

Copyright
by
Anirudh Srinivasan
2023

The Thesis Committee for Anirudh Srinivasan
certifies that this is the approved version of the following thesis:

**Computational Modeling of Politeness
Across Diverse Languages**

SUPERVISING COMMITTEE:

Eunsol Choi, Supervisor

Kyle Mahowald

**Computational Modeling of Politeness
Across Diverse Languages**

by
Anirudh Srinivasan

Thesis

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

Master of Science in Computer Science

**The University of Texas at Austin
May 2023**

Abstract

Computational Modeling of Politeness Across Diverse Languages

Anirudh Srinivasan, MSCS
The University of Texas at Austin, 2023

SUPERVISOR: Eunsol Choi

We study politeness phenomena in nine typologically diverse languages. Politeness is an important facet of communication and is sometimes argued to be cultural-specific, yet existing computational linguistic study is limited to English. We create TyDiP, a dataset containing three-way politeness annotations for 500 examples in each language, totaling 4.5K examples. We evaluate how well multilingual models can identify politeness levels – they show a fairly robust zero-shot transfer ability, yet fall short of estimated human accuracy significantly. We further study mapping the English politeness strategy lexicon into nine languages via automatic translation and lexicon induction, analyzing whether each strategy’s impact stays consistent across languages. Lastly, we empirically study the complicated relationship between formality and politeness through transfer experiments. We hope our dataset will support various research questions and applications, from evaluating multilingual models to constructing polite multilingual agents. The data and code is publicly available at on GitHub <https://github.com/Genius1237/TyDiP> and HuggingFace <https://huggingface.co/datasets/Genius1237/TyDiP>.

Table of Contents

Chapter 1: Introduction	6
Chapter 2: Related Work	9
Chapter 3: TYDiP: Multilingual Politeness Dataset	11
3.1 TYDiP: Multilingual Politeness Dataset	11
3.1.1 Annotation Process	12
3.1.2 Inter-annotator Agreement	14
3.1.3 Final Dataset	16
Chapter 4: Predicting Politeness	20
4.1 Predicting Politeness	20
Chapter 5: Building and Analyzing Politeness Strategies in Nine Languages	24
5.1 Building and Analyzing Politeness Strategies in Nine Languages	24
Chapter 6: Transfer between Formality and Politeness	29
6.1 Transfer between Formality and Politeness	29
Chapter 7: Conclusion	33
Appendix A: Appendix	35
A.1 Annotation UI	35
Works Cited	37

Chapter 1: Introduction

Whether politeness phenomena and strategies are universal across languages or not have been controversial among sociologists and linguists. While Brown and Levinson (1978) claimed their universality, other followup work (Korac-Kakabadse et al., 2001) claimed how communication patterns can differ based on cultures and other social constructs such as gender (Mills, 2003) and domains.

To contribute to the linguistic study of cross-cultural politeness, we collect politeness labels on nine typologically and culturally diverse languages, Hindi, Korean, Spanish, Tamil, French, Vietnamese, Russian, Afrikaans, and Hungarian. This language set covers five scripts and eight language families. We follow the seminal work (Danescu-Niculescu-Mizil et al., 2013) closely, focusing on politeness exhibited in requests as they involve the speaker imposing on the listener, requiring them to employ various politeness techniques. To capture rich linguistic strategies that can be lost in translation (Lembersky et al., 2011), we collect sentences written in each target language. To minimize the domain shift among languages, we collect examples in each language from their respective Wikipedia User talk pages, where editors make requests about administrative and editorial decisions.

Crowdsourcing labels in low-resource languages is challenging as recruiting and validating crowdworkers is nontrivial. Thus, we carefully design an annotation process that includes a translation task to evaluate annotator’s language proficiency and a model-in-the-loop qualification task which filters workers whose labels diverge from highly confident predictions from multilingual models. After this process, we observe high agreement among the annotators in our dataset despite the subjectivity of the task. Interestingly, the annotators agree with each other more when assigning politeness score on requests in their native languages compared to assigning politeness score on requests in English, which is their second language.

Equipped with our new multilingual politeness dataset, we evaluate zero-shot transfer ability of existing multilingual models in predicting politeness – subjective and pragmatic language interpretation task. Pretrained language models (Conneau et al., 2020) fine-tuned on annotated English politeness data (Danescu-Niculescu-Mizil et al., 2013) show competitive performances on all languages, weighing in the universality of politeness phenomena across languages. We also witness impressive zero-shot performance of a high-capacity pretrained language model (Brown et al., 2020). We observe a degradation in classification performances when we translate the target language (via Google Translate API) to English, suggesting politeness might not be preserved in the current machine translation model. Despite the simplicity of classification task, we report a substantial difference between the estimated human accuracy and the best model accuracy (over 10% difference in accuracy in six out of nine languages).

Lastly, we provide two studies delving into politeness phenomena. We map English politeness strategy lexicon to create politeness strategy lexicon in nine languages by using tools like automatic translation, lexicon alignment (Dou and Neubig, 2021) and large-scale corpora in the same domain. Despite the limitations of automatic lexicon mapping, we largely observe consistent correlation with politeness score for each politeness strategy in nine languages we study, with some interesting exceptions. We then compare the notion of politeness and formality which has been studied in multilingual setting (Briakou et al., 2021). Our empirical results supports that notions of politeness and formality cannot be used interchangeably. However, when we control for semantics, politeness classifier judges the formal version of the same sentence as more polite than its informal variant.

We release our annotated data and aligned politeness lexicon to support future work. Our dataset can support various end applications, such as building multilingual agents optimized for politeness (Silva et al., 2022), developing a translation model that preserves politeness level (Fu et al., 2020), evaluating the impact of different pretraining corpus and modeling architecture for modeling subjective tasks in a wide range

of languages (Hu et al., 2020), understanding cultural-specific politeness strategies, and many more.

Chapter 2: Related Work

Politeness & Formality Danescu-Niculescu-Mizil et al. (2013) presents the first quantitative, linguistic study of politeness, annotating two types of corpora – requests extracted from conversations between users on Wikipedia User Talk Pages and user comments from Stackoverflow. Followup work explored interpreting neural networks’ politeness predictions (Aubakirova and Bansal, 2016) and controllable text generation with target politeness level (Sennrich et al., 2016; Niu and Bansal, 2018; Fu et al., 2020). While these work considers politeness phenomena in English, we expand it to study the phenomena in nine languages. A related concept to politeness is formality, studied in multiple prior work (Lahiri, 2016; Pavlick and Tetreault, 2016; Rao and Tetreault, 2018; Briakou et al., 2021).

Multilingual Models Recent progresses in pretrained language models have brought better representation for multitude of languages. Multilingual language models like mBERT (Devlin et al., 2019), XLMR (Conneau et al., 2020), based on the transformer architecture, are pretrained with the masked language modeling objective on a large amount of corpora (El-Kishky et al., 2020; Suarez et al., 2019) spanning over 100 languages. While the community also recognizes the varying quality of unlabeled data in a range languages (Caswell et al., 2022), such multilingual models provide improved representations for modeling low resource languages. When finetuned on downstream task data in a single language, these models make reasonable predictions in multiple languages (Wu and Dredze, 2019). Multilingual models have also been evaluated in a prompting setup for different tasks like Machine Translation (Tan et al., 2022) and different Multilingual NLU tasks (Zhao and Schütze, 2021; Lin et al., 2021; Winata et al., 2021).

Multilingual Benchmarks Despite recent progresses in NLP resources and benchmarks, partially powered by affordable crowdsourcing (Snow et al., 2008), linguistic resources in low resource languages are still severely limited to compared to resources in English (Joshi et al., 2020). Many existing datasets are translated from English data (Conneau et al., 2018; Longpre et al., 2021). While translating approach for dataset construction have advantage of ensuring similar data distribution across languages, data collected in such fashion will not reflect the language usages of diverse population, introducing translationese which can be different from purely native text (Lembersky et al., 2011). We provide resources for nine typologically diverse languages, capturing a subtle phenomena of politeness.

Chapter 3: TyDiP: Multilingual Politeness Dataset

3.1 TyDiP: Multilingual Politeness Dataset

Motivation Our goal is to construct high-quality multilingual evaluation data with native content, covering a wide range of languages on the task of politeness prediction. Following prior work (Danescu-Niculescu-Mizil et al., 2013), we focus on identifying politeness in requests, where requests involve speaker imposing on the listener. This scenario elicit speakers to employ diverse strategies to minimize the imposition of requests, or apologizing for the imposition (Lakoff, 1977). For each request text, we aim to collect a graded politeness score (between -3 and 3, with 0.5 increment).

Language Selection We chose Hindi, Korean, Spanish, Tamil, French, Vietnamese, Russian, Afrikaans, and Hungarian. Our criteria for selecting languages were (1) covering low resource language when possible, (2) languages with rich discussion on Wikipedia editor forum and (3) languages where we can recruit native speaker annotators on a crowdsourcing platform, Prolific.¹

Source Sentence Collection We source requests from Wikipedia user talk pages from target language Wikipedia dumps.² Each request is a part of a conversation between editors on Wikipedia. We follow the pre-processing step from prior work (Danescu-Niculescu-Mizil et al., 2013), extracting each request as a sequence of two successive sentences where the second sentence ends with a question mark (?). We present one example here: "I'm somewhat puzzled by your recent edits on the Harper page, which have left two different sets of footnotes. Could you please explain your rationale for the change?"

¹<https://prolific.co/>

²<https://dumps.wikimedia.org/backup-index.html>

Language	Family	Script	Total # Requests	% Target / Eng / Other	Avg len (in bytes)
Hindi	Indo-Aryan	Devanagari	4,412	71 / 26 / 3	351
Korean	Korean	Hangul	43,219	96 / 3 / 1	183
Spanish	Romance	Latin	180,832	97 / 2 / 1	181
Tamil	Dravidian	Tamil	5,590	92 / 8 / 0	325
French	Romance	Latin	354,544	98 / 1 / 1	179
Vietnamese	Austroasiatic	Latin	22,070	96 / 4 / 0	210
Russian	Slavic	Cyrillic	291,220	98 / 1 / 1	254
Afrikaans	Germanic	Latin	3,399	85 / 11 / 4	134
Hungarian	Uralic	Latin	80,825	98 / 1 / 1	132

Table 3.1: Languages chosen for our study and their data statistics. We report the number of available requests in Wikipedia User talk pages after pre-processing step, the distribution of languages after language identification, and the average length in bytes for each request.

3.1.1 Annotation Process

Collecting annotations for non-English data for a wide range of languages is non-trivial in all aspects, from source text collection, annotator recruiting to annotation validation. We describe our annotation process here and hope that our collection strategy can provide insights for future multilingual data collection efforts for other tasks and domains.

Pre-processing We observe that a sizable portion of the requests is written in language other than its own. Thus, we filter sentences not belonging to the target language with a language identification with `langdetect` (Nakatani, 2010).

Table 3.1 shows data statistics, including the language distribution among these requests. We use the `Polyglot` tokenizer³ for preprocessing.

Annotator Recruiting We collect our annotation on a crowdsourcing platform, Prolific, which allows us to find workers based on their first language. Instead of developing separate guidelines for each language, we recruit bilingual annotators who can understand English guidelines. We also filter the workers by their task approval rate ($> 98\%$).

³<https://github.com/aboSamoor/polyglot>

To annotators who meet these criteria, we perform qualification process which involves translation task and the target task, which we describe below.

Target Task Qualification Inspired by strong zero-shot transfer performances of multilingual models on a variety of tasks (Conneau et al., 2018; Wu and Dredze, 2019), we use a multilingual classifier trained on existing English politeness dataset (Danescu-Niculescu-Mizil et al., 2013) to select sentences for the qualification task.⁴ We sample examples where the classifier assigned very high or very low politeness score for each language. Language-proficient researchers verified the correctness of model predictions on a subset (four) of languages. While the model was not always correct, their highly confident predictions were mostly correct. These requests, paired with the predicted politeness label, were used to filter crowdworkers.

Translation Qualification Task Inspired by prior work (Pavlick et al., 2014) which employed a translation task to assess the language proficiency of crowdworkers, we estimate their language proficiency by evaluating their translation skills.⁵

We present crowdworkers with a set of five requests (assigned either very polite or very impolite rating by the model) in the target language, and ask them to translate into English as well as to label a politeness score. We first compared the annotator’s translation with the output from Google Translate API.⁶ If the edit distance between their translation and the output from Google translate, we remove them from the annotator pool as they could be using this service. We also computed the distance between the user’s politeness score and the model’s predicted labels, and pruned workers who provided scores that varies significantly from model predictions.

⁴We describe this model (XLMR-target) in detail in Section 4.

⁵In original politeness dataset collection, instead of translation task they introduced paraphrasing task to ensure linguistic attentiveness.

⁶<https://cloud.google.com/translate>

The qualification is not completely automatic, with constant monitoring on four languages on which language-proficient researchers continuously provide sanity checks. Fifteen workers per language took our qualifier task, and after this filtering we ended up with 7 Afrikaans, 9 Spanish, 9 Hungarian, 10 Tamil, 10 Russian, 11 Hindi, 11 Korean, 11 French and 11 Vietnamese workers.

Final Data Collection / Postprocessing The annotators annotated 5 English requests and 15 target language requests instances per task. The annotation interface can be found in the appendix. We collect 3-way annotations for each request. Annotating 20 examples took approximately seven minutes and annotators were paid \$3 for it, translating to \$25.43/hr.

3.1.2 Inter-annotator Agreement

Ensuring data quality is challenging, especially when we do not have in-house native speaker to inspect for all languages we study. Following prior work (Pavlick et al., 2014; Danescu-Niculescu-Mizil et al., 2013), we estimate the annotation quality by comparing inter-annotator agreement with agreement between randomly assigned labels according to the data distribution⁷. As we study continuous rather than categorical value, we compute pairwise spearman correlation to measure agreement score instead of Cohen’s Kappa.

As each annotator provided scores for both English sentences and sentences of their native languages, we report both agreement numbers, split by language in Table 3.2. We consistently observe a positive correlation among the annotators’ scores. Interestingly, we observe substantially higher agreement when annotators were labeling their own language compared to labeling English across all nine languages. This suggests the interpreting politeness of foreign language can be less precise and more variable compared to interpreting that of native language. As our main goal

⁷Specifically, we shuffle the politeness scores of each annotation set (20 examples).

Language	en		target	
hi	0.31	(0.4)	0.39	(0.2)
ko	0.34	(0.34)	0.6	(0.12)
es	0.28	(0.12)	0.52	(0.16)
ta	0.38	(0.32)	0.33	(0.17)
fr	0.45	(0.3)	0.53	(0.21)
vi	0.38	(0.31)	0.41	(0.17)
ru	0.43	(0.34)	0.51	(0.16)
af	0.35	(0.34)	0.37	(0.2)
hu	0.38	(0.3)	0.52	(0.19)
average	0.36	(0.31)	0.46	(0.17)

Table 3.2: Pairwise correlation (mean and standard deviation (in brackets)) for each language annotator, on English data and their native language data.

is collecting target language annotations, this would not impact the quality of our dataset, which studies how native speakers perceive native contents. We plot the averaged pairwise spearman correlation of annotations and that of random assignments in Figure 3.3. In both English and their native languages, annotator correlation is substantially higher than correlation from random label assignments, which hovers around zero as expected.

Figure 3.4 compares the overall IRR metrics on our annotations with the IRR on the annotations released by Danescu-Niculescu-Mizil et al. (2013) on English request data. They release the 5 way annotation done on their data and also a single score for each sentence after averaging and normalization. We report two scores in Table 3.3: a correlation with raw annotations and a correlation with the final aggregated scores.

Figure 3.1 shows the distribution of the pairwise correlation metric over different HITs for each language. Each subplot has the distribution over the english and target language parts of each HIT, as well as a baseline method where the scores are shuffled before computing the correlation.

The correlations in the random baseline are close to 0 and the correlations on

	Raw Score		Aggregate Score	
hi	0.27	(0.48)	0.41	(0.44)
ko	0.27	(0.48)	0.45	(0.46)
es	0.24	(0.48)	0.4	(0.42)
ta	0.25	(0.48)	0.49	(0.39)
fr	0.25	(0.46)	0.52	(0.35)
vi	0.29	(0.48)	0.46	(0.39)
ru	0.21	(0.52)	0.45	(0.4)
af	0.28	(0.48)	0.44	(0.44)
hu	0.22	(0.5)	0.48	(0.38)

Table 3.3: Agreement with original English labels (Danescu-Niculescu-Mizil et al., 2013) annotated by speakers of various languages. Mean and std deviation (in brackets).

the annotations are significantly higher. The correlations on the English annotations do show more variance in their distribution.

3.1.3 Final Dataset

We collect three way annotations for 500 randomly sampled requests for each language. We normalize each annotator’s score to a normal distribution with a mean of zero and standard deviation of 1, and then average the score of three annotators to get a final score for each item, which ranges from -3 (very impolite) to +3 (very polite). We plot the final politeness distribution per language in Figure 3.2.

Table 3.4 summarizes the distribution of scores across languages. All the languages have a mean close to 0, with similarly shaped distribution of scores. The minimum and maximum scores seem to vary a bit across languages. Some languages like Spanish have a higher median score and a higher number of sentences with a positive scores. Examples of annotated sentences are in Table 3.5.

We split these examples into 4 quartiles based on their politeness scores, and consider sentences from the top and bottom 25 percentile of politeness scores only (corresponding to positive and negative politeness), following prior work (Danescu-

Pairwise Correlation Metric and Random Baseline across Languages

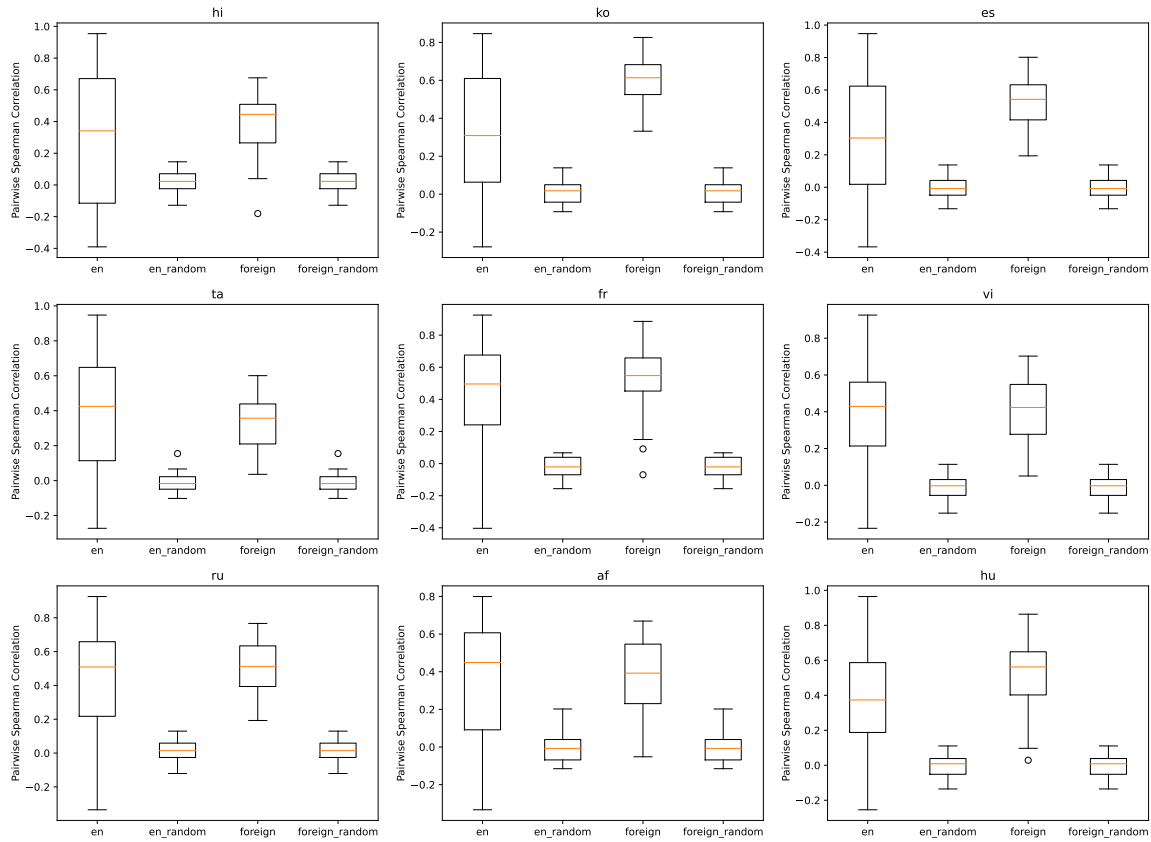


Figure 3.1: Pairwise Correlation Metric over different HITs, along with Random Baseline value for different Languages

Niculescu-Mizil et al., 2013; Aubakirova and Bansal, 2016). This results in a balanced binary politeness prediction task, while reducing the number of examples by half. We refer to this dataset (containing half of the total TyDiP dataset) as TyDiP evaluation dataset.

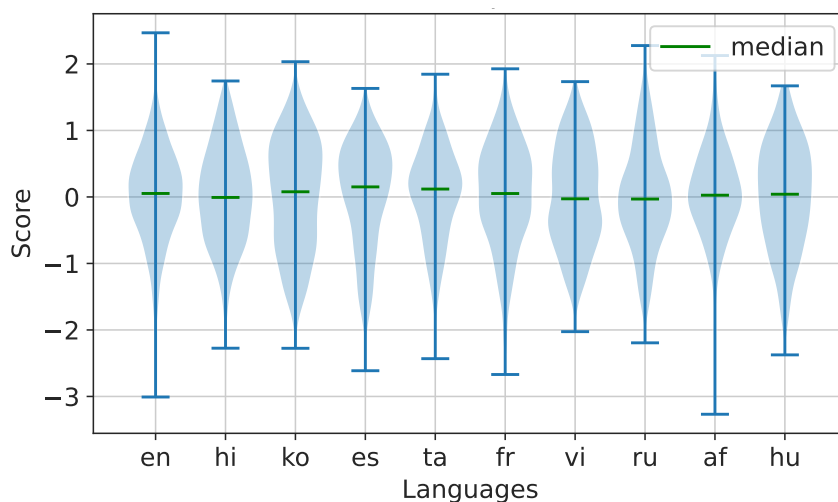


Figure 3.2: Distribution of final politeness scores per language, with mean and median highlighted.

Lang	Request	Translation	Score
hi	इस आर्टिकिल में हम आपकी उन सभी दुवधियों का हल ले कर आये हैं . नेपाल कहाँ और क्यों जाएँ?	In this article, we have brought the solution of all those dilemmas of yours. Where and why to visit Nepal?	-0.65
ko	위백엔 url 이 있습니다. 왜님은그것을따르지않으세요?	Wieback has a url. Why don't you follow it?	-0.95
es	Hola Cosconiano. ¿Has entendido lo que te han puntualizado o haces caso omiso?	Hello Cosconiano. Have you understood what they have pointed out to you or do you ignore it?	-1.24
ta	ஆங்கிலத்தில் எழுதியது நீ தானே? நடுவால் கடுமம் பட்ட காணாம ஆக்கியபோ?	Did you write in English? A little lost in the middle?	-1.53
fr	Pour la première, c'est possible et serait faisable à ce jour. Cependant estce souhaitable ?	For the first, it is possible and would be feasible today. However, is it desirable?	0.62
vi	Nếu bạn nằm trong trường hợp của tôi, thì bạn sẽ nghĩ gì về Minh Huy? BQV sẽ làm gì với Minh Huy?	If you were in my situation, what would you think of Minh Huy? What will BQV do with Minh Huy?	0.63
ru	Никакого желания редактировать занбсплёные тексты нет. Ну как вам этот абзац?	There is no desire to edit the written texts. Well, how do you like this paragraph?	-0.96
af	Jy maak aantygings van "veldtogte". Dalk is daar 'n balk in jou oog?	You make allegations of "campaigns". Maybe there is a beam in your eye?	-1.30
hu	Szia! Te jobban értesz a halakhoz, átnéznéd a következő cikkeket?	Hi! You know more about fish, would you like to review the following articles?	0.82

Table 3.5: Examples from the final TyDIP dataset. The politeness scale is from -3 (very impolite) to +3 (very polite).

Lang	# of Examples	Scores				% Positive
		<i>Mean</i>	<i>Std</i>	<i>Min</i>	<i>Max</i>	
hi	500	-0.0005	0.7756	-2.2745	1.7437	0.4980
ko	500	-0.0016	0.8659	-2.2760	2.0335	0.5320
es	500	-0.0007	0.8241	-2.6132	1.6320	0.5740
ta	500	0.0035	0.7546	-2.4319	1.8454	0.5540
fr	500	0.0003	0.8406	-2.6700	1.9268	0.5340
vi	500	-0.0019	0.7810	-2.0258	1.7338	0.4920
ru	500	0.0047	0.8234	-2.1938	2.2763	0.4880
af	500	0.0049	0.7845	-3.2678	2.1264	0.5200
hu	500	0.0027	0.8322	-2.3750	1.6705	0.5200

Table 3.4: Statistics on Final Politeness Scores

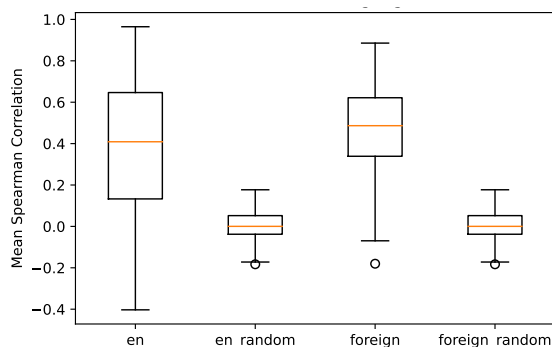


Figure 3.3: Spearman correlation. The first and third graph represents our annotated data in English and target languages respectively, and the second and the fourth shows correlation for random assignments, which hovers around zero as expected.




Figure 3.4: Pairwise correlation metric on our annotations compared to the annotations released by Danescu-Niculescu-Mizil et al. (2013)

Chapter 4: Predicting Politeness

4.1 Predicting Politeness

Equipped with politeness data for nine languages, we evaluate cross-lingual transfer performance of multilingual language models (Conneau et al., 2020). We are interested in following research questions:

1. Can a **multilingual** model trained on English politeness data predict politeness of different languages?
2. Can we use a **monolingual** model trained on English politeness data by translating target languages into English?

Models We study two fine-tuned pretrained language models, one English model (RoBERTa (Liu et al., 2019)) and one multilingual model (XLM-RoBERTa (Conneau et al., 2020)) which supports all nine languages we study.

We randomly split data from Danescu-Niculescu-Mizil et al. (2013) to yield 1,926 training and 251 evaluation examples in English. With this training dataset, we fine-tuned each model for five epochs with a batch size of 32 and an learning rate of 5e-6 on a Quadro RTX 6000 machine. We use the large variants for both models.

At inference time, we translate the target language requests into English using Google Translate API (optional for XLMR model, necessary for RoBERTa model).

We use one large-scale language model, GPT3 (Brown et al., 2020) Davinci-002, in a zero-shot prompting setup with the following prompt:

```
Is this request polite?  
<input example>
```

Model	Input Lang.	en	hi	ko	es	ta	fr	vi	ru	af	hu	Avg
Majority	-	0.537	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
XLMR	target	0.892	0.868	0.784	0.84	0.78	0.82	0.844	0.668	0.856	0.812	0.808
XLMR	en	0.892	0.884	0.752	0.848	0.748	0.84	0.816	0.688	0.836	0.8	0.801
RoBERTa	en	0.912	0.868	0.692	0.836	0.768	0.812	0.796	0.684	0.856	0.768	0.786
GPT3	target	0.808	0.732	0.708	0.732	0.596	0.764	0.692	0.688	0.688	0.76	0.706
GPT3	en	0.808	0.668	0.62	0.732	0.652	0.72	0.664	0.612	0.7	0.652	0.668

Table 4.1: Accuracy on TYDiP evaluation dataset. The XLMR and RoBERTa models are finetuned in English politeness data from Danescu-Niculescu-Mizil et al. (2013), while GPT3 model is prompted in a zero-shot fashion. When Input Lang. column is “en”, we use Google Translate API to translate the target language into English.

Then, we compute the probabilities for two options for next token – “yes” and “no” respectively, which map to “polite” and “impolite” labels respectively. Designing prompts for each language is non-trivial, so in this initial study we use this exact same English template for all languages.

Results Table 4.1 reports the model performances. Following recent question answering benchmark (Clark et al., 2020), we only aggregate the scores on non-English languages to focus on transfer performances. Both finetuned language models (XLMR and RoBERTa) boast strong performance in English, reaching an accuracy hovering 90%. Even zero-shot GPT model performs competitively, reporting an accuracy of 80.8%.

In terms of XLMR model, the results were fairly split on whether it is better to use automatically translated English input, matching the training data, or using the target language input as is. Using the text in English showed better performances in and four (Hindi, Spanish, French, Russian) and using the target language input was better in five languages (Korean, Tamil, Vietnamese, Afrikaans, and Hungarian). Using the target language yields a slightly better performance, questioning whether automatic translation maintain the politeness level.

Large-scale language model, GPT3, even used in a zero-shot fashion without

much prompt engineering (Gao et al., 2021) shows competitive performances, significantly outperforming the majority baseline. Similar to XLMR, using target language as is showed better performance than using translated text (70.6 vs. 66.8) on average, and in seven out of nine languages.

Comparing performances across languages is tricky as the annotation was done by different sets of annotators on different items for each language. To put these numbers in context, we provide a comparison between estimated human performance and model performance in the next section. Would human agreement be lower on languages with weaker model performance?

Comparison with human agreement To compute a comparable number between the annotators and models, we use our original 3-way annotated data before aggregating politeness score. We treat one annotator’s label as the human prediction and consider the other two as references, taking their mean to get the gold politeness score. We repeat this random sampling process for each example in test set 1,000 times and plot the distribution of accuracy scores in Figure 4.1.

Annotators shows varying degree of agreements – we notice a particularly stronger agreement in Korean and Hungarian, but overall we observe strong agreement, hovering around 90%. Interestingly, models significantly underperform in these languages with high human agreement, making the gap between human and model performance large. Six out of the nine languages have a gap of at least 10%, and two of them being greater than 15%.




Figure 4.1: Comparing the our best model accuracy (XLMR-target) vs. annotator accuracy on politeness prediction.

Chapter 5: Building and Analyzing Politeness Strategies in Nine Languages

5.1 Building and Analyzing Politeness Strategies in Nine Languages

In this section, we develop a set of linguistic politeness strategies based on existing English strategies (Danescu-Niculescu-Mizil et al., 2013), and see how can explaining politeness phenomena in nine diverse languages we study. While politeness strategies are not necessary for building a high-performing classifier, it can be helpful to understand politeness phenomena.

The original English study presents a list of politeness strategies along with each strategies relation to assigned politeness score. They found many statistically significant correlations between politeness strategies and human perception, such as words belonging to gratitude lexicon (appreciate), counterfactual modal (could/would) correlates with being polite, and starting the sentence with first person pronoun correlating with being impolite.

Developing such a politeness lexicon for each language requires expert annotation, which can be infeasible for low-resource languages with a fewer language-proficient researchers (Joshi et al., 2020). Thus, we aim to automatically generate politeness strategies for other languages from the English ones. For this initial study, we focus on lexicon-based strategies (15 out of 20 strategies), excluding strategies involving dependency parsing.

Mapping English Lexicon to Target Languages To build a politeness lexicon in nine languages, we use two NLP tools – translation and word alignments.

We sample 5000 Wikipedia editor requests that are not included in our anno-

tated data for each of nine languages.¹ We first automatically translate target language sentence into English (with Google Translate API) and then align the words in the translated English sentence to the words in original sentence in the target language.

Aligning words in parallel corpora has been long-standing task in NLP. Traditionally, alignments can be obtained as a byproduct of training statistical MT systems (Och and Ney, 2003; Dyer et al., 2013). Yet, this typically require a large parallel corpus, which we lack for nine languages we study. We instead use alignment method using the similarity between token representations from multilingual pretrained language models (mBERT (Devlin et al., 2019)), fast-align (Dou and Neubig, 2021).

For each word in English politeness lexicon, we collect their aligned word in the target language. As the alignments maps a sequence of words to a sequence of words, sometimes a single word English lexicon is mapped to multiple words in the target language. For each word in the English lexicon, we consider up to top five target language word sequences as its matching lexicon. We show examples of induced lexicon in Table 5.1 and full lexicon in the repository.

As automatically generated lexicon can be imprecise for either incorrect translation or alignments, we manually inspected the generated lexicon in four languages for which we have language-proficient researchers. We found that the alignments were mostly reasonable, but erroneous and imprecise for words with multiple senses. Not every lexicon was mapped to foreign words either, we show the coverage statistics (average % of words in lexicon mapped to foreign language words), which hovers around 60-70%, at the bottom of Figure 5.1.

Analysis with Induced Lexicon Using automatically induced lexicon, we analyze our multilingual politeness data, mirroring the analysis from Danescu-Niculescu-Mizil

¹For languages (Afrikaans, Hindi) with less than 5K requests, we used all available data.

Strategy	Hindi	Korean	Spanish	Tamil
Please please	कृपया	바랍니다	por favor	முகபிமுமா
HASHEDGE think apparently postulate	लगता	생각하시나요 아무래도	creo al parecer postular	நினகைக்கிறீர்கள் முன்வகைக்க
Deference great	महान्	대단한	gran	படுபு
1st prsn pl we	हम	우리가	nos	நாம்
Indirect hello	नमस्ते	안녕하세요	hola	வணக்கம்
Direct so	तो	그럼		அப்படி

Table 5.1: Examples of politeness strategies lexicon gathered from our alignment method and then cherry picked by language proficient researcher.

et al. (2013). We report the average politeness score of sentences exhibiting each each strategy in Figure 5.1. The baseline value here would be 0. We observe that the average politeness score for each strategy across languages are somewhat consistent (e.g., PLEASE strategy being positively correlated in all languages except Spanish). The diverging patterns can be an error with strategy mapping and needs further investigation. Interestingly, in languages with lower model performance (Korean, Tamil), we observe more diverging patterns (e.g., indirect greeting having positive implications in these two languages while mildly negative in English).

In the Figure 5.2, we include the occurrence of different strategies in the most polite (4th quartile) and least polite (1st quartile) subsections of our data. Cells shaded in light orange represent a baseline value of 0.25 and anything deviating from this appear in Dark Green or Red. We can clearly see difference across the the 2 quartiles for some of these strategies.

The aligned politeness strategy lexicon (Section 5.1) relies on multiple automatic toolkits (machine translation system and word alignments), thus analysis

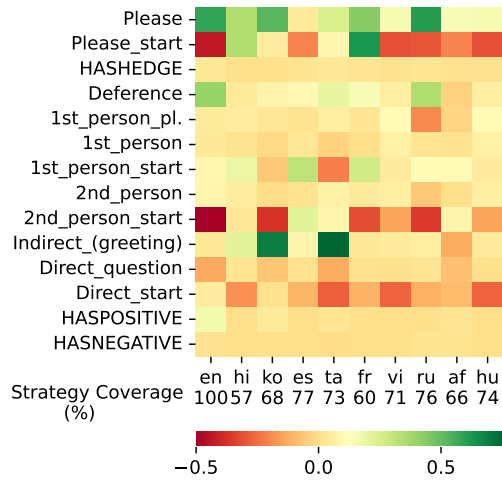


Figure 5.1: Induced politeness strategies and their relation to politeness scores in nine languages. We plot the average politeness score of a set of sentences containing corresponding strategy. Here, the baseline value is 0. The number of strategies covered by the induced lexicon is also mentioned for each language.

should be interpreted with caution.

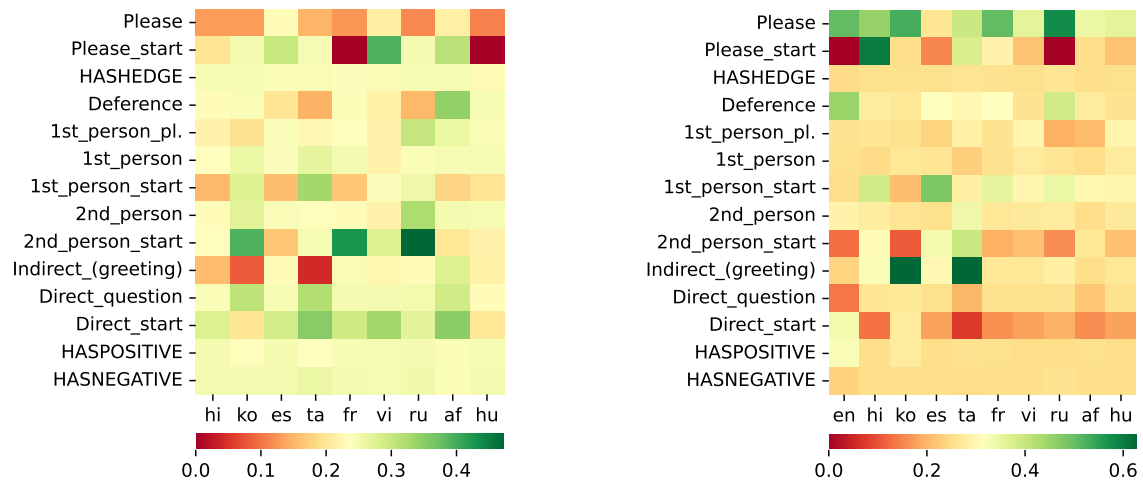


Figure 5.2: Presence of Politeness Strategies across the 1st and 4th quartiles (left and right plot) of data. Any cells deviating from the baseline value of 0.25 represent significant results

Chapter 6: Transfer between Formality and Politeness

6.1 Transfer between Formality and Politeness

While we are not aware of computational linguistic studies in politeness covering multiple languages, prior work (Briakou et al., 2021; Rao and Tetreault, 2018) has explored formality in four languages (English, French, Italian and Portuguese). In this section, we study the connections between the formality and politeness. Would formally written sentences perceived as more polite by our classifier?

We use GYAFC (Rao and Tetreault, 2018) and X-FORMAL (Briakou et al., 2021), two datasets containing informal sentences from the L6 Yahoo Answers Corpus¹ and four formal rewrites for each sentence (dataset statistics can be found in the Table 6.1).

In Table 6.2, we report zero-shot transfer results from politeness classifier to formality classification. We will use our best multilingual politeness classifier (XLMR-target) from Section 4.1. We calibrate the threshold of our politeness classifier to account for the different data distribution of positive and negative examples using the dev set. We get the probability for the 80th percentile of scores from the English dev set and use this on the test sets. Somewhat surprisingly, the classifier performs worse than the majority baseline. Table 6.3 shows performance numbers of transfer in the reverse direction, i.e from formality to politeness. We similarly finetune XLMR model on the English train set from GYAFC (Rao and Tetreault, 2018), and evaluated it on TyDiP evaluation dataset, using target language as an input. After the threshold calibration, the model performs better than the majority baseline, but substantially underperforms the in-domain performance reported in Table 4.1.

¹<https://webscope.sandbox.yahoo.com/catalog.php?datatype=1>

Dataset	# Informal	# Formal
Rao and Tetreault (2018)		
English	2478	10992
Briakou et al. (2021)		
French	1000	4000
Italian	1000	4000
Portuguese	1000	4000

Table 6.1: Statistics of Formality Data used for Evaluation (as test sets)

Model	en	fr	it	pt
Majority Baseline	80			
Before Calib.	40.12	37.74	38.10	36.72
After Calib.	73.48	74.08	74.44	73.76

Table 6.2: Transfer from politeness to formality. Formality classification accuracy on X-FORMAL dataset.

Does this mean formality and politeness are not linked? Upon inspection (see Table 6.4 for examples), we find that politeness prediction for the the informal and formal rewrites of the same sentence often stay consistent. Looking into the model’s prediction on (informal, formal) sentence pairs, we find that almost 80% of pairs in English have the same politeness prediction for both sentences. The left column in Table 6.5 depicts this across four languages, suggesting that politeness could be further linked to the *content*, not just *style* of the writing.

In their original work, Rao and Tetreault (2018) report that commonly used techniques to make sentences formal include phrasal paraphrases, punctuation changes, expansions, contractions, capitalization and normalization which are fairly stylistic. Would such rewriting make sentences to be perceived more polite? We investigate this by further looking into (informal, formal) sentence pairs – for each version of the sentence in the pair, we compute their politeness probability (as assigned by the classifier) and report percentage of pairs where formal version of the same sentence were viewed as more *polite* than its impolite counterpart. The right column in Ta-

	en	hi	ko	es	ta
Before Calib.	0.537	0.5	0.5	0.5	0.5
After Calib.	0.557	0.612	0.588	0.644	0.564
	fr	vi	ru	af	hu
Before Calib.	0.5	0.5	0.5	0.5	0.5
After Calib.	0.6	0.624	0.644	0.564	0.528

Table 6.3: Transfer from formality to politeness. Politeness classification accuracy on TyDiP evaluation dataset.

Sentence	Formality	Politeness
Hey, I'm in NYC I'll help you out if your around!	Informal	Polite
I am in New York City. I will help you if you are nearby.	Formal	Polite
why do they try to sound british ?	Informal	Impolite
Why do they attempt to sound British?	Formal	Impolite

Table 6.4: Four example with annotated formality label and predicted politeness label. The formality labels are from Rao and Tetreault (2018) and the politeness labels assigned by our classifier.

Table 6.5 presents these results – for about 70% examples, such rewriting indeed made the sentence perceived as more polite, despite often not enough to flip the politeness decision.

Language	$\mathbb{1}(\text{polite} \text{ — formal}) = \mathbb{1}(\text{polite} \text{ — informal})$	$\text{p}(\text{polite} \text{ — formal}) > \text{p}(\text{polite} \text{ — informal})$
English	0.811	0.682
French	0.775	0.702
Italian	0.764	0.734
Portuguese	0.779	0.696

Table 6.5: Analysing politeness predictions on (informal, formal) sentence pairs. The left column represents the fraction of pairs for which the same politeness label is assigned to both sentences. The right column represents the fraction of pairs for which the classifier’s probability of being polite for the formal sentence is higher than that of its informal counterparts.

Chapter 7: Conclusion

In this thesis, we present TYDIP, a corpus of requests paired with its perceived politeness score spanning nine languages along with rich analysis on the dataset. We evaluate multiple multilingual models in zero-shot politeness prediction and find that they are able to perform well without being trained on data from the same language, while not reaching human-level performances yet.

Future Work One of the potential uses for our dataset is politeness-controlled text generation translation with multilingual models. Existing work in this area (Sennrich et al., 2016; Fu et al., 2020) is limited in scope due to a lack of data across many languages. We wish to explore this area, define concrete tasks using the data from our dataset and evaluate different multilingual generation models like on this task. This will become useful as an evaluation task as there is more use of language assistants and translators in cross-language communication in this global age.

Another direction for future work is generating politeness strategies for diverse sets of languages. Our current method makes use of imperfect automated tools and requires human supervision at some stages. Improving this stage will allow more accurate analysis of politeness phenomena. Further, they can also be used for politeness-controlled translation tasks to produce outputs more grounded in fine-grained politeness strategies.

Our dataset is moderately sized (250 examples per language in the evaluation portion, and a total of 500 examples per language) and still covers a limited number of languages. We had intended to cover more languages (one example being Japanese), but this were hindered by the number of annotators we could recruit for each language. Expanding to different domains beyond Wikipedia would be fruitful too, as Wikipedia editor forum is fairly constrained domain.

Ethical Considerations

The data we annotate comes from Wikipedia User Talk pages, which is an online forum for communication between editors on Wikipedia. This data spans nine different languages and contains speakers from different countries and demographics. The annotation is done by crowdworkers recruited from the online platform Prolific. These workers aren't restricted to a particular country. They are paid a wage of \$25.43/hr which is higher than the average pay stipulated on the platform. We use this data to evaluate an existing model across multiple languages, and do not use it for training as such.

Acknowledgements

We would like to thank Yasumasa Onoe, Bernardo Oviedo, Gokul Anandaraman for their help in the earlier phase of the project development and inspecting data quality. We would also like to thank Cristian Danescu-Niculescu-Mizil for answering questions about his work. We'd like to thank Joel Tetreault and Yahoo for providing access to the formality datasets. We also thank Akari Asai for providing feedback on the paper. We'd like to thank Anuj Diwan for the thoughtful discussions on this topic and providing helpful feedback along the way.

Appendix A: Appendix

A.1 Annotation UI

Figure A.1 contains the user interface used for the final annotation process.

Politeness Annotation

Please rate how polite the request comes across for you. Each request is a part of a conversation between editors on wikipedia. There can be multiple factors affecting the politeness level of a sentence, such as style and content.

- **Style:** How did the requester phrase their request? The same content could be conveyed in both an impolite and polite manner. After reading the sentence, if you're able to think of a more polite way of saying the same thing, the style is not very polite
- **Content:** What is being asked in this request? For example, if the request is assuming the other editor made a mistake or violated rules intentionally, it's less polite than one where editors are discussing what changes to make to a page. The content score will reflect the nature of the request

You'll assign a single politeness score after considering all factors.

A small fraction of the sentences may not be actual valid text. Simply mark those as invalid on the rightmost column.

The task is spread over 4 pages. The first page contains 5 sentences in English and the remaining 3 pages contain 5 sentences each in a Foreign language.

You can navigate between the pages using the Previous and Next Buttons.

Before we start, please enter your Prolific ID

Prolific ID:

<code>\$(en_sentence0)</code>	Very Impolite	Impolite	Somewhat Impolite	Neutral	Somewhat Polite	Polite	Very Polite	<input type="checkbox"/>	Invalid
<code>\$(en_sentence1)</code>	Very Impolite	Impolite	Somewhat Impolite	Neutral	Somewhat Polite	Polite	Very Polite	<input type="checkbox"/>	Invalid
<code>\$(en_sentence2)</code>	Very Impolite	Impolite	Somewhat Impolite	Neutral	Somewhat Polite	Polite	Very Polite	<input type="checkbox"/>	Invalid
<code>\$(en_sentence3)</code>	Very Impolite	Impolite	Somewhat Impolite	Neutral	Somewhat Polite	Polite	Very Polite	<input type="checkbox"/>	Invalid
<code>\$(en_sentence4)</code>	Very Impolite	Impolite	Somewhat Impolite	Neutral	Somewhat Polite	Polite	Very Polite	<input type="checkbox"/>	Invalid

Page 1/4

If you have any comments about any particular sentence or about the set of factors (style/content) for a particular sentence, please leave them below.

Click the button below once you're done. If there are any missing fields, you'll be notified and you have to check the form again

Figure A.1: Annotation Interface

Works Cited

Malika Aubakirova and Mohit Bansal. Interpreting neural networks to improve politeness comprehension. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2035–2041, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1216. URL <https://aclanthology.org/D16-1216>.

Eleftheria Briakou, Di Lu, Ke Zhang, and Joel Tetreault. Olá, bonjour, salve! XFORMAL: A benchmark for multilingual formality style transfer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3199–3216, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.256. URL <https://aclanthology.org/2021.naacl-main.256>.

Penelope Brown and Stephen C Levinson. Universals in language usage: Politeness phenomena. In *Questions and politeness: Strategies in social interaction*, pages 56–311. Cambridge University Press, 1978.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran As-

sociates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf>.

Isaac Caswell, Julia Kreutzer, Lisa Wang, Ahsan Wahab, Daan van Esch, Nasanbayar Ulzii-Orshikh, Allahsera Auguste Tapo, Nishant Subramani, Artem Sokolov, Claytone Sikasote, Monang Setyawan, Supheakmungkol Sarin, Sokhar Samb, Benoît Sagot, Clara Rivera, Annette Rios Gonzales, Isabel Papadimitriou, Salomey Osei, Pedro Ortiz Suarez, Iroro Orife, Kelechi Ogueji, Rubungo Andre Niyongabo, Toan Q. Nguyen, Mathias Muller, Andr’e Muller, Shamsuddeen Hassan Muhammad, Nanda Firdausi Muhammad, Ayanda Mnyakeni, Jamshidbek Mirzakhlov, Tapiwanashe Matangira, Colin Leong, Nze Lawson, Sneha Kudugunta, Yacine Jernite, M. Jenny, Orhan Firat, Bonaventure F. P. Dossou, Sakhile Dlamini, Nisansa de Silva, Sakine cCabuk Balli, Stella Rose Biderman, Alessia Battisti, Ahmed Baruwa, Ankur Bapna, Pallavi N. Baljekar, Israel Abebe Azime, Ayodele Awokoya, Duygu Ataman, Orevaoghene Ahia, Oghenefego Ahia, Sweta Agrawal, and Mofetoluwa Adeyemi. Quality at a glance: An audit of web-crawled multilingual datasets. *Transactions of the Association for Computational Linguistics*, 10:50–72, 2022.

J. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. Tydi qa: A benchmark for information-seeking question answering in typologically diverse languages. *Transactions of the Association for Computational Linguistics*, 8:454–470, 2020.

Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. XNLI: Evaluating cross-lingual sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1269. URL <https://aclanthology.org/D18-1269>.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.747. URL <https://aclanthology.org/2020.acl-main.747>.

Cristian Danescu-Niculescu-Mizil, Moritz Sudhof, Dan Jurafsky, Jure Leskovec, and Christopher Potts. A computational approach to politeness with application to social factors. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 250–259, Sofia, Bulgaria, August 2013. Association for Computational Linguistics. URL <https://aclanthology.org/P13-1025>.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.

Zi-Yi Dou and Graham Neubig. Word alignment by fine-tuning embeddings on parallel corpora. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2112–2128, Online, April 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.eacl-main.181. URL <https://aclanthology.org/2021.eacl-main.181>.

Chris Dyer, Victor Chahuneau, and Noah A. Smith. A simple, fast, and effective

reparameterization of IBM model 2. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648, Atlanta, Georgia, June 2013. Association for Computational Linguistics. URL <https://aclanthology.org/N13-1073>.

Ahmed El-Kishky, Vishrav Chaudhary, Francisco Guzmán, and Philipp Koehn. A massive collection of cross-lingual web-document pairs. In *EMNLP*, 2020.

Liye Fu, Susan Fussell, and Cristian Danescu-Niculescu-Mizil. Facilitating the communication of politeness through fine-grained paraphrasing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5127–5140, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.416. URL <https://aclanthology.org/2020.emnlp-main.416>.

Tianyu Gao, Adam Fisch, and Danqi Chen. Making pre-trained language models better few-shot learners. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.295. URL <https://aclanthology.org/2021.acl-long.295>.

Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization. *ArXiv*, abs/2003.11080, 2020.

Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. The state and fate of linguistic diversity and inclusion in the NLP world. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online, July 2020. Association for Com-

putational Linguistics. doi: 10.18653/v1/2020.acl-main.560. URL <https://aclanthology.org/2020.acl-main.560>.

Nada Korac-Kakabadse, Alexander Kouzmin, Andrew Korac-Kakabadse, and Lawson Savery. Low-and high-context communication patterns: towards mapping cross-cultural encounters. *Cross cultural management: An international journal*, 2001.

Shibamouli Lahiri. Squinky! a corpus of sentence-level formality, informativeness, and implicature, 2016.

Robin Lakoff. What you can do with words: Politeness, pragmatics and performatives. In *Proceedings of the Texas conference on performatives, presuppositions and implicatures*, pages 79–106. ERIC, 1977.

Gennadi Lembersky, Noam Ordan, and Shuly Wintner. Language models for machine translation: Original vs. translated texts. *Computational Linguistics*, 38:799–825, 2011.

Xi Victoria Lin, Todor Mihaylov, Mikel Artetxe, Tianlu Wang, Shuohui Chen, Daniel Simig, Myle Ott, Naman Goyal, Shruti Bhosale, Jingfei Du, Ramakanth Pasunuru, Sam Shleifer, Punit Singh Koura, Vishrav Chaudhary, Brian O’Horo, Jeff Wang, Luke Zettlemoyer, Zornitsa Kozareva, Mona Diab, Veselin Stoyanov, and Xian Li. Few-shot learning with multilingual language models, 2021. URL <https://arxiv.org/abs/2112.10668>.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach, 2019. URL <https://arxiv.org/abs/1907.11692>.

S. Longpre, Yi Lu, and Joachim Daiber. Mkqa: A linguistically diverse benchmark for multilingual open domain question answering. *Transactions of the Association for Computational Linguistics*, 9:1389–1406, 2021.

Sara Mills. Gender and politeness. 2003.

Shuyo Nakatani. Language detection library for java, 2010. URL <https://github.com/shuyo/language-detection>.

Tong Niu and Mohit Bansal. Polite dialogue generation without parallel data. *Transactions of the Association for Computational Linguistics*, 6:373–389, 2018. doi: 10.1162/tacl_a_00027. URL <https://aclanthology.org/Q18-1027>.

Franz Josef Och and Hermann Ney. A systematic comparison of various statistical alignment models. *Computational Linguistics*, 29(1):19–51, 2003. doi: 10.1162/089120103321337421. URL <https://aclanthology.org/J03-1002>.

Ellie Pavlick and Joel Tetreault. An empirical analysis of formality in online communication. *Transactions of the Association for Computational Linguistics*, 4:61–74, 2016. doi: 10.1162/tacl_a_00083. URL <https://aclanthology.org/Q16-1005>.

Ellie Pavlick, Matt Post, Ann Irvine, Dmitry Kachaev, and Chris Callison-Burch. The language demographics of Amazon Mechanical Turk. *Transactions of the Association for Computational Linguistics*, 2:79–92, 2014. doi: 10.1162/tacl_a_00167. URL <https://aclanthology.org/Q14-1007>.

Sudha Rao and Joel Tetreault. Dear sir or madam, may I introduce the GYAFC dataset: Corpus, benchmarks and metrics for formality style transfer. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume*

1 (*Long Papers*), pages 129–140, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1012. URL <https://aclanthology.org/N18-1012>.

Rico Sennrich, Barry Haddow, and Alexandra Birch. Controlling politeness in neural machine translation via side constraints. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 35–40, San Diego, California, June 2016. Association for Computational Linguistics. doi: 10.18653/v1/N16-1005. URL <https://aclanthology.org/N16-1005>.

Diogo Silva, David Semedo, and João Magalhães. Polite task-oriented dialog agents: To generate or to rewrite? In *WASSA, 2022*.

Rion Snow, Brendan T. O’Connor, Dan Jurafsky, and A. Ng. Cheap and fast – but is it good? evaluating non-expert annotations for natural language tasks. In *EMNLP*, 2008.

Pedro Ortiz Suarez, Benoît Sagot, and Laurent Romary. Asynchronous pipeline for processing huge corpora on medium to low resource infrastructures. 2019.

Zhixing Tan, Xiangwen Zhang, Shuo Wang, and Yang Liu. MSP: Multi-stage prompting for making pre-trained language models better translators. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6131–6142, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.424. URL <https://aclanthology.org/2022.acl-long.424>.

Genta Indra Winata, Andrea Madotto, Zhaojiang Lin, Rosanne Liu, Jason Yosinski, and Pascale Fung. Language models are few-shot multilingual learners. In *Proceedings of the 1st Workshop on Multilingual Representation Learning*, pages 1–15, Punta Cana, Dominican Republic, November 2021. Association

for Computational Linguistics. doi: 10.18653/v1/2021.mrl-1.1. URL <https://aclanthology.org/2021.mrl-1.1>.

Shijie Wu and Mark Dredze. Beto, bentz, becas: The surprising cross-lingual effectiveness of bert. In *EMNLP*, 2019.

Mengjie Zhao and Hinrich Schütze. Discrete and soft prompting for multilingual models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8547–8555, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.672. URL <https://aclanthology.org/2021.emnlp-main.672>.