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013, pages 81-90. https://arxiv.org/pdf/1207.6365, 2013. 49, 52, 13, 770, 1037, 1045, 1057, 1083, 1085, 1105, 1109, 1119, 1337, 1338, 1370, 1401, 1426, 1468, 1489, 1490, 1499, 1511, 1512, 1535, 1555, 1566, 1606, 1609, 1610, 1758, 1781, 1786, 1791, 1922, 1927, 1966, 1999, 2027, 2029, 2037, 2042, 2052, 2075, 2076, 2078, 2082
[CW15a] Kenneth L Clarkson and David P Woodruff. Input sparsity and hardness for robust subspace approximation. In 2015 IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS), pages 310-329. IEEE, https: //arxiv.org/pdf/1510.06073, 2015. 54, 56, 246, 993, 1046, 1112, 1131, \(1428,1474,1506,1652,1740,1741,1743,1744,1966,1967,1971,2027\)
[CW15b] Kenneth L Clarkson and David P Woodruff. Sketching for m-estimators: A unified approach to robust regression. In Proceedings of the Twenty-Sixth

Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 921939. SIAM, 2015. 52, 246, 993, 1109, 1338, 1455, 1456, 1740, 1747, 1748, 1749, 1966, 1967, 1971, 2027
[CWW19] Kenneth L. Clarkson, Ruosong Wang, and David P. Woodruff. Dimensionality reduction for tukey regression. In \(I C M L\), 2019. 1431
\(\left[\mathrm{CWZ}^{+} 17\right]\) Changyou Chen, Wenlin Wang, Yizhe Zhang, Qinliang Su, and Lawrence Carin. A convergence analysis for a class of practical variance-reduction stochastic gradient mcmc. arXiv preprint arXiv:1709.01180, 2017. 181
[CYYM14] Kai-Wei Chang, Scott Wen-tau Yih, Bishan Yang, and Chris Meek. Typed tensor decomposition of knowledge bases for relation extraction. In Empirical Methods in Natural Language Processing (EMNLP), pages 1568-1579, 2014. 57, 1469
[Dal17] Arnak S Dalalyan. Theoretical guarantees for approximate sampling from smooth and log-concave densities. Journal of the Royal Statistical Society: Series \(B\) (Statistical Methodology), 79(3):651-676, 2017. 16, 17, 18, 21, 171, \(172,174,251,253,254,262,351,356\)
[Dan16] Amit Daniely. Complexity theoretic limitations on learning halfspaces. In Proceedings of the forty-eighth annual ACM symposium on Theory of Computing (STOC), pages 105-117. ACM, 2016. 258
[DB13] Marco F Duarte and Richard G Baraniuk. Spectral compressive sensing.

Applied and Computational Harmonic Analysis, 35(1):111-129, 2013. 41, 676, 686, 763
[DB14] Alexandre Défossez and Francis Bach. Constant step size least-mean-square: Bias-variance trade-offs and optimal sampling distributions. arXiv preprint arXiv:1412.0156, 2014. 13, 74
[DB16] Aymeric Dieuleveut and Francis Bach. Nonparametric stochastic approximation with large step-sizes. The Annals of Statistics, 44(4):1363-1399, 2016. 13, 74
[DBLJ14] Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In Advances in neural information processing systems, pages 1646-1654, 2014. 13, 74
[DCWY18] Raaz Dwivedi, Yuansi Chen, Martin J Wainwright, and Bin Yu. Log-concave sampling: Metropolis-hastings algorithms are fast! In COLT. arXiv preprint arXiv:1801.02309, 2018. 16, 18, 171, 174
[ \(\left.\mathrm{DDH}^{+} 09\right]\) Anirban Dasgupta, Petros Drineas, Boulos Harb, Ravi Kumar, and Michael W Mahoney. Sampling algorithms and coresets for \(\ell_{p}\) regression. SIAM Journal on Computing, 38(5):2060-2078, 2009. 14, 74, 1117, 1124, 1134, 1190, 1406, 1407, 1408, 1409, 1516, 1717, 1966, 1971, 1973, 2033, 2034, 2035, 2044, 2045, 2050, 2051, 2066, 2067
[DDT \({ }^{+}\)08] Marco F Duarte, Mark A Davenport, Dharmpal Takhar, Jason N Laska, Ting Sun, Kevin F Kelly, and Richard G Baraniuk. Single-pixel imaging via compressive sampling. IEEE signal processing magazine, 25(2):83-91, 2008. 33, 617, 630
[DDTS06] David Leigh Donoho, Iddo Drori, Yaakov Tsaig, and Jean-Luc Starck. Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit. Department of Statistics, Stanford University, 2006. 38, 918
[DFK \(\left.{ }^{+} 04\right]\) Petros Drineas, Alan M. Frieze, Ravi Kannan, Santosh Vempala, and V. Vinay. Clustering large graphs via the singular value decomposition. Machine Learning, 56(1-3):9-33, 2004. 48, 1036
[DFS16] Amit Daniely, Roy Frostig, and Yoram Singer. Toward deeper understanding of neural networks: The power of initialization and a dual view on expressivity. In Advances in Neural Information Processing Systems (NeurIPS), pages 2253-2261, 2016. 358
[DG94] Jean-Pierre Dedieu and R. J. Gregorac. Corrigendum: "Obreschkoff's theorem revisited: what convex sets are contained in the set of hyperbolic polynomials?" [J. Pure Appl. Algebra 81 (1992), no. 3, 269-278; MR1179101 (93g:12001)] by Dedieu. J. Pure Appl. Algebra, 93(1):111-112, 1994. 600
[DG04] Jeffrey Dean and Sanjay Ghemawat. MapReduce: Simplified data processing on large clusters. In \(O S D I, 2004.2315\)
[DG08] Jeffrey Dean and Sanjay Ghemawat. MapReduce: simplified data processing on large clusters. Communications of the ACM, 51(1):107-113, 2008. 2315
[DGL+99] Yevgeniy Dodis, Oded Goldreich, Eric Lehman, Sofya Raskhodnikova, Dana Ron, and Alex Samorodnitsky. Improved testing algorithms for monotonicity. In Randomization, Approximation, and Combinatorial Optimization. Algorithms and Techniques, pages 97-108. Springer, 1999. 2218
[DGR96] A Dutt, M Gu, and V Rokhlin. Fast algorithms for polynomial interpolation, integration, and differentiation. SIAM Journal on Numerical Analysis, 33(5):1689-1711, 1996. 198
[DGRS15] Alexandros G Dimakis, Anna Gál, Ankit Singh Rawat, and Zhao Song. Batch codes through dense graphs without short cycles. In ISIT, 2015. xxix
[DHKP97] Martin Dietzfelbinger, Torben Hagerup, Jyrki Katajainen, and Martti Penttonen. A reliable randomized algorithm for the closest-pair problem. Journal of Algorithms, 25(1):19-51, 1997. 1334
[DJS \({ }^{+}\)19] Huaian Diao, Rajesh Jayaram, Zhao Song, Wen Sun, and David P. Woodruff. Optimal sketching for kronecker product regression and low rank approximation. Manuscript, 2019. xxix, 2024, 2514, 2515, 2516
[DK17] Arnak S Dalalyan and Avetik G Karagulyan. User-friendly guarantees for the langevin monte carlo with inaccurate gradient. arXiv preprint arXiv:1710.00095, 2017. 16, 17, 171, 172, 181
[DKM06] Petros Drineas, Ravi Kannan, and Michael W Mahoney. Fast monte carlo algorithms for matrices ii: Computing a low-rank approximation to a matrix. SIAM Journal on computing, 36(1):158-183, 2006. 1338
[DKP \({ }^{+}\)17] David Durfee, Rasmus Kyng, John Peebles, Anup B. Rao, and Sushant Sachdeva. Sampling random spanning trees faster than matrix multiplication. In Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC), pages 730-742, 2017. 29, 550, 555, 567, 568
[DKR02] Petros Drineas, Iordanis Kerenidis, and Prabhakar Raghavan. Competitive recommendation systems. In STOC, 2002. 48, 1036
[DLDM98] Lieven De Lathauwer and Bart De Moor. From matrix to tensor: Multilinear algebra and signal processing. In Institute of Mathematics and Its Applications Conference Series, volume 67, pages 1-16. Citeseer, 1998. 57, 1469
[DLL \(\left.{ }^{+} 19\right]\) Simon S Du, Jason D Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In ICML. https://arxiv.org/pdf/1811.03804, 2019. 256, 1909, 2140
[DLT \(\left.{ }^{+} 18\right]\) Simon S. Du, Jason D. Lee, Yuandong Tian, Barnabás Póczos, and Aarti Singh. Gradient descent learns one-hidden-layer CNN: don't be afraid of spurious local minima. In ICML, 2018. 21, 250, 257, 351, 2140
[DLTY06] Ke-xue Dai, Guo-hui Li, Dan Tu, and Jian Yuan. Prospects and current studies on background subtraction techniques for moving objects detection from surveillance video. Journal of Image and Graphics, 7(002), 2006. 630
[DM08] Wei Dai and Olgica Milenkovic. Subspace pursuit for compressive sensing: Closing the gap between performance and complexity. Technical report, ILLINOIS UNIV AT URBANA-CHAMAPAIGN, 2008. 38, 918
[DM16] Alain Durmus and Eric Moulines. High-dimensional bayesian inference via the unadjusted langevin algorithm. arXiv preprint arXiv:1605.01559, 2016. 205, 209
[DMIMW12] Petros Drineas, Malik Magdon-Ismail, Michael W Mahoney, and David P Woodruff. Fast approximation of matrix coherence and statistical leverage. Journal of Machine Learning Research, 13(Dec):3475-3506, 2012. 1338, 1468, 1568, 1593, 1907, 1922, 1926
[DMM06a] Petros Drineas, Michael W Mahoney, and S Muthukrishnan. Sampling algorithms for \(\ell_{2}\) regression and applications. In Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm, pages 1127-1136. Society for Industrial and Applied Mathematics, 2006. 1907
[DMM06b] Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Subspace sampling and relative-error matrix approximation: Column-based methods. In Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 9th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX 2006 and 10th International Workshop on Randomization and Computation, RANDOM 2006, Barcelona, Spain, August 28-30 2006, Proceedings, pages 316326, 2006. 52, 1109, 1433, 1468
[DMM06c] Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Subspace sampling and relative-error matrix approximation: Column-row-based methods. In Algorithms - ESA 2006, 14th Annual European Symposium, Zurich, Switzerland, September 11-13, 2006, Proceedings, pages 304-314, 2006. 52, 1109, 1433, 1468, 1922
[DMM08] Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Relativeerror CUR matrix decompositions. SIAM J. Matrix Analysis Applications, \(30(2): 844-881,2008.52,1109,1114,1433,1468,1476,1604\)
\(\left[\mathrm{DMM}^{+}\right.\)17] Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. arXiv preprint arXiv:1710.01688, 2017. 257
[DMMS11] Petros Drineas, Michael W Mahoney, S Muthukrishnan, and Tamás Sarlós. Faster least squares approximation. Numerische Mathematik, 117(2):219249, 2011. 1337, 1907, 1922, 1931
[Don06] David L. Donoho. Compressed sensing. IEEE Trans. Information Theory, \(52(4): 1289-1306,2006\). lxvi, lxviii, \(9,33,34,35,617,619,624,917\)
[DPFJ16] Kévin Degraux, Gabriel Peyré, Jalal Fadili, and Laurent Jacques. Sparse support recovery with non-smooth loss functions. In Advances in Neural Information Processing Systems, pages 4269-4277, 2016. 988
[DPPR17] David Durfee, John Peebles, Richard Peng, and Anup B. Rao. Determinantpreserving sparsification of SDDM matrices with applications to counting
and sampling spanning trees. In Proceedings of the 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), 2017. 29, 550, 555
[DR10] Amit Deshpande and Luis Rademacher. Efficient volume sampling for row/column subset selection. In 2010 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS), pages 329-338. IEEE, https: //arxiv.org/pdf/1004.4057, 2010. 1114, 1476, 1604
[DRD18] Arnak S Dalalyan and Lionel Riou-Durand. On sampling from a log-concave density using kinetic langevin diffusions. arXiv preprint arXiv:1807.09382, 2018. 16, 17, 171, 172
[DRW \(\left.{ }^{+} 16\right]\) Kumar Avinava Dubey, Sashank J Reddi, Sinead A Williamson, Barnabas Poczos, Alexander J Smola, and Eric P Xing. Variance reduction in stochastic gradient langevin dynamics. In Advances in neural information processing systems, pages 1154-1162, 2016. 181
[DS13] Alexander Munro Davie and Andrew James Stothers. Improved bound for complexity of matrix multiplication. Proceedings of the Royal Society of Edinburgh Section A: Mathematics, 143(2):351-369, 2013. 12, 4
[DSL08] Vin De Silva and Lek-Heng Lim. Tensor rank and the ill-posedness of the best low-rank approximation problem. SIAM Journal on Matrix Analysis and Applications, 30(3):1084-1127, 2008. 59, 1471, 1497, 1652
[DSM \(\left.{ }^{+} 16\right]\) Alain Durmus, Umut Simsekli, Eric Moulines, Roland Badeau, and Gaël Richard. Stochastic gradient richardson-romberg markov chain monte carlo.

In Advances in Neural Information Processing Systems, pages 2047-2055, 2016. 181
[DSS16] Amit Daniely and Shai Shalev-Shwartz. Complexity theoretic limitations on learning dnf's. In Conference on Learning Theory (COLT), pages 815-830, 2016. 258
[DSSW18] Huaian Diao, Zhao Song, Wen Sun, and David P. Woodruff. Sketching for kronecker product regression and p-splines. In AISTATS. https://arxiv. org/pdf/1712.09473, 2018. xxix, 2024, 2025, 2026, 2028, 2038, 2039, 2040, 2041, 2066
[DSWY19] Huaian Diao, Zhao Song, David P. Woodruff, and Xin Yang. Total least squares regression in input sparsity time. Manuscript, 2019. xxix
[DTMR18] Sarah Dean, Stephen Tu, Nikolai Matni, and Benjamin Recht. Safely learning to control the constrained linear quadratic regulator. arXiv preprint arXiv:1809.10121, 2018. 257
[Dun10] Mark Dunster. Legendre and Related Functions. Handbook of Mathematical Functions, Cambridge University Press, 2010. 785
[DV06] Amit Deshpande and Santosh Vempala. Adaptive sampling and fast low-rank matrix approximation. In Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, pages 292-303. Springer, 2006. 1338
[DV07] Amit Deshpande and Kasturi R. Varadarajan. Sampling-based dimension reduction for subspace approximation. In Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007, pages 641-650, 2007. 54, 1112, 1474
[DVTV09] Amit Deshpande, Kasturi R. Varadarajan, Madhur Tulsiani, and Nisheeth K. Vishnoi. Algorithms and hardness for subspace approximation. CoRR, abs/0912.1403, 2009. 56, 1428
[DW17] Huaian Diao and David P. Woodruff. Kronecker product and spline regression. manuscript, 2017. 1518
[DWG \(\left.{ }^{+} 13\right]\) Denisa Duma, Mary Wootters, Anna C Gilbert, Hung Q Ngo, Atri Rudra, Matthew Alpert, Timothy J Close, Gianfranco Ciardo, and Stefano Lonardi. Accurate decoding of pooled sequenced data using compressed sensing. In International Workshop on Algorithms in Bioinformatics, pages 70-84. Springer, 2013. 33, 617
[DZHZ06] Chris Ding, Ding Zhou, Xiaofeng He, and Hongyuan Zha. R1-pca: rotational invariant \(\ell_{1}\)-norm principal component analysis for robust subspace factorization. In Proceedings of the 23rd international conference on Machine learning, pages 281-288. ACM, 2006. l, lxxv, 1286, 1287, 1288, 1298
[DZPS19] Simon S Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. In \(I C L R\). https: //arxiv.org/pdf/1810.02054, 2019. 257, 352, 1909, 2140, 2142, 2143, 2144, \(2145,2146,2150,2152,2154,2157,2159,2161,2174\)
[ECG \(\left.{ }^{+} 09\right]\) Yaniv Erlich, Kenneth Chang, Assaf Gordon, Roy Ronen, Oron Navon, Michelle Rooks, and Gregory J Hannon. DNA sudoku-harnessing highthroughput sequencing for multiplexed specimen analysis. Genome research, 19(7):1243-1253, 2009. 33, 617
[EFW10] Alexandre V. Evfimievski, Ronald Fagin, and David P. Woodruff. Epistemic privacy. J. ACM, 58(1):2, 2010. 1046
[EIM11] Alina Ene, Sungjin Im, and Benjamin Moseley. Fast clustering using MapReduce. In Proceedings of the 17 th ACM SIGKDD international conference on Knowledge discovery and data mining, pages 681-689. ACM, 2011. 2315
[EKK \(\left.{ }^{+} 00\right]\) Funda Ergün, Sampath Kannan, S Ravi Kumar, Ronitt Rubinfeld, and Mahesh Viswanathan. Spot-checkers. Journal of Computer and System Sciences, 60(3):717-751, 2000. 2218
[Elm90] Jeffrey L Elman. Finding structure in time. Cognitive science, 14(2):179-211, 1990. 22, 352
[ELMM19] Yonina Eldar, Jerry Li, Cameron Musco, and Christopher Musco. The sample complexity of toeplitz covariance estimation. In manuscript, 2019. 1906
[EM06] Paul HC Eilers and Brian D Marx. Multidimensional density smoothing with p-splines. In Proceedings of the 21st international workshop on statistical modelling, 2006. 2024
[End10] Joachim HG Ender. On compressive sensing applied to radar. Signal Processing, 90(5):1402-1414, 2010. 622
[ES09] Lars Eldén and Berkant Savas. A newton-grassmann method for computing the best multilinear rank-(r1,r2,r3) approximation of a tensor. SIAM J. Matrix Analysis Applications, 31(2):248-271, 2009. 57, 1469
[ESAZ09] Yaniv Erlich, Noam Shental, Amnon Amir, and Or Zuk. Compressed sensing approach for high throughput carrier screen. In Communication, Control, and Computing, 2009. Allerton 2009. 47th Annual Allerton Conference on, pages 539-544. IEEE, 2009. 33, 617
[EV03] Cristian Estan and George Varghese. New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice. ACM Transactions on Computer Systems (TOCS), 21(3):270-313, 2003. 34, 618
[Exc13] Stack Exchange. Low-rank matrix approximation in terms of entry-wise \(\ell_{1}\) norm, 2013. 53, 1110
[FB17] C. Daniel Freeman and Joan Bruna. Topology and geometry of half-rectified network optimization. In ICLR, 2017. 351
[FEGK13] Ahmed K Farahat, Ahmed Elgohary, Ali Ghodsi, and Mohamed S Kamel. Distributed column subset selection on mapreduce. In 2013 IEEE 13th International Conference on Data Mining (ICDM), pages 171-180. IEEE, 2013. 1318, 1433, 1476, 1784, 1788
[Fei02] Uriel Feige. Relations between average case complexity and approximation complexity. In Proceedings of the thiry-fourth annual ACM symposium on Theory of computing(STOC), pages 534-543. ACM, 2002. 1041, 1081, 1671
[Fel80] H. J. Fell. On the zeros of convex combinations of polynomials. Pacific J. Math., 89(1):43-50, 1980. 600
[FFSS07] Dan Feldman, Amos Fiat, Micha Sharir, and Danny Segev. Bi-criteria lineartime approximations for generalized k-mean/median/center. In Proceedings of the 23rd ACM Symposium on Computational Geometry, Gyeongju, South Korea, June 6-8, 2007, pages 19-26, 2007. 1474
[FGKS15] Roy Frostig, Rong Ge, Sham M Kakade, and Aaron Sidford. Competing with the empirical risk minimizer in a single pass. In Conference on learning theory (COLT), pages 728-763, 2015. 13, 15, 74, 77
[FGRW12] Vitaly Feldman, Venkatesan Guruswami, Prasad Raghavendra, and Yi Wu. Agnostic learning of monomials by halfspaces is hard. SIAM Journal on Computing, 41(6):1558-1590, 2012. 13, 74
[FGV17] V. Feldman, C. Guzmán, and S. S. Vempala. Statistical query algorithms for mean vector estimation and stochastic convex optimization. In \(S O D A\), pages 1265-1277, 2017. 1967
[FHHP11] Wai Shing Fung, Ramesh Hariharan, Nicholas JA Harvey, and Debmalya Panigrahi. A general framework for graph sparsification. In Proceedings of the forty-third annual ACM symposium on Theory of computing (STOC), pages 71-80. ACM, 2011. 551, 552, 556, 559, 561
[Fis04] Eldar Fischer. The art of uninformed decisions: A primer to property testing.

Current Trends in Theoretical Computer Science: The Challenge of the New Century, 1:229-264, 2004. 2218
[FIS05] Gereon Frahling, Piotr Indyk, and Christian Sohler. Sampling in dynamic data streams and applications. In Symposium on Computational Geometry 2005, pages 142-149, 2005. 2135
[FKG71] Cees M Fortuin, Pieter W Kasteleyn, and Jean Ginibre. Correlation inequalities on some partially ordered sets. Communications in Mathematical Physics, 22(2):89-103, 1971. 28, 549
[FKM \(\left.{ }^{+} 05\right]\) Joan Feigenbaum, Sampath Kannan, Andrew McGregor, Siddharth Suri, and Jian Zhang. On graph problems in a semi-streaming model. Theoretical Computer Science, 348(2-3):207-216, 2005. 2094
[FKV04] Alan M. Frieze, Ravi Kannan, and Santosh Vempala. Fast monte-carlo algorithms for finding low-rank approximations. J. ACM, 51(6):1025-1041, 2004. 1468
[FL11] Dan Feldman and Michael Langberg. A unified framework for approximating and clustering data. In Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011, pages 569-578, 2011. 54, 1112, 2089, 2090, 2094, 2098
[FL12] Albert Fannjiang and Wenjing Liao. Coherence pattern-guided compressive sensing with unresolved grids. SIAM Journal on Imaging Sciences, 5(1):179202, 2012. 763
[FLP15] Dimitris Fotakis, Michael Lampis, and Vangelis Th Paschos. Sub-exponential approximation schemes for csps: From dense to almost sparse. arXiv preprint arXiv:1507.04391, 2015. 1269, 1274
[FM92] Tomás Feder and Milena Mihail. Balanced matroids. In Proceedings of the twenty-fourth annual ACM symposium on Theory of computing (STOC), pages 26-38. ACM, 1992. 555
[FMMN11] Shmuel Friedland, V Mehrmann, A Miedlar, and M Nkengla. Fast low rank approximations of matrices and tensors. Electron. J. Linear Algebra, 22(10311048):462, 2011. 58, 1470, 1476
[FMPS13] Shmuel Friedland, Volker Mehrmann, Renato Pajarola, and Susanne K. Suter. On best rank one approximation of tensors. Numerical Lin. Alg. with Applic., 20(6):942-955, 2013. 57, 1469
[FMS07] Dan Feldman, Morteza Monemizadeh, and Christian Sohler. A ptas for kmeans clustering based on weak coresets. In Proceedings of the twenty-third annual symposium on Computational geometry, pages 11-18. ACM, 2007. 2094
[FMS \({ }^{+}\)10] Jon Feldman, S. Muthukrishnan, Anastasios Sidiropoulos, Clifford Stein, and Zoya Svitkina. On distributing symmetric streaming computations. ACM Transactions on Algorithms, 6(4), 2010. Previously in SODA'08. 2315
[FMSW10] Dan Feldman, Morteza Monemizadeh, Christian Sohler, and David P. Woodruff. Coresets and sketches for high dimensional subspace approxima-
tion problems. In Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2010, Austin, Texas, USA, January 17-19, 2010, pages 630-649, 2010. 54, 1112
[Fou11] Simon Foucart. Hard thresholding pursuit: an algorithm for compressive sensing. SIAM Journal on Numerical Analysis, 49(6):2543-2563, 2011. 38, 918
[Fre75] Michael L Fredman. On computing the length of longest increasing subsequences. Discrete Mathematics, 11(1):29-35, 1975. 2218
[FRS16] Zachary Friggstad, Mohsen Rezapour, and Mohammad R Salavatipour. Local search yields a ptas for k -means in doubling metrics. In Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on, pages 365-374. IEEE, 2016. 2094
[FS97] Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. Journal of computer and system sciences, 55(1):119-139, 1997. 14, 74
[FS99] Roger Fischlin and Jean-Pierre Seifert. Tensor-based trapdoors for cvp and their application to public key cryptography. Cryptography and Coding, pages 801-801, 1999. 57, 1469
[FS05] Gereon Frahling and Christian Sohler. Coresets in dynamic geometric data streams. In Proceedings of the thirty-seventh annual ACM symposium on Theory of computing (STOC), pages 209-217. ACM, 2005. 2090, 2094, 2136
[FS12] Dan Feldman and Leonard J Schulman. Data reduction for weighted and outlier-resistant clustering. In Proceedings of the twenty-third annual ACMSIAM symposium on Discrete Algorithms, pages 1343-1354. Society for Industrial and Applied Mathematics, 2012. 2094, 2136, 2137
[FSGM \(\left.{ }^{+} 99\right]\) Min Fang, Narayanan Shivakumar, Hector Garcia-Molina, Rajeev Motwani, and Jeffrey D Ullman. Computing iceberg queries efficiently. In Internaational Conference on Very Large Databases (VLDB'98), New York, August 1998. Stanford InfoLab, 1999. 33, 617
[FSS13] Dan Feldman, Melanie Schmidt, and Christian Sohler. Turning big data into tiny data: Constant-size coresets for k-means, pca and projective clustering. In Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1434-1453. Society for Industrial and Applied Mathematics, 2013. 48, 1036
[FT07] Shmuel Friedland and Anatoli Torokhti. Generalized rank-constrained matrix approximations. SIAM Journal on Matrix Analysis and Applications, 29(2):656-659, 2007. 1135, 1508
[FT15] Shmuel Friedland and Venu Tammali. Low-rank approximation of tensors. In Numerical Algebra, Matrix Theory, Differential-Algebraic Equations and Control Theory, pages 377-411. Springer, 2015. 57, 58, 1469, 1470, 1476
[Gan05] Sumit Ganguly. Counting distinct items over update streams. In International Symposium on Algorithms and Computation, pages 505-514. Springer, 2005. 2092, 2117
[Gao08] Junbin Gao. Robust \(\ell_{1}\) principal component analysis and its bayesian variational inference. Neural computation, 20(2):555-572, 2008. 53, 516, 1111
[GBT84] Harold N Gabow, Jon Louis Bentley, and Robert E Tarjan. Scaling and related techniques for geometry problems. In Proceedings of the sixteenth annual ACM symposium on Theory of computing (STOC), pages 135-143. ACM, 1984. 1334
[GG11] Nicolas Gillis and François Glineur. Low-rank matrix approximation with weights or missing data is np-hard. SIAM Journal on Matrix Analysis and Applications, 32(4):1149-1165, 2011. 51, 1038, 1081, 1334
[GGH14] Quanquan Gu, Huan Gui, and Jiawei Han. Robust tensor decomposition with gross corruption. In Advances in Neural Information Processing Systems(NIPS), pages 1422-1430, 2014. 58, 1470
[GGI \({ }^{+} 02\) ] Anna C Gilbert, Sudipto Guha, Piotr Indyk, S Muthukrishnan, and Martin Strauss. Near-optimal sparse Fourier representations via sampling. In Proceedings of the thiry-fourth annual ACM symposium on Theory of computing, pages 152-161. ACM, 2002. 38, 40, 46, 623, 636, 676, 682, 688, 759, 762, 919
[GHK15] Rong Ge, Qingqing Huang, and Sham M Kakade. Learning mixtures of gaussians in high dimensions. In Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing (STOC), pages 761-770. ACM, https://arxiv.org/pdf/1503.00424, 2015. 57, 1469
[GI10] Anna Gilbert and Piotr Indyk. Sparse recovery using sparse matrices. Proceedings of the IEEE, 98(6):937-947, 2010. 8, 34, 618
[Gi198] David Gillman. A chernoff bound for random walks on expander graphs. SIAM Journal on Computing, 27(4):1203-1220, 1998. 26, 27, 502, 503
[GK09] Rahul Garg and Rohit Khandekar. Gradient descent with sparsification: an iterative algorithm for sparse recovery with restricted isometry property. In ICML, volume 9, pages 337-344, 2009. 38, 918
[GKK12] Ashish Goel, Michael Kapralov, and Sanjeev Khanna. On the communication and streaming complexity of maximum bipartite matching. In Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms, pages 468-485. SIAM, 2012. 2094
[GKKT17] Surbhi Goel, Varun Kanade, Adam Klivans, and Justin Thaler. Reliably learning the ReLU in polynomial time. In Conference on Learning Theory (COLT), 2017. 258
[GKM18] Surbhi Goel, Adam Klivans, and Raghu Meka. Learning one convolutional layer with overlapping patches. In \(I C M L\). arXiv preprint arXiv:1802.02547, 2018. 257, 1451
[GKP12] Ashish Goel, Michael Kapralov, and Ian Post. Single pass sparsification in the streaming model with edge deletions. arXiv preprint arXiv:1203.4900, 2012. 2094
[GKU19] Mohsen Ghaffari, Fabian Kuhn, and Jara Uitto. Conditional hardness results for massively parallel computation from distributed lower bounds. In 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS). IEEE, 2019. 1917
[GL89] Oded Goldreich and Leonid A Levin. A hard-core predicate for all oneway functions. In Proceedings of the twenty-first annual ACM symposium on Theory of computing, pages 25-32. ACM, 1989. 38, 919
[GL04] Andreas Goerdt and André Lanka. An approximation hardness result for bipartite clique. In Electronic Colloquium on Computational Complexity, Report, volume 48. https://eccc.weizmann.ac.il/report/2004/048/, 2004. 1081, 1672
[GLM17] Rong Ge, Jason D. Lee, and Tengyu Ma. Learning one-hidden-layer neural networks with landscape design. In \(I C L R, 2017.21,251,351\)
[GLM18] Rong Ge, Jason D. Lee, and Tengyu Ma. Learning one-hidden-layer neural networks with landscape design. In ICLR. http://arxiv.org/pdf/1711. 00501, 2018. 2140
[GLPS10] Anna C Gilbert, Yi Li, Ely Porat, and Martin J Strauss. Approximate sparse recovery: optimizing time and measurements. SIAM Journal on Computing 2012 (A preliminary version of this paper appears in STOC 2010), 41(2):436453, 2010. xliii, \(34,35,620,623,624,625,626,627,628,632,688,690,989\)
[GLPS17] Anna C Gilbert, Yi Li, Ely Porat, and Martin J Strauss. For-all sparse recovery in near-optimal time. ACM Transactions on Algorithms (TALG), 13(3):32, 2017. 623
[GLSS18] Ankit Garg, Yin-Tat Lee, Zhao Song, and Nikhil Srivastava. A matrix expander chernoff bound. In STOC. https://arxiv.org/pdf/1704.03864, 2018. xxv, 5, 32, 553, 2141, 2513
[GM15] Rong Ge and Tengyu Ma. Decomposing overcomplete 3rd order tensors using sum-of-squares algorithms. In The 18th. International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'2015), and the 19th. International Workshop on Randomization and Computation (RANDOM'2015). https://arxiv.org/pdf/1504.05287, 2015. 58, 1470
[GMH13] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 6645-6649. IEEE, 2013. 20, 250, 351
[GMMO00] Sudipto Guha, Nina Mishra, R. Motwani, and L. O'Callaghan. Clustering data streams. In FOCS, pages 359-366, 2000. 2094
[GMS05] Anna C Gilbert, S Muthukrishnan, and Martin Strauss. Improved time bounds for near-optimal sparse Fourier representations. In Optics \& Photonics 2005, pages 59141A-59141A. International Society for Optics and Pho-
tonics, 2005. 37, 38, 40, 623, 636, 676, 682, 688, 689, 691, 762, 919, 923, 1904
[GNP \(\left.{ }^{+} 13\right]\) Anna C Gilbert, Hung Q Ngo, Ely Porat, Atri Rudra, and Martin J Strauss. \(\ell_{2} / \ell_{2}\)-foreach sparse recovery with low risk. In International Colloquium on Automata, Languages, and Programming, pages 461-472. Springer, 2013. 623
[GO16] Inc. Gurobi Optimization. Gurobi optimizer reference manual, 2016. 2041
[Gol65] Sidney Golden. Lower bounds for the Helmholtz function. Physical Review, 137(4B):B1127, 1965. 28, 506
[Goo99] Michael T Goodrich. Communication-efficient parallel sorting. SIAM Journal on Computing, 29(2):416-432, 1999. 2324, 2413
[Goo05] Joseph W Goodman. Introduction to Fourier optics. Roberts and Company Publishers, 2005. 36, 917
[GP14] Mina Ghashami and Jeff M Phillips. Relative errors for deterministic lowrank matrix approximations. In Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 707-717. Society for Industrial and Applied Mathematics, https://arxiv.org/pdf/ 1307.7454, 2014. 1310, 1779
[GQ14] Donald Goldfarb and Zhiwei Qin. Robust low-rank tensor recovery: Models and algorithms. SIAM Journal on Matrix Analysis and Applications, \(35(1): 225-253,2014.58,1470\)
[GRV09] Navin Goyal, Luis Rademacher, and Santosh Vempala. Expanders via random spanning trees. In Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms (SODA), pages 576-585. Society for Industrial and Applied Mathematics, 2009. 550, 552, 555
[GS \(\left.{ }^{+} 99\right]\) Antoine Guitton, William W Symes, et al. Robust and stable velocity analysis using the huber function. In 69th SEG meeting, Houston, USA, Expanded Abstracts, pages 1166-1169, 1999. 990
[GS05] Alex Graves and Jürgen Schmidhuber. Framewise phoneme classification with bidirectional LSTM and other neural network architectures. Neural Networks, 18(5-6):602-610, 2005. 353
[GS12] Surya Ganguli and Haim Sompolinsky. Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. Annual review of neuroscience, 35:485-508, 2012. 33, 617
[GSS11] Shayan Oveis Gharan, Amin Saberi, and Mohit Singh. A randomized rounding approach to the traveling salesman problem. In Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science (FOCS), pages 550-559, 2011. 550, 555
[GSS17] Alon Gonen and Shai Shalev-Shwartz. Fast rates for empirical risk minimization of strict saddle problems. arXiv preprint arXiv:1701.04271, 2017. 13, 74
[GSZ11] Michael T Goodrich, Nodari Sitchinava, and Qin Zhang. Sorting, searching, and simulation in the mapreduce framework. In ISAAC, volume 7074, pages 374-383. Springer, 2011. 2315, 2324, 2413, 2414, 2415
[Gue83] Alain Guenoche. Random spanning tree. Journal of Algorithms, 4(3):214220, 1983. 29, 550
[GV15] Nicolas Gillis and Stephen A Vavasis. On the complexity of robust pca and \(\ell_{1}\)-norm low-rank matrix approximation. arXiv preprint arXiv:1509.09236, 2015. 53, 1110, 1268, 1270, 1274, 1276, 1334, 1370
[GVL13] Gene H. Golub and Charles F. Van Loan. Matrix computations. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, 2013. 2024
[GVS03] David Guillamet, Jordi Vitri, and Bernt Schiele. Introducing a weighted nonnegative matrix factorization for image classification. Pattern Recognition Letters, 24(14):2447-2454, 2003. 1044
[GVS15] Ian J Goodfellow, Oriol Vinyals, and Andrew M Saxe. Qualitatively characterizing neural network optimization problems. In \(I C L R, 2015.20,250,254\), 351, 356, 364, 366
[GW95] Michel X Goemans and David P Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. Journal of the ACM (JACM), 42(6):1115-1145, 1995. 1269
[GWD14] Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. CoRR, abs/1410.5401, 2014. 354
[GZ16] Venkatesan Guruswami and David Zuckerman. Robust Fourier and polynomial curve fitting. Manuscript, 2016. 765
[GZD10] Quanquan Gu, Jie Zhou, and Chris H. Q. Ding. Collaborative filtering: Weighted nonnegative matrix factorization incorporating user and item graphs. In SDM, 2010. 1044
[Haa81] Uffe Haagerup. The best constants in the khintchine inequality. Studia Mathematica, 70:231-283, 1981. 2505
[Hah98] Richard LT Hahnloser. On the piecewise analysis of networks of linear threshold neurons. Neural Networks, 11(4):691-697, 1998. 23, 353
[Har70] Richard A Harshman. Foundations of the parafac procedure: Models and conditions for an "explanatory" multi-modal factor analysis. ., 1970. 57, 1469, 1797
[Har14] Sariel Har-Peled. Low rank matrix approximation in linear time. CoRR, abs/1410.8802, 2014. 49, 1037
[Hås90] Johan Håstad. Tensor rank is np-complete. Journal of Algorithms, 11(4):644654, 1990. 60, 1472, 1669, 1681, 1686, 1689
[Hås00] Johan Håstad. On bounded occurrence constraint satisfaction. Information Processing Letters, 74(1-2):1-6, 2000. 1670
[Hås01] Johan Håstad. Some optimal inapproximability results. Journal of the ACM (JACM), 48(4):798-859, 2001. 1269, 1670, 1671
[Haz01] Michiel Hazewinkel. Gram matrix. Encyclopedia of Mathematics, Springer, 2001. 786
[HD08] Heng Huang and Chris Ding. Robust tensor factorization using r 1 norm. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 1-8. IEEE, 2008. 57, 58, 1469, 1470
[Hea08] Alexander D Healy. Randomness-efficient sampling within nc. Computational Complexity, 17(1):3-37, 2008. 26, 501, 502, 509, 531, 537, 553
[HHW18] Bernhard Haeupler, D. Ellis Hershkowitz, and David Wajc. Round- and message-optimal distributed graph algorithms. In arXiv preprint. http:// arxiv.org/pdf/1801.05127, 2018. 2317
[HIKP12a] Haitham Hassanieh, Piotr Indyk, Dina Katabi, and Eric Price. Nearly optimal sparse Fourier transform. In Proceedings of the forty-fourth annual ACM symposium on Theory of computing, pages 563-578. ACM, 2012. lxvii, lxviii, \(37,38,40,46,164,623,625,636,676,682,688,689,691,694,700,701,715\), 719, 759, 762, 769, 803, 822, 919, 923, 924, 926, 965, 977, 978, 979, 996, 1015, 1016, 1017, 1018, 1904
[HIKP12b] Haitham Hassanieh, Piotr Indyk, Dina Katabi, and Eric Price. Simple and practical algorithm for sparse Fourier transform. In Proceedings of the twenty-
third annual ACM-SIAM symposium on Discrete Algorithms, pages 11831194. SIAM, 2012. 38, 164, 623, 636, 701, 762, 803, 919, 924, 1904
[HJLS13] David W Hosmer Jr, Stanley Lemeshow, and Rodney X Sturdivant. Applied logistic regression, volume 398. John Wiley \& Sons, 2013. 14, 16, 17, 74, 171
[HK13] Daniel Hsu and Sham M Kakade. Learning mixtures of spherical gaussians: moment methods and spectral decompositions. In Proceedings of the 4 th conference on Innovations in Theoretical Computer Science(ITCS), pages 1120. ACM, https://arxiv.org/pdf/1206.5766, 2013. 57, 1469
[HK15] Qingqing Huang and Sham M Kakade. Super-resolution off the grid. In Advances in Neural Information Processing Systems, pages 2647-2655, 2015. 763
[HKT16] Fumio Hiai, Robert Koenig, and Marco Tomamichel. Generalized log-majorization and multivariate trace inequalities. arXiv preprint arXiv:1609.01999, September 2016. 507
[HKZ11] Daniel Hsu, Sham M. Kakade, and Tong Zhang. Robust matrix decomposition with sparse corruptions. ITIT, 57(11):7221-7234, 2011. 1043
[HL13] Christopher J Hillar and Lek-Heng Lim. Most tensor problems are np-hard. In Journal of the \(A C M\) (JACM), volume 60(6), page 45. https://arxiv. org/pdf/0911.1393, 2013. 60, 1472, 1673, 1676
[HLS \({ }^{+}\)18] Elad Hazan, Holden Lee, Karan Singh, Cyril Zhang, and Yi Zhang. Spectral filtering for general linear dynamical systems. In Advances in Neural Information Processing Systems (NIPS), 2018. 24, 257, 354
[HM17] Moritz Hardt and Tengyu Ma. Identity matters in deep learning. In ICLR, 2017. 21, 251, 257, 258, 277, 351
[HMR18] Moritz Hardt, Tengyu Ma, and Benjamin Recht. Gradient descent learns linear dynamical systems. Journal of Machine Learning Research (JMLR), 19(29):1-44, 2018. 23, 257, 354
[HMRW14] Moritz Hardt, Raghu Meka, Prasad Raghavendra, and Benjamin Weitz. Computational limits for matrix completion. In COLT, 2014. 51, 1038
[ \(\mathrm{HMT}^{+}\)13] Hong Huang, Satyajayant Misra, Wei Tang, Hajar Barani, and Hussein AlAzzawi. Applications of compressed sensing in communications networks. arXiv preprint arXiv:1305.3002, 2013. 620
[ \(\mathrm{HNH}^{+}\)13] Furong Huang, Niranjan U. N, Mohammad Umar Hakeem, Prateek Verma, and Animashree Anandkumar. Fast detection of overlapping communities via online tensor methods on gpus. \(C o R R\), abs/1309.0787, 2013. 1797
[HO14] Nicholas JA Harvey and Neil Olver. Pipage rounding, pessimistic estimators and matrix concentration. In Proceedings of the twenty-fifth annual ACMSIAM symposium on Discrete algorithms (SODA), pages 926-945. Society for Industrial and Applied Mathematics, 2014. 553
[Hoe63] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. Journal of the American Statistical Association, 58(301):13-30, 1963. 2506
[HPK05] Sariel Har-Peled and Akash Kushal. Smaller coresets for k-median and kmeans clustering. In Proceedings of the twenty-first annual symposium on Computational geometry, pages 126-134. ACM, 2005. 2094
[HPM04] Sariel Har-Peled and Soham Mazumdar. On coresets for k-means and kmedian clustering. In Proceedings of the thirty-sixth annual ACM symposium on Theory of computing, pages 291-300. ACM, 2004. 2094
[HPS05] Tamir Hazan, Simon Polak, and Amnon Shashua. Sparse image coding using a 3d non-negative tensor factorization. In Tenth IEEE International Conference on Computer Vision(ICCV), volume 1, pages 50-57. IEEE, 2005. 57, 1469
[HR69] Josef Hadar and William R Russell. Rules for ordering uncertain prospects. The American economic review, 59(1):25-34, 1969. 654
[HR16] Ishay Haviv and Oded Regev. The restricted isometry property of subsampled Fourier matrices. In \(S O D A\), pages 288-297. arXiv preprint arXiv:1507.01768, 2016. lxvii, lxviii, 37, 38, 623, 762, 918, 919, 920, 923, 925, 1922
[HRRS11] Frank R Hampel, Elvezio M Ronchetti, Peter J Rousseeuw, and Werner A Stahel. Robust statistics: the approach based on influence functions, volume 196. John Wiley \& Sons, 2011. 56, 1428
[HS90] Yingbo Hua and Tapan K Sarkar. Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise. Acoustics, Speech and Signal Processing, IEEE Transactions on, 38(5):814-824, 1990. 687
[HS97] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. Neural computation, 9(8):1735-1780, 1997. 23, 353
[HSE12] Bo Hu, Zhao Song, and Martin Ester. User features and social networks for topic modelling in online social media. In The 2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, (ASONAM 2012), 26-29 August, 2012, Kadir Has University, Istanbul, Turkey., 2012. xxix
[HSS15] Samuel B Hopkins, Jonathan Shi, and David Steurer. Tensor principal component analysis via sum-of-square proofs. In 28th Annual Conference on Learning Theory (COLT), pages 956-1006. https://arxiv.org/pdf/1507. 03269, 2015. 58, 1470
[HSS19] MohammadTaghi Hajiaghayi, Saeed Seddighin, and Xiaorui Sun. Massively parallel approximation algorithms for edit distance and longest common subsequence. In Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1654-1672. SIAM, 2019. 1917
[HSSS16] Samuel B Hopkins, Tselil Schramm, Jonathan Shi, and David Steurer. Fast spectral algorithms from sum-of-squares proofs: tensor decomposition and
planted sparse vectors. In Proceedings of the 48th Annual Symposium on the Theory of Computing. ACM, https://arxiv.org/pdf/1512.02337, 2016. \(25,58,548,1470\)
[HSSS19] MohammadTaghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, and Xiaorui Sun. Approximating lcs in linear time: Beating the barrier. In Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1181-1200. SIAM, 2019. 2192, 2214, 2216
[HSYZ18] Wei Hu, Zhao Song, Lin F. Yang, and Peilin Zhong. Nearly optimal dynamic \(k\)-means clustering for high-dimensional data. arXiv preprint arXiv:1802.00459, 2018. 2087
[HSZ17] Elad Hazan, Karan Singh, and Cyril Zhang. Learning linear dynamical systems via spectral filtering. In Advances in Neural Information Processing Systems (NeurIPS), pages 6702-6712, 2017. 24, 257, 354
[HT16] Daniel Hsu and Matus Telgarsky. Greedy bi-criteria approximations for \(k\) medians and \(k\)-means. arXiv preprint arXiv:1607.06203, 2016. 1474
[Hub64] Peter J. Huber. Robust estimation of a location parameter. The Annals of Mathematical Statistics, 35(1):73-101, 1964. 56, 246, 989, 992, 1007, 1370, 1428
[HW71] David Lee Hanson and Farroll Tim Wright. A bound on tail probabilities for quadratic forms in independent random variables. The Annals of Mathematical Statistics, 42(3):1079-1083, 1971. 2505
[HX16] Nicholas J. A. Harvey and Keyulu Xu. Generating random spanning trees via fast matrix multiplication. In LATIN 2016: Theoretical Informatics, volume 9644, pages 522-535, 2016. 29, 550
[HZRS16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770-778, 2016. 20, 22, 250, 252, 351
[IBY \({ }^{+}\)07] Michael Isard, Mihai Budiu, Yuan Yu, Andrew Birrell, and Dennis Fetterly. Dryad: distributed data-parallel programs from sequential building blocks. ACM SIGOPS Operating Systems Review, 41(3):59-72, 2007. 2315
[IK14] Piotr Indyk and Michael Kapralov. Sample-optimal Fourier sampling in any constant dimension. In Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on, pages 514-523. IEEE, 2014. lxvii, lxviii, 37, 38, 40, 164, 623, 636, 676, 682, 683, 691, 762, 919, 920, 923, 924, 926, 927, 937, 1904
[IKP14] Piotr Indyk, Michael Kapralov, and Eric Price. (Nearly) Sample-optimal sparse Fourier transform. In Proceedings of the Twenty-Fifth Annual ACMSIAM Symposium on Discrete Algorithms, pages 480-499. SIAM, 2014. 38, \(46,164,623,636,682,759,762,919,1904\)
[IMS17a] Sungjin Im, Benjamin Moseley, and Xiaorui Sun. Efficient massively parallel methods for dynamic programming. In Proceedings of the 49 th Annual ACM

SIGACT Symposium on Theory of Computing, pages 798-811. ACM, 2017. 1917
[IMS17b] Sungjin Im, Benjamin Moseley, and Xiaorui Sun. Efficient massively parallel methods for dynamic programming. In STOC, pages 798-811, 2017. 2315
[Ind04] Piotr Indyk. Algorithms for dynamic geometric problems over data streams. In Proceedings of the thirty-sixth annual ACM symposium on Theory of computing, pages 373-380. ACM, 2004. 2088, 2135
[Ind06] Piotr Indyk. Stable distributions, pseudorandom generators, embeddings, and data stream computation. Journal of the ACM (JACM), 53(3):307-323, 2006. 666, 1118, 1170, 1186, 2030, 2051, 2070, 2074, 2086
[Ind07] Piotr Indyk. Sketching, streaming and sublinear-space algorithms. Graduate course notes, available at, 2007. 33, 617
[IP11] Piotr Indyk and Eric Price. K-median clustering, model-based compressive sensing, and sparse recovery for earth mover distance. In Proceedings of the forty-third annual ACM symposium on Theory of computing, pages 627-636. ACM, 2011. 1262, 2094
[IPW11] Piotr Indyk, Eric Price, and David P Woodruff. On the power of adaptivity in sparse recovery. In Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on, pages 285-294. IEEE, 2011. 623
[IPZ98] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? In Proceedings. 39th Annual Sympo-
sium on Foundations of Computer Science (FOCS), pages 653-662. IEEE, 1998. 1273, 1670
[IR08a] Piotr Indyk and Milan Ruzic. Near-optimal sparse recovery in the 11 norm. In Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on, pages 199-207. IEEE, 2008. 34, 617
[IR08b] Piotr Indyk and Milan Ruzic. Near-optimal sparse recovery in the 11 norm. In Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on, pages 199-207. IEEE, 2008. 989
[IR13] Piotr Indyk and Ilya Razenshteyn. On model-based RIP-1 matrices. In International Colloquium on Automata, Languages, and Programming, pages 564-575. Springer, 2013. 622
[Ise09] Arieh Iserles. A first course in the numerical analysis of differential equations. Number 44 in . Cambridge university press, 2009. 19, 182
[IT03] Piotr Indyk and Nitin Thaper. Fast image retrieval via embeddings. In Workshop on Statistical and Computational Theories of Vision (at ICCV), 2003. 1262, 1264
[IVWW19] Piotr Indyk, Ali Vakilian, Tal Wagner, and David P. Woodruff. Sampleoptimal low-rank approximation of distance matrices. In COLT, 2019. 2027
[IW97] Russell Impagliazzo and Avi Wigderson. \(\mathrm{P}=\mathrm{BPP}\) if E requires exponential circuits: Derandomizing the XOR lemma. In Proceedings of the twenty-ninth
annual ACM symposium on Theory of computing (STOC), pages 220-229. ACM, 1997. 1704
[IW05] P. Indyk and David P. Woodruff. Optimal approximations of the frequency moments of data streams. In STOC, pages 202-208, 2005. 1976
[Iwe08] Mark A Iwen. A deterministic sub-linear time sparse Fourier algorithm via non-adaptive compressed sensing methods. In Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms, pages 20-29. Society for Industrial and Applied Mathematics, 2008. 38, 919, 1904
[Iwe10] Mark A Iwen. Combinatorial sublinear-time Fourier algorithms. Foundations of Computational Mathematics, 10(3):303-338, 2010. 38, 919, 920, 1904
[Iwe13] Mark A Iwen. Improved approximation guarantees for sublinear-time Fourier algorithms. Applied And Computational Harmonic Analysis, 34(1):57-82, 2013. 38, 623, 636, 762, 919, 920, 1904
[Jai10] Anil K Jain. Data clustering: 50 years beyond k-means. Pattern recognition letters, 31(8):651-666, 2010. 2088
[JGH18] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In Advances in neural information processing systems, pages 8571-8580, 2018. 21, 251, 253, 254, 269, 328
[Jia06] Tiefeng Jiang. How many entries of a typical orthogonal matrix can be approximated by independent normals? The Annals of Probability, 34(4):14971529, 2006. 1929, 1941
[JLGJ18] Chi Jin, Lydia T Liu, Rong Ge, and Michael I Jordan. On the local minima of the empirical risk. In Advances in Neural Information Processing Systems (NeurIPS), pages 4901-4910, 2018. 13, 74
[JMZ15] Bo Jiang, Shiqian Ma, and Shuzhong Zhang. Tensor principal component analysis via convex optimization. Mathematical Programming, 150(2):423457, 2015. 58, 1470
[JO14a] Prateek Jain and Sewoong Oh. Learning mixtures of discrete product distributions using spectral decompositions. In 27th Annual Conference on Learning Theory (COLT), pages 824-856. https://arxiv.org/pdf/1311.2972, 2014. 57, 1469
[JO14b] Prateek Jain and Sewoong Oh. Provable tensor factorization with missing data. In Advances in Neural Information Processing Systems (NIPS), pages 1431-1439. https://arxiv.org/pdf/1406.2784, 2014. 58, 1470
[JPT13] Gabriela Jeronimo, Daniel Perrucci, and Elias Tsigaridas. On the minimum of a polynomial function on a basic closed semialgebraic set and applications. SIAM Journal on Optimization, 23(1):241-255, 2013. 1053, 1507
[JSA15] Majid Janzamin, Hanie Sedghi, and Anima Anandkumar. Beating the perils of non-convexity: Guaranteed training of neural networks using tensor meth-
ods. In arXiv preprint. https://arxiv.org/pdf/1506.08473, 2015. 57, 358, 1469
[JST11] Hossein Jowhari, Mert Sağlam, and Gábor Tardos. Tight bounds for lp samplers, finding duplicates in streams, and related problems. In Proceedings of the thirtieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems, pages 49-58. ACM, 2011. 619
[JW18] Rajesh Jayaram and David P Woodruff. Perfect lp sampling in a data stream. In 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS), pages 544-555. IEEE, 2018. 2072, 2073
[JZ13] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In Advances in neural information processing systems, pages 315-323, 2013. 13, 15, 74, 77
[KABO10] Alexandros Karatzoglou, Xavier Amatriain, Linas Baltrunas, and Nuria Oliver. Multiverse recommendation: n -dimensional tensor factorization for context-aware collaborative filtering. In Proceedings of the fourth ACM conference on Recommender systems, pages 79-86. ACM, 2010. 57, 1469
[ \(\mathrm{KAC}^{+}\)08] Fredrik Kahl, Sameer Agarwal, Manmohan Krishna Chandraker, David Kriegman, and Serge Belongie. Practical global optimization for multiview geometry. International Journal of Computer Vision, 79(3):271-284, 2008. 53, 1111
[Kah97] Nabil Kahale. Large deviation bounds for Markov chains. Combinatorics, Probability and Computing, 6(04):465-474, 1997. 26, 502
[Kap16] Michael Kapralov. Sparse Fourier transform in any constant dimension with nearly-optimal sample complexity in sublinear time. In Symposium on Theory of Computing Conference, STOC'16, Cambridge, MA, USA, June 19-21, 2016, 2016. lxvii, lxviii, \(37,38,164,623,636,762,919,923,925,926,1903\), 1904
[Kap17] Michael Kapralov. Sample efficient estimation and recovery in sparse fft via isolation on average. In Foundations of Computer Science, 2017. FOCS'17. IEEE 58th Annual IEEE Symposium on. https://arxiv.org/pdf/1708. 04544, 2017. lxvii, lxviii, 37, 38, 164, 623, 636, 637, 919, 923, 925, 926, 964, 965, 993, 1904
[Kar84] Narendra Karmarkar. A new polynomial-time algorithm for linear programming. In Proceedings of the sixteenth annual ACM symposium on Theory of computing(STOC), pages 302-311. ACM, 1984. 10, 13, 3, 15
[Kar93] David R Karger. Global min-cuts in rnc, and other ramifications of a simple min-cut algorithm. In \(S O D A\), volume 93, pages \(21-30,1993.555,556\)
[Kaw10] Akitoshi Kawamura. Lipschitz continuous ordinary differential equations are polynomial-space complete. Computational Complexity, 19(2):305-332, 2010. 20, 177
[Kaw16] Kenji Kawaguchi. Deep learning without poor local minima. In Advances in Neural Information Processing Systems, pages 586-594, 2016. 21, 251, 257
[KB06] Tamara Kolda and Brett Bader. The tophits model for higher-order web link analysis. In Workshop on link analysis, counterterrorism and security, volume 7, pages 26-29, 2006. 57, 1469
[KB09] Tamara G. Kolda and Brett W. Bader. Tensor decompositions and applications. SIAM Review, 51(3):455-500, 2009. 57, 1469
[KB13] Nal Kalchbrenner and Phil Blunsom. Recurrent continuous translation models. In Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, pages 1700-1709, 2013. 22, 352
[ \(\left.\mathrm{KBG}^{+} 10\right]\) Raghunandan M Kainkaryam, Angela Bruex, Anna C Gilbert, John Schiefelbein, and Peter J Woolf. poolmc: Smart pooling of mrna samples in microarray experiments. BMC bioinformatics, 11(1):299, 2010. 33, 617
[KBJ78] Roger Koenker and Gilbert Bassett Jr. Regression quantiles. Econometrica: journal of the Econometric Society, pages 33-50, 1978. 1430
[KC07] Yong-Deok Kim and Seungjin Choi. Nonnegative tucker decomposition. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR)., pages 1-8. IEEE, 2007. 1706
[KC12] Akitoshi Kawamura and Stephen Cook. Complexity theory for operators in analysis. ACM Transactions on Computation Theory (TOCT) (A preliminary version of this paper appeared in STOC 2010), 4(2):5, 2012. 20, 177
[KDS08] Wim P Krijnen, Theo K Dijkstra, and Alwin Stegeman. On the nonexistence of optimal solutions and the occurrence of "degeneracy" in the candecomp/parafac model. Psychometrika, 73(3):431-439, 2008. 59, 1471
[KF82] Ker-I Ko and Harvey Friedman. Computational complexity of real functions. Theoretical Computer Science, 20(3):323-352, 1982. 20, 177
[KG12] Serap Kirbiz and Bilge Günsel. Perceptually weighted non-negative matrix factorization for blind single-channel music source separation. In ICPR, 2012. 1044
[KH01] Roger Koenker and Kevin F Hallock. Quantile regression. Journal of economic perspectives, 15(4):143-156, 2001. 14, 74
[KHL89] JB Kruskal, RA Harshman, and ME Lundy. How 3-mfa data can cause degenerate parafac solutions, among other relationships. Multiway data analysis, pages 115-121, 1989. 59, 1471
[Kir47] Gustav Kirchhoff. Uber die auflosung der gliechungen, auf welche man bei der undersuchung der linearen vertheilung galvanischer strome gefuhrt wird. Poggendorgs Ann. Phys. Chem., pages 497-508, 1847. 29, 550
[KK03] Qifa Ke and Takeo Kanade. Robust subspace computation using \(\ell_{1}\) norm. Technical Report CMU-CS-03-172, Carnegie Mellon University, Pittsburgh, PA., 2003. 53, 1110, 1370
[KK05] Qifa Ke and Takeo Kanade. Robust \(\ell_{1}\) norm factorization in the presence of outliers and missing data by alternative convex programming. In 2005 IEEE

Computer Society Conference on Computer Vision and Pattern Recognition (CVPR), volume 1, pages 739-746. IEEE, 2005. l, lxxv, 53, 1110, 1286, 1295, 1296, 1298, 1299, 1334, 1370
[KKMO07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for max-cut and other 2-variable csps? SIAM Journal on Computing, 37(1):319-357, 2007. 1269
[KL11] J. Kelner and A. Levin. Spectral sparsification in the semi-streaming setting. In Symposium on Theoretical Aspects of Computer Science (STACS), 2011. 1310, 1779, 2094
\(\left[\mathrm{KLC}^{+} 15\right]\) Eunwoo Kim, Minsik Lee, Chong-Ho Choi, Nojun Kwak, and Songhwai Oh. Efficient-norm-based low-rank matrix approximations for large-scale problems using alternating rectified gradient method. IEEE transactions on neural networks and learning systems, 26(2):237-251, 2015. 53, 1110, 1370
[Kle99] Jon M. Kleinberg. Authoritative sources in a hyperlinked environment. J. ACM, 46(5):604-632, 1999. 48, 1036
[KLM \({ }^{+}\)14a] Michael Kapralov, Yin Tat Lee, Cameron Musco, Christopher Musco, and Aaron Sidford. Single pass spectral sparsification in dynamic streams. In 2014 IEEE 55th Annual Symposium on Foundations of Computer Science (FOCS), pages 561-570. IEEE, https://arxiv.org/pdf/1407.1289, 2014. 1310, 1779
\(\left[K_{L M}{ }^{+} 14 b\right]\) Raimondas Kiveris, Silvio Lattanzi, Vahab Mirrokni, Vibhor Rastogi, and Sergei Vassilvitskii. Connected components in mapreduce and beyond. In Proceedings of the ACM Symposium on Cloud Computing, pages 1-13. ACM, 2014. 2324
\(\left[\mathrm{KLM}^{+} 17\right]\) Michael Kapralov, Yin Tat Lee, CN Musco, CP Musco, and Aaron Sidford. Single pass spectral sparsification in dynamic streams. SIAM Journal on Computing, 46(1):456-477, 2017. 2094
[KLP \(\left.{ }^{+} 16\right]\) Rasmus Kyng, Yin Tat Lee, Richard Peng, Sushant Sachdeva, and Daniel A Spielman. Sparsified cholesky and multigrid solvers for connection laplacians. In Proceedings of the forty-eighth annual ACM symposium on Theory of Computing, pages 842-850. ACM, 2016. 6
[KLS19] Rasmus Kyng, Kyle Luh, and Zhao Song. Four deviations suffice for rank 1 matrices. In Manuscript. https://arxiv.org/pdf/1901.06731, 2019. xxvi, 33, 1913, 2513
[KLT \(\left.{ }^{+} 16\right]\) Kuldeep Kulkarni, Suhas Lohit, Pavan Turaga, Ronan Kerviche, and Amit Ashok. Reconnet: Non-iterative reconstruction of images from compressively sensed measurements. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 449-458, 2016. 988
[KM93] Eyal Kushilevitz and Yishay Mansour. Learning decision trees using the Fourier spectrum. SIAM Journal on Computing, 22(6):1331-1348, 1993. 38, 919
[KM09] Jonathan Kelner and Aleksander Madry. Faster generation of random spanning trees. In Proceedings of the 50th annual Symposium on Foundations of Computer Science (FOCS), pages 13-21, 2009. Available at https://arxiv.org/abs/0908.1448. 29, 550
[KM11] Tamara G Kolda and Jackson R Mayo. Shifted power method for computing tensor eigenpairs. SIAM Journal on Matrix Analysis and Applications, 32(4):1095-1124, 2011. 58, 1469
[KMN \(\left.{ }^{+} 02\right]\) Tapas Kanungo, David M Mount, Nathan S Netanyahu, Christine D Piatko, Ruth Silverman, and Angela Y Wu. A local search approximation algorithm for k-means clustering. In Proceedings of the eighteenth annual symposium on Computational geometry, pages 10-18. ACM, 2002. 2093
[KMP10] Ioannis Koutis, Gary L Miller, and Richard Peng. Approaching optimality for solving sdd linear systems. In 2010 IEEE 51st Annual Symposium on Foundations of Computer Science. IEEE, 2010. 6
[KMP11] Ioannis Koutis, Gary L Miller, and Richard Peng. A nearly-m log n time solver for sdd linear systems. In 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science, pages 590-598. IEEE, 2011. 6
[KN14] Daniel M Kane and Jelani Nelson. Sparser johnson-lindenstrauss transforms. In Journal of the ACM (JACM), volume 61(1), page 4. https://arxiv.org/ pdf/1012.1577, 2014. 1781, 1786, 1791, 2076, 2079
[KNO18] Valentin Khrulkov, Alexander Novikov, and Ivan Oseledets. Expressive power of recurrent neural networks. In International Conference on Learning Representations, 2018. 22, 352
[KNPW11] Daniel M Kane, Jelani Nelson, Ely Porat, and David P Woodruff. Fast moment estimation in data streams in optimal space. In Proceedings of the forty-third annual ACM symposium on Theory of computing, pages 745-754. ACM, 2011. 2029
[Knu69] Donald E Knuth. The art of computer programming. vol. 2: Seminumerical algorithms. addisonwesley. Reading, MA, pages 229-279, 1969. 1807, 1819, 1820
[Knu72] D. F. Knuth. Selected combinatorial research problems. Computer Science Department, Stanford University, 1972. 1915, 2191
[Knu98] Donald E. Knuth. The art of computer programming, vol. 2 : seminumerical algorithms, 1998. 59, 1471
[KNW10a] Daniel M Kane, Jelani Nelson, and David P Woodruff. On the exact space complexity of sketching and streaming small norms. In Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms, pages 1161-1178. Society for Industrial and Applied Mathematics, 2010. 1303, \(1304,1305,1307,1308\)
[KNW10b] Daniel M Kane, Jelani Nelson, and David P Woodruff. On the exact space complexity of sketching and streaming small norms. In Proceedings of the
twenty-first annual ACM-SIAM symposium on Discrete Algorithms, pages 1161-1178. SIAM, 2010. 2030, 2051, 2069
[Ko83] Ker-I Ko. On the computational complexity of ordinary differential equations. Information and control, 58(1-3):157-194, 1983. 20, 177
[Ko10] Ker-I Ko. Polynomial-time computability in analysis: A survey. 2010. 20, 177
[Koe00] Roger Koenker. Galton, edgeworth, frisch, and prospects for quantile regression in econometrics. Journal of Econometrics, 95(2):347-374, 2000. 14, 74
[Koe05] Roger Koenker. Quantile Regression. Cambridge University Press, 2005. 14, 74
[KOSZ13] Jonathan A Kelner, Lorenzo Orecchia, Aaron Sidford, and Zeyuan Allen Zhu. A simple, combinatorial algorithm for solving sdd systems in nearly-linear time. In Proceedings of the forty-fifth annual ACM symposium on Theory of computing, pages 911-920. ACM, 2013. 6
[Kow75] Sophie Kowalevsky. Zur theorie der partiellen differentialgleichungen. In Journal für die reine und angewandte Mathematik, 80, pages 1-32. (German spelling of her lastname used at that time), 1875. 233
[KPHF12] U. Kang, Evangelos E. Papalexakis, Abhay Harpale, and Christos Faloutsos. Gigatensor: scaling tensor analysis up by 100 times - algorithms and discoveries. In \(K D D\), pages 316-324, 2012. 1797
[KPPS17] Rasmus Kyng, Jakub Pachocki, Richard Peng, and Sushant Sachdeva. A framework for analyzing resparsification algorithms. In Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 2032-2043. SIAM, 2017. 26, 549, 554
[KPSZ18] Rasmus Kyng, Richard Peng, Robert Schwieterman, and Peng Zhang. Incomplete nested dissection. In Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, pages 404-417. ACM, 2018. 6
[Kro83] Pieter M Kroonenberg. Three-mode principal component analysis: Theory and applications, volume 2. DSWO press, 1983. 57, 233, 1469
[KS96] David R Karger and Clifford Stein. A new approach to the minimum cut problem. Journal of the ACM (JACM), 43(4):601-640, 1996. 555, 556
[KS08] Tamara G Kolda and Jimeng Sun. Scalable tensor decompositions for multiaspect data mining. In Eighth IEEE International Conference on Data Mining (ICDM), pages 363-372. IEEE, 2008. 57, 1469
[KS09] Adam R Klivans and Alexander A Sherstov. Cryptographic hardness for learning intersections of halfspaces. Journal of Computer and System Sciences, 75(1):2-12, 2009. 258
[KS16] Rasmus Kyng and Sushant Sachdeva. Approximate gaussian elimination for laplacians-fast, sparse, and simple. In 2016 IEEE 57 th Annual Symposium on Foundations of Computer Science (FOCS), pages 573-582. IEEE, 2016. 26, 6, 548, 554, 559
[KS18] Rasmus Kyng and Zhao Song. A matrix chernoff bound for strongly rayleigh distributions and spectral sparsifiers from a few random spanning trees. In FOCS. https://arxiv.org/pdf/1810.08345, 2018. xxvi, 5, 32, 1911, 2141, 2513
[KS19] Michal Koucky and Michael Saks. Constant factor approximations to edit distance on far input pairs in nearly linear time. Manuscript, 2019. 2193
[KSH12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems, pages 1097-1105, 2012. 20, 22, 250, 251, 351
[KSS10] Amit Kumar, Yogish Sabharwal, and Sandeep Sen. Linear-time approximation schemes for clustering problems in any dimensions. Journal of the ACM (JACM), 57(2):5, 2010. 2094
[KST54] Tamás Kovári, Vera Sós, and Pál Turán. On a problem of k. zarankiewicz. In Colloquium Mathematicum, volume 1, pages 50-57, 1954. 2284
[KSV08] Ravindran Kannan, Hadi Salmasian, and Santosh Vempala. The spectral method for general mixture models. SIAM J. Comput., 38(3):1141-1156, 2008. 48, 1036
[KSV10] Howard Karloff, Siddharth Suri, and Sergei Vassilvitskii. A model of computation for mapreduce. In Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms, pages 938-948. Society for Industrial and Applied Mathematics, 2010. 2315, 2324, 2328
[KSV13] Peter Keevash, Benny Sudakov, and Jacques Verstraëte. On a conjecture of erdős and simonovits: Even cycles. Combinatorica, 33(6):699-732, 2013. 2284
[KSW16] Karin Knudson, Rayan Saab, and Rachel Ward. One-bit compressive sensing with norm estimation. IEEE Transactions on Information Theory, 62(5):2748-2758, 2016. 636
[KSZC03] Balachander Krishnamurthy, Subhabrata Sen, Yin Zhang, and Yan Chen. Sketch-based change detection: methods, evaluation, and applications. In Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement, pages 234-247. ACM, 2003. 33, 34, 617, 618
[Kul90] Vidyadhar G. Kulkarni. Generating random combinatorial objects. Journal of Algorithms, 11(2):185-207, 1990. 29, 550
[KV09] Ravi Kannan and Santosh Vempala. Spectral algorithms. Foundations and Trends in Theoretical Computer Science, 4(3-4):157-288, 2009. 53, 1110
[KV17] Ravindran Kannan and Santosh Vempala. Randomized algorithms in numerical linear algebra. Acta Numerica, 26:95-135, 2017. 2043
[KVW14] Ravindran Kannan, Santosh S Vempala, and David P Woodruff. Principal component analysis and higher correlations for distributed data. In Proceedings of The 27th Conference on Learning Theory (COLT), pages 1040-1057, 2014. 1318, 1781, 1784, 1786, 1788, 1789, 1791
[KVZ19] Michael Kapralov, Ameya Velingker, and Amir Zandieh. Dimensionindependent sparse Fourier transform. In Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 2709-2728. SIAM, 2019. lxvii, lxviii, 37, 39, 623, 919, 923, 925, 1902
[KW11] Felix Krahmer and Rachel Ward. New and improved Johnson-Lindenstrauss embeddings via the restricted isometry property. SIAM Journal on Mathematical Analysis, 43(3):1269-1281, 2011. 1922
[Kwa08] Nojun Kwak. Principal component analysis based on \(\ell_{1}\)-norm maximization. IEEE transactions on pattern analysis and machine intelligence, 30(9):16721680, 2008. l, lxxv, 53, 1110, 1286, 1291, 1298, 1299, 1370
[KYFD15] Liwei Kuang, Laurence Yang, Jun Feng, and Mianxiong Dong. Secure tensor decomposition using fully homomorphic encryption scheme. IEEE Transactions on Cloud Computing, 2015. 57, 1469
[LA03] Wu-Sheng Lu and Andreas Antoniou. New method for weighted low-rank approximation of complex-valued matrices and its application for the design of 2-d digital filters. In ISCAS, 2003. 51, 1038
[LAKD08] Song Li, Shaikh S. Ahmed, Gerhard Klimeck, and Eric Darve. Computing entries of the inverse of a sparse matrix using the FIND algorithm. J. Comput. Physics, 227(22):9408-9427, 2008. 1925
[Lan06] J Landsberg. The border rank of the multiplication of \(2 \times 2\) matrices is seven.

In Journal of the American Mathematical Society, volume 19(2), pages 447459, 2006. 59, 1471
[Lan12] Joseph M Landsberg. Tensors: geometry and applications, volume 128. American Mathematical Society Providence, RI, USA., http://www.math.tamu. edu/~joseph.landsberg/Tbookintro.pdf, 2012. 59, 1471
[Lat97] Rafal Latala. Estimation of moments of sums of independent real random variables. The Annals of Probability, pages 1502-1513, 1997. 1390
[LBKW14] Yingyu Liang, Maria-Florina Balcan, Vandana Kanchanapally, and David P. Woodruff. Improved distributed principal component analysis. In NIPS, 2014. 48, 1036
\(\left[\mathrm{LBN}^{+} 17\right]\) Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S. Schoenholz, Jeffrey Pennington, and Jascha Sohl-Dickstein. Deep neural networks as Gaussian processes. arXiv:1711.00165, 2017. 358
[LDFU13] Yichao Lu, Paramveer Dhillon, Dean P Foster, and Lyle Ungar. Faster ridge regression via the subsampled randomized hadamard transform. In Advances in neural information processing systems, pages 369-377, 2013. 164, 165, 168, 1907, 1922, 1937
[LDP07a] Michael Lustig, David Donoho, and John M Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. Magnetic resonance in medicine, 58(6):1182-1195, 2007. 33, 617
[LDP07b] Michael Lustig, David Donoho, and John M Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. Magnetic resonance in medicine, 58(6):1182-1195, 2007. 988
[LDSP08] Michael Lustig, David L Donoho, Juan M Santos, and John M Pauly. Compressed sensing MRI. IEEE signal processing magazine, 25(2):72-82, 2008. 621
\(\left[L^{+}{ }^{+} 18\right]\) Xiaotong Lu, Weisheng Dong, Peiyao Wang, Guangming Shi, and Xuemei Xie. Convcsnet: A convolutional compressive sensing framework based on deep learning. arXiv preprint arXiv:1801.10342, 2018. 988
[Lee12] Jeff Leek. Prediction: the lasso vs. just using the top 10 predictors. http://simplystatistics.org/2012/02/23/ prediction-the-lasso-vs-just-using-the-top-10/, 2012. 1925
[Lee17] Yin Tat Lee. Uniform sampling and inverse maintenance. In Talk at Michael Cohen Memorial Symposium. Available at: https://simons.berkeley .edu/ talks/welcome-andbirds-eye-view-michaels-work., 2017. 119
[Lee18] Yin Tat Lee. Personal communication. UW, 2018. 1909, 2140
[Lez98] Pascal Lezaud. Chernoff-type bound for finite Markov chains. Annals of Applied Probability, pages 849-867, 1998. 26, 502
\(\left[\mathrm{LFC}^{+} 16\right]\) Canyi Lu, Jiashi Feng, Yudong Chen, Wei Liu, Zhouchen Lin, and Shuicheng Yan. Tensor robust principal component analysis: Exact recovery of corrupted low-rank tensors via convex optimization. In Proceedings of the IEEE

Conference on Computer Vision and Pattern Recognition (CVPR), pages 5249-5257, 2016. 57, 58, 1469, 1470
[LG14] François Le Gall. Powers of tensors and fast matrix multiplication. In Proceedings of the 39th international symposium on symbolic and algebraic computation(ISSAC), pages 296-303. ACM, 2014. 12, 4, 75, 89, 158, 159, 1426, 1468, 1921, 2466
[LGU18] Francois Le Gall and Florent Urrutia. Improved rectangular matrix multiplication using powers of the coppersmith-winograd tensor. In Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms(SODA), pages 1029-1046. SIAM, 2018. 15, 67, 75, 89, 158, 159
\(\left[L^{\prime} P^{+} 15\right]\) Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971, 2015. 20, 250, 351
[LHW17] Xingguo Li, Jarvis Haupt, and David Woodruff. Near optimal sketching of low-rank tensor regression. In Advances in Neural Information Processing Systems, pages 3466-3476, 2017. 2025, 2027
[Lib13] Edo Liberty. Simple and deterministic matrix sketching. In Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD), pages 581-588. ACM, 2013. 1310, 1779
[Lin66] Joram Lindenstrauss. A short proof of Liapounoff's convexity theorem. J. Math. Mech., 15:971-972, 1966. 593
[LJCJ17] Lihua Lei, Cheng Ju, Jianbo Chen, and Michael I Jordan. Non-convex finitesum optimization via scsg methods. In Advances in Neural Information Processing Systems, pages 2348-2358, 2017. 15, 77
[LL18] Yuanzhi Li and Yingyu Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. In Advances in Neural Information Processing Systems (NeurIPS), 2018. 21, 251, 257, 259, 273, 351, 1909, 2140
[Llo82] Stuart Lloyd. Least squares quantization in pcm. IEEE transactions on information theory, 28(2):129-137, 1982. 2093
[LLR16] Yuanzhi Li, Yingyu Liang, and Andrej Risteski. Recovery guarantee of weighted low-rank approximation via alternating minimization. CoRR, abs/1602.02262, 2016. 51, 1038
[LM00] Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. Annals of Statistics, pages 1302-1338, 2000. 390, 2506
[LM15] Shachar Lovett and Raghu Meka. Constructive discrepancy minimization by walking on the edges. In SIAM Journal on Computing (A preliminary version of paper appeared in FOCS 2012), volume 44(5), pages 1573-1582. https://arxiv.org/pdf/1203.5747, 2015. 596
[LMH15] Hongzhou Lin, Julien Mairal, and Zaid Harchaoui. A universal catalyst for first-order optimization. In Advances in Neural Information Processing Systems, pages 3384-3392, 2015. 15, 77
[LMH17] Hongzhou Lin, Julien Mairal, and Zaid Harchaoui. Catalyst acceleration for first-order convex optimization: from theory to practice. arXiv preprint arXiv:1712.05654, 2017. 15, 77
[LMP13] Mu Li, Gary L. Miller, and Richard Peng. Iterative row sampling. In 54 th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA, pages 127-136. https://arxiv. org/pdf/1211.2713, 2013. 52, 1109, 1926, 1966
[LMS98] Gad M Landau, Eugene W Myers, and Jeanette P Schmidt. Incremental string comparison. SIAM Journal on Computing, 27(2):557-582, 1998. 1914, 2214, 2215, 2221
[LMS11] Daniel Lokshtanov, Dániel Marx, and Saket Saurabh. Lower bounds based on the exponential time hypothesis. In Bull. EATCS 105, pages 41-72, 2011. 1671
[LMSV11] Silvio Lattanzi, Benjamin Moseley, Siddharth Suri, and Sergei Vassilvitskii. Filtering: a method for solving graph problems in MapReduce. In Proceedings of the 23rd ACM symposium on Parallelism in algorithms and architectures, pages 85-94. ACM, 2011. 2315, 2316
[LMV00a] Lieven De Lathauwer, Bart De Moor, and Joos Vandewalle. A multilinear singular value decomposition. SIAM J. Matrix Analysis Applications, 21(4):1253-1278, 2000. 57, 1469
[LMV00b] Lieven De Lathauwer, Bart De Moor, and Joos Vandewalle. On the best rank1 and rank- \(\left(R_{1}, R_{2}, \cdots, R_{n}\right)\) approximation of higher-order tensors. SIAM J. Matrix Analysis Applications, 21(4):1324-1342, 2000. 57, 1469
[LMWY13] Ji Liu, Przemyslaw Musialski, Peter Wonka, and Jieping Ye. Tensor completion for estimating missing values in visual data. IEEE Trans. Pattern Anal. Mach. Intell., 35(1):208-220, 2013. 57, 58, 1469, 1470
[LN18] Yi Li and Vasileios Nakos. Sublinear-time algorithms for compressive phase retrieval. In ISIT. arXiv preprint arXiv:1709.02917, 2018. 635
[LN19] Yi Li and Vasileios Nakos. Deterministic sparse Fourier transform with an \(\ell_{\infty}\) guarantee. arXiv preprint arXiv:1903.00995, 2019. lxvii, lxviii, 37, 38, 919, 923, 926, 1904
[LNN15] Kasper Green Larsen, Jelani Nelson, and Huy L Nguyên. Time lower bounds for nonadaptive turnstile streaming algorithms. In Proceedings of the fortyseventh annual ACM symposium on Theory of computing, pages 803-812. ACM, 2015. 621
[LNNT16] Kasper Green Larsen, Jelani Nelson, Huy L Nguyên, and Mikkel Thorup. Heavy hitters via cluster-preserving clustering. In Foundations of Computer

Science (FOCS), 2016 IEEE 57th Annual Symposium on, pages 61-70. IEEE, https://arxiv.org/pdf/1604.01357, 2016. 618, 620, 631, 1779, 2029, 2053
[LNRW19] J. Li, A. Nikolov, I. Razenshteyn, and E. Waingarten. On mean estimation for general norms with statistical queries. In COLT, 2019. 1967
[LNW18] Yi Li, Vasileios Nakos, and David P. Woodruff. On low-risk heavy hitters and sparse recovery schemes. In \(R A N D O M / A P P R O X\). arXiv preprint arXiv:1709.02919, 2018. 623
[LOG16] Wuchen Li, Stanley Osher, and Wilfrid Gangbo. Fast algorithm for earth mover's distance based on optimal transport and \(\ell_{1}\) type regularization i. arXiv preprint arXiv:1609.07092, 2016. 1262
[Loj63] S Lojasiewicz. A topological property of real analytic subsets. Coll. du CNRS, Les équations aux dérivées partielles, 117:87-89, 1963. 364
[LP04] C. A. Leon and F. Perron. Optimal Hoeffding bounds for discrete reversible Markov chains. Annals of Applied Probability, 14(2), 2004. 26, 502
[LPS88] Alexander Lubotzky, Ralph Phillips, and Peter Sarnak. Ramanujan graphs. Combinatorica, 8(3):261-277, 1988. 27, 503, 595
[LPS19] Yin Tat Lee, Swati Padmanabhan, and Zhao Song. Tbd. Manuscript, 2019. xxix
[LPSY18] David Liau, Eric Price, Zhao Song, and Ger Yang. Stochastic multi-armed bandits in constant space. In AISTATS, 2018. xxix
[LPW97] W.-S Lu, S.-C Pei, and P.-H Wang. Weighted low-rank approximation of general complex matrices and its application in the design of 2-d digital filters. In IEEE Transactions on Circuits and Systems, volume 44, pages 650-655, 1997. 51, 1038
[LPW09] David Asher Levin, Yuval Peres, and Elizabeth Lee Wilmer. Markov chains and mixing times. American Mathematical Soc., 2009. 1219
[LRHG13] Ben London, Theodoros Rekatsinas, Bert Huang, and Lise Getoor. Multirelational learning using weighted tensor decomposition with modular loss. In arXiv preprint. https://arxiv.org/abs/1303.1733, 2013. 58, 1470
[LRSB12] Nicolas Le Roux, Mark W Schmidt, and Francis R Bach. A stochastic gradient method with an exponential convergence rate for finite training sets. In NIPS, pages 2672-2680, 2012. 13, 74
[LS90] László Lovász and Miklós Simonovits. The mixing rate of markov chains, an isoperimetric inequality, and computing the volume. In 31st Annual Symposium on Foundations of Computer Science (FOCS), pages 346-354. IEEE, 1990. 16, 171
[LS92] László Lovász and Miklós Simonovits. On the randomized complexity of volume and diameter. In 33rd Annual Symposium on Foundations of Computer Science (FOCS), pages 482-492. IEEE, 1992. 16, 171
[LS93] László Lovász and Miklós Simonovits. Random walks in a convex body and
an improved volume algorithm. Random structures \(\mathcal{E}\) algorithms, 4(4):359412, 1993. 16, 171
[LS00] Daniel D. Lee and H. Sebastian Seung. Algorithms for non-negative matrix factorization. In NIPS, 2000. 49, 1036
[LS13] Yin Tat Lee and Aaron Sidford. Path finding I: Solving linear programs with \(\widetilde{O}(\sqrt{r a n k})\) linear system solves. arXiv preprint arXiv:1312.6677, 2013. 7
[LS14] Yin Tat Lee and Aaron Sidford. Path finding methods for linear programming: Solving linear programs in \(O(\sqrt{\text { rank }})\) iterations and faster algorithms for maximum flow. In 2014 IEEE 55th Annual Symposium on Foundations of Computer Science (FOCS), pages 424-433. IEEE, 2014. 10, 11, 13, 3, 5, 15, 21, 1895
[LS15] Yin Tat Lee and Aaron Sidford. Efficient inverse maintenance and faster algorithms for linear programming. In Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on, pages 230-249. IEEE, 2015. 10, 13, 3, 15, 1046, 1895
[LS17] Yin Tat Lee and He Sun. An sdp-based algorithm for linear-sized spectral sparsification. In Proceedings of the 49 th annual acm sigact symposium on theory of computing, pages 678-687. ACM, 2017. 29, 596
[LS18] Yin Tat Lee and He Sun. Constructing linear-sized spectral sparsification in almost-linear time. SIAM Journal on Computing (A Preliminary version of this paper appeared in FOCS 15), 47(6):2315-2336, 2018. 29, 596
[LSS19] Jerry Li, Ruoqi Shen, and Zhao Song. Tbd. Manuscript, 2019. xxix, 2308, 2515
[LSSS14] Roi Livni, Shai Shalev-Shwartz, and Ohad Shamir. On the computational efficiency of training neural networks. In Advances in Neural Information Processing Systems (NeurIPS), pages 855-863, 2014. 258
[LSV18] Yin Tat Lee, Zhao Song, and Santosh S Vempala. Algorithmic theory of odes and sampling from well-conditioned logconcave densities. arXiv preprint arXiv:1812.06243, 2018. xxv, 18, 24, 174, 1897, 2512
[LSW15] Yin Tat Lee, Aaron Sidford, and Sam Chiu-wai Wong. A faster cutting plane method and its implications for combinatorial and convex optimization. In Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on, pages 1049-1065. IEEE, 2015. 21
[LSW \({ }^{+}\)19] Yingyu Liang, Zhao Song, Mengdi Wang, Lin Yang, and Xin Yang. Sketching transformed matrices with applications to natural language processing. Manuscript, 2019. xxix
[LSY19] Yibo Lin, Zhao Song, and Lin F Yang. Towards a theoretical understanding of hashing-based neural nets. In AISTATS, 2019. xxix
[LSZ19] Yin Tat Lee, Zhao Song, and Qiuyi Zhang. Solving empirical risk minimization in the current matrix multiplication time. In COLT. https: //arxiv.org/pdf/1905.04447.pdf, 2019. xxiv, 3, 8, 24, 1896, 2512, 2513
[LT13] Michel Ledoux and Michel Talagrand. Probability in Banach Spaces: isoperimetry and processes. Springer Science \& Business Media, 2013. 511
[LV06a] László Lovász and Santosh Vempala. Fast algorithms for logconcave functions: Sampling, rounding, integration and optimization. In 47 th Annual IEEE Symposium on Foundations of Computer Science (FOCS)., pages 5768. IEEE, 2006. 16, 171
[LV06b] László Lovász and Santosh Vempala. Hit-and-run from a corner. SIAM Journal on Computing, 35(4):985-1005, 2006. 16, 17, 18, 171, 174
[LV13] Xiaodong Li and Vladislav Voroninski. Sparse signal recovery from quadratic measurements via convex programming. SIAM Journal on Mathematical Analysis, 45(5):3019-3033, 2013. 635
[LV17] Yin Tat Lee and Santosh S Vempala. Geodesic walks in polytopes. In Proceedings of the 49 th Annual ACM SIGACT Symposium on Theory of Computing (STOC), pages 927-940. ACM, https://arxiv.org/pdf/1606.04696.pdf, 2017. 19, 177, 182, 197
[LV18] Yin Tat Lee and Santosh S Vempala. Convergence rate of riemannian hamiltonian monte carlo and faster polytope volume computation. In STOC. https//arxiv.org/pdf/1710.06261.pdf, 2018. 19, 173, 182
[LWC13] David Lawlor, Yang Wang, and Andrew Christlieb. Adaptive sub-linear time Fourier algorithms. Advances in Adaptive Data Analysis, 5(01):1350003, 2013. 38, 919
[LY17] Yuanzhi Li and Yang Yuan. Convergence analysis of two-layer neural networks with ReLU activation. In Advances in Neural Information Processing Systems (NeurIPS). http://arxiv.org/abs/1705.09886, 2017. 21, 250, 257, 351, 2140
[Lya40] AA Lyapunov. On completely additive vector functions. Izv. Akad. Nauk SSSR, 4:465-478, 1940. 593
[LZBJ14] Tao Lei, Yuan Zhang, Regina Barzilay, and Tommi Jaakkola. Low-rank tensors for scoring dependency structures. In Association for Computational Linguistics(ACL), Best student paper award, 2014. 57, 1469
[LZMB15] Tao Lei, Yuan Zhang, Alessandro Moschitti, and Regina Barzilay. High-order low-rank tensors for semantic role labeling. In In Proceedings of the 2015 Conference of the North America Chapter of the Association For Computational Linguistics-Human Language Technologies (NAACLHLT 2015. Citeseer, 2015. 57, 1469
[LZWW16] Gui-Fu Lu, Jian Zou, Yong Wang, and Zhongqun Wang. L1-norm-based principal component analysis with adaptive regularization. Pattern Recognition, 60:901-907, 2016. 1334
[Mad13] Aleksander Madry. Navigating central path with electrical flows: From flows to matchings, and back. In 2013 IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS), pages 253-262. IEEE, 2013. 11, 12, 3, 4, 1088, 1089, 1094, 1095, 1100, 1131, 1506
[Mad16] Aleksander Madry. Computing maximum flow with augmenting electrical flows. In 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS), pages 593-602. IEEE, 2016. 11, 12, 3, 4
[Mah11a] M. W. Mahoney. Randomized algorithms for matrices and data. Foundations and Trends \(®\) in Machine Learning, 3(2):123-224, 2011. 1974
[Mah11b] Michael W. Mahoney. Randomized algorithms for matrices and data. Foundations and Trends in Machine Learning, 3(2):123-224, 2011. 2025
[Man92] Yishay Mansour. Randomized interpolation and approximation of sparse polynomials stpreliminary version. In International Colloquium on Automata, Languages, and Programming, pages 261-272. Springer, 1992. 38, 919
[Mar73] GA Margulis. Explicit construction of expanders. Problemy Peredaci Informacii, 9(4):71-80, 1973. 595
[Mas69] James L Massey. Shift-register synthesis and bch decoding. Information Theory, IEEE Transactions on, 15(1):122-127, 1969. 46, 759, 763
[Mau03] Andreas Maurer. A bound on the deviation probability for sums of nonnegative random variables. J. Inequalities in Pure and Applied Mathematics, 4(1):15, 2003. 1389
[MB11] Eric Moulines and Francis R Bach. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In Advances in Neural Information Processing Systems, pages 451-459, 2011. 13, 74
[MBZ10] Sergio V Macua, Pavle Belanovic, and Santiago Zazo. Consensus-based distributed principal component analysis in wireless sensor networks. In Signal Processing Advances in Wireless Communications (SPAWC), 2010 IEEE Eleventh International Workshop on, pages 1-5. IEEE, 2010. 1318, 1784, 1788
[McG09] Andrew McGregor. Graph mining on streams. Encyclopedia of Database Systems, pages 1271-1275, 2009. 2316
[Meg12] Nimrod Megiddo. Progress in Mathematical Programming: Interior-Point and Related Methods. Springer Science \& Business Media, 2012. 6
[Mek04] Raghu Meka. discrepancy-and-beating-the-union-bound/. 2004. 1912
[MES08] Julien Mairal, Michael Elad, and Guillermo Sapiro. Sparse representation for color image restoration. Image Processing, IEEE Transactions on, 17(1):5369, 2008. 51, 1038
[MH09] Morten Mørup and Lars Kai Hansen. Sparse coding and automatic relevance determination for multi-way models. In SPARS'09-Signal Processing with Adaptive Sparse Structured Representations, 2009. 1706
[MHG15] Cun Mu, Daniel Hsu, and Donald Goldfarb. Successive rank-one approximations for nearly orthogonally decomposable symmetric tensors. SIAM Journal on Matrix Analysis and Applications, 36(4):1638-1659, 2015. 58, 1469
[MHWG14] Cun Mu, Bo Huang, John Wright, and Donald Goldfarb. Square deal: Lower bounds and improved relaxations for tensor recovery. In The Thirty-first

International Conference on Machine Learning (ICML), pages 73-81. https: //arxiv.org/pdf/1307.5870, 2014. 58, 1470
[ \(\left.M K B^{+} 10\right]\) Tomáš Mikolov, Martin Karafiát, Lukáš Burget, Jan Černockỳ, and Sanjeev Khudanpur. Recurrent neural network based language model. In Eleventh Annual Conference of the International Speech Communication Association, 2010. 22, 352
\(\left[M K B^{+}{ }^{+11] ~ T o m a ́ s ̌ ~ M i k o l o v, ~ S t e f a n ~ K o m b r i n k, ~ L u k a ́ s ̌ ~ B u r g e t, ~ J a n ~ C ̌ e r n o c k y ̀, ~ a n d ~ S a n j e e v ~}\right.\) Khudanpur. Extensions of recurrent neural network language model. In Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on, pages 5528-5531. IEEE, 2011. 22, 352
[MKCP16] P. P. Markopoulos, S. Kundu, S. Chamadia, and D. A. Pados. Efficient \(\ell_{1}-\) Norm Principal-Component Analysis via Bit Flipping. ArXiv e-prints, 2016. 53, 1110, 1370
[MKP13] Panos P. Markopoulos, George N. Karystinos, and Dimitrios A. Pados. Some options for \(\ell_{1}\)-subspace signal processing. In \(I S W C S\) 2013, The Tenth International Symposium on Wireless Communication Systems, Ilmenau, TU Ilmenau, Germany, August 27-30, 2013, pages 1-5, 2013. 53, 1110, 1370
[MKP14] Panos P. Markopoulos, George N. Karystinos, and Dimitrios A. Pados. Optimal algorithms for \(\ell_{1}\)-subspace signal processing. IEEE Trans. Signal Processing, 62(19):5046-5058, 2014. 53, 1110, 1370
[MLF15] Zhuang Ma, Yichao Lu, and Dean Foster. Finding linear structure in large datasets with scalable canonical correlation analysis. In International Conference on Machine Learning, pages 169-178, 2015. 15, 77
[MM13] Xiangrui Meng and Michael W Mahoney. Low-distortion subspace embeddings in input-sparsity time and applications to robust linear regression. In Proceedings of the forty-fifth annual ACM symposium on Theory of computing, pages 91-100. ACM, https://arxiv.org/pdf/1210.3135, 2013. 49, 52, 1037, 1109, 1119, 1134, 1185, 1198, 1253, 1254, 1260, 1336, 1337, 1338, 1341, 1344, 1345, 1370, 1401, 1426, 1468, 1489, 1516, 1535, 1717, 1922, 1966, 1971, 2027, 2029, 2030, 2031, 2032
[MMD08] Michael W Mahoney, Mauro Maggioni, and Petros Drineas. Tensor-cur decompositions for tensor-based data. SIAM Journal on Matrix Analysis and Applications, 30(3):957-987, 2008. 58, 1469, 1470
[MMN18] Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. Proceedings of the National Academy of Sciences, 115(33):E7665-E7671, 2018. 358
[MMSW15] Konstantin Makarychev, Yury Makarychev, Maxim Sviridenko, and Justin Ward. A bi-criteria approximation algorithm for \(k\) means. arXiv preprint arXiv:1507.04227, 2015. 1474
[MNV09] Meena Mahajan, Prajakta Nimbhorkar, and Kasturi Varadarajan. The planar k-means problem is np-hard. In International Workshop on Algorithms and Computation, pages 274-285. Springer, 2009. 2093
[Moi13] Ankur Moitra. An almost optimal algorithm for computing nonnegative rank. In \(S O D A, 2013.1044,1046,1050,1899\)
[Moi14a] Ankur Moitra. Algorithmic Aspects of Machine Learning. Cambridge University Press, 2014. 58, 1470, 2094
[Moi14b] Ankur Moitra. Tensor decompositions and their applications, 2014. 1797
[Moi15] Ankur Moitra. The threshold for super-resolution via extremal functions. In STOC, 2015. lxxi, 9, 40, 44, 46, 47, 622, 678, 683, 687, 759, 762, 764
[Mør11] Morten Mørup. Applications of tensor (multiway array) factorizations and decompositions in data mining. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 1(1):24-40, 2011. 57, 1469
[Mou10] Nima Mousavi. How tight is chernoff bound, 2010. 2506, 2507
[MP80] William J. Masek and Mike Paterson. A faster algorithm computing string edit distances. J. Comput. Syst. Sci., 20(1):18-31, 1980. 1915, 2191, 2214
[MP14] Gregory T Minton and Eric Price. Improved concentration bounds for countsketch. In Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 669-686. Society for Industrial and Applied Mathematics, 2014. 793, 794
[MPB15] Ali Mousavi, Ankit B Patel, and Richard G Baraniuk. A deep learning approach to structured signal recovery. In Communication, Control, and Computing (Allerton), 2015 53rd Annual Allerton Conference on, pages 13361343. IEEE, 2015. 988
[MPS14] Andrew M McDonald, Massimiliano Pontil, and Dimitris Stamos. Spectral ksupport norm regularization. In Advances in Neural Information Processing Systems, pages 3644-3652, 2014. 1968, 1971, 1984
[MR05] Elchanan Mossel and Sébastien Roch. Learning nonsingular phylogenies and hidden markov models. In Proceedings of the thirty-seventh annual ACM symposium on Theory of computing (STOC), pages 366-375. ACM, https: //arxiv.org/pdf/cs/0502076, 2005. 57, 1469
[MR10] Dana Moshkovitz and Ran Raz. Two-query pcp with subconstant error. In Journal of the \(A C M\) (JACM), volume \(57(5)\), page 29. A preliminary version appeared in the Proceedings of The 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS 08), FOCS 08 Best paper award, https://eccc.weizmann.ac.il/eccc-reports/2008/ TR08-071/, 2010. 1670, 1671, 1702
[MR18] Pasin Manurangsi and Daniel Reichman. The computational complexity of training ReLU(s). arXiv preprint arXiv:1810.04207, 2018. 258
[MS17] Tomoya Murata and Taiji Suzuki. Doubly accelerated stochastic variance reduced dual averaging method for regularized empirical risk minimization. In Advances in Neural Information Processing Systems, pages 608-617, 2017. \(13,15,18,19,74,77,173,174\)
[MSS15a] Adam W. Marcus, Daniel A. Spielman, and Nikhil Srivastava. Interlacing families I: Bipartite Ramanujan graphs of all degrees. Ann. of Math. (2), 182(1):307-325, 2015. 595, 597, 599, 600
[MSS15b] Adam W. Marcus, Daniel A. Spielman, and Nikhil Srivastava. Interlacing families II: Mixed characteristic polynomials and the Kadison-Singer problem. Ann. of Math. (2), 182(1):327-350, 2015. 31, 554, 563, 591, 594, 597, 601, 607, 1912, 1913
[MSS16] Tengyu Ma, Jonathan Shi, and David Steurer. Polynomial-time tensor decompositions with sum-of-squares. In Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on, pages 438-446. IEEE, https://arxiv.org/pdf/1610.01980, 2016. 58, 1470
[MSS18] A. Marcus, D. Spielman, and N. Srivastava. Interlacing families IV: Bipartite ramanujan graphs of all sizes. SIAM Journal on Computing, 47(6):2488-2509, 2018. 595
[MST15] Aleksander Madry, Damian Straszak, and Jakub Tarnawski. Fast generation of random spanning trees and the effective resistance metric. In Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 2019-2036, 2015. Available at http://arxiv.org/pdf/1501.00267v1.pdf. 29, 550
[MT18] Jakub Marecek and Tigran Tchrakian. Robust spectral filtering and anomaly detection. arXiv preprint arXiv:1808.01181, 2018. 257
[Mut05a] S. Muthukrishnan. Data Streams: Algorithms and Applications. Foundations and Trends in Theoretical Computer Science, 1(2):117-236, 2005. 33, 617
[Mut05b] Shanmugavelayutham Muthukrishnan. Data streams: Algorithms and applications. Foundations and Trends \({ }^{\Omega}\) in Theoretical Computer Science, 1(2):117-236, 2005. 2088
[MV18] Oren Mangoubi and Nisheeth K Vishnoi. Dimensionally tight running time bounds for second-order hamiltonian monte carlo. arXiv preprint arXiv:1802.08898, 2018. 18, 19, 173, 174
[MW10] Morteza Monemizadeh and David P Woodruff. 1-pass relative-error lpsampling with applications. In Proceedings of the twenty-first annual ACMSIAM symposium on Discrete Algorithms, pages 1143-1160. SIAM, 2010. 1568, 1593, 1799
[MW17] Cameron Musco and David P Woodruff. Sublinear time low-rank approximation of positive semidefinite matrices. In 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), pages 672-683. IEEE, 2017. 2027
[MXZZ13] Deyu Meng, Zongben Xu, Lei Zhang, and Ji Zhao. A cyclic weighted median method for \(\ell_{1}\) low-rank matrix factorization with missing entries. In \(A A A I\), volume 4, page 6, 2013. 53, 1110, 1370
[MZIC17] Sami Merhi, Ruochuan Zhang, Mark A Iwen, and Andrew Christlieb. A new class of fully discrete sparse Fourier transforms: Faster stable implementations with guarantees. Journal of Fourier Analysis and Applications, pages 1-34, 2017. 38, 919
\(\left[\mathrm{N}^{+} 03\right]\) Yurii Nesterov et al. Random walk in a simplex and quadratic optimization over convex polytopes. CORE, 2003. 1673
[Nac10] Mergen Nachin. Lower bounds on the column sparsity of sparse recovery matrices. UAP: MIT Undergraduate Thesis, 2010. 622
[Nad64] Elizbar A Nadaraya. On estimating regression. Theory of Probability \& Its Applications, 9(1):141-142, 1964. 14, 74
[Nak94] A.M. Nakhushev. Cauchy-kovaleskaya theorem. Hazewinkel, Michiel, Encyclopedia of Mathematics, 1994. 233
[Nak17a] Vasileios Nakos. Almost optimal phaseless compressed sensing with sublinear decoding time. In 2017 IEEE International Symposium on Information Theory, ISIT 2017, Aachen, Germany, June 25-30, 2017, pages 1142-1146, 2017. 635
[Nak17b] Vasileios Nakos. On fast decoding of high-dimensional signals from one-bit measurements. In 44 th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland, pages 61:1-61:14, 2017. 636
[Nak19] Vasileios Nakos. One-bit expandersketch for one-bit compressed sensing. In ISIT. arXiv preprint arXiv:1711.04049, 2019. 636
[ \(\mathrm{NDH}^{+}\)17] Tigran Nagapetyan, Andrew B Duncan, Leonard Hasenclever, Sebastian J Vollmer, Lukasz Szpruch, and Konstantinos Zygalakis. The true cost of
stochastic gradient langevin dynamics. arXiv preprint arXiv:1706.02692, 2017. 181
[NDTTJ18] Aleksandar Sasho Nikolov, Daniel Dadush, Kunal Talwar, and Nicole Tomczak-Jaegermann. Balancing vectors in any norm. FOCS, 2018. 596
[Nes83] Yurii E Nesterov. A method for solving the convex programming problem with convergence rate \(O\left(1 / k^{2}\right)\). In Dokl. Akad. Nauk \(S S S R\), volume 269, pages 543-547, 1983. 13, 74
[Nes98] Yurii Nesterov. Introductory lectures on convex programming volume i: Basic course. Lecture notes, 1998. 15, 77, 81, 82, 160
[Nes04] Yurii Nesterov. Introductory lectures on convex optimization: A basic course, volume 87. Springer Science \& Business Media, 2004. 13, 74, 267, 332
[NJ02] Andrew Y Ng and Michael I Jordan. On discriminative vs. generative classifiers: A comparison of logistic regression and naive bayes. In Advances in neural information processing systems, pages 841-848, 2002. 16, 171
[NJLS09] Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. SIAM Journal on optimization, 19(4):1574-1609, 2009. 13, 74
[NN89] Yu Nesterov and Arkadi Nemirovsky. Self-concordant functions and polynomial-time methods in convex programming. Report, Central Economic and Mathematic Institute, USSR Acad. Sci, 1989. 10, 3
[NN91] Yu Nesterov and Arkadi Nemirovsky. Acceleration and parallelization of the path-following interior point method for a linearly constrained convex quadratic problem. SIAM Journal on Optimization, 1(4):548-564, 1991. 15
[NN94] Yurii Nesterov and Arkadii Nemirovskii. Interior-point polynomial algorithms in convex programming, volume 13. Siam, 1994. 13, 15
[NN13a] Jelani Nelson and Huy L Nguyên. Osnap: Faster numerical linear algebra algorithms via sparser subspace embeddings. In 2013 IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS), pages 117-126. IEEE, https://arxiv.org/pdf/1211.1002, 2013. 49, 52, 13, 770, 1037, 1109, 1337, 1338, 1370, 1426, 1468, 1489, 1535, 1758, 1922, 1966, 2027, 2029
[NN13b] Jelani Nelson and Huy L Nguyên. Sparsity lower bounds for dimensionality reducing maps. In Proceedings of the forty-fifth annual ACM symposium on Theory of computing, pages 101-110. ACM, 2013. 622
[NN14] Jelani Nelson and Huy L Nguyên. Lower bounds for oblivious subspace embeddings. In International Colloquium on Automata, Languages, and Programming, pages 883-894. Springer, 2014. 1242, 1922
[Nol07] John P Nolan. Stable distributions. 2007. 2030, 2054
[NPSS06] Assaf Naor, Yuval Peres, Oded Schramm, and Scott Sheffield. Markov chains in smooth banach spaces and gromov-hyperbolic metric spaces. Duke Mathematical Journal, 134(1):165-197, 2006. 509
[NS17] Yurii Nesterov and Sebastian U Stich. Efficiency of the accelerated coordinate descent method on structured optimization problems. SIAM Journal on Optimization, 27(1):110-123, 2017. 13, 74
[NS19] Vasileios Nakos and Zhao Song. Stronger L2/L2 compressed sensing; without iterating. In Proceedings of the 51st Annual ACM Symposium on Theory of Computing (STOC), 2019. xxvi, 9, 35, 48, 2029, 2514
[NSW18] Vasileios Nakos, Zhao Song, and Zhengyu Wang. The power of (careful) iterative loop analysis in compressed sensing. In manuscript, 2018. xxvii, 48, 964, 1904, 2514
[NSW19a] Vasileios Nakos, Zhao Song, and Zhengyu Wang. (Nearly) sample-optimal sparse Fourier transform in any dimension; RIPless and Filterless. In FOCS, 2019. xxvii, 9, 37, 48, 164, 246, 1902, 1904, 2514
[NSW19b] Vasileios Nakos, Zhao Song, and Zhengyu Wang. Robust sparse recovery via m-estimators. Manuscript, 2019. 1967, 2514
[NSWZ18] Vasileios Nakos, Xiaofei Shi, David P. Woodruff, and Hongyang Zhang. Improved algorithms for adaptive compressed sensing. In 45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic, pages 90:1-90:14, 2018. 623
[NT09a] Deanna Needell and Joel A Tropp. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. Applied and computational harmonic analysis, 26(3):301-321, 2009. lxvi, lxviii, 34, 35, 617, 624
[NT09b] Deanna Needell and Joel A Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. Applied and computational harmonic analysis, 26(3):301-321, 2009. 38, 918, 919, 925
[NV09] Deanna Needell and Roman Vershynin. Uniform uncertainty principle and signal recovery via regularized orthogonal matching pursuit. Foundations of computational mathematics, 9(3):317-334, 2009. 38, 918, 919
[NV10] D Needel and R Vershynin. Signal recovery from inaccurate and incomplete measurements via regularized orthogonal matching pursuit. IEEE Journal of Selected Topics in Signal Processing, pages 310-316, 2010. 38, 918
[NW14] Jelani Nelson and David P. Woodruff. Personal communication. ., 2014. 1563, 1566
[O'M08] Owen O'Malley. Terabyte sort on apache hadoop. Yahoo Tech. Rep, 2008. 2331
[OO18] Samet Oymak and Necmiye Ozay. Non-asymptotic identification of LTI systems from a single trajectory. arXiv preprint arXiv:1806.05722, 2018. 257
[OPS17] Junier B Oliva, Barnabás Póczos, and Jeff Schneider. The statistical recurrent unit. In International Conference on Machine Learning (ICML), pages 26712680, 2017. 353
[OS14] Sewoong Oh and Devavrat Shah. Learning mixed multinomial logit model from ordinal data. In Advances in Neural Information Processing Systems (NIPS), pages 595-603. https://arxiv.org/pdf/1411.0073, 2014. 57, 1469
[Ose11] Ivan V. Oseledets. Tensor-train decomposition. SIAM J. Scientific Computing, 33(5):2295-2317, 2011. 58, 1469, 1711
[OST08] Ivan V Oseledets, DV Savostianov, and Eugene E Tyrtyshnikov. Tucker dimensionality reduction of three-dimensional arrays in linear time. SIAM Journal on Matrix Analysis and Applications, 30(3):939-956, 2008. 58, 1470, 1476
[OT09] Ivan V Oseledets and Eugene E Tyrtyshnikov. Breaking the curse of dimensionality, or how to use svd in many dimensions. SIAM Journal on Scientific Computing, 31(5):3744-3759, 2009. 58, 1469, 2510
[OTZ11] Ivan Oseledets, Eugene Tyrtyshnikov, and Nickolai Zamarashkin. Tensortrain ranks for matrices and their inverses. Computational Methods in Applied Mathematics Comput. Methods Appl. Math., 11(3):394-403, 2011. 1711
[OY05] S. Oh, S. Kwon and J. Yun. A method for structured linear total least norm on blind deconvolution problem. Applied Mathematics and Computing, 19:151-164, 2005. 2024
[OYDS11] Henrik Ohlsson, Allen Y Yang, Roy Dong, and S Shankar Sastry. Compressive phase retrieval from squared output measurements via semidefinite programming. arXiv preprint arXiv:1111.6323, pages 1-27, 2011. 635
[Oym18] Samet Oymak. Learning compact neural networks with regularization. arXiv preprint arXiv:1802.01223, 2018. 257
[Paa97] Pentti Paatero. A weighted non-negative least squares algorithm for threeway "parafac" factor analysis. Chemometrics and Intelligent Laboratory Systems, 38(2):223-242, 1997. 58, 1470
[Paa00] Pentti Paatero. Construction and analysis of degenerate parafac models. Journal of chemometrics, 14(3):285-299, 2000. 16, 57, 59, 171, 1469, 1471
[Pag13] Rasmus Pagh. Compressed matrix multiplication. ACM Transactions on Computation Theory (TOCT), 5(3):9, 2013. 1482, 1486, 1518
[PB17] Jeffrey Pennington and Yasaman Bahri. Geometry of neural network loss surfaces via random matrix theory. In International Conference on Machine Learning (ICML), International Convention Centre, Sydney, Australia, 2017. 358
[PBLJ15] Anastasia Podosinnikova, Francis Bach, and Simon Lacoste-Julien. Rethinking lda: moment matching for discrete ica. In Advances in Neural Information Processing Systems(NIPS), pages 514-522. https://arxiv.org/pdf/1507. 01784, 2015. 57, 1469
[PC08] Anh Phan and Andrzej Cichocki. Fast and efficient algorithms for nonnegative tucker decomposition. Advances in Neural Networks-ISNN 2008, pages 772-782, 2008. 1706
[Pee96] René Peeters. Orthogonal representations over finite fields and the chromatic number of graphs. Combinatorica, 16(3):417-431, 1996. 51, 1038
[Phi] Phillips. Compressed sense. https://www.philips.com/healthcare/ resources/landing/compressed-sense. 36, 917
[Pin60] Mark S Pinsker. Information and information stability of random variables and processes. San Francisco: Holden-Day, 1964, originally published in Russian in 1960, 1960. 1219
[PJ92] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. SIAM Journal on Control and Optimization, 30(4):838855, 1992. 13, 74
[PK16] Young Woong Park and Diego Klabjan. Iteratively reweighted least squares algorithms for \(\ell_{1}\)-norm principal component analysis. arXiv preprint arXiv:1609.02997, 2016. 53, 1110, 1370
[PKB14] Dimitris Papailiopoulos, Anastasios Kyrillidis, and Christos Boutsidis. Provable deterministic leverage score sampling. In Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining, pages 997-1006. ACM, 2014. 1926
\(\left[\mathrm{PLR}^{+} 16\right]\) Ben Poole, Subhaneil Lahiri, Maithreyi Raghu, Jascha Sohl-Dickstein, and Surya Ganguli. Exponential expressivity in deep neural networks through transient chaos. In Advances In Neural Information Processing Systems, pages 3360-3368, 2016. 00047. 358
[PLY10] Yanwei Pang, Xuelong Li, and Yuan Yuan. Robust tensor analysis with 11-
norm. IEEE Transactions on Circuits and Systems for Video Technology, 20(2):172-178, 2010. 57, 58, 1469, 1470
[PMvdG \(\left.{ }^{+} 13\right]\) Jack Poulson, Bryan Marker, Robert A van de Geijn, Jeff R Hammond, and Nichols A Romero. Elemental: A new framework for distributed memory dense matrix computations. ACM Transactions on Mathematical Software (TOMS), 39(2):13, 2013. 1318, 1784, 1788
[Pol63] Boris Teodorovich Polyak. Gradient methods for minimizing functionals. Zhurnal Vychislitel'noi Matematiki i Matematicheskoi Fiziki, 3(4):643-653, 1963. 364
[ \(\left.\mathrm{PP}^{+} 08\right]\) Kaare Brandt Petersen, Michael Syskind Pedersen, et al. The matrix cookbook. Technical University of Denmark, 7:15, 2008. 1221, 1223
[PP13] Ninh Pham and Rasmus Pagh. Fast and scalable polynomial kernels via explicit feature maps. In Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining(KDD), pages 239-247. ACM, 2013. 1482, 1486, 1518, 1797
[PP14] Robin Pemantle and Yuval Peres. Concentration of lipschitz functionals of determinantal and other strong rayleigh measures. Combinatorics, Probability and Computing, 23(1):140-160, 2014. 28, 29, 549, 557, 566, 567, 569
[PR14] Sameer Pawar and Kannan Ramchandran. A robust R-FFAST framework for computing a k -sparse n -length DFT in \(\mathrm{O}(\mathrm{k} \log \mathrm{n})\) sample complexity using
sparse-graph codes. In Information Theory (ISIT), 2014 IEEE International Symposium on, pages 1852-1856. IEEE, 2014. 38, 919
[Pri11] Eric Price. Efficient sketches for the set query problem. In Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms, pages 41-56. Society for Industrial and Applied Mathematics, 2011. 633, 964, 965, 993
[Pri13] Eric C. Price. Sparse recovery and Fourier sampling. PhD thesis, Massachusetts Institute of Technology, 2013. 164
[PRS16] Gopal Pandurangan, Peter Robinson, and Michele Scquizzato. Fast distributed algorithms for connectivity and mst in large graphs. In Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures, pages 429-438. ACM, 2016. 2317
[PRSZ18] Rina Panigrahy, Ali Rahimi, Sushant Sachdeva, and Qiuyi Zhang. Convergence results for neural networks via electrodynamics. In ITCS, 2018. 21, 251, 351
[PRT02] Jiming Peng, Cornelis Roos, and Tamás Terlaky. Self-regular functions and new search directions for linear and semidefinite optimization. Mathematical Programming, 93(1):129-171, 2002. 12, 4
[PRTV00] Christos H. Papadimitriou, Prabhakar Raghavan, Hisao Tamaki, and Santosh Vempala. Latent semantic indexing: A probabilistic analysis. J. Comput. Syst. Sci., 61(2):217-235, 2000. 48, 1036
[Prü18] Heinz Prüfer. Neuer beweis eines satzes uber permutationen. Arch. Math. Phys., 27:742-744, 1918. 580
[PS85] Franco P Preparata and Michael Ian Shamos. Computational geometry: an introduction. Springer Science \& Business Media, 1985. 1334
[PS12] Ely Porat and Martin J Strauss. Sublinear time, measurement-optimal, sparse recovery for all. In Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms, pages 1215-1227. Society for Industrial and Applied Mathematics, 2012. 623
[PS15] Eric Price and Zhao Song. A robust sparse Fourier transform in the continuous setting. In Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on, pages 583-600. IEEE, 2015. xxvi, lxxiii, 9, 10, 40, \(46,47,164,622,759,764,765,769,803,806,822,920,965,977,996,1016\), 1901, 1904, 1905, 1908, 2514
[PS17] Aaron Potechin and David Steurer. Exact tensor completion with sum-ofsquares. In arXiv preprint. https://arxiv.org/pdf/1702.06237, 2017. 58, 1470
[PSG17] Jeffrey Pennington, Samuel S. Schoenholz, and Surya Ganguli. Resurrecting the sigmoid in deep learning through dynamical isometry: theory and practice. arXiv:1711.04735, November 2017. 358
[PSL \(\left.{ }^{+} 12\right]\) Mark S Pearce, Jane A Salotti, Mark P Little, Kieran McHugh, Choonsik Lee, Kwang Pyo Kim, Nicola L Howe, Cecile M Ronckers, Preetha Rajaraman,

Alan W Craft, et al. Radiation exposure from ct scans in childhood and subsequent risk of leukaemia and brain tumours: a retrospective cohort study. The Lancet, 380(9840):499-505, 2012. 620
[PSW17] Eric Price, Zhao Song, and David P. Woodruff. Fast regression with an \(\ell_{\infty}\) guarantee. In International Colloquium on Automata, Languages, and Programming (ICALP), 2017. xxix, 164, 1242, 1920, 2515, 2516
[PTBD16] Ho N Phien, Hoang D Tuan, Johann A Bengua, and Minh N Do. Efficient tensor completion: Low-rank tensor train. In arXiv preprint. https: //arxiv.org/pdf/1601.01083, 2016. 1711
[PTC13] Anh Huy Phan, Petr Tichavský, and Andrzej Cichocki. Fast alternating LS algorithms for high order CANDECOMP/PARAFAC tensor factorizations. IEEE Transactions on Signal Processing, 61(19):4834-4846, 2013. 1797
[PW11] Eric Price and David P Woodruff. ( \(1+\) eps)-approximate sparse recovery. In Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on, pages 295-304. IEEE, 2011. 620, 633, 634, 989, 991, 996, 1334
[PW16] Mert Pilanci and Martin J Wainwright. Iterative hessian sketch: Fast and accurate solution approximation for constrained least-squares. The Journal of Machine Learning Research, 17(1):1842-1879, 2016. 86
[PW17] Mert Pilanci and Martin J Wainwright. Newton sketch: A near linear-time optimization algorithm with linear-quadratic convergence. SIAM Journal on Optimization, 27(1):205-245, 2017. 86, 358
[PYLR17] Ramtin Pedarsani, Dong Yin, Kangwook Lee, and Kannan Ramchandran. Phasecode: Fast and efficient compressive phase retrieval based on sparsegraph codes. IEEE Transactions on Information Theory, 63(6):3663-3691, 2017. 635
[QBI \(\left.{ }^{+} 13\right]\) Saad Qaisar, Rana Muhammad Bilal, Wafa Iqbal, Muqaddas Naureen, and Sungyoung Lee. Compressive sensing: From theory to applications, a survey. Journal of Communications and networks, 15(5):443-456, 2013. 620
[QOSG02] Yongming Qu, George Ostrouchov, Nagiza Samatova, and Al Geist. Principal component analysis for dimension reduction in massive distributed data sets. In Proceedings of IEEE International Conference on Data Mining (ICDM), 2002. 1318, 1784, 1788
[Rao19] Sharavas Rao. Improved lower bounds for the restricted isometry property of subsampled fourier matrices. In arXiv preprint. https://arxiv.org/pdf/ 1903.12146.pdf, 2019. 920, 925
[Ren88] James Renegar. A polynomial-time algorithm, based on newton's method, for linear programming. Mathematical Programming, 40(1-3):59-93, 1988. \(10,12,13,3,4,7\)
[Ren92a] James Renegar. On the computational complexity and geometry of the firstorder theory of the reals, part I: introduction. preliminaries. the geometry of semi-algebraic sets. the decision problem for the existential theory of the reals. J. Symb. Comput., 13(3):255-300, 1992. 1046, 1052, 1053, 1088, 1131, 1339, 1483, 1506
[Ren92b] James Renegar. On the computational complexity and geometry of the firstorder theory of the reals, part II: the general decision problem. preliminaries for quantifier elimination. J. Symb. Comput., 13(3):301-328, 1992. 1046, \(1052,1053,1088,1131,1339,1506\)
[Ren92c] James Renegar. On the computational complexity and geometry of the firstorder theory of the reals. part iii: quantifier elimination. Journal of Symbolic Computation, 13(3):329-352, 1992. 1088
[Ren92d] James Renegar. On the computational complexity of approximating solutions for real algebraic formulae. SIAM J. Comput., 21(6):1008-1025, 1992. 1046
[Ren01] James Renegar. A mathematical view of interior-point methods in convex optimization, volume 3. Siam, 2001. 6
[Rey89] George O Reynolds. The New Physical Optics Notebook: Tutorials in Fourier Optics. ERIC, 1989. 36, 917
[RHS \(\left.{ }^{+} 16\right]\) Sashank J Reddi, Ahmed Hefny, Suvrit Sra, Barnabas Poczos, and Alex Smola. Stochastic variance reduction for nonconvex optimization. In International conference on machine learning, pages 314-323, 2016. 15, 77
[Rip04] BD Ripley. Robust statistics. Course notes, 2004. 988
[RM14] Emile Richard and Andrea Montanari. A statistical model for tensor pca. In Advances in Neural Information Processing Systems, pages 2897-2905. https://arxiv.org/pdf/1411.1076, 2014. 58, 1470
[RMCS13] Vibhor Rastogi, Ashwin Machanavajjhala, Laukik Chitnis, and Anish Das Sarma. Finding connected components in map-reduce in logarithmic rounds. In Data Engineering (ICDE), 2013 IEEE 29th International Conference on, pages 50-61. IEEE, 2013. lxxviii, 2317, 2323, 2324, 2334, 2469, 2502
[RNSS16] Avik Ray, Joe Neeman, Sujay Sanghavi, and Sanjay Shakkottai. The search problem in mixture models. In arXiv preprint. https://arxiv.org/pdf/ 1610.00843, 2016. 57, 1469
[Rot13] Thomas Rothvoß. Approximating bin packing within O(logOPTloglogOPT) bins. In Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on, pages 20-29. IEEE, https://arxiv.org/pdf/1301.4010, 2013. 596
[Rot17] Thomas Rothvoss. Constructive discrepancy minimization for convex sets. SIAM Journal on Computing (A preliminary of version of this paper appeared in FOCS 2014), 46(1):224-234, 2017. 596
[RPK86] Robert Roy, Arogyaswami Paulraj, and Thomas Kailath. Esprit-a subspace rotation approach to estimation of parameters of cisoids in noise. Acoustics, Speech and Signal Processing, IEEE Transactions on, 34(5):1340-1342, 1986. 762
[RR17] Shravas Rao and Oded Regev. A sharp tail bound for the expander random sampler. arXiv preprint arXiv:1703.10205, 2017. 26, 502
[RR19] Victor Reis and Thomas Rothvoss. Linear size sparsifier and the geometry of the operator norm ball. arXiv preprint arXiv:1907.02145, 2019. 1912
[RRT17] Maxim Raginsky, Alexander Rakhlin, and Matus Telgarsky. Non-convex learning via stochastic gradient langevin dynamics: a nonasymptotic analysis. In COLT. arXiv preprint arXiv:1702.03849, 2017. 17, 172
[RS67] A Rényi and G Szekeres. On the height of trees. J. Austral. Math. Soc, 7(4):497-5, 1967. 560
[RS19a] Victor Reis and Zhao Song. Tbd. Manuscript, 2019. xxix
[RS19b] Aviad Rubinstein and Zhao Song. Reducing approximate longest common subsequence to approximate edit distance. arXiv preprint arXiv:1904.05451, 2019. xxix, 1915, 2190, 2216, 2515
[RSSS19] Aviad Rubinstein, Saeed Seddighin, Zhao Song, and Xiaorui Sun. Approximation algorithms for lcs and lis with truly improved running times. In FOCS, 2019. xxix, 1914, 1915, 2192, 2213, 2515
[RST10] Steffen Rendle and Lars Schmidt-Thieme. Pairwise interaction tensor factorization for personalized tag recommendation. In Proceedings of the third ACM international conference on Web search and data mining(WSDM), pages 8190. ACM, 2010. 57, 1469
[RSW16] Ilya Razenshteyn, Zhao Song, and David P Woodruff. Weighted low rank approximations with provable guarantees. In Proceedings of the 48 th Annual

Symposium on the Theory of Computing (STOC), 2016. xxvii, 6, 56, 60, 1131, 1339, 1428, 1506, 1649, 1652, 1654, 1655, 1657, 1658, 1664, 1665, 1672, 1898, 2027, 2514
[RTP16] Thomas Reps, Emma Turetsky, and Prathmesh Prabhu. Newtonian program analysis via tensor product. In Proceedings of the 43 rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages(POPL), volume 51:1, pages 663-677. ACM, 2016. 57, 1469
[RTV05] Cornelis Roos, Tamás Terlaky, and J-Ph Vial. Interior point methods for linear optimization. Springer Science \& Business Media, 2005. 6
[Rub18] Aviad Rubinstein. Approximating edit distance. https://theorydish.blog/2018/07/20/approximating-edit-distance/, 2018. 1914
[Rud99] Mark Rudelson. Random vectors in the isotropic position. Journal of Functional Analysis, 164(1):60-72, 1999. 25, 26, 30, 503, 548, 553, 590, 2141
[RV08] Mark Rudelson and Roman Vershynin. On sparse reconstruction from Fourier and gaussian measurements. Communications on Pure and Applied Mathematics, 61(8):1025-1045, 2008. lxvii, lxviii, 37, 38, 623, 762, 918, 919, 923, 925
[RV09] Mark Rudelson and Roman Vershynin. Smallest singular value of a random rectangular matrix. Communications on Pure and Applied Mathematics, 62(12):1707-1739, 2009. 1545
[RV10] Mark Rudelson and Roman Vershynin. Non-asymptotic theory of random matrices: extreme singular values. arXiv preprint arXiv:1003.2990, 2010. 1226
[RVW00] Omer Reingold, Salil Vadhan, and Avi Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders and extractors. In Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on, pages 3-13. IEEE, 2000. 27, 503
[RVW16] Tim Roughgarden, Sergei Vassilvitskii, and Joshua R. Wang. Shuffles and circuits: (on lower bounds for modern parallel computation). In Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2016, Asilomar State Beach/Pacific Grove, CA, USA, July 11-13, 2016, pages 1-12, 2016. 2316
[RVW18] Tim Roughgarden, Sergei Vassilvitskii, and Joshua R Wang. Shuffles and circuits (on lower bounds for modern parallel computation). Journal of the ACM (JACM), 65(6):41, 2018. 1916
[RYP03] Ryan Rifkin, Gene Yeo, and Tomaso Poggio. Regularized least-squares classification. Nato Science Series Sub Series III Computer and Systems Sciences, 190:131-154, 2003. 1923
[SA96] Emilio Salinas and Laurence F. Abbott. A model of multiplicative neural responses in parietal cortex. Proceedings of the National Academy of Sciences, 93(21):11956-11961, 1996. 23, 353
[SA15] Hanie Sedghi and Anima Anandkumar. Provable methods for training neural networks with sparse connectivity. In \(I C L R\). arXiv preprint arXiv:1412.2693, 2015. 358
[SAH \({ }^{+}\)13] Lixin Shi, O Andronesi, Haitham Hassanieh, Badih Ghazi, Dina Katabi, and Elfar Adalsteinsson. Mrs sparse-FFT: Reducing acquisition time and artifacts for in vivo 2d correlation spectroscopy. In ISMRM'13, Int. Society for Magnetic Resonance in Medicine Annual Meeting and Exhibition, 2013. 41, 676, 687
[Sar06] Tamás Sarlós. Improved approximation algorithms for large matrices via random projections. In 47 th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 21-24 October 2006, Berkeley, California, USA, Proceedings, pages 143-152, 2006. 52, 13, 1109, 1117, 1337, 1338, 1370, 1426, 1468, 1907, 1920, 1922, 1924, 1927, 1966
[SAZ09] Noam Shental, Amnon Amir, and Or Zuk. Rare-allele detection using compressed se (que) nsing. arXiv preprint arXiv:0909.0400, 2009. 33, 617
[SBG04] Age K. Smilde, Rasmus Bro, and Paul Geladi. Multi-way Analysis with Applications in the Chemical Sciences. Wiley, 2004. 57, 1469
[SBS \({ }^{+}\)17] Hojjat Salehinejad, Julianne Baarbe, Sharan Sankar, Joseph Barfett, Errol Colak, and Shahrokh Valaee. Recent advances in recurrent neural networks. arXiv preprint arXiv:1801.01078, 2017. 23, 353, 357
[SBT17] David Sutter, Mario Berta, and Marco Tomamichel. Multivariate trace inequalities. Communications in Mathematical Physics, 352(1):37-58, 2017. arXiv:1604.03023. 501, 507, 508, 517, 518, 522, 524
[SC15] Jimin Song and Kevin C Chen. Spectacle: fast chromatin state annotation using spectral learning. Genome biology, 16(1):33, 2015. 57, 1469
[SC16] Daniel Soudry and Yair Carmon. No bad local minima: Data independent training error guarantees for multilayer neural networks. arXiv preprint arXiv:1605.08361, 2016. 21, 251
[Sch61] Craige Schensted. Longest increasing and decreasing subsequences. Canadian Journal of Mathematics, 13:179-191, 1961. 2218
[Sch81] Ralph Otto Schmidt. A signal subspace approach to multiple emitter location spectral estimation. Ph. D. Thesis, Stanford University, 1981. 762
[Sch12] Leonard J Schulman. Cryptography from tensor problems. In IACR Cryptology ePrint Archive, volume 2012, page 244. https://eprint.iacr.org/ 2012/244, 2012. 57, 1469
[Sch18] Aaron Schild. An almost-linear time algorithm for uniform random spanning tree generation. In Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC), 2018. 29, 550, 555
[SGGSD17] Samuel S. Schoenholz, Justin Gilmer, Surya Ganguli, and Jascha SohlDickstein. Deep information propagation. In \(I C L R\), 2017. 358
[SGS15] Rupesh K Srivastava, Klaus Greff, and Jürgen Schmidhuber. Training very deep networks. In Advances in neural information processing systems (NeurIPS), pages 2377-2385, 2015. 255
[SH05] Amnon Shashua and Tamir Hazan. Non-negative tensor factorization with applications to statistics and computer vision. In Proceedings of the 22nd international conference on Machine learning(ICML), pages 792-799. ACM, 2005. 57, 1469
[SH11] Martin Slawski and Matthias Hein. Sparse recovery by thresholded nonnegative least squares. In Advances in Neural Information Processing Systems, pages 1926-1934, 2011. 988
[Sha11] Ohad Shamir. A variant of azuma's inequality for martingales with subgaussian tails. ArXiv e-prints, abs/1110.2392, 10 2011. 283
[SHM \(\left.{ }^{+} 16\right]\) David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of Go with deep neural networks and tree search. Nature, 529(7587):484-489, 2016. 20, 250, 351
[Shp90] D. Shpak. A weighted-least-squares matrix decomposition method with applications to the design of two-dimensional digital filters. In IEEE Thirty Third Midwest Symposium on Circuits and Systems, 1990. 51, 1038
[SHW \({ }^{+}\)16] Mao Shaowu, Zhang Huanguo, Wu Wanqing, Zhang Pei, Song Jun, and Liu Jinhui. Key exchange protocol based on tensor decomposition problem. China Communications, 13(3):174-183, 2016. 57, 1469
[Sie] Siemens. Compressed sensing beyond speed. https://www.healthcare. siemens.com/magnetic-resonance-imaging/clinical-specialities/ compressed-sensing. 36, 917
[SJ03] Nathan Srebro and Tommi S. Jaakkola. Weighted low-rank approximations. In ICML, 2003. 50, 51, 1038
[SK18] Eric Price Sushrut Karmalkar. Compressed sensing with adversarial sparse noise via 11 regression. arXiv preprint arXiv:1809.08055, 2018. 990
[SL09] Roman Sandler and Michael Lindenbaum. Nonnegative matrix factorization with earth mover's distance metric. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 1873-1880. IEEE, 2009. 1262, 1263
\(\left[\mathrm{SLC}^{+} 17\right]\) Fanhua Shang, Yuanyuan Liu, James Cheng, KW Ng, and Yuichi Yoshida. Variance reduced stochastic gradient descent with sufficient decrease. arXiv preprint arXiv:1703.06807, 2017. 13, 74
[SLJ \({ }^{+}\)15] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 1-9, 2015. 22, 251
[SLRB17] Mark Schmidt, Nicolas Le Roux, and Francis Bach. Minimizing finite sums with the stochastic average gradient. Mathematical Programming, 162(1-2):83-112, 2017. 15, 77
[Slu77] Eric V Slud. Distribution inequalities for the binomial law. The Annals of Probability, 5(3):404-412, 1977. 2506, 2507
[Sma98] Steve Smale. Mathematical problems for the next century. The mathematical intelligencer, 20(2):7-15, 1998. 1895
[SMH11] Ilya Sutskever, James Martens, and Geoffrey Hinton. Generating text with recurrent neural networks. In International Conference on Machine Learning (ICML), pages 1017-1024, 2011. 23, 353
[SMT \({ }^{+}\)18] Max Simchowitz, Horia Mania, Stephen Tu, Michael I Jordan, and Benjamin Recht. Learning without mixing: Towards a sharp analysis of linear system identification. In Conference on Learning Theory (COLT). arXiv preprint arXiv:1802.08334, 2018. 257
[Son17] Zhao Song. High dimensional Fourier transform in the continuous setting. In Manuscript, 2017. 21, 250, 351, 1901, 2140, 2309
[SP97] Mike Schuster and Kuldip K Paliwal. Bidirectional recurrent neural networks. IEEE Transactions on Signal Processing, 45(11):2673-2681, 1997. 353
[Spe85] Joel Spencer. Six standard deviations suffice. Trans. Amer. Math. Soc., 289(2):679-706, 1985. 30, 590
[SS91] Hava T Siegelmann and Eduardo D Sontag. Turing computability with neural nets. Applied Mathematics Letters, 4(6):77-80, 1991. 354
[SS10] Michael Saks and C Seshadhri. Estimating the longest increasing sequence in polylogarithmic time. In Proceedings of the Fifty-First Annual IEEE Symposium on Foundations of Computer Science, 2010. 2214, 2218, 2219
[SS11] Daniel A Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. SIAM Journal on Computing, 40(6):1913-1926, 2011. 25, 29, 13, 548, 551, 552, 559, 581, 596
[SS14] Zhao Song and Wen Sun. Probabilistic recharging model in uncertain environments. In \(A A M A S\), pages 1343-1344, 2014. xxix
[SS16] Shai Shalev-Shwartz. SDCA without duality, regularization, and individual convexity. In International Conference on Machine Learning, pages 747-754, 2016. 15, 77
[SS17] Tselil Schramm and David Steurer. Fast and robust tensor decomposition with applications to dictionary learning. manuscript, 2017. 58, 1470
[SS18] Itay Safran and Ohad Shamir. Spurious local minima are common in twolayer ReLU neural networks. In International Conference on Machine Learning (ICML). http://arxiv.org/abs/1712.08968, 2018. 21, 251, 351
[SS19] Zhao Song and Wen Sun. Efficient model-free reinforcement learning in metric spaces. arXiv preprint arXiv:1905.00475, 2019. xxix
[SSB14] Haşim Sak, Andrew Senior, and Françoise Beaufays. Long short-term memory recurrent neural network architectures for large scale acoustic modeling. In Fifteenth annual conference of the international speech communication association, 2014. 22, 352
[SSBD14] Shai Shalev-Shwartz and Shai Ben-David. Understanding machine learning: From theory to algorithms. Cambridge university press, 2014. 14, 74
[SSN12] Martin Sundermeyer, Ralf Schlüter, and Hermann Ney. Lstm neural networks for language modeling. In Thirteenth annual conference of the international speech communication association, 2012. 22, 352
[SSS \({ }^{+}\)17] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of Go without human knowledge. Nature, 550(7676):354, 2017. 20, 250, 351
[SSV12] Zhao Song, Seyed Abbas Sadat, and Richard T. Vaughan. Mo-lost: Adaptive ant trail untangling in multi-objective multi-colony robot foraging. In Eleventh International Conference on Autonomous Agents and Multiagent Systems(AAMAS2012), June 4-8 2012, Universitat Politecnica de Valencia, Valencia, Spain., 2012. xxix
[SSZ13] Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. Journal of Machine Learning Research, 14(Feb):567-599, 2013. 13, 74
[SSZ14] Shai Shalev-Shwartz and Tong Zhang. Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. In International Conference on Machine Learning, pages 64-72, 2014. 15, 22, 77, 251
[ST04] Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing (STOC), pages 81-90, New York, NY, USA, 2004. ACM. 25, 6, 372, 381, 464, 548, 551, 596
[ST11] Daniel A Spielman and Shang-Hua Teng. Spectral sparsification of graphs. SIAM Journal on Computing, 40(4):981-1025, 2011. 29
[Ste06] Alwin Stegeman. Degeneracy in candecomp/parafac explained for \(\mathrm{p} \times \mathrm{p} \times 2\) arrays of rank p+1 or higher. Psychometrika, 71(3):483-501, 2006. 59, 1471
[Ste08] Alwin Stegeman. Low-rank approximation of generic \(\mathrm{p} \times \mathrm{q} \times 2\) arrays and diverging components in the candecomp/parafac model. SIAM Journal on Matrix Analysis and Applications, 30(3):988-1007, 2008. 59, 1471
[STLS14] Marco Signoretto, Dinh Quoc Tran, Lieven De Lathauwer, and Johan A. K. Suykens. Learning with tensors: a framework based on convex optimization and spectral regularization. Machine Learning, 94(3):303-351, 2014. 57, 1469
[Str69] Volker Strassen. Gaussian elimination is not optimal. Numerische Mathematik, 13(4):354-356, 1969. 89, 761, 775, 792, 1133, 1426, 1468, 1515
[SV82] Yossi Shiloach and Uzi Vishkin. An \(O(\log n)\) parallel connectivity algorithm. J. Algorithms, 3(1):57-67, 1982. 2315
[SV12a] Nariankadu D. Shyamalkumar and Kasturi R. Varadarajan. Efficient subspace approximation algorithms. Discrete \& Computational Geometry, 47(1):44-63, 2012. 54, 1112
[SV12b] Zhao Song and Richard T. Vaughan. Multi-robot, multi-patch foraging with maximum sustainable yield. In Artificial Life(ALIFE XIII), July 19-22, 2012, Michigan State University, East Lansing, Michigan, USA, 2012. xxix
[SV13] Zhao Song and Richard T. Vaughan. Sustainable robot foraging: Adaptive fine-grained multi-robot task allocation for maximum sustainable yield of biological resources. In IROS, pages 3309-3316, 2013. xxix
[SV16] Damian Straszak and Nisheeth K Vishnoi. Irls and slime mold: Equivalence and convergence. arXiv preprint arXiv:1601.02712, 2016. 34, 617
[SVL14] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In Advances in neural information processing systems (NeurIPS), pages 3104-3112, 2014. 22, 352
[SVWX17] Le Song, Santosh Vempala, John Wilmes, and Bo Xie. On the complexity of learning neural networks. In Advances in Neural Information Processing Systems (NeurIPS), pages 5514-5522, 2017. 258
[SW11] Christian Sohler and David P Woodruff. Subspace embeddings for the \(\ell_{1}\) norm with applications. In Proceedings of the forty-third annual ACM sym-
posium on Theory of computing, pages 755-764. ACM, 2011. 52, 1109, 1117, \(1118,1170,1185,1191,1253,1336,1338,1341,1344,1345,1975\)
[SW15] Tselil Schramm and Benjamin Weitz. Low-rank matrix completion with adversarial missing entries. \(C o R R\), abs/1506.03137, 2015. 51, 1038, 1043
[SW18] Christian Sohler and David P. Woodruff. Strong coresets for k-median and subspace approximation, goodbye dimension. In FOCS. IEEE, 2018. 2094
[SWY \(\left.{ }^{+} 19\right]\) Zhao Song, Ruosong Wang, Lin F. Yang, Hongyang Zhang, and Peilin Zhong. Efficient symmetric norm regression via linear sketching. Manuscript, 2019. xxix, 1965, 2515, 2516
[SWZ16] Zhao Song, David P. Woodruff, and Huan Zhang. Sublinear time orthogonal tensor decomposition. In Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems (NIPS) 2016, December 5-10, 2016, Barcelona, Spain, pages 793-801, 2016. xxviii, \(3,4,58,61,1469,2025,2514\)
[SWZ17] Zhao Song, David P Woodruff, and Peilin Zhong. Low rank approximation with entrywise \(\ell_{1}\)-norm error. In Proceedings of the 49 th Annual Symposium on the Theory of Computing (STOC). ACM, https://arxiv.org/pdf/ 1611.00898, 2017. xxvii, 3, 6, 56, 60, 1145, 1334, 1335, 1336, 1337, 1338, \(1339,1342,1345,1357,1370,1373,1399,1428,1429,1433,1476,1488,1506\), 1604, 1617, 1619, 1623, 1624, 1628, 1629, 1645, 1646, 1647, 1717, 1718, 1720, \(1725,1727,1728,1729,1736,1737,1738,1739,1779,1784,1788,2027,2082\), 2094, 2514
[SWZ18] Zhao Song, David P Woodruff, and Peilin Zhong. Towards a zero-one law for entrywise low rank approximation. arXiv preprint arXiv:1811.01442, 2018. xxviii, 60, 1007, 1967, 2027, 2082, 2083, 2514
[SWZ19a] Zhao Song, David P Woodruff, and Peilin Zhong. Average case column subset selection for entrywise \(\ell_{1}\)-norm loss. Manuscript, 2019. 2514
[SWZ19b] Zhao Song, David P Woodruff, and Peilin Zhong. Relative error tensor low rank approximation. In \(S O D A\). arXiv preprint arXiv:1704.08246, 2019. xxviii, \(3,6,7,60,246,1341,1357,1433,1899,2025,2027,2042,2094,2514\)
[SY19] Zhao Song and Xin Yang. Quadratic suffices for over-parametrization via matrix chernoff bound. arXiv preprint arXiv:1906.03593, 2019. xxix, 25, 548, 1908, 1909, 2139, 2516
[SZ13] Zhao Song and Yuke Zhu. Graphical model-based learning in high dimensional feature spaces. In Twenty-Seventh AAAI Conference on Artificial Intelligence, 2013. xxix
[TBR15] Gongguo Tang, Badri Narayan Bhaskar, and Benjamin Recht. Near minimax line spectral estimation. Information Theory, IEEE Transactions on, 61(1):499-512, 2015. 763
[TBSR13] Gongguo Tang, Badri Narayan Bhaskar, Parikshit Shah, and Benjamin Recht. Compressed sensing off the grid. Information Theory, IEEE Transactions on, 59(11):7465-7490, 2013. 41, 676, 686, 763
[TD99] Françoise Tisseur and Jack Dongarra. A parallel divide and conquer algorithm for the symmetric eigenvalue problem on distributed memory architectures. SIAM Journal on Scientific Computing, 20(6):2223-2236, 1999. 1318, 1784, 1788
[Ter13] Tamás Terlaky. Interior point methods of mathematical programming, volume 5. Springer Science \& Business Media, 2013. 6
[TG07] Joel Tropp and Anna C Gilbert. Signal recovery from partial information via orthogonal matching pursuit. IEEE Trans. Inform. Theory, 53(12):46554666, 2007. 38, 918
[Tho65] Colin J Thompson. Inequality with applications in statistical mechanics. Journal of Mathematical Physics, 6(11):1812-1813, 1965. 28, 506
[Tia17] Yuandong Tian. An analytical formula of population gradient for twolayered ReLU network and its applications in convergence and critical point analysis. In International Conference on Machine Learning (ICML). http://arxiv.org/abs/1703.00560, 2017. 21, 250, 351, 358, 2140
[Tib96] Robert Tibshirani. Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society. Series B (Methodological), pages 267-288, 1996. 14, 74, 1334
[Tij] H Tijms. Understanding probability: Chance rules in everyday life. 2004. 25, 548
[TK11] Petr Tichavsky and Zbyněk Koldovsky. Weight adjusted tensor method for blind separation of underdetermined mixtures of nonstationary sources. IEEE Transactions on Signal Processing, 59(3):1037-1047, 2011. 58, 1470
[TLW \({ }^{+}\)06] Dharmpal Takhar, Jason N Laska, Michael B Wakin, Marco F Duarte, Dror Baron, Shriram Sarvotham, Kevin F Kelly, and Richard G Baraniuk. A new compressive imaging camera architecture using optical-domain compression. In Computational Imaging IV, volume 6065, page 606509. International Society for Optics and Photonics, 2006. 33, 617
[TM17] Davoud Ataee Tarzanagh and George Michailidis. Fast monte carlo algorithms for tensor operations. In arXiv preprint. https://arxiv.org/pdf/ 1704.04362, 2017. 1476, 1504
[Tre01] Luca Trevisan. Non-approximability results for optimization problems on bounded degree instances. In Proceedings of the thirty-third annual ACM symposium on Theory of computing (STOC), pages 453-461. ACM, 2001. 1671, 1703
[Tro11a] Joel A Tropp. Freedman's inequality for matrix martingales. Electronic Communications in Probability, 16:262-270, 2011. 25, 548
[Tro11b] Joel A Tropp. Improved analysis of the subsampled randomized hadamard transform. Advances in Adaptive Data Analysis, 3(01n02):115-126, 2011. 1907, 1922
[Tro11c] Joel A Tropp. User-friendly tail bounds for matrix martingales. Technical report, CALIFORNIA INST OF TECH PASADENA, 2011. 553, 558, 2510
[Tro12] Joel A Tropp. User-friendly tail bounds for sums of random matrices. Foundations of computational mathematics, 12(4):389-434, 2012. 25, 26, 30, 31, 503, 511, 548, 551, 553, 557, 559, 590, 591, 2141
[Tro15] Joel A. Tropp. An introduction to matrix concentration inequalities. Foundations and Trends® in Machine Learning, 8(1-2):1-230, 2015. 504, 2141, 2510
[TS19] Ewin Tang and Zhao Song. Tbd. Manuscript, 2019. xxix
[TSHK11] Ryota Tomioka, Taiji Suzuki, Kohei Hayashi, and Hisashi Kashima. Statistical performance of convex tensor decomposition. In Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems (NIPS). Proceedings of a meeting held 12-14 December 2011, Granada, Spain., pages 972-980, 2011. 58, 1470
[Tso10] Charalampos E. Tsourakakis. MACH: fast randomized tensor decompositions. In \(S D M\), pages 689-700, 2010. 1797, 1798, 1804
[TSSW00] Luca Trevisan, Gregory B Sorkin, Madhu Sudan, and David P Williamson. Gadgets, approximation, and linear programming. SIAM Journal on Computing, 29(6):2074-2097, 2000. 1269
[Tsy09] Alexandre B Tsybakov. Introduction to nonparametric estimation. revised and extended from the 2004 french original. translated by vladimir zaiats, 2009. 1219
[Tuk60] John W Tukey. A survey of sampling from contaminated distributions. Contributions to probability and statistics, pages 448-485, 1960. 992, 993, 1007
[Tur41] Paul Turán. On an external problem in graph theory. Mat. Fiz. Lapok, 48:436-452, 1941. 2284
[TWZS15] Hsiao-Yu Fish Tung, Chao-Yuan Wu, Manzil Zaheer, and Alexander J Smola. Spectral methods for the hierarchical dirichlet process. 2015. 1797, 1799, 1814, 1887
[TZ12] Mikkel Thorup and Yin Zhang. Tabulation-based 5-independent hashing with applications to linear probing and second moment estimation. SIAM Journal on Computing, 41(2):293-331, 2012. 621
[UHZ \(\left.{ }^{+} 16\right]\) Madeleine Udell, Corinne Horn, Reza Zadeh, Stephen Boyd, et al. Generalized low rank models. Foundations and Trends \(®\) in Machine Learning, \(9(1): 1-118,2016.55,56,1342,1427,1428,1450\)
[Vai87] Pravin M Vaidya. An algorithm for linear programming which requires \(O\left(\left((m+n) n^{2}+(m+n)^{1.5} n\right) L\right)\) arithmetic operations. In FOCS. IEEE, 1987. \(10,3,11\)
[Vai89a] Pravin M Vaidya. A new algorithm for minimizing convex functions over convex sets. In 30th Annual Symposium on Foundations of Computer Science(FOCS), pages 338-343. IEEE, 1989. 10, 3
[Vai89b] Pravin M Vaidya. Speeding-up linear programming using fast matrix multiplication. In FOCS. IEEE, 1989. 10, 13, 15, 3, 20, 66, 77
[Val90] Leslie G. Valiant. A bridging model for parallel computation. Commun. ACM, 33(8):103-111, 1990. 2315
[Vap92] Vladimir Vapnik. Principles of risk minimization for learning theory. In Advances in neural information processing systems, pages 831-838, 1992. 13, 14, 74
[Vap13] Vladimir Vapnik. The nature of statistical learning theory. Springer science \& business media, 2013. 13, 74
[Vas09] M Alex O Vasilescu. A multilinear (tensor) algebraic framework for computer graphics, computer vision, and machine learning. PhD thesis, Citeseer, 2009. 57, 1469
[vdBLSS19] Jan van den Brand, Yin Tat Lee, Aaron Sidford, and Zhao Song. Tbd. Manuscript, 2019. xxix
[vdH03] Joris van der Hoeven. Majorants for formal power series. Citeseer, 2003. 232
[Ver10] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. arXiv preprint arXiv:1011.3027, 2010. 1942
[Ver14] Sergio Verdú. Total variation distance and the distribution of relative information. In ITA, pages 1-3. Citeseer, 2014. 1219
[VGT12] Gaël Varoquaux, Alexandre Gramfort, and Bertrand Thirion. Small-sample brain mapping: sparse recovery on spatially correlated designs with randomization and clustering. arXiv preprint arXiv:1206.6447, 2012. 988
[VL92] Charles F Van Loan. Computational frameworks for the fast Fourier transform, volume 10 of Frontiers in Applied Mathematics. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1992. 2024
[VL00] Charles F Van Loan. The ubiquitous kronecker product. Journal of computational and applied mathematics, 123(1-2):85-100, 2000. 2029, 2042, 2065, 2080
[VLP93] Charles F Van Loan and N. Pitsianis. Approximation with Kronecker products. In Linear algebra for large scale and real-time applications (Leuven, 1992), volume 232 of NATO Adv. Sci. Inst. Ser. E Appl. Sci., pages 293-314. Kluwer Acad. Publ., Dordrecht, 1993. 2024
[Voe11] David G Voelz. Computational Fourier Optics: A MATLAB Tutorial (SPIE Tutorial Texts Vol. TT89). SPIE press, 2011. 36, 917
[VT01] H. L. Van Trees. Detection, estimation, and modulation theory. Wiley, 1 edition, September 2001. 53, 1111
[VT02] M Alex O Vasilescu and Demetri Terzopoulos. Multilinear analysis of image ensembles: Tensorfaces. In European Conference on Computer Vision, pages 447-460. Springer, 2002. 57, 1469
[VT04] M Alex O Vasilescu and Demetri Terzopoulos. Tensortextures: Multilinear image-based rendering. In ACM Transactions on Graphics (TOG), volume 23:3, pages 336-342. ACM, 2004. 57, 1469
[VW18] Santosh Vempala and John Wilmes. Polynomial convergence of gradient descent for training one-hidden-layer neural networks. arXiv preprint arXiv:1805.02677, 2018. 351
[WA03] Hongcheng Wang and Narendra Ahuja. Facial expression decomposition. In Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on, pages 958-965. IEEE, 2003. 57, 1469
[WA16] Yining Wang and Animashree Anandkumar. Online and differentially-private tensor decomposition. In Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems (NIPS) 2016, December 5-10, 2016, Barcelona, Spain. https://arxiv.org/pdf/ 1606.06237, 2016. 58, 1469, 1809, 1810, 1826, 1850, 1855, 1861
[Wag08] R. Wagner. Tail estimates for sums of variables sampled by a random walk. Combinatorics, Probability and Computing, 17(2), 2008. 26, 502
[Wag11] David Wagner. Multivariate stable polynomials: theory and applications. Bulletin of the American Mathematical Society, 48(1):53-84, 2011. 599
[Wai19] M.J. Wainwright. High-Dimensional Statistics: A Non-Asymptotic Viewpoint. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019. 2068
[Wal77] Alastair J Walker. An efficient method for generating discrete random variables with general distributions. ACM Transactions on Mathematical Software (TOMS), 3(3):253-256, 1977. 1807
[Wan16] Yining Wang. Personal communication. 2016. 1808
[Wan19a] Lan Wang. A new tuning-free approach to high-dimensional regression. ., 2019. 2026
[Wan19b] Lan Wang. Personal communication. ., 2019. 2026
[Wat64] Geoffrey S Watson. Smooth regression analysis. Sankhy \(\bar{a}\) : The Indian Journal of Statistics, Series A, pages 359-372, 1964. 14, 74
[Wea04] Nik Weaver. The Kadison-Singer problem in discrepancy theory. Discrete Math., 278(1-3):227-239, 2004. 593
[Wes94] Carl-Fredrik Westin. A tensor framework for multidimensional signal processing. PhD thesis, Linköping University Electronic Press, 1994. 57, 1469
[Whi12] Tom White. Hadoop: the definitive guide. O'Reilly, 2012. 2315
[Wik16] Wikipedia. Random permutation. In https://en.wikipedia.org/wiki/Random_permutation, 2016. 1819
[Wil96] David Bruce Wilson. Generating random spanning trees more quickly than the cover time. In Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing (STOC), pages 296-303, New York, NY, USA, 1996. ACM. 29, 550
[Wil05] Herbert S Wilf. Generating functionology. AK Peters/CRC Press, 2005. 666
[Wil12] Virginia Vassilevska Williams. Multiplying matrices faster than coppersmithwinograd. In Proceedings of the forty-fourth annual ACM symposium on Theory of computing (STOC), pages 887-898. ACM, 2012. 12, 4, 75, 89, 158, 159, 761, 775, 792, 881, 1133, 1426, 1468, 1515, 1921
[WKL09] Lan Wang, Bo Kai, and Runze Li. Local rank inference for varying coefficient models. Journal of the American Statistical Association, 104(488):1631-1645, 2009. 2026
[WL09] Lan Wang and Runze Li. Weighted wilcoxon-type smoothly clipped absolute deviation method. Biometrics, 65(2):564-571, 2009. 2026
[WLSH14] Chi Wang, Xueqing Liu, Yanglei Song, and Jiawei Han. Scalable momentbased inference for latent dirichlet allocation. In \(E C M L-P K D D\), pages 290305, 2014. 1797
\(\left[W^{+}{ }^{+}{ }^{15]}\right.\) Ye Wang, Meng Li, Xinyang Yi, Zhao Song, Michael Orshansky, and Constantine Caramanis. Novel power grid reduction method based on 11 regularization. In \(D A C, 2015\). xxix
[WM01] B. Walczak and DL Massart. Dealing with missing data: Part i. Chemometrics and Intelligent Laboratory Systems, 58(1):15-27, 2001. 58, 1469
[Woj96] P. Wojtaszczyk. Banach spaces for analysts, volume 25. Cambridge University Press, 1996. 1987
[Woo50] Max A Woodbury. Inverting modified matrices. Memorandum report, 42(106):336, 1950. 14
[Woo14a] David P Woodruff. Low rank approximation lower bounds in row-update streams. In Advances in Neural Information Processing Systems (NIPS), pages 1781-1789, 2014. 1922, 1928, 1934, 1946
[Woo14b] David P. Woodruff. Sketching as a tool for numerical linear algebra. Foundations and Trends in Theoretical Computer Science, 10(1-2):1-157, 2014. 52, 53, 770, 1056, 1109, 1110, 1116, 1134, 1377, 1479, 1482, 1516, 1561, 1767, 1781, 1786, 1791, 1992, 2000, 2025, 2027, 2031, 2037, 2042, 2077
[WPB \(\left.{ }^{+} 18\right]\) Lan Wang, Bo Peng, Jelena Bradic, Runze Li, and Yunan Wu. A tuningfree robust and efficient approach to high-dimensional regression. Technical report, School of Statistics, University of Minnesota, 2018. 2026
[Wri73] Farrol Tim Wright. A bound on tail probabilities for quadratic forms in independent random variables whose distributions are not necessarily symmetric. The Annals of Probability, 1(6):1068-1070, 1973. 2505
[Wri97] Stephen J Wright. Primal-dual interior-point methods, volume 54. Siam, 1997. 6
[WS15] Yining Wang and Aarti Singh. Column subset selection with missing data via active sampling. In The 18th International Conference on Artificial Intelligence and Statistics (AISTATS), pages 1033-1041, 2015. 1433, 1476
[WTSA15] Yining Wang, Hsiao-Yu Tung, Alexander J Smola, and Anima Anandkumar. Fast and guaranteed tensor decomposition via sketching. In Advances in Neural Information Processing Systems (NIPS), pages 991-999. https:// arxiv.org/pdf/1506.04448, 2015. 58, 1469, 1518, 1797, 1799, 1800, 1803, 1808, 1809, 1810, 1811, 1813, 1814, 1824, 1826, 1828, 1887
[WW19] Ruosong Wang and David P. Woodruff. Tight bounds for \(\ell_{p}\) oblivious subspace embeddings. In \(S O D A\), pages 1825-1843, 2019. 1966, 1971, 1972, 2031, 2032
[WWS \({ }^{+}\)05] Hongcheng Wang, Qing Wu, Lin Shi, Yizhou Yu, and Narendra Ahuja. Out-of-core tensor approximation of multi-dimensional matrices of visual data. ACM Transactions on Graphics (TOG), 24(3):527-535, 2005. 57, 1469
[WX05] Avi Wigderson and David Xiao. A randomness-efficient sampler for matrixvalued functions and applications. In Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on, pages 397-406. IEEE, 2005. 26, 27, 502, 503, 504, 505, 508
[WX06] Avi Wigderson and David Xiao. Retraction of Wigderson-Xiao: A randomness-efficient sampler for matrix valued functions, eccc tr05-107. 2006. 28, 501, 504
[WX08] Avi Wigderson and David Xiao. Derandomizing the AW matrix-valued Chernoff bound using pessimistic estimators and applications. Theory of Computing, 4(3):53-76, 2008. 505
[WY13] Naiyan Wang and Dit-Yan Yeung. Bayesian robust matrix factorization for image and video processing. In Proceedings of the IEEE International Conference on Computer Vision, pages 1785-1792, 2013. 1334
[WYWY12] Naiyan Wang, Tiansheng Yao, Jingdong Wang, and Dit-Yan Yeung. A probabilistic approach to robust matrix factorization. In European Conference on Computer Vision, pages 126-139. Springer, 2012. 1334
[WZ13] David P. Woodruff and Qin Zhang. Subspace embeddings and \(\ell_{p}\)-regression using exponential random variables. In COLT 2013-The 26th Annual Conference on Learning Theory, June 12-14, 2013, Princeton University, NJ, USA, pages 546-567, 2013. 52, 1109, 1117, 1119, 1456, 1966, 1971
[WZ16] David P Woodruff and Peilin Zhong. Distributed low rank approximation of implicit functions of a matrix. In 32nd IEEE International Conference on Data Engineering (ICDE). https://arxiv.org/pdf/1601.07721, 2016. 1318, 1784, 1788
[WZC \(\left.{ }^{+} 18\right]\) Tsui-Wei Weng, Huan Zhang, Hongge Chen, Zhao Song, Cho-Jui Hsieh, Duane Boning, Inderjit S Dhillon, and Luca Daniel. Towards fast computation of certified robustness for relu networks. In ICML. https://arxiv.org/pdf/ 1804.09699, 2018. xxix
[XBSD \(\left.{ }^{+} 18\right]\) Lechao Xiao, Yasaman Bahri, Jascha Sohl-Dickstein, Samuel S. Schoenholz, and Jeffrey Pennington. Dynamical isometry and a mean field theory of CNNs: How to train 10,000-layer vanilla convolutional neural networks. In International Conference on Machine Learning (ICML), 2018. 358
[XCS10] Huan Xu, Constantine Caramanis, and Sujay Sanghavi. Robust pca via outlier pursuit. In Advances in Neural Information Processing Systems, pages 2496-2504, 2010. 1334, 1370
[XCS11] Liang Xiong, Xi Chen, and Jeff Schneider. Direct robust matrix factorizatoin for anomaly detection. In Data Mining (ICDM), 2011 IEEE 11th International Conference on, pages 844-853. IEEE, 2011. 1334
[XJR02] Eric P Xing, Michael I Jordan, and Stuart Russell. A generalized mean field algorithm for variational inference in exponential families. In Proceedings of the Nineteenth conference on Uncertainty in Artificial Intelligence, pages 583-591. Morgan Kaufmann Publishers Inc., 2002. 14, 74
[XZ14] Lin Xiao and Tong Zhang. A proximal stochastic gradient method with progressive variance reduction. SIAM Journal on Optimization, 24(4):20572075, 2014. 15, 77
[Yan18] Wei Yang. Classification on CIFAR-10/100 and ImageNet with PyTorch, 2018. Accessed: 2018-04. lxxi, lxxii, 265, 344, 345, 346, 347, 348, 349
[Yao77] Andrew Chi-Chin Yao. Probabilistic computations: Toward a unified measure
of complexity. In Foundations of Computer Science, 1977., 18th Annual Symposium on, pages 222-227. IEEE, 1977. 1242
[Yao82] Andrew Chi-Chih Yao. On constructing minimum spanning trees in \(k\) dimensional spaces and related problems. SIAM Journal on Computing, 11(4):721-736, 1982. 1334
[YC09] Jiho Yoo and Seungjin Choi. Weighted nonnegative matrix co-trifactorization for collaborative prediction. In \(A C M L\), 2009. 1044
[YC14] Tatsuya Yokota and Andrzej Cichocki. Multilinear tensor rank estimation via sparse tucker decomposition. In Soft Computing and Intelligent Systems (SCIS), 2014 Joint 7th International Conference on and Advanced Intelligent Systems (ISIS), 15th International Symposium on, pages 478-483. IEEE, 2014. 1706
[YCRM16] Jiyan Yang, Yin-Lam Chow, Christopher Ré, and Michael W Mahoney. Weighted sgd for \(\ell_{p}\) regression with randomized preconditioning. In Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 558-569. Society for Industrial and Applied Mathematics, https://arxiv.org/pdf/1502.03571, 2016. 1717
[YCS11] Yusuf Kenan Yilmaz, Ali Taylan Cemgil, and Umut Simsekli. Generalised coupled tensor factorisation. In Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011. Proceedings of a meeting held 12-14 December 2011, Granada, Spain., pages 2151-2159, 2011. 57, 1469
[YCS16] Xinyang Yi, Constantine Caramanis, and Sujay Sanghavi. Solving a mixture of many random linear equations by tensor decomposition and alternating minimization. In arXiv preprint. https://arxiv.org/pdf/1608.05749, 2016. 57, 1469
[Ye97] Yinyu Ye. Interior point algorithms: theory and analysis. Springer, 1997. 6
[Yel14] Yelp. Yelp dataset. http://www.yelp.com/dataset_challenge, 2014. 1426
[YFS16] Yuning Yang, Yunlong Feng, and Johan AK Suykens. Robust low-rank tensor recovery with regularized redescending m-estimator. IEEE transactions on neural networks and learning systems, 27(9):1933-1946, 2016. 58, 1470
[YLPR15] Dong Yin, Kangwook Lee, Ramtin Pedarsani, and Kannan Ramchandran. Fast and robust compressive phase retrieval with sparse-graph codes. In 2015 IEEE International Symposium on Information Theory (ISIT), pages 2583-2587. IEEE, 2015. 635
[YMM13] J. Yang, X. Meng, and M. Mahoney. Quantile regression for large-scale applications. In \(I C M L\), pages 881-887, 2013. 1966, 1971
[YMM14] Jiyan Yang, Xiangrui Meng, and Michael W. Mahoney. Quantile regression for large-scale applications. SIAM J. Scientific Computing, 36(5), 2014. 1450
[You40] G. Young. Maximum likelihood estimation and factor analysis. Psychometrika, 6(1):49-53, 1940. 50, 1038
[YS17] Greg Yang and Samuel Schoenholz. Mean field residual networks: On the edge of chaos. In Advances in Neural Information Processing Systems (NeurIPS), pages 7103-7114, 2017. 358
[YS18] Greg Yang and Sam S. Schoenholz. Deep mean field theory: Layerwise variance and width variation as methods to control gradient explosion. ICLR open review, 2018. 358
[YTM94] Yinyu Ye, Michael J Todd, and Shinji Mizuno. An \(O(\sqrt{n L})\)-iteration homogeneous and self-dual linear programming algorithm. Mathematics of Operations Research, 19(1):53-67, 1994. 9
[YV18] Grigory Yaroslavtsev and Adithya Vadapalli. Massively parallel algorithms and hardness for single-linkage clustering under \(\ell_{p}\) distances. In International Conference on Machine Learning, pages 5596-5605, 2018. 1917
[YX15] Zai Yang and Lihua Xie. Achieving high resolution for super-resolution via reweighted atomic norm minimization. In Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on, pages 3646-3650. IEEE, 2015. 686, 763
[YZD12] Linbin Yu, Miao Zhang, and Chris Ding. An efficient algorithm for \(\ell_{1}\)-norm principal component analysis. In 2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 1377-1380. IEEE, 2012. 53, 1111, 2506, 2507
[ZBH \({ }^{+}\)17] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In International Conference on Learning Representations (ICLR), 2017. 20, 250, 351
[ZBSG05] Yunhui Zheng, David J Brady, Michaell E Sullivan, and Bob D Guenther. Fiber-optic localization by geometric space coding with a two-dimensional gray code. Applied optics, 44(20):4306-4314, 2005. 33, 617
[ZCF \({ }^{+}\)10] Matei Zaharia, Mosharaf Chowdhury, Michael J Franklin, Scott Shenker, and Ion Stoica. Spark: Cluster computing with working sets. HotCloud, 10(1010):95, 2010. 2315
[ZCS \(\left.{ }^{+} 19\right]\) Huan Zhang, Hongge Chen, Zhao Song, Duane Boning, Inderjit S Dhillon, and Cho-Jui Hsieh. The limitations of adversarial training and the blind-spot attack. In ICLR. arXiv preprint arXiv:1901.04684, 2019. xxix
[ZCZG18] Difan Zou, Yuan Cao, Dongruo Zhou, and Quanquan Gu. Stochastic gradient descent optimizes over-parameterized deep relu networks. In arXiv preprint. https://arxiv.org/pdf/1811.08888, 2018. 1909, 2140
[ZCZJ14] Yuchen Zhang, Xi Chen, Denny Zhou, and Michael I Jordan. Spectral methods meet em: A provably optimal algorithm for crowdsourcing. In \(A d-\) vances in Neural Information Processing Systems (NIPS), pages 1260-1268. https://arxiv.org/pdf/1406.3824, 2014. 57, 1469
[ZF13] Anastasios Zouzias and Nikolaos M. Freris. Randomized extended kaczmarz for solving least squares. SIAM J. Matrix Analysis Applications, 34(2):773793, 2013. 1927
[ZG01] Tong Zhang and Gene H. Golub. Rank-one approximation to high order tensors. SIAM J. Matrix Analysis Applications, 23(2):534-550, 2001. 57, 1469
[ZH05] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 67(2):301-320, 2005. 14, 74
[Zha97] Zhengyou Zhang. Parameter estimation techniques: A tutorial with application to conic fitting. Image and vision Computing, 15(1):59-76, 1997. 56, 1428, 1430
[ZK16] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. arXiv preprint arXiv:1605.07146, 2016. 255
[ZL17] Yi Zhou and Yingbin Liang. Critical points of neural networks: Analytical forms and landscape properties. arXiv preprint arXiv:1710.11205, 2017. 351
[ZLC17] Yuchen Zhang, Percy Liang, and Moses Charikar. A hitting time analysis of stochastic gradient langevin dynamics. In COLT. arXiv preprint arXiv:1702.05575, 2017. 17, 172
[ZLS \(\left.{ }^{+} 12\right]\) Yinqiang Zheng, Guangcan Liu, Shigeki Sugimoto, Shuicheng Yan, and Masatoshi Okutomi. Practical low-rank matrix approximation under robust
\(\ell_{1}\)-norm. In 2012 IEEE Conference on Computer Vision and Pattern Recognition, Providence, RI, USA, June 16-21, 2012, pages 1410-1417, 2012. 53, 1110, 1370
[ZLSD18] Jiong Zhang, Yibo Lin, Zhao Song, and Inderjit S Dhillon. Learning long term dependencies via fourier recurrent units. In International Conference on Machine Learning (ICML), 2018. xxix, 339, 353, 988
[ZMKL05] Cai-Nicolas Ziegler, Sean M McNee, Joseph A Konstan, and Georg Lausen. Improving recommendation lists through topic diversification. In Proceedings of the 14 th international conference on World Wide Web, pages 22-32. ACM, 2005. 1426
[Zol86] Vladimir M Zolotarev. One-dimensional stable distributions, volume 65. American Mathematical Soc., 1986. 667
[Zou12] Anastasios Zouzias. A matrix hyperbolic cosine algorithm and applications. In International Colloquium on Automata, Languages, and Programming, pages 846-858. Springer, 2012. 29
[ZP04] Peng Zhang and Jing Peng. Svm vs regularized least squares classification. In ICPR (1), pages 176-179, 2004. 1923
[ZPB06] Yunhui Zheng, Nikos P Pitsianis, and David J Brady. Nonadaptive group testing based fiber sensor deployment for multiperson tracking. IEEE Sensors Journal, 6(2):490-494, 2006. 33, 617
[ZS16] Ruohan Zhang and Zhao Song. Maximum sustainable yield problem for robot foraging and construction system. In IJCAI, pages 2725-2731, 2016. xxix
[ZSB15] Ruohan Zhang, Zhao Song, and Dana H Ballard. Global policy construction in modular reinforcement learning. In Twenty-Ninth AAAI Conference on Artificial Intelligence, 2015. xxix
[ZSD17] Kai Zhong, Zhao Song, and Inderjit S Dhillon. Learning non-overlapping convolutional neural networks with multiple kernels. arXiv preprint arXiv:1711.03440, 2017. xxix, 21, 25, 251, 257, 351, 548, 2140
[ZSJ \({ }^{+}\)17] Kai Zhong, Zhao Song, Prateek Jain, Peter L. Bartlett, and Inderjit S. Dhillon. Recovery guarantees for one-hidden-layer neural networks. In ICML, 2017. xxix, 21, 25, 57, 251, 257, 351, 358, 548, 1451, 1469, 2140
[ZSJD18] Kai Zhong, Zhao Song, Prateek Jain, and Inderjit S Dhillon. Nonlinear inductive matrix completion based on one-layer neural networks. arXiv preprint arXiv:1805.10477, 2018. xxix
[ZW13] Syed Zubair and Wenwu Wang. Tensor dictionary learning with sparse tucker decomposition. In Digital Signal Processing (DSP), 2013 18th International Conference on, pages 1-6. IEEE, 2013. 1706
[ZWX \({ }^{+}\)17] Shun Zheng, Jialei Wang, Fen Xia, Wei Xu, and Tong Zhang. A general distributed dual coordinate optimization framework for regularized loss minimization. The Journal of Machine Learning Research, 18(1):4096-4117, 2017. 13, 74
[ZWZ16] Junyu Zhang, Zaiwen Wen, and Yin Zhang. Subspace methods with local refinements for eigenvalue computation using low-rank tensor-train format. Journal of Scientific Computing, pages 1-22, 2016. 1711
[ZX17] Yuchen Zhang and Lin Xiao. Stochastic primal-dual coordinate method for regularized empirical risk minimization. The Journal of Machine Learning Research, 18(1):2939-2980, 2017. 13, 58, 74, 1470
[ZYJ17] Lijun Zhang, Tianbao Yang, and Rong Jin. Empirical risk minimization for stochastic convex optimization: \(O(1 / n)\)-and \(O\left(1 / n^{2}\right)\)-type of risk bounds. arXiv preprint arXiv:1702.02030, 2017. 13, 74
[ZZGM15] Xiaodong Zheng, Shanfeng Zhu, Junning Gao, and Hiroshi Mamitsuka. Instance-wise weighted nonnegative matrix factorization for aggregating partitions with locally reliable clusters. In IJCAI, 2015. 1044

\section*{Vita}

Zhao Song was born in Xi'an, Shaanxi, China, the son of Li Song (father) and Mei Sun (mother), the grandson of Kexin Song (father's father), Weifeng Guan (father's mother), Yongzhe Sun (mother's father), Fenglin Yang (mother's mother).

Permanent address: Jinhua North Road 280, C-22-06
Xi'an, Shaanxi, China 710032

This dissertation was typeset with \(\mathrm{AT}_{\mathrm{E}} \mathrm{X}^{\dagger}\) by the author.

\footnotetext{
\({ }^{\dagger}{ }^{1} T_{E} \mathrm{X}\) is a document preparation system developed by Leslie Lamport as a special version of Donald Knuth's TEX Program.
}```


[^0]:    ${ }^{1}$ Throughout this thesis, we assume there is no redundant constraints and hence $n \geq d$. Note that papers in different communities uses different symbols to denote the number of variables and constraints in a linear program.
    ${ }^{2}$ We use $O^{*}$ to hide $n^{o(1)}$ and $\log ^{O(1)}(1 / \delta)$ factors and $\widetilde{O}$ to hide $\log ^{O(1)}(n / \delta)$ factors.

[^1]:    ${ }^{3}$ The dual exponent of matrix multiplication $\alpha$ is the supremum among all $a \geq 0$ such that it takes $n^{2+o(1)}$ time to multiply an $n \times n$ matrix by an $n \times n^{a}$ matrix.

[^2]:    ${ }^{4}$ [Dal17] works in a parameter regime where the weight changes of all layers except the last one make negligible contribution to the final output.

[^3]:    ${ }^{5}$ Intuitively, an RNN recurrently applies the same network unit for $L$ times if the input sequence is of length $L$. When this unit has "operator norm" larger than one or smaller than one, the final output can possibly exponentially explode or vanish in $L$. More importantly, when one back propagates through time - which intuitively corresponds to applying the reverse unit multiple times- the gradient can also vanish or explode. Controlling the operator norm of a non-linear operator can be quite challenging.

[^4]:    ${ }^{6}$ That is the first vertex $v_{1}$ is chosen uniformly at random - which is the stationary distribution of the graph $G$.

[^5]:    ${ }^{7}$ But not the only one: pseudorandom permutations for sparse FT in high dimensions also incur an exponential loss, and it is not known whether this can be avoided.

[^6]:    ${ }^{8}$ We use $f \lesssim g$ to denote that $f \leq C g$ for some universal constant $C$.

[^7]:    ${ }^{9}$ https://en.wikipedia.org/wiki/Tensor_rank_decomposition\#Border_rank

[^8]:    ${ }^{1}$ Throughout this paper, we assume there is no redundant constraints and hence $n \geq d$. Note that papers in different communities uses different symbols to denote the number of variables and constraints in a linear program.
    ${ }^{2}$ We use $O^{*}$ to hide $n^{o(1)}$ and $\log { }^{O(1)}(1 / \delta)$ factors and $\widetilde{O}$ to hide $\log ^{O(1)}(n / \delta)$ factors.

[^9]:    ${ }^{3}$ The dual exponent of matrix multiplication $\alpha$ is the supremum among all $a \geq 0$ such that it takes $n^{2+o(1)}$ time to multiply an $n \times n$ matrix by an $n \times n^{a}$ matrix.

[^10]:    ${ }^{4}$ The classical view of central path method is to take a Newton step on the system (2.2), which turns out to be same as taking a gradient step on the energy defined in (2.3). However, our main algorithm will choose a different energy and this gradient descent view is crucial for designing our algorithm.

[^11]:    ${ }^{5}$ If the input is deterministic, so is the output and the runtime.

[^12]:    ${ }^{1}$ The logistic function is $g(t)=\frac{1}{1+e^{-t}}$ and the logistic loss is $-\log (g(t))=\log \left(1+e^{-t}\right)$.

[^13]:    ${ }^{1}$ [Dal17] works in a parameter regime where the weight changes of all layers except the last one make negligible contribution to the final output.

[^14]:    ${ }^{2}$ A recurrent neural network executed on input sequences with time horizon $L$ is very similar to a feedforward neural network with $L$ layers. The main difference is that in a feedforward network, weight matrices are different across layers, and thus independently randomly initialized; in contrast, in an RNN, the same weight matrix is applied across the entire time horizon, so we do not have fresh new randomness for proofs that involve in induction. In other words, the over-parameterized convergence theory of DNN is much simpler than that of RNN.

[^15]:    ${ }^{3}$ Without loss of generality, one can re-scale and assume $\left\|x_{i}\right\| \leq 1 / \sqrt{2}$ for every $i \in[n]$. Again, without loss of generality, one can pad each $x_{i}$ by an additional coordinate to ensure $\left\|x_{i}\right\|=1 / \sqrt{2}$. Finally, without loss of generality, one can pad each $x_{i}$ by an additional coordinate $\frac{1}{\sqrt{2}}$ to ensure $\left\|x_{i}\right\|=1$. This last coordinate $\frac{1}{\sqrt{2}}$ is equivalent to introducing a (random) bias term, because $A\left(\frac{y}{\sqrt{2}}, \frac{1}{\sqrt{2}}\right)=\frac{A}{\sqrt{2}}(y, 0)+b$ where $b \sim \mathcal{N}\left(0, \frac{1}{m} I\right)$. In our proofs, the specific constant $\frac{1}{\sqrt{2}}$ does not matter.
    ${ }^{4}$ If $\left\|y_{i}^{*}\right\| \leq \Omega$ for some parameter $\Omega$, our complexities shall also grow in poly $(\Omega)$.

[^16]:    ${ }^{5}$ We note that if one jointly trains all the layers, in certain parameter regimes, it may be equivalent to as if only the last layer is trained [Dal17]. We therefore choose to fix the last layer $B$ to avoid such confusion.

[^17]:    ${ }^{6}$ This spectral norm bound seems small, but is in fact quite large: it can totally change the outputs and fit the training data, because weights are randomly initialized (per entry) at around $\frac{1}{\sqrt{m}}$ for $m$ being large.

[^18]:    ${ }^{7}$ Although Theorem 5.15 .1 was not explicitly stated until version 5 of this paper, its proof was fully contained in the proofs of Theorem 5.13.1 and 5.14.1. Since some readers cannot find it, we state it here as a separate theorem.

[^19]:    ${ }^{8}$ See (5.19) and (5.21) in the proofs.

[^20]:    ${ }^{9}$ Namely, if the number of sign changes is bounded in all layers, then $h_{i, \ell}$ and $\widetilde{h}_{i, \ell}$ cannot be too far away by applying matrix concentration; and reversely, if $h_{i, \ell}$ is not far from $\widetilde{h}_{i, \ell}$ in all layers, then the number of sign changes per layer must be small.

[^21]:    ${ }^{10}$ This is the only technical idea that we borrowed from [LL18], which is the over-parameterization theory for 2-layer neural networks.

[^22]:    ${ }^{11}$ In principle, each such proof may require a careful rewriting of the main body of this paper. We choose to sketch only the proof difference (in the appendix) in order to keep this paper short. If there is sufficient interest from the readers, we can consider adding the full proofs in the future revision of this paper.
    ${ }^{12}$ Note that the loss function when combined with the neural network together $f\left(B h_{i, L}\right)$ is not gradient dominant. Therefore, one cannot apply classical theory on gradient dominant functions to derive our same result.

[^23]:    ${ }^{13}$ Again, this cannot be derived from classical theory of finding approximate saddle points for non-convex functions, because weights $\vec{W}$ with small $\left\|\nabla f\left(B h_{i, L}\right)\right\|$ is a very different (usually much harder) task comparing to having small gradient with respect to $\vec{W}$ for the entire composite function $f\left(B h_{i, L}\right)$.
    ${ }^{14}$ This is because attaining constant objective error $\epsilon=1 / 4$ for the cross-entropy loss suffices to imply perfect training accuracy.

[^24]:    ${ }^{15}$ More specifically, we can let $X_{k}=m\left(\phi\left(W_{\ell} h_{i, \ell-1}\right)-\phi\left(W_{\ell} h_{j, \ell-1}\right)\right)_{k}^{2}$ which is $O(1)$-subgaussian and let $X=X_{1}+\cdots+X_{m}$. We have $\operatorname{Pr}\left[X \geq \mathbb{E}[X]\left(1-\delta_{\ell-1}\right)\right] \geq 1-e^{-\Omega\left(\delta_{\ell-1}^{2} \mathbb{E}[X]\right)}$.

[^25]:    ${ }^{16}$ Alternatively, one can fully specify all the constants without using the big- $O$ notation. This was done in our prior work [AZLS18] but is notation-heavy. We refrain from doing so in this simplified paper.

[^26]:    ${ }^{17}$ That is, $f\left(z+z^{\prime} ; y\right) \leq f(z)+\left\langle\nabla f(z ; y), z^{\prime}\right\rangle+\frac{1}{2}\left\|z^{\prime}\right\|^{2}$.

[^27]:    ${ }^{18}$ Recall $\frac{\partial f(z ; y)}{z_{y}}=p_{y}\left(1-p_{y}\right)$ where $p_{j}=\frac{e^{z_{j}}}{\sum_{i=1}^{d} e^{z_{i}}}$. If $p_{y}>1 / 2$, then $z$ correctly predicts the target label $y$ because $p_{y}>p_{j}$ for $j \neq z$.

[^28]:    ${ }^{19}$ We acknowledge the existence of more careful modifications to avoid loosing too many such factors, but do not present such result for the simplicity of this paper.
    ${ }^{20}$ We note that in all of our applications of Lemma 5.8.1, the minimal choice of $\epsilon$ is around $\delta^{3}$ from the proof of $\delta$-separateness. Therefore, choosing $\tau=\frac{\delta^{2}}{100 L}$ is safe. We are aware of slightly more involved proofs that are capable of handling much larger values of $\tau$.

[^29]:    ${ }^{21}$ For simplicity, we only show how to modify Lemma 5.8 .5 with success probability $1-e^{-\Omega\left(\log ^{2} m\right)}$ because that is all we need to the downstream application of Lemma 5.8.5. If one is willing to be more careful, the success probability can be much higher.

[^30]:    ${ }^{1}$ Intuitively, an RNN recurrently applies the same network unit for $L$ times if the input sequence is of length $L$. When this unit has "operator norm" larger than one or smaller than one, the final output can possibly exponentially explode or vanish in $L$. More importantly, when one back propagates through time -which intuitively corresponds to applying the reverse unit multiple times - the gradient can also vanish or explode. Controlling the operator norm of a non-linear operator can be quite challenging.

[^31]:    ${ }^{2}$ Most of the technical lemmas of this paper remain to hold (and become much simpler) once $W$ is replaced with $W_{\ell}$. This is carefully treated by [AZLS19].

[^32]:    ${ }^{3}$ At a first glance, one may question how it is possible for SGD to enjoy a logarithmic time dependency in $\epsilon^{-1}$; after all, even when minimizing strongly-convex and Lipschitz-smooth functions, the typical convergence rate of SGD is $T \propto 1 / \epsilon$ as opposed to $T \propto \log (1 / \epsilon)$. We quickly point out there is no contradiction here if the stochastic pieces of the objective enjoy a common global minimizer. In math terms, suppose we want to minimize some function $f(x)=\frac{1}{n} \sum_{i=1}^{n} f_{i}(x)$, and suppose $x^{*}$ is the global minimizer of convex functions $f_{1}(x), \ldots, f_{n}(x)$. Then, if $f(x)$ is $\sigma$-strongly convex, and each each $f_{i}(x)$ is $L$-Lipschitz smooth, then SGD -moving in negative direction of $\nabla f_{i}(x)$ for a random $i \in[n]$ per step-can find $\epsilon$-minimizer of this function in $O\left(\frac{L^{2}}{\sigma^{2}} \log \frac{1}{\epsilon}\right)$ iterations.
    ${ }^{4}$ Our theorem holds even when $A, B$ are at random initialization and only the hidden weight matrix $W$ is trained. This is much more difficult to analyze than the convex task of training only the last layer $B$ [Dal17]. Training only the last layer can significantly reduce the learning power of (recurrent or not) neural networks in practice.

[^33]:    ${ }^{5}$ If it only satisfies $\left\|x_{i, 1}\right\| \leq 1$ one can pad it with an additional coordinate to make $\left\|x_{i, 1}\right\|=1$ hold. As for the assumption $\left[x_{i, 1}\right]_{d_{x}}=\frac{1}{\sqrt{2}}$, this is equivalent to adding a bias term $\mathcal{N}\left(0, \frac{1}{m}\right)$ for the first layer.

[^34]:    ${ }^{6}$ The index $\ell$ starts from 2, because $B h_{i, 1}=B \phi\left(A x_{i, 1}\right)$ remains constant if we are not optimizing over $A$ and $B$.

[^35]:    ${ }^{7}$ We remark here that the $\widetilde{O}$ notation may hide additional polynomial dependency in $\log \log \epsilon^{-1}$. This is not necessary, at the expense of slightly complicating the proofs, as shown by follow up [AZLS19].

[^36]:    ${ }^{8}$ For instance, proving gradient norm lower bound in Theorem 6.15 .2 for a single neuron $k \in[m]$ is easy, but how to apply concentration across neurons? Crucially, due to the recurrent structure these quantities are never independent, so we have to build necessary probabilistic tools to tackle this. If one is willing to ignore such subtleties, then our sketched proof is sufficiently short and gives a good overview.

[^37]:    ${ }^{9}$ More precisely, letting $v=\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}$, we have $W\left(I-U_{\ell-1} U_{\ell-1}^{\top}\right) h_{i, \ell}=\left(W \frac{v}{\|v\|}\right)\|v\|$. Here, $W \frac{v}{\|v\|}$ is a random Gaussian vector in $\mathcal{N}\left(0, \frac{2}{m} I\right)$ and is independent of all $\left\{h_{i, a} \mid i \in[n], a<\ell\right\}$.

[^38]:    ${ }^{10}$ This is the only place that we rely on Assumption 6.2.1. This assumption is somewhat necessary in the following sense. If $x_{i, \ell}=x_{j, \ell}$ for some pair $i \neq j$ for all the first ten layers $\ell=1,2, \ldots, 10$, and if $y_{i, \ell}^{*} \neq y_{i, \ell}^{*}$ for even just one of these layers, then there is no hope in having the training objective decrease to zero. Of course, one can make more relaxed assumption on the input data, involving both $x_{i, \ell}$ and $y_{i, \ell}^{*}$. While this is possible, it complicates the statements so we do not present such results in this paper.

[^39]:    ${ }^{12}$ We use $(\operatorname{poly}(\rho))^{2}$ to emphasize that the first polynomial needs to be bigger than the second poly $(\rho)$.

[^40]:    ${ }^{13}$ Indeed, after Gram-Schmidt we can write $U_{\ell-1}=\left[\widehat{h}_{1}, \ldots, \widehat{h}_{n(\ell-1)}\right]$, where each $\widehat{h}_{j}$ only depends on the randomness of $A$ and $W\left[\widehat{h}_{1}, \ldots, \widehat{h}_{j-1}\right]$. In other words, conditioning on any choice of $A$ and $W\left[\widehat{h}_{1}, \ldots, \widehat{h}_{j-1}\right]$, we still have $W \widehat{h}_{j}$ is an independent Gaussian vector from $\mathcal{N}\left(0, \frac{2}{m} I\right)$. Similarly, $v_{i, \ell}$ may depend on the randomness of $A$ and $W U_{\ell-1}$, but conditioning on any choice of $A$ and $W U_{\ell-1}$, we still have $W v_{i, \ell}$ follows from $\mathcal{N}\left(0, \frac{2}{m} I\right)$. This proves that the entries of $M_{1}, M_{2}, M_{3}$ are independent.

[^41]:    ${ }^{14}$ Note that if $z_{1}$ and $z_{2}$ are random, then they satisfy such constraints by Lemma 6.11.1.

[^42]:    ${ }^{15}$ The proof of the base case is a replication of Claim 6.11.7 and only simpler. Indeed, for the base case of $\ell=0$, one can view $h_{i, \ell}=x_{i, 1}$ and $h_{j, \ell}=x_{j, 1}$, and view $U_{\ell}$ as an empty matrix. Then, the same analysis of Claim 6.11 .7 but with slightly different notations will apply.

[^43]:    ${ }^{16}$ Note that if $z_{1}, z_{2}$ and $z_{4}$ are random, then they satisfy such constraints by Lemma 6.11.1.

[^44]:    ${ }^{17}$ This follows from a similar argument as Footnote 13, taking into account the additional fact that each $z_{j}$ may only depend on the randomness of $A, W U_{\ell-1}$, and $W z_{b-1}, \ldots W z_{j-1}$.

[^45]:    ${ }^{18}$ Indeed,

    $$
    \widetilde{g}_{\ell-1}=\widetilde{W} \widetilde{h}_{\ell-2}+A x_{\ell-1}=M_{1} z_{1}+M_{2} z_{2}
    $$

[^46]:    ${ }^{19}$ Since we have done this too many times, let us quickly point out the calculation. By Corollary 6.11.16,

[^47]:    ${ }^{1}$ That is the first vertex $v_{1}$ is chosen uniformly at random - which is the stationary distribution of the graph $G$.

[^48]:    ${ }^{2}$ The similar version is an exercise 4 in page 163 (Section VII) of [Con78], and also can be found here, https://math.stackexchange.com/questions/882147/ find-a-conformal-map-from-semi-disk-onto-unit-disk

[^49]:    ${ }^{1}$ [SS11] analyzed sampling with replacement, but based on [Tro12], a folklore result shows the same behavior can be obtained by doing independent coin flips for every edge with low leverage score, again with inclusion probabilities proportional to leverage scores.

[^50]:    ${ }^{2}$ Their full statement is more general, see [AG15] Corollary 1.9.

[^51]:    ${ }^{3}$ Note, however, that we are able to prove deterministic bounds on the predictable quadratic variation

[^52]:    ${ }^{1}$ This was pointed out to us by Tarun Kathuria.

[^53]:    ${ }^{2}$ Which approximates the complete bipartite graph.
    ${ }^{3}$ More precisely, it is the largest non-trivial eigenvalue that is being controlled. The largest eigenvalue of the adjacency matrix of a $d$-regular graph is the trivial eigenvalue of $d$.

[^54]:    ${ }^{4}$ Gil Kalai suggested the connection between that result and the Kadison-Singer problem.

[^55]:    ${ }^{1}$ the lower bounds in [LNN15] also depend heavily on the streaming model, so they do not transfer necessarily to all scenarios where sparse recovery finds application.

[^56]:    ${ }^{2}$ The authors do not specifically address the exponent in the poly $(\log n)$, but we estimate it to be $\geq 4$.

[^57]:    ${ }^{3}$ We note that this term is exactly $k$, not $O(k)$. Although not important for our main result, it will be crucial for some of our applications of our techniques.

[^58]:    ${ }^{1}$ We use $f \lesssim g$ to denote that $f \leq C g$ for some universal constant $C$.

[^59]:    ${ }^{1}$ There is a catch here, that the coefficients of the exponentials are potentially unbounded, if the frequencies are arbitrarily close together. We first use Gram determinants to show that the signal is $\delta$-close to one

[^60]:    ${ }^{2}$ Outputting $Q=\underset{\text { degree- } d}{\arg \min }\left\|Q^{\prime}-x\right\|_{S, w}$ is not good enough, because it only gives constant success probability.

[^61]:    ${ }^{3}$ The proofs in this section are identical to [HIKP12a] and [PS15].

[^62]:    ${ }^{4}$ We will set $B$ to be $O(k), D$ to be poly $(k)$ and $\sigma$ to be $T / \operatorname{poly}(k)$.

[^63]:    ${ }^{5}$ We omitted the proofs here, because the proofs are identical to the one-cluster situation.

[^64]:    ${ }^{1}$ We can also generalize its CUR decomposition algorithms by using our techniques, but it is not the main focus in this work.

[^65]:    ${ }^{1}$ We use the notation $\widetilde{O}(f):=O\left(f \cdot \log { }^{O(1)} f\right)$.

[^66]:    ${ }^{1}$ Recall the Frobenius norm $\|A\|_{F}$ of a matrix $A$ is $\left(\sum_{i=1}^{n} \sum_{j=1}^{d} A_{i, j}^{2}\right)^{1 / 2}$.

[^67]:    ${ }^{2}$ https://en.wikipedia.org/wiki/Tensor_rank_decomposition\#Border_rank

[^68]:    ${ }^{3}$ The entries of $A$ are assumed to fit in $n^{\delta}$ words.

[^69]:    ${ }^{4} \mathrm{~T}$ denotes the tube which is the column in 3rd dimension of tensor.

[^70]:    ${ }^{5}$ For simplicity, we define $U \otimes V \otimes W=\sum_{i=1}^{k} U_{i} \otimes V_{i} \otimes W_{i}$, where $U_{i}$ is the $i$-th column of $U$.
    ${ }^{6} \operatorname{vec}\left(V_{i}^{*} \otimes W_{i}^{*}\right)$ denotes a row vector that has length $n_{1} n_{2}$ where $V_{i}^{*}$ has length $n_{1}$ and $W_{i}^{*}$ has length $n_{2}$.
    ${ }^{7}\left(V^{* \top} \odot W^{* \top}\right)$ denotes a $k \times n_{1} n_{2}$ matrix where the $i$-th row is $\operatorname{vec}\left(V_{i}^{*} \otimes W_{i}^{*}\right)$, where length $n_{1}$ vector $V_{i}^{*}$ is the $i$-th column of $n_{1} \times k$ matrix $V^{*}$, and length $n_{2}$ vector $W_{i}^{*}$ is the $i$-th column of $n_{2} \times k$ matrix $W^{*}, \forall i \in[k]$.

[^71]:    ${ }^{8}$ https://arxiv.org/pdf/1504.06729v1.pdf

[^72]:    ${ }^{9}$ The entries of $A$ are assumed to fit in $n^{\delta}$ words.

[^73]:    ${ }^{10}$ The entries of $A$ and $W$ are assumed to fit in $n^{\delta}$ words.

[^74]:    ${ }^{11}$ The entries of $A$ and $W$ are assumed to fit in $n^{\delta}$ words.

[^75]:    ${ }^{12}$ Personal communication with Russell Impagliazzo and Ryan Williams.

[^76]:    ${ }^{13}$ The first two parts are accomplished by personal communication with Dana Moshkovitz and Govind Ramnarayan.

[^77]:    ${ }^{1}$ http://yining-wang.com/fftlda-code.zip

[^78]:    ${ }^{2}$ https://github.com/huanzhang12/sampling_tensor_decomp/

[^79]:    ${ }^{3}$ http://www.cs.cmu.edu/~enron/
    ${ }^{4}$ http://www.cs.princeton.edu/~blei/lda-c/
    ${ }^{5}$ https://archive.ics.uci.edu/ml/machine-learning-databases/bag-of-words
    ${ }^{6}$ https://archive.ics.uci.edu/ml/machine-learning-databases/bag-of-words
    ${ }^{7}$ http://archive.ics.uci.edu/ml/datasets/NSF+Research+Award+Abstracts+1990-2003

[^80]:    ${ }^{1}$ Personal communication with James Renegar

[^81]:    ${ }^{1}$ https://en.wikipedia.org/wiki/Moore-Penrose_pseudoinverse

[^82]:    ${ }^{2}$ https://en.wikipedia.org/wiki/Hypergeometric_distribution

[^83]:    ${ }^{1}$ Throughout the paper, we use $\widetilde{O}(f)$ to denote $f$ polylog $f$.

[^84]:    ${ }^{2}$ https://sublinear.info/index.php?title=Open_Problems:30.

[^85]:    ${ }^{3}$ Even for $\ell_{p}$ norms with $p<2$, embeddings based on random projections will necessarily induce a distortion factor polynomial in $d$, as shown in [WW19].

[^86]:    ${ }^{4}$ See, e.g., [Mah11a], for a survey on leverage scores.

[^87]:    ${ }^{5}$ Alternatively, we can use the oblivious subspace embedding in $[A L S+18]$ for this step. However, as we have discussed, the oblivious subspace embedding in [ALS $\left.{ }^{+} 18\right]$ requires stronger assumptions on the function $G: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ than those in Assumption 25.1.2, which restricts the class of Orlicz norms to which our algorithm can be applied.

[^88]:    ${ }^{1}$ We remark that while the $\operatorname{nnz}(b)$ term is not written in the Theorem of [DSSW18], their approach of leverage score sampling from a well-conditioned basis requires one to sample from a well conditioned basis of $\left[A_{1} \otimes A_{2}, b\right]$ for a subspace embedding. As stated, their algorithm only sampled from $\left[A_{1} \otimes A_{2}\right]$. To fix this omission, their algorithm would require an additional $n n z(b)$ time to leverage score sample from the augmented matrix.

[^89]:    ${ }^{2}$ For a function $f(n, d, \epsilon, \delta), \widetilde{O}(f)=O(f \cdot \operatorname{poly}(\log n))$, where we assume $n \geq d, \epsilon^{-1}, \delta^{-1}$

[^90]:    ${ }^{3}$ In discussion with the authors of these works, the original $O\left((d \log d)^{1 / p}\right)$ distortion factors stated in these papers should be replaced with $O(d \log d)$; as we do not optimize the poly $(d)$ factors in our analysis, this does not affect our bounds.

[^91]:    ${ }^{4}$ To amplify the probability, we can sketch $A$ and $A U^{\top} U$ with a sparse JL matrix (e.g., Lemma 26.10.1 with $k_{i}=\Theta\left(q k^{2} /\left(\delta \epsilon^{2}\right)\right)$ for each $\left.i\right)$ in input sparsity time to estimate the cost of a given solution. We can then repeat $\log (1 / \delta)$ times and take the minimum to get failure probability $1-\delta$.

[^92]:    ${ }^{1}$ It is easy to see that $k$ points are needed in a coreset - when there are only $k$ points in the dataset, the optimal $k$-means cost is 0 , so a coreset has to contain all $k$ points.
    ${ }^{2}$ It is also possible to obtain an $\widetilde{O}\left(k^{2} \cdot \operatorname{poly}(d)\right)$ space algorithm for dynamic streams by combining the techniques from [Che09] and $\left[\mathrm{BFL}^{+} 17\right]$.

[^93]:    ${ }^{3}$ We assume for now that we know OPT. Our actual algorithm uses exponential search to guess the value of OPT.

[^94]:    ${ }^{4}$ At any time during the stream, for any point $p \in[\Delta]^{d}$, the number of deletions of $p$ so far is always no more than the number of insertions of $p$.

[^95]:    ${ }^{5}$ At the end of the stream, for any point $p \in[\Delta]{ }^{d}$, the number of insertions of $p$ is at most one more than the number of deletions of $p$.

[^96]:    ${ }^{6}$ The proof is slightly different since Theorem 27.3 .1 only claims a constant success probability. See Section 27.3.3 for the detailed proof of Theorem 27.3.1.

[^97]:    ${ }^{1}$ For simplicity, let us assume $\gamma=0$.

[^98]:    ${ }^{1}$ Since the definitions are equivalent up to a factor of 2 (each substitution is an insertion and a deletion), this difference is irrelevant as we consider constant factor approximations of edit distance.
    ${ }^{2}$ One may also consider the case of intermediate alphabet size, e.g. $\Sigma=\{A, C, G, T\}$. It is plausible that

[^99]:    ${ }^{3}$ The near-linear time approximation algorithms for edit distance [KS19, BR19] also incur a sublinear additive error term, but that is OK for our reduction.

[^100]:    ${ }^{4}$ This is wlog since $|1(A)-0(B)|=|0(A)-1(B)|$, so the premise of this lemma is symmetric.

[^101]:    ${ }^{5}$ In particular, in Eq. (29.3) we require that $\beta^{\prime}+\gamma<\frac{1}{4 c+2}$, where $c$ is the approximation factor of the edit distance algorithm.

[^102]:    ${ }^{6}$ Note that we do not define a $\widehat{M_{B}}$.

[^103]:    ${ }^{1}$ Alternatively, the cost of a substitution is doubled as it requires a deletion and an insertion.
    ${ }^{2}$ By truly sub-quadratic we mean $O\left(n^{2-\epsilon}\right)$, for any constant $\epsilon>0$

[^104]:    ${ }^{3}$ We define $\widetilde{O}(f)$ to be $f \cdot \log { }^{O(1)}(f)$.

[^105]:    ${ }^{4}$ For example, $A_{1}=0^{n / 2} 0^{n / 2}, A_{2}=0^{n / 2} 1^{n / 2}, A_{3}=1^{n / 2} 1^{n / 2}$.
    ${ }^{5}$ When we use our solution to approximate edit distance, we can safely assume that $\|\operatorname{lcs}(A, B)\|=\Omega(1)$ since otherwise the edit distance of the two strings is very close to $2 n$.

[^106]:    ${ }^{6}$ Note that in every instance in each string there is a character that appears with frequency at least $1 /|\Sigma|$, but in general that may not be the same character.
    ${ }^{7}$ Their algorithm obtains an additive error of $\delta n$ in time $2^{\widetilde{O}(1 / \delta)}$. When the solution size is bounded by $\lambda n$, one needs to set $\delta<\lambda$ in order to guarantee a multiplicative factor approximation.

[^107]:    ${ }^{8}$ There is a term $(1 / \lambda)^{1 / \lambda}$ in the running time.

[^108]:    ${ }^{9}$ The equality holds if we assume $\delta=\Omega(1)$.

[^109]:    ${ }^{1}$ The spectral gap of a graph $G$ is the second smallest eigenvalue of the normalized Laplacian of $G$.
    ${ }^{2}$ We can construct $G$ as the following: a bridge connects two 3-regular expanders where each expander has $n / 2$ vertices.

[^110]:    ${ }^{3}$ The construction of $S_{v}$ for spanning forest algorithm is slightly different from that described in the connectivity algorithm. $S_{v}$ in spanning forest algorithm has a stronger property: $\forall u \in V \backslash S_{v}, \operatorname{dist}_{G}(u, v)$ must be at least $\operatorname{dist}_{G}\left(u^{\prime}, v\right)$ for any $u^{\prime} \in S_{v}$.

