

Copyright
by
Christopher Scott Hammel
2015

The Thesis Committee for Christopher Scott Hammel
certifies that this is the approved version of the following thesis:

Standardization for Intelligent Detection and
Autonomous Operation of Non-Structured Hardware,
and its Application on Railcar Brake Release Operation

APPROVED BY

SUPERVISING COMMITTEE:

Delbert Tesar, Supervisor

Pradeep Ashok

**Standardization for Intelligent Detection and
Autonomous Operation of Non-Structured Hardware,
and its Application on Railcar Brake Release Operation**

by

Christopher Scott Hammel, B.S.M.E.

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 ENGINEERING

THE UNIVERSITY OF TEXAS AT AUSTIN

May 2015

This work is dedicated to my grandpa, Dr. David John Hammel, who was the first engineer I ever met, and whose health was too poor to attend my master's graduation, and to my grandma, Genie Hammel, who has provided more than enough support for both of them even when he couldn't be present. This work is also dedicated to my parents, Scott and Lisa Hammel, who have always been a constant source of unwavering love and support and the best role models I could ask for. I would also like to dedicate this work to my girlfriend, Maranda Moody, who has supported me lovingly through the many long, absent nights and weekends of work.

Acknowledgments

I would like to thank Dr. Tesar for his guidance during graduate school. I would like to thank Dr. Benito Fernandez for re-igniting my love of learning and engineering and my desire to enter graduate school. I would like to thank Dr. Pradeep Ashok and Andrew Boddiford for believing in me and aiding my acceptance into the graduate program, and for their continued guidance and friendship.

I would also like to thank the other graduate students of the RRG, Dhiral Chheda, Timothy Woodard, Scott Hamill, and Gurtej Singh for their friendship and camaraderie through the many long nights spent working toward a common goal.

Standardization for Intelligent Detection and Autonomous Operation of Non-Structured Hardware, and its Application on Railcar Brake Release Operation

Christopher Scott Hammel, M.S.E.
The University of Texas at Austin, 2015

Supervisor: Delbert Tesar

This thesis introduces a standard framework for evaluating and planning for desired autonomous (or semi-autonomous) operations, then applies the framework, in detail, to the task of automating emergency brake release before rail-car decoupling. A significant hurdle to be accounted for is the lack of standardization of much of the hardware of interest in industry. Non-standardized rail car components must be formally structured as fully as possible to improve the reliability of the robotic automation. This brake release task requires either pushing or pulling a “bleed rod” that protrudes from the side of each rail car. The requirements for each step of the evaluation and planning process will be laid out in this thesis, as an example of how it should be applied to future automation tasks.

Table of Contents

Acknowledgments	v
Abstract	vi
List of Tables	xv
List of Figures	xvi
Chapter 1. Introduction	1
1.1 Objective	1
1.2 Background	2
1.3 Motivation	4
1.3.1 Mobile Robotics	4
1.3.2 Reliability	5
1.3.3 Time Efficiency	5
1.4 Generic Steps	6
1.4.1 Structuring the Hardware	6
1.4.1.1 Data Mining	7
1.4.2 Modify/Optimize Object Detection	7
1.4.3 Determine Object Detection Potential	8
1.4.4 Determine Mobile Robot Requirements	9
1.5 Results of Steps as Applied to Brake Release	9
1.5.1 Structuring the Brake Rod Hardware	9
1.5.2 Brake Rod Detection	10
1.5.3 Brake Rod Detection Potential	11
1.5.4 Mobile Robot Requirements	13
1.5.4.1 Manipulator	13
1.5.4.2 Mobile Platform	14

Chapter 2. San Antonio Yard Visit - Understanding the Problem	16
2.1 “Service Valve” Housing	16
2.1.1 Valve Handle	17
2.2 Bleed Rod	19
2.2.1 Variations	19
2.3 Reservoir	20
2.4 Measurements and Analysis	21
2.4.1 Force Measurements	24
2.5 System Assumptions	24
2.5.1 Expected Layout/Procedure	25
2.6 Necessary Components for Lab Testing System	27
2.6.1 Necessary Variability	27
2.7 Vision	28
2.7.1 Requirements	28
2.7.1.1 Different Heights	28
2.7.2 Camera Configuration	28
2.7.2.1 Multiple Cameras	28
2.7.2.2 Multiple Camera Pairs	29
2.7.2.3 Vertical Axis	30
Chapter 3. Laboratory Test-Bed Modifications	31
3.1 Service Valve	31
3.1.1 Mounting Location	32
3.2 Bleed Rods	35
3.2.1 Requirements	35
3.2.2 Vision Testing	36
3.2.3 Rod Operation Testing	37
3.3 Pneumatic Reservoir	38
3.3.1 Supporting Calculations	39

Chapter 4. Robot Modifications	43
4.1 End-Effector	43
4.2 Camera Mount	44
4.2.1 Background	44
4.2.2 Angle Testing	47
Chapter 5. Robot/Vision Hardware Requirements	48
5.1 Rod Height	48
5.1.1 Robot and Mobile Platform During Actuation	48
5.1.2 Mobile Platform Stability Requirements	50
5.1.3 Vision	50
5.2 Detection Camera Position	52
5.2.1 Degrees of Freedom	53
5.2.1.1 Camera Depth	54
5.2.1.2 Distance Along Rail	54
5.2.1.3 Mounting Angle	55
5.2.2 Simulation Layout	55
5.2.3 Viewing Angle	56
5.2.3.1 Performance Metric	56
5.2.4 Viewing Distance	57
5.2.4.1 Performance Metric	58
5.2.5 Detection Potential Metric	60
5.2.5.1 Fusion Methods	61
5.2.5.2 Normalization	62
Chapter 6. Primary Operational Steps	65
6.1 Bleed Rod Detection	65
6.1.1 Wayside Sensing	66
6.1.1.1 Speed Issues	68
6.1.1.2 Distance Issues	68
6.1.2 Searching Scheme Without Wayside Sensing	68
6.1.2.1 Blind Searching (No Start-of-Car Reference)	68
6.1.2.2 Optional Reference Point Detection	69

6.1.3	Searching Scheme With Wayside Sensing	70
6.2	Bleed Rod Alignment	72
6.2.1	Coordinate Determination	72
6.2.2	Ideal Alignment	72
6.2.2.1	Robot-Favored Alignment	72
6.2.2.2	Camera-Favored Alignment	72
6.3	Approaching the Rod	73
6.3.1	F/T Sensor	73
6.4	Bleed Rod Actuation	73
6.5	Valve-Release Confirmation	74
Chapter 7. Bleed Rod Detection Algorithm		76
7.1	Getting Started with Bleed Rod Detection	78
7.2	Comparing the Vision Problems	79
7.2.1	Parameter Modifications	79
7.2.1.1	Window Dimensions	79
7.2.1.2	ANN Output Threshold	79
7.2.1.3	“Hit Count” Threshold	80
7.2.1.4	Window Step-Sizes	80
7.2.2	Heuristic Parameter Modifications	81
7.2.2.1	Populated Pixel Ratio	81
7.2.2.2	Centroid Location	82
7.2.3	Additions (Heuristics)	83
7.2.3.1	Image Height	83
7.2.3.2	Window Edge Criteria	84
7.2.3.3	Rod vs. Loop Width	84
7.3	Results	85
Chapter 8. Robot Alignment/Positioning		87
8.1	Robot Pose and Branching Avoidance	87
8.2	Position Determination	87
8.2.1	Fully Static Robot	88
8.2.1.1	Detection Accuracy	88

8.2.1.2	Alignment Accuracy	89
8.2.2	Fully Moving Robot	89
8.2.2.1	Ground Inconsistencies	90
8.2.2.2	Position Error (Slip)	90
8.2.2.3	Filtering Latency	91
8.3	Rod-Centered Reference Frame	91
Chapter 9. Initial Rod Contact		92
9.1	Contact Confirmation	92
9.1.1	F/T Sensor Threshold	92
9.1.2	Safety Concerns	92
9.2	Safe Rod-End Placement	93
9.2.1	Robot Wrist Torque Limits	93
9.2.2	Expected Initial Alignment Accuracy	94
Chapter 10. Rod Actuation		95
10.1	Actuation Methods Considered	95
10.1.1	Pneumatic Piston	95
10.1.2	Robotic Arm Manipulation	96
10.1.3	Conclusion	97
10.1.3.1	Arm Selection	97
10.2	Key Factors for Rod Actuation	98
10.2.1	Travel	98
10.2.1.1	Different Bleed Rod Geometries	99
10.2.2	Possible Problems	101
10.2.2.1	Notched Rod-End	101
10.2.2.2	Insufficient Exposed Rod Length	102
10.2.2.3	Empty Reservoir Tank	102
10.2.3	Notched Rod Solution	105
10.2.3.1	Case 1	105
10.2.3.2	Case 2	106

Chapter 11. Valve-Release Confirmation	109
11.1 Considered Techniques	109
11.1.1 Auditory (microphone)	109
11.1.2 Force Monitoring	110
11.1.3 Piston-Extension Monitoring	110
11.2 Auditory Feature Recognition	111
11.2.1 Signal Frequency	114
11.2.2 Signal Amplitude	115
11.2.3 Time Threshold	116
11.2.4 Sound Signature Classification	118
11.2.5 Characterized Noise Rejection	118
11.3 Microphone Placement	119
11.3.1 Experiments	119
11.3.2 Results	121
11.3.3 Practical Application	123
Chapter 12. Adaptation for Mobile Manipulator	127
12.1 Manipulator Requirements	127
12.1.1 Modularity Between Tasks	127
12.1.1.1 Workspace	127
12.1.1.2 End-Effector Versatility/Interchangeability	128
12.1.2 Field-Repairability/Modularity	128
12.2 Static Vehicle Stability	130
12.2.1 Base Stability Through Workspace	130
12.2.2 Outriggers	130
12.3 Vehicle Maneuverability	132
12.3.1 Steering	132
12.3.1.1 Skid-Steer/Differential Steering	132
12.3.1.2 Ackerman Steering	133
12.3.1.3 Independently Steered Wheels	134
12.3.2 Tires	135
12.3.2.1 Custom for Rail-Traversing	135

12.4	Sensor Preservation	136
12.4.1	Problems to Anticipate	136
12.4.1.1	Camera Vibration	136
12.4.1.2	Microphone Noise	138
Chapter 13.	Phase II Recommendations	139
13.1	Further Development Needs	139
13.1.1	Vision Algorithm	139
13.1.1.1	ANN Retraining	139
13.1.1.2	3D Camera Choice	140
13.1.2	Multiple On-Board Camera feeds	140
13.1.3	Pulling Rather than Pushing	141
13.2	Mobile Manipulator Yard Navigation	141
13.3	Mobile Robot Recommendations	143
13.3.1	Mobile Platform	143
13.3.1.1	Recommendation	144
13.3.1.2	Others Considerations	146
13.3.1.3	Future Applications	149
13.3.2	Manipulator	154
13.3.2.1	Recommended Development	154
13.3.2.2	Recommended Actuators	157
Chapter 14.	Conclusions & Results	161
14.1	Introduction	161
14.2	San Antonio Yard Visit - Understanding the Problem	163
14.3	Laboratory Test-Bed Modifications	166
14.4	Robot Modifications	167
14.5	Robot/Vision Hardware Requirements	169
14.5.1	Mobile Platform	169
14.5.2	Vision Hardware	170
14.6	Primary Operational Steps	176
14.7	Bleed Rod Detection Algorithm	180
14.8	Robot Alignment/Positioning	181

14.9 Initial Rod Contact & Rod Actuation	183
14.10 Valve-Release Confirmation	185
14.11 Adaptation for Mobile Manipulator	192
Bibliography	200

List of Tables

2.1	Bleed Rod position measurements taken on cars during the San Antonio Yard Visit	22
7.1	ANN Performance Results through 4 stages of development . .	86
10.1	Summary of potential fault events during rod operation, and associated sensor responses	104
11.1	Tabulated results from the experimental microphone conditions	122
12.1	Mobile Robot criteria and static stability evaluations	130
13.1	Manipulator Requirements [8]	156
13.2	Manipulator Comparison [8]	160
14.1	Features of brake release valve	165
14.2	Required Vision System Features	173
14.3	Sound Recognition Features	190
14.4	Mobile Manipulator Features	197

List of Figures

1.1	“True” rod training images	11
1.2	Camera workspace evaluation	13
1.3	Mobile Platform static stability schematic	15
2.1	Bleed Rod with worn notch	20
2.2	Plot of F-Values vs. number of clusters for bleed rod position data	24
2.3	Bleed Rod Position 1	25
2.4	Bleed Rod Position 2	25
2.5	Procedure schematic	26
3.1	Unmodified rolling stock test-bed	31
3.2	“Service Portion” schematic	32
3.3	Rear view of test chassis without service valve	33
3.4	Service valve mounting plate schematic	34
3.5	Rear view of test chassis with installed service valve	34
3.6	Installed mounting plate and service valve (front/back views)	35
3.7	Bleed Rod showing length variability	36
3.8	SolidWorks model and installed bleed rod end showing internal temporary mounting components	36
3.9	Left View (SolidWorks) of first Bleed Rod design	37
3.10	Previous pressure manifold and the hole limiting air flow	38
3.11	Custom pressure manifold for air tank	39
3.12	Moody diagram used to find the Darcy friction factor	41
4.1	Conceptual conical contact plate	44
4.2	Individual camera mount	45
4.3	Individual camera mount with T-strut	45
4.4	Cameras and mounts placed in dual camera mounting frame	47

5.1	Mobile Manipulator platform loads at different rod heights . . .	49
5.2	Top view schematic of key parameters that affect detection potential	53
5.3	Ratio of visible image size to actual size; the camera is at the origin, and locations on the graph represent rod locations . . .	57
5.4	Decreasing image signal-to-noise ratio sketch	59
5.5	Ratio of rod width in the image (in pixels) to edge noise; the camera is at the origin, and locations on the graph represent rod locations	60
5.6	Potential for rod detection (scaled 0 - 1); camera is at the origin, and locations on the graph represent rod locations; non-zero values represent feasible rod locations	64
6.1	Top view of dual camera configuration; cameras facing away from each other	66
6.2	ASUS Xtion camera, rod view impeded from right side by a vertical bar	67
6.3	Rod location schematic	69
6.4	Notched Rod	74
7.1	“True” rod training images	78
7.2	Window step-size schematic (for clarity, the labeled step is much larger than the actual step size, which gives the spacing between the “true positive detections”)	81
7.3	Horizontal centroid heuristic	82
7.4	Horizontal Centroid heuristic	82
7.5	Image height heuristic schematic	83
7.6	Window edge heuristic criteria	84
7.7	Width comparison schematic for pre-ANN heuristic filtering	85
9.1	Top view schematic of the resisting moment required by an alignment error	94
10.1	Valve handle with rod connection	99
10.2	Rod geometry variations	100
10.3	Actual and replicated worn rod notches	101
10.4	Free-body diagram for “Case 1”	105

10.5	Sketch of end-effector contact plate with lower lip to lift the rod in “Case 2”	108
11.1	RODE VideoMic Pro directional microphone	112
11.2	“Cone” for the RODE VideoMic Pro	113
11.3	On-line audio processing flow chart	114
11.4	Example of audio signal without before (top) and after (bottom) bandpass filter	115
11.5	Audio detections from amplitude threshold, before (top) and after (bottom) minimum duration threshold	117
11.6	Overhead view of experimental setup, showing the concrete cylinder and thin-walled structure obstructions and positions A, B, C, and D	121
11.7	Lower position valve, line-of-sight is completely obstructed	124
11.8	Lower position valve with no direct obstructions	125
11.9	Upper position valve with obstructions	125
12.1	Mobile manipulator platform loads with rear outrigger	131
12.2	Ackerman steering schematic [1]	133
13.1	Mobile platform offset wheel configuration	145
13.2	Low impulse grouser wheel	147
13.3	24" grouser wheel climbing a 6" rail	147
13.4	Off-the-shelf mobile robot system components	148
13.5	Versatile MMP for brake release & air hose lacing	151
13.6	Required manipulator workspace [8]	156
13.7	Schunk LWA 4D robotic arm	158
13.8	UR10 & UR5	159
14.1	Left View (SolidWorks) of bleed rod mockup	167
14.2	Cameras and mounts placed in dual camera mounting frame	169
14.3	Mobile manipulator platform loads at different rod heights	170
14.4	Potential for rod detection (scaled 0 - 1); camera is at the origin, and locations on the graph represent rod locations; non-zero values represent feasible rod locations	172
14.5	Rod location schematic	178

14.6 Notched rod	180
14.7 Audio signal before (top) and after (bottom) bandpass filter	187

Chapter 1

Introduction

1.1 Objective

Union Pacific Railroad (UP) is interested in performing repetitive physical tasks on stationary cars in switch yards using mobile robots. This thesis introduces a standard framework for evaluating and planning for desired autonomous (or semi-autonomous) operations, then applies, in detail, the framework to the task of automated brake release before car decoupling.

A significant hurdle to be accounted for is the lack of standardization of much of the hardware of interest. Non-standardized rail car components must be formally structured as fully as possible to improve the reliability of the robotic automation. The first demonstrated application of robotics for UP's switchyards was the automated handle pull to decouple moving car consists in the hump location.

The second application demonstration of mobile robotics UP is focusing on is releasing the emergency brakes on cars, which must take place prior to uncoupling (which was demonstrated at UTexas). This brake release task requires either pushing or pulling a "bleed rod" that protrudes from the side of

each rail car. The requirements for each step of the process will be laid out in this thesis, and hardware recommendations given based on these requirements.

1.2 Background

Numerous tasks regularly take place in a switch yard on stationary cars (before or after uncoupling/re-coupling has taken place). These tasks include those necessary for standard operations (brake release, hose de-lacing/lacing), inspection tasks (looking for wheel cracks, truck bolt inspection, checking brake pad wear, etc.), and maintenance tasks (i.e. brake pad replacement). All of these tasks are currently done by humans.

Most of these tasks share some common requirements, while certain tasks may have additional, unique requirements. The specifics of the common requirements as they apply to the brake release task will be detailed throughout this thesis. Common requirements for a mobile platform include the following:

1. Capability, both physically and intelligently, to maneuver throughout a yard
2. Adequate platform localization performance within the RR car environment
3. Sufficient power reserve to support necessary running time between charges
4. Modularity and robustness to allow for in-house repairs/replacements

5. Sufficient static stability for task performance

Railroad switch yards are very complex, but can be made into well structured environments. In general, the greater the level of known structure beforehand, the more robust an operation will be. The knowledge of both yard layout and common switch yard obstacles can, and should, be utilized as much as possible.

Common requirements pertaining to robotic arms for mobile switch yard use are the following:

1. Low power requirement
2. Lightweight (also to minimize the power requirement)
3. Weather-proofing/ruggedization
4. Modularity to allow for in-house replacement from a minimum set of on-hand hardware
5. Internal controller and drives, preventing the need for tethering

Requirements for on-board vision development common to most tasks are:

1. Ability of cameras to function in changing conditions (varying lighting, dust, snow, possibly rain)

2. Object detection algorithm adaptation for each piece of hardware
3. Mapping of detection potential at different camera positions

1.3 Motivation

Thousands of rail cars go through hump yards and/or switch yards each day, and each and every car must be operated on at least once. Because the cars are stationary during the tasks suggested to be done by a mobile robot, safety is not as great a factor as for other dynamic operations (e.g. pin-pulling). The three major factors motivating a move away from human operators in this case are safety, reliability, and time efficiency.

These tasks are monotonous and repetitive, and are burdens for human beings, who are capable of higher-level thinking and decision making, especially given the proper information.

1.3.1 Mobile Robotics

Mobile robots have been targeted, at the request of UP, as the best option for these tasks primarily due to the minimal need for yard modification. On the whole, a mobile robot is largely self-contained. Installation/acquisition of additional hardware may be necessary to support their operation, but a properly designed/chosen mobile robot should be able to run in an existing, unmodified switch yard.

1.3.2 Reliability

Reliability is an issue with regard to inspection tasks. A common problem, according to UP, is inaccurate inspection records. Employees either neglect to look closely enough at the components they're inspecting, or don't do the inspection at all but record that they have. Specifically, the Truck Bolt, which has to be inspected from underneath the rail car, is often ignored during routine inspection. Neglect like this results in failures that could easily have been prevented, had the routine in place been properly followed.

1.3.3 Time Efficiency

Time Efficiency can be improved due to multiple aspects of this sort of automation; it should be noted that while the physical tasks would be done by robots instead of the current human operators, human operators would ideally be providing intelligent oversight of one or more robots at a time.

If neither the reliability nor speed of operation improve while moving from humans to robots, the fact still remains that multiple robots can operate simultaneously with oversight from a single human, thereby increasing the relative speed of the overall task. However, because automated performance is, in general, more repeatable and characterizable than human performance, improvements in both speed and reliability of a single operation are a desired result.

1.4 Generic Steps

1.4.1 Structuring the Hardware

The human brain easily, often subconsciously, accounts for unexpected variation or imprecision. Because these tasks have always been performed by humans, neither the need to establish and maintain standardized parameters (nor to document the distributions of these parameters) has existed. The distributions of non-standardized hardware must be understood as fully as possible, as the automation task should utilize prior knowledge to reduce uncertainty. Variations may exist in hardware position, size, shape, material, stiffness; or in task operating parameters, such as required travel, actuation force, actuation angle. It is important to identify the parameter variations that are relevant to the problem.

For positions, some reference point, or datum, must be established. This reference point should be something that is relevant to the location of the hardware of interest (components on RR cars), and that is either more easily identifiable than the hardware of interest or known some other way with respect to the mobile robot (e.g. the ground).

As much as possible, grouping of these different variations of the hardware should be done manually, to reduce processing time needed for further data mining.

1.4.1.1 Data Mining

Data mining refers to any technique used to aid in extracting information/correlations/relationships from seemingly incomprehensible, random, or uncorrelated data. Common data mining techniques include clustering (e.g. k-means clustering, hierarchical clustering), Principal Component Analysis (PCA), and Non-Linear PCA. One primary purpose of data mining is to find groupings of distinct “classes” of variations that represent significantly tighter distributions (smaller standard deviations).

To provide a basic, 1-dimensional example, assume a sample set has a mean of 5 with a standard deviation of 3. There might actually be two subsets with means of 1 and 9 whose standard deviations are both 0.5. Now, instead of trying to account for one value somewhere in the range [-1,11], two separate “types” would be planned for, in the regions [0,2] and [8,10].

A grouping with these characteristics would be easily identified by plotting the data. However, if a similar split were present across a set of 10-dimensional data, which is much more difficult to visualize, it would likely not be recognizable without the aid of data mining tools.

1.4.2 Modify/Optimize Object Detection

An easily adaptable object-detection algorithm was developed for the Automated Pin-Pulling System (APPS) [3]; the algorithm utilizes a depth camera (3D camera) and an artificial neural network (ANN). The most taxing modification that is needed for any new hardware is collecting images of the

object of interest for retraining the ANN. Other parameters in the algorithm should be tuned for the expected images.

The next, more open-ended task is developing heuristic requirements. The heuristics are checked during pre-processing (before running the computationally expensive ANN); if the conditions aren't met, further processing is not required. Heuristics serve to speed up the search process, as well as increase the algorithm's robustness to false-positive detections.

1.4.3 Determine Object Detection Potential

Potential for detection in any scene relies on a clear view of the image of interest. Any object will have an ideal viewing direction (most distinguishable, most visible/unimpeded, etc). This view should first be determined. Once the desired view is chosen, the positioning of the camera to obtain this view will be analyzed.

While a mobile robot would strive toward placing the camera in the ideal position to obtain the desired view, perfection is not often feasible, so a range within which detection can occur must be defined (two of the most common barriers to detection potential are described and mapped in Section 5.2). In addition to mapping detection potential just due to camera constraints, there may be constraints on feasible camera location due to any number of other factors (safety concerns, blocked views, insufficient reach/dexterity, etc). These additional constraints may decrease the size of the reliable workspace for camera position. Understanding the regions of expected

failure is paramount to a consistently reliable operation, and to the ability to safely modify operating schemes in the future.

1.4.4 Determine Mobile Robot Requirements

A mobile robot, for any task in a switch yard, must embody the features described in Chapter 12. However, each individual task has its own requirements, including maximum payload, size of workspace, variation of motion required for each iteration, speed during operation, end-effector precision, and dexterity during operation.

To plan the operating scheme, one should start by examining the method used by a human operator.

1.5 Results of Steps as Applied to Brake Release

1.5.1 Structuring the Brake Rod Hardware

The brake release valve handle is a pin on what is effectively a sprung ball joint, which releases when pushed to any direction by $\sim 30^\circ$. The bleed rods (one for each side of the car) are pinned to the end of the valve pin. So, when the rod on one side is pushed or pulled, the rod on the other side moves with it. This is only a problem if the opposing side impedes proper travel.

Measurements were taken, in all 3 dimensions, from the rail and from the uncoupling lever. The measurements from the rail can provide constraints that the mobile platform must meet, and both the rail measurements and the

lever measurements can be used to focus in on a search area intelligently. The last measurements that were taken are the size of the loop in the rod-ends (for replication and possible end-effector design purposes), and the distance to the nearest rigid face on the car behind the rod-end (to establish a reliable depth-filter range and required end-effector clearance). The measurements taken for 10 cars (Car 1 was taken for repairs just after measuring began) are in Table 2.1, in Chapter 2.

This data ideally is grouped into 4 clusters, but clusters 3 and 4 have small enough deviations from the two major clusters that they don't need to be accounted for separately with hardware (the clustering evaluation is shown in Figure 2.2). A schematic of the possible rod positions is shown in Figure 2.5.

1.5.2 Brake Rod Detection

After adequate bleed rod training images were gathered and used for training, the detection was tested on recorded videos from the San Antonio yard visit to establish a performance baseline. From here, additional training images were taken, both True images (bleed rod images, shown in Figure 1.1) and False images (images of objects frequently giving false-positive detections), and the ANN was retrained to improve performance.

Object detection parameters that were tuned using the initial training set include sliding window size, Neural Network output threshold, "hit count" threshold, and sliding-window step size. The existing heuristic threshold values

that were modified for the rod image are the populated pixel ratio and centroid location. New heuristics that were added are rod height (as it appears within the window), window edge crossing, and rod vs. loop width ratio.

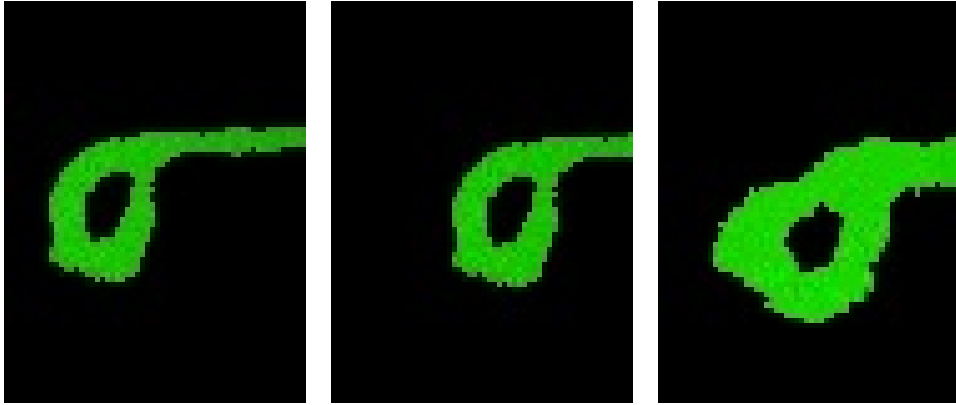


Figure 1.1: “True” rod training images

Using the videos taken during the San Antonio yard visit, testing the rod detection algorithm shows that 99.52% of the detections (i.e. “potential rod” windows within a frame) are true positives, and 100% of the full recognitions (i.e. confirmed rods within a frame) are true positives.

1.5.3 Brake Rod Detection Potential

The easily distinguishable side view must be directly in view, and that image must be sufficiently clear. Some views will be obstructed, but never from both sides of a rod; therefore, a pair of cameras must be used to search both forward and backward along a rail car. Less than ideal images can potentially still be used for detection, but, once degraded past a certain point, proper detection becomes impossible. A useful workspace can be defined for camera

position, relative to the bleed rod, such that detection is possible anywhere within the workspace. Additionally, a metric for detection potential has been defined to aid in optimizing camera position.

The two parameters that affect the image requirements are viewing angle and viewing depth (see Figure 5.2). These parameters affect the detection potential equally, as either a poor viewing angle or a poor viewing depth can make detection impossible.

The metric used to describe the detection potential based on the viewing angle is the **ratio of projected rod length to actual rod length**. This value is 1 at 0° (the full length of the rod is visible), and is 0 at 90° . The most descriptive metric related to viewing depth is the ratio of the rod width to the camera noise (w/ϵ), which is equivalent to a standard **signal-to-noise ratio**.

A single metric is necessary to evaluate the overall “goodness” of a camera position with respect to the rod. A linear weighting scheme is used to combine normalized metrics. The allowable workspace for the camera, as determined using these metrics, is shown in Figure 1.2, where any value greater than 0 denotes a viable camera position. In the figure, the camera is positioned at the origin, and the plot refers to “goodness” of rod locations relative to the camera.

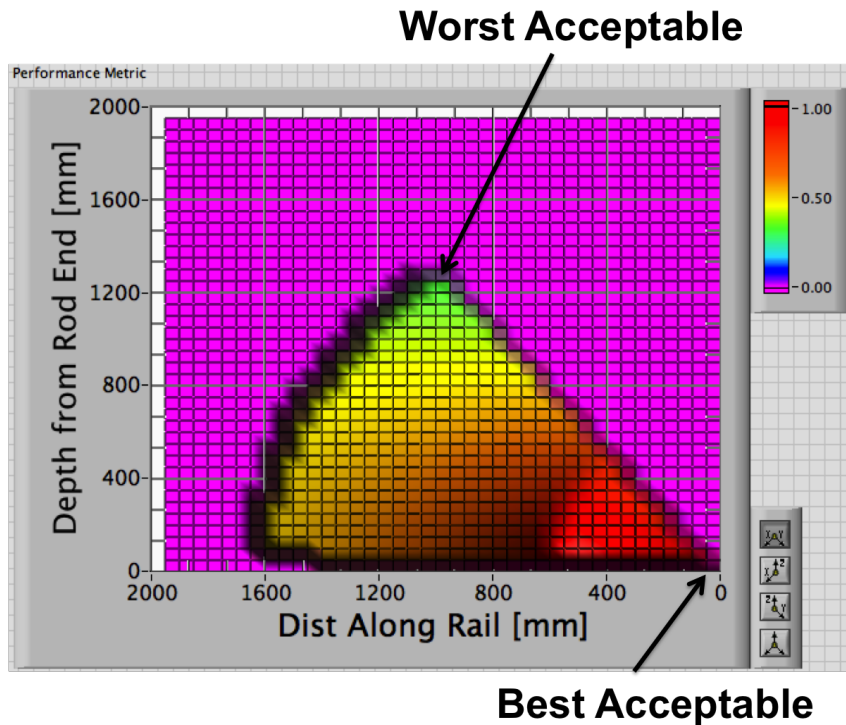


Figure 1.2: Camera workspace evaluation

1.5.4 Mobile Robot Requirements

1.5.4.1 Manipulator

The requirements for the manipulator intended for the brake release task include allowable payload, range of motion (max-to-min end-effector height, i.e. workspace), and available travel at the desired heights.

The maximum payload must account for end-effector weight and required rod-operation force (~ 30 lbf). The workspace must allow the end-effector to reach both the high and low rod configurations (offset by appx. 45"). Each valve requires the same operating travel to be opened, but

rods have inherent compliance (which adds to the necessary travel) based on geometry. The maximum possible travel must be available to the manipulator at each of the two heights.

There are very few lightweight, low-power manipulators on the market that meet the payload requirement. However, within such a small, well-defined workspace, careful planning can ensure success without exceeding manufacturer ratings.

1.5.4.2 Mobile Platform

The requirements for a mobile platform for the brake release task must be considered separately for the static rod-pushing operation and for the driving portions of the task. The driving portion requires the capability to traverse the known terrain of a switch yard, and to somewhat precisely control steering. The static portion of the brake release task requires that a mobile platform have sufficient stability to push the rod without sliding or falling over.

More complex tasks will certainly require Independent All-Wheel Steering, but this brake release task should be feasible using a mobile platform with Ackerman Steering. In no case should skid-steer be used for a mobile platform in a switch yard.

Figure 1.3 shows the parameters considered when evaluating the geometry-based static stability of a mobile platform for the brake release task.

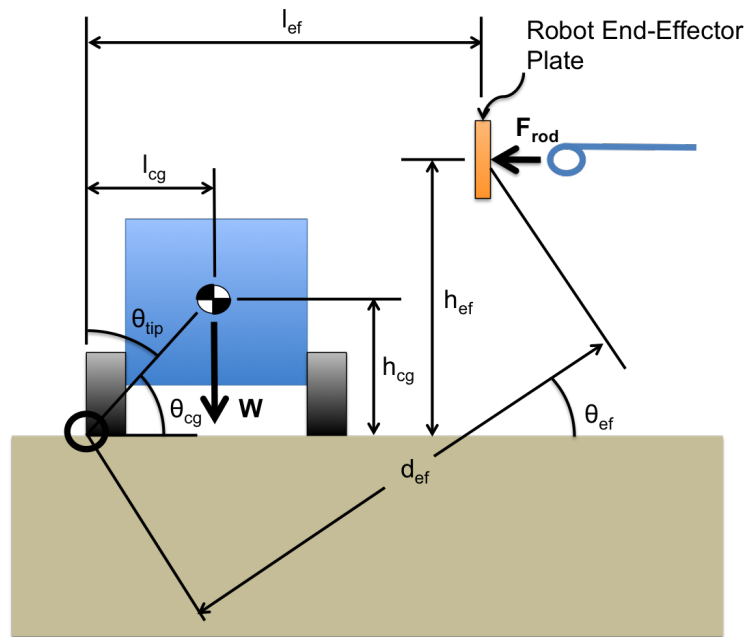


Figure 1.3: Mobile Platform static stability schematic

The static stability evaluation shows that only one mobile platform on the market is immediately suitable for continued development of an Automated Brake Release System.

Chapter 2

San Antonio Yard Visit - Understanding the Problem

The primary goal of visiting this rail yard in San Antonio was to gather information about the brake release rods (bleed rods). The information needed to establish the constraints of this task includes rod positions, mounts, clearances, travels, forces, and any other still unknown variations or constants.

This chapter describes the findings during the San Antonio Yard Visit, the conclusions drawn from them, and outlines the plan for operation and initial testing. The plan for testing described here was carried out in a controlled laboratory environment, which is discussed in following chapters.

2.1 “Service Valve” Housing

Each housing is very similar, with a few different valve models, but all can be assumed to be equivalent for this purposes. The only documentation that could be located was a valve instruction manual [5]. The housings seem to all be standardized, and have the same mounting configurations (bolt-hole placement, etc). Each housing has a release valve, a release reservoir input, a bleed output, and a brake piston pressure output. When the valve is opened,

the release reservoir (which contains air pressurized to ~ 70 psi) is opened to the brake piston release valve. This pressurized air opens the brake piston release valve, which locks open and bleeds the brake cylinder, allowing the brake piston to retract over a period of about 30 seconds.

2.1.1 Valve Handle

The valve handle is a pin on what is effectively a sprung ball joint, which releases when pushed to any direction by $\sim 30^\circ$. The bleed rods (one for each side of the car) are pinned to the end of the valve pin. So, when the rod on one side is pushed or pulled, the rod on the other side moves with it. This is only a problem if the opposing side impedes proper travel.

Bleed Rod



Service Valve Housing



Bleed Rod

Valve Handle

2.2 Bleed Rod

The rods themselves are made from standard 0.5" steel rod stock. Each rod is custom-bent for the particular car, because valve-mounting positions are not consistent enough for standardization. While the valve connections always requires the same amount of travel, there may be much additional travel at the rod-end due to compliance along the rod (bending due to large moment-arms). This variation for each rod position should be considered.

2.2.1 Variations

The most typical bleed rod end is a standard loop. One variation of this is a loop end with a notch rubbed into the shaft from the holding plate contact, shown below. The other variation is an L-shaped end instead of a loop. None of these rod ends were seen in the yard.



Figure 2.1: Bleed Rod with worn notch

2.3 Reservoir

The pressurized air in the reservoir serves only as the piston valve release force. When the bleed rod is operated and its valve is open, the air blowing off (audibly) is from this tank, not from the brake piston itself. The reservoir is usually good for multiple brake piston releases before needing to be re-pressurized. As its supply is depleted, the pressure decreases. As the pressure decreases, the pitch of the blow-off sound decreases. For this reason, testing different reservoir pressures in the lab will be necessary.

2.4 Measurements and Analysis

There is little to no documentation or specifications regarding the bleed rod position, shape, or size. In order to automate a search process as robustly as possible, the problem needs to be standardized and constrained wherever applicable. Measurements were taken with respect to all standard, documented, recognizable baselines. From this point, the data can be analyzed to find any patterns or groupings that can be utilized.

Measurements were taken, in all 3 dimensions, from the rail and from the uncoupling lever. The measurements from the rail can provide constraints that the mobile platform must meet, and both the rail measurements and the lever measurements can be used to focus in on a search area intelligently. The last measurements that were taken are the size of the loop in the rod-ends (for replication and possible end-effector design purposes), and the distance to the nearest rigid face on the car behind the rod-end (to establish a reliable depth-filter range and required end-effector clearance). The measurements taken for 10 cars (Car 1 was taken for repairs just after measuring began) are in Table 2.1

Table 2.1: Bleed Rod position measurements taken on cars during the San Antonio Yard Visit

Car #	Car Color	Car Serial #	Car Serial #	Rod Position	Rod Height from ground [in]	x-distance from handle [ft]	rod-end depth from handle [in]	depth to nearest face/object [in]	rod-end depth from rail [in]	Rod-ID [in]	Rod-End OD [in]	Rod Diameter [in]
1	Gray	36081	43									
2	Black	400549	31	Center	29.08	-6	4.5	28.5	0.5	1.6	0.5	0.5
3	Gray	86059	56	End	4.25	-5	6.25	32	0.45	1.47	0.53	0.53
4	Gray	705463	56	End	3.92	-11	3.5	30	0.65	1.65	0.53	0.53
5	Gray/Rust	84033	44	Center	34.50	-9	7.5	32	0.3	1.3	0.5	0.5
6	Red/Brown	793007	15	Center	28.75	1	3.75	29	0.51	1.54	0.52	0.52
7	Brown	4049	22	Center	22.50	-6	5.25	28	0.45	1.41	0.45	0.45
8	Brown	4039	24.5	Center	22.50	-6	7	29.5	0.4	1.35	0.5	0.5
9	Yellow	866069	26	Center	35.75	-4	4.5	26	0.58	1.57	0.5	0.5
10	Yellow	2035	57	End	3.75	-13	5.5	28.5	0.57	1.6	0.5	0.5
11	Gray	9707	57	End	3.75	-8	4.5	33	0.33	1.3	0.53	0.53

Because humans don't always readily see possible groupings or relationships, k-means clustering was performed on the data. The two obvious groupings of rod position are low and centered along the length of the car, or high near one end of the car. However, the clustering extracted 2 clusters within the already known groupings.

When looking for the most relevant number of clusters to explain data groupings, an F-value can be used to evaluate how well the additional clusters account for variation. F-values are a ratio of variation within the clusters to variation between the clusters. The smaller this ratio, the more informative the clusters are, relative to the un-clustered data set. Additional clusters will decrease the F-value, but eventually adding more clusters results in little to no improvement in the resulting information (see Figure 2.2).

This data ideally is grouped into 4 clusters, but clusters 3 and 4 have small enough deviations from the two major clusters that they don't need to be accounted for separately with hardware.



Figure 2.2: Plot of F-Values vs. number of clusters for bleed rod position data

2.4.1 Force Measurements

2.5 System Assumptions

- The rod will always be in one of three positions, two of which data has been collected for that is representative of the entire population
- If the rod does not travel far enough to open the valve, the robot should be able to detect it, but there is no requirement to make it release
- The rod ends will always be visible (view not impeded) either from the right or the left



Figure 2.3: Bleed Rod Position 1 Figure 2.4: Bleed Rod Position 2

2.5.1 Expected Layout/Procedure

At least one valve housing will be affixed to the existing rolling-stock chassis. Bleed rods will be made to route to both standard positions, or positions that will be sufficiently representative of them. The robot will, first, find the uncoupling lever at the beginning of a car, which will be used as a reference point for the start of a car. Then, with cameras at the proper starting height, the robot will begin traversing the car looking for the brake release rod-end. If it is not found within the expected distance from the uncoupling lever, then the cameras will move to the expected height of the next possible position and continue searching.

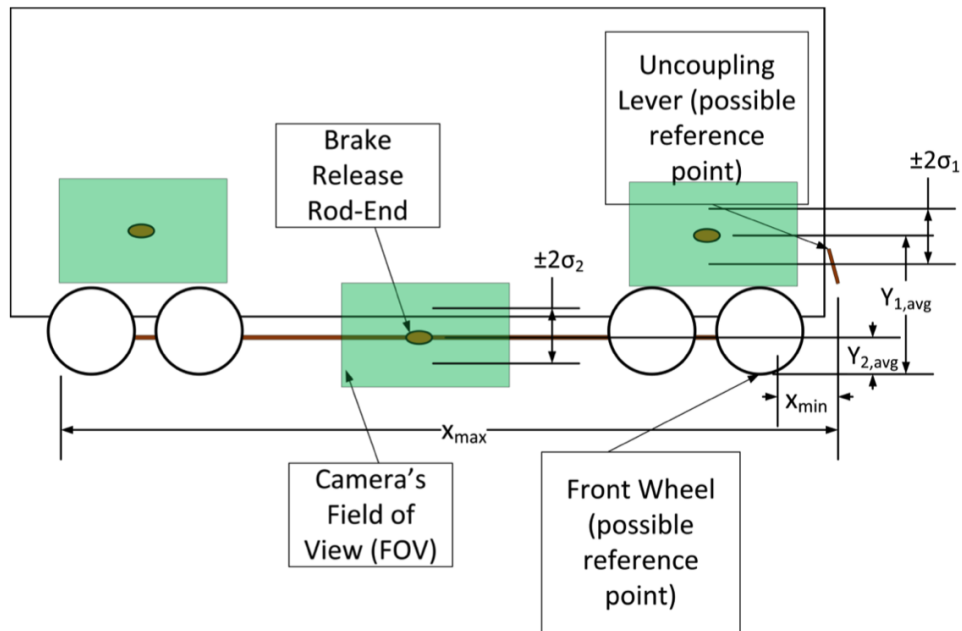


Figure 2.5: Procedure schematic

Once a rod-end is found, in either expected region, the robot will receive its coordinates, align its end-effector with the rod, move toward it until contact is made, then actuate a pneumatic cylinder to push the rod. When the rod is pushed, it will be held while the system verifies that the valve has properly opened and is blowing off the pressurized air (microphone, LabVIEW); the rod should be held 0-3 seconds, depending on the valve model, so it will be assumed that the valve should be held open for at least 3 seconds.

2.6 Necessary Components for Lab Testing System

- Service Valve Housing
- 2-3 custom bleed rods
- Pressurized Air Reservoir

2.6.1 Necessary Variability

- Bleed rods should have vastly different geometries
 - If different actuation travels are to be tested, this variation is necessary, as travel depends only upon rod geometry
- Bleed rod mount should be mobile between 3 and 8 inches
 - The distance from the rod end to the guide plate face may change the view of the depth camera; this change is not anticipated to be detrimental, but it should be planned for in case the change is significant and the current assumption is incorrect
- Easily swappable view impedance to test backward-facing camera
 - Both impeded and unimpeded cases (missed by one camera, found by the other) need to be testable

2.7 Vision

2.7.1 Requirements

We need to be able to see every rod end in the Field of View (FOV), ideally as close to the center of the frame as possible. Any view that is impeded by one camera angle should be visible from the second camera angle.

2.7.1.1 Different Heights

Based on the averages and possible ranges of the two different rod positions, rods will be too far separated to all be within the FOV of a stationary camera. This, then, requires two distinct camera locations.

2.7.2 Camera Configuration

2.7.2.1 Multiple Cameras

At least 2 cameras will be required, possibly 3, in the horizontal plane of the bleed rod. The first camera will be facing the same direction as the robot is moving along the track, about 30° from parallel to the rail. This view of the rod seemed to produce the best results from the video taken in the yard. Getting a clear side-view of the rod is imperative, as that is where its most distinct features are visible.

One car was encountered on which this view of the bleed rod was impeded by a pipe. In this case, a camera at the same angle, facing the other direction, picked up the rod-end clearly. As stated in the assumptions,

it is assumed that at least one of these two views will always be clear, as was the case with all bleed rods encountered in the yard.

A third camera could point perpendicular to the rail, to be used to find the uncoupling lever. The first two cameras should be mounted close to the cars (need to check railroad standards and get as close as possible), but the third will be mounted at the same distance as it was for the uncoupling lever detection developed in the Automated Pin Pulling Project [3] (report submitted to UP, Jan., 2014). More details are given regarding placement of the first two cameras to maximize detection potential in Section 5.2.

2.7.2.2 Multiple Camera Pairs

One option to address the two different rod heights is having separate sets of cameras. The RRG has successfully run two ASUS Xtion cameras simultaneously, but the initialization of both has not been successfully automated. Having 4 or more cameras running simultaneously should be possible, but will likely require multiple separate USB-Hubs and may require driver customization outside of the standard open-source functionality.

Separate pairs of cameras running simultaneously allows both rod heights to be searched at all times, eliminating the need for any knowledge of location along a car that is being searched. However, the additional computational overhead required to run the detection algorithm on 4 camera feeds simultaneously may overwhelm an on-board computer's capabilities.

2.7.2.3 Vertical Axis

An alternative to multiple sets of cameras is having one set be mobile between the two required heights. However, because only one possible rod height can be searched at any given time with this configuration, current location along a car must be known to make this option viable. Research needs to be done to find a cost-effective, controllable vertical axis that fits these requirements. Pneumatics would be a robust, simple option if fine adjustment or encoding is not required, but would require on-board pressurized air.

- Approximately 750mm of travel
- Robust encoder
- Low power requirement
- Lightweight

A mobile set of cameras would minimize the computational requirement, but would add weight, required power, and physical complexity.

Chapter 3

Laboratory Test-Bed Modifications

The Rolling Stock Test-Bed, built for the Automated Pin-Pulling System (APPS), is used as the base for the Brake Release mock-up. Figure 3.1 shows the unmodified chassis in the RRG laboratory.

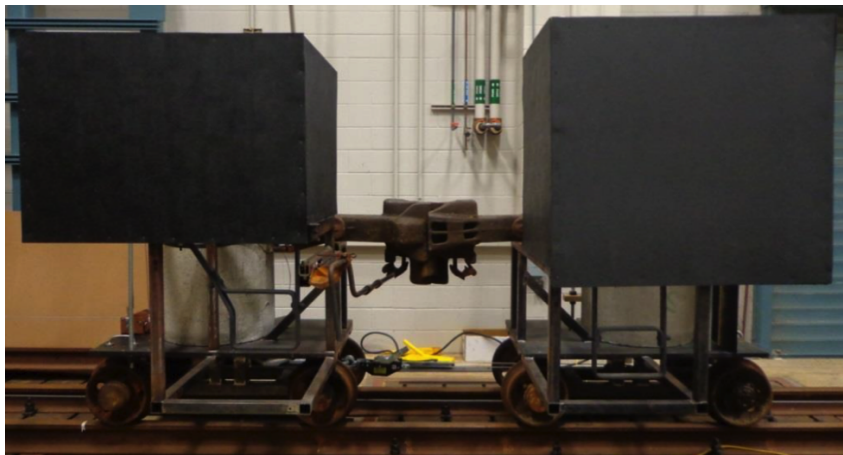


Figure 3.1: Unmodified rolling stock test-bed

3.1 Service Valve

The valve acquired from the train yard in San Antonio is the service valve used in most train cars. The bleed rod is attached to the service valve via the release valve handle. On the train car the service valve is attached to all

the other components via the pipe bracket. When the bleed rod is pressed, air is released from a pressure reservoir and a loud bleed off sound can be heard. This sound will later be used to verify that the rod has, indeed, been pressed and the brake release initiated.

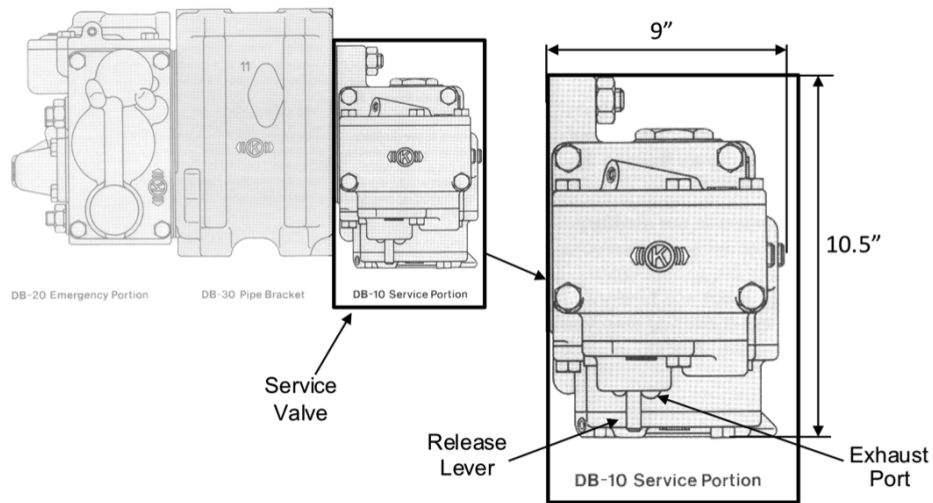


Figure 3.2: “Service Portion” schematic

3.1.1 Mounting Location

In choosing a mounting location on the test chassis for the Service Valve, the two primary considerations were possible rod configurations and reliable structure for support. If the service valve were mounted too close to the front of the chassis, there would be too little room to allow for multiple rod configurations. A location was chosen in the rear of the chassis on the lower portion of the platform, with the Service Valve inverted (release lever pointing upward instead of downward). Being placed at the back of the chassis provides

sufficient space for rod variations, and being inverted allows the gusseted bars to be used for support, while still making the release lever accessible. A rod can connect to the release lever from either the top or bottom of the range of rod locations, with extra space for extraneous bends (discussed below, in Section 3.2).

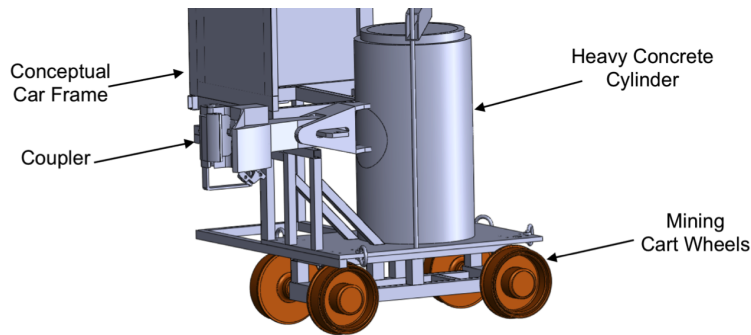


Figure 3.3: Rear view of test chassis without service valve

When designing the mounting plate, the pneumatic layout of the service valve had to be established. Through testing (given minimal literature regarding valve layout/operation), the pneumatic ports of interest on the service valve were determined [5]. All of the pneumatic ports on the Service Valve are on the side that normally is bolted to the Pipe Bracket. The circles labeled as Pneumatic Ports in Figure 3.2 are bores to hold rubber seals used to plug the ports. The hole labeled “Air Supply,” an $1/8$ " NPT-threaded hole for a push-to-connect pneumatic fitting, corresponds to the only port requiring pressure for testing purposes; this pressure provides the blow-off sound. Figure 3.6 shows the Air Supply port connected to the Air Supply Hose. The service

valve already had bolt holes which were also laid out on the mounting plate.

The mounting plate is then secured to the steel chassis frame.

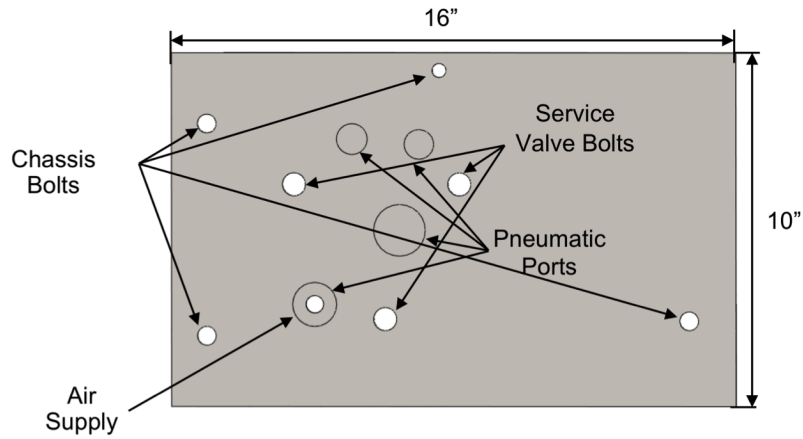


Figure 3.4: Service valve mounting plate schematic

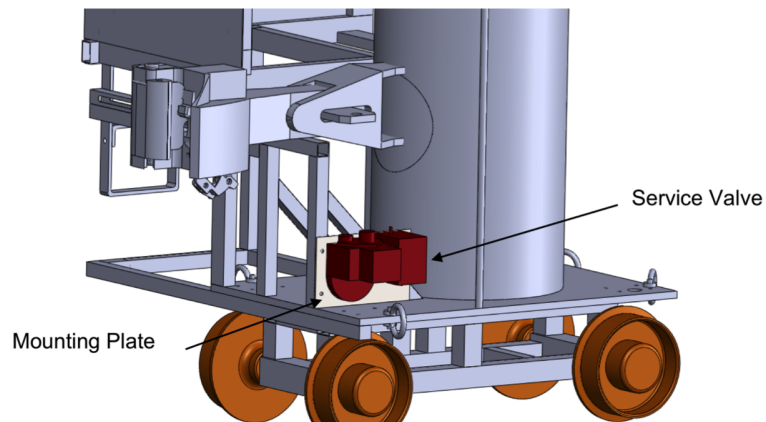


Figure 3.5: Rear view of test chassis with installed service valve

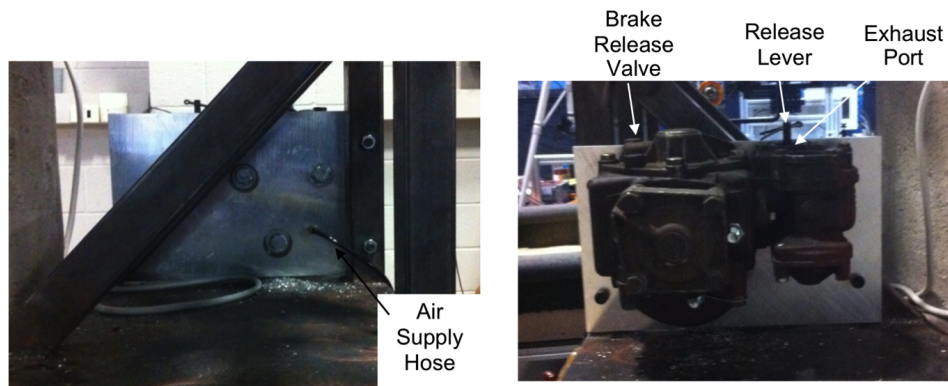


Figure 3.6: Installed mounting plate and service valve (front/back views)

3.2 Bleed Rods

3.2.1 Requirements

The bleed rods (internal rods and rod-ends) were separated using a modular design, to maximize reconfigurability. Each rod-end and internal rod are threaded and connected using a turnbuckle; a turnbuckle is used, instead of a rigid, fixed connection, to provide a variable length. The bleed rods seen in the San Antonio yard were all in a certain range from the face of the rail cars, between 3.5" and 7.5". The bleed rod in the lab must be able to be set to any distance within this range.

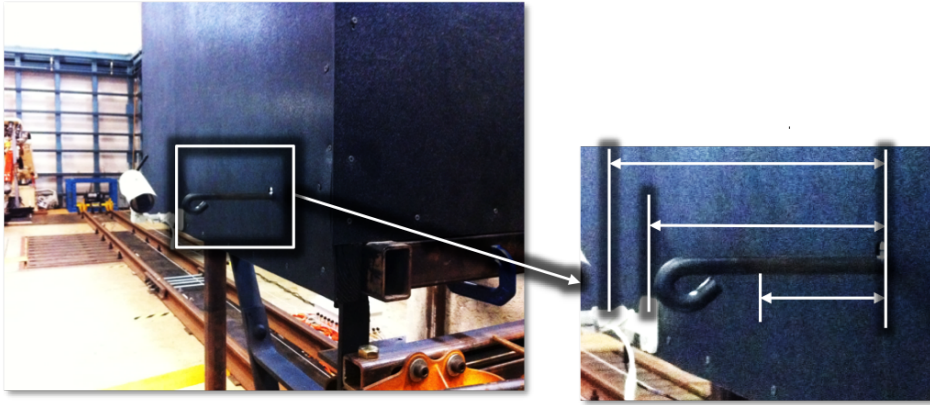


Figure 3.7: Bleed Rod showing length variability

3.2.2 Vision Testing

Before the internal rods were designed and fabricated, the bleed rod end needed to be attached to the chassis for vision testing. Steel plates were designed and fabricated to hold the rod in place via a clamp which allows for manually setting the rod itself within the desired range of lengths.

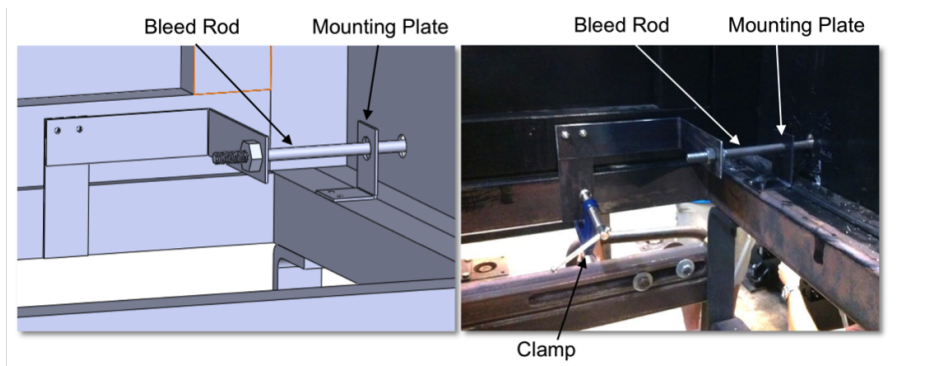


Figure 3.8: SolidWorks model and installed bleed rod end showing internal temporary mounting components

3.2.3 Rod Operation Testing

The various rod designs are needed to test the range of conditions observed in the San Antonio yard. The full range of possible travels to start the air release on the brake release valve needed to be tested. In order to do this, it was decided to use the extra space to provide for multiple internal rod configurations; one being the most direct route to the release valve, and the rest with varying degrees of bends (internal bending moments). The internal rod attachments will also be threaded on the end and attached to the external portion using a turnbuckle, retaining the ability to change the bleed rod length at any time.

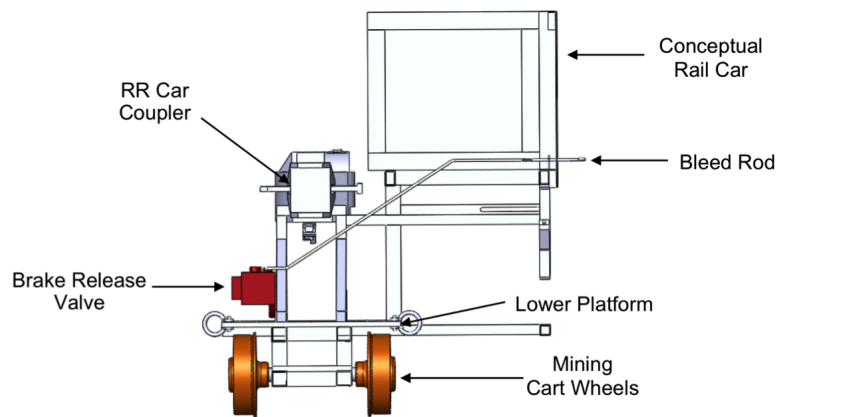


Figure 3.9: Left View (SolidWorks) of first Bleed Rod design

3.3 Pneumatic Reservoir

Prior to the introduction of the pneumatic reservoir to the test set-up, the pneumatics of the service valve were directly connected to the air source located along the wall of the lab. This air source is held at a constant 80 psi which is in the desired range ($\sim 80 - 90$ psi), but whenever the release valve is opened an almost immediate pressure drop was observed resulting in a short loud burst of sound followed by an extended, quiet, bleed-off noise. However the desired sound, which emulates the sound heard in the train yard, is a loud, extended, bleed-off noise at about the volume of the short initial burst that was obtained. To achieve the desired bleed-off noise the wall air source was connected to a 10 gallon air tank which was then connected to the service valve. The pressure manifold that came with the air tank still limited the flow too much, so a custom manifold was designed.

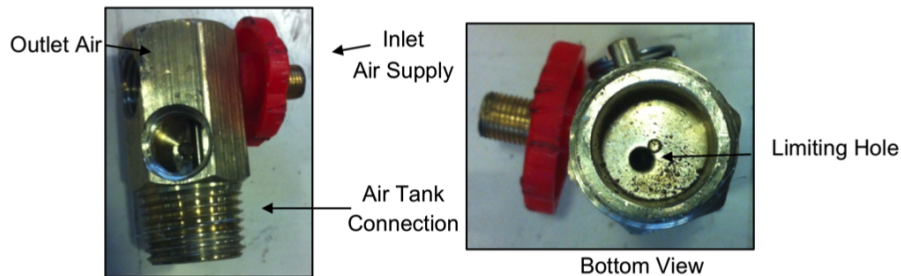


Figure 3.10: Previous pressure manifold and the hole limiting air flow

The custom manifold eliminates the limiting hole found in the standard manifold and also increased the diameter of the tube connecting the reservoir

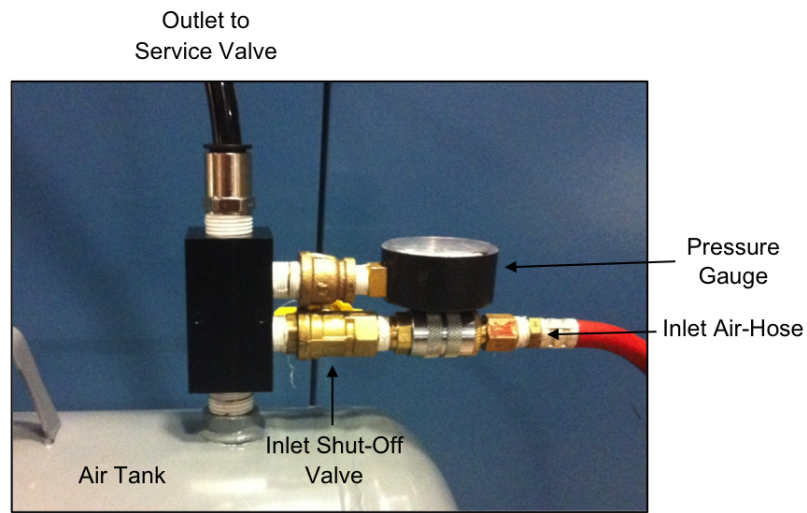


Figure 3.11: Custom pressure manifold for air tank

to the service valve to $\frac{1}{2}$ " from $\frac{1}{4}$ ". The service valve mounting plate air supply hole was increased from $\frac{1}{8}$ " NPT to $\frac{1}{2}$ " NPT.

3.3.1 Supporting Calculations

By increasing the diameter of all the connections, eliminating the hole in the previous pressure manifold, and introducing a pneumatic reservoir to the service valve greatly diminished the observed supply pressure drop. This was accomplished by reducing the head losses between the air source and the valve. The pressure loss, Δp , due to head loss can be expressed as:

$$\Delta p = \rho g h_f \quad (3.1)$$

where:

- ρ is the fluid density
- g is the acceleration due to gravity
- h_f is the head loss due to friction (major loss)

The head loss, h_f , can then be calculated using the Darcy-Weisbach equation:

$$h_f = f_D \frac{L V^2}{D 2g} \quad (3.2)$$

where:

- f_D is the Darcy friction factor, found via a Moody diagram (Figure 3.12)
- L is the length of the pipe
- D is the internal diameter of the pipe
- V is the average fluid velocity

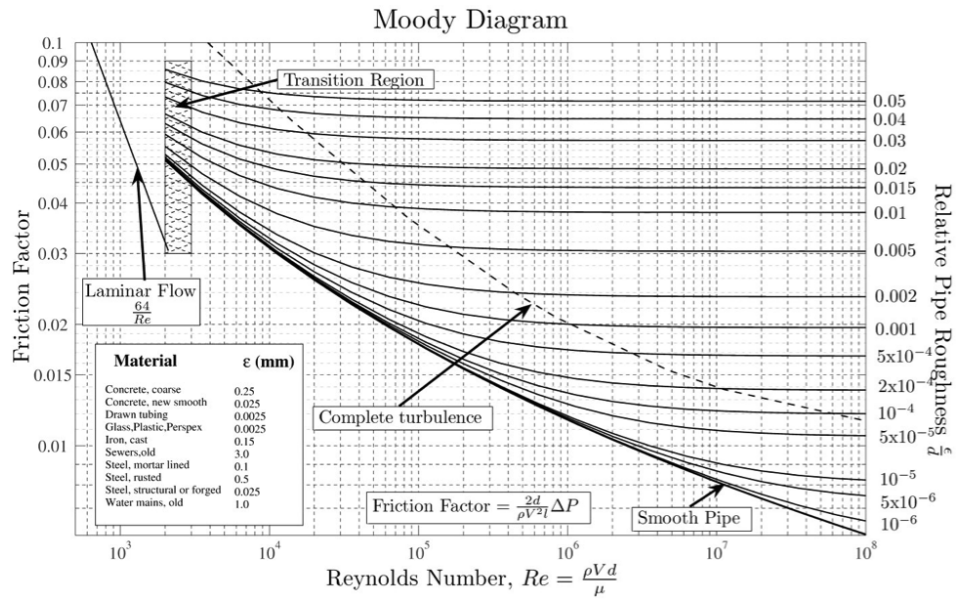


Figure 3.12: Moody diagram used to find the Darcy friction factor

For the conditions used in this situation (80 psi reservoir pressure, room temperature, and a pipe diameter of $\frac{1}{4}$ " before and $\frac{1}{2}$ " after) it is fair to assume turbulent flow. The relative pipe roughness ($\frac{\epsilon}{d}$) of the $\frac{1}{2}$ " pipe will be lower than that of the $\frac{1}{4}$ " pipe and the Reynolds Number ($Re = \frac{\rho V d}{\mu}$) will be higher for the $\frac{1}{2}$ " pipe than for the $\frac{1}{4}$ " pipe. Knowing this and referring to the Moody Diagram, Figure 3.12, the Darcy friction factor (f_D) of the $\frac{1}{2}$ " pipe will be lower than that of the $\frac{1}{4}$ " pipe. If the length (L), gravity (g), fluid density (ρ), and velocity (V) of the two pipes can be considered close enough to one another, the lower Darcy friction factor and larger diameter of the $\frac{1}{2}$ " pipe will result in a smaller head loss (h_f), thus creating a smaller pressure drop (Δp). This smaller pressure drop means a greater pressure at the exhaust outlet,

which will result in a higher exhaust flow rate, able to sustain the desired blow-off sound.

Chapter 4

Robot Modifications

4.1 End-Effector

The bleed rod requires only a linear force, in most cases, applied in a direction directly perpendicular to the rails. The simplest contact piece that can be used is a flat plate. In the implemented version, a Force/Torque sensor should be used, possibly in conjunction with a more complex contact plate (see Figure 4.1, and Section 6.3.1 in [8] for other possible contact plate shapes), but here, to save time and money, a simple flat contact plate was used. The plate installed is HDPE, which offers a bit of compliance during contact, preventing wear on the rod-end.

A shaped contact plate, like the one shown in Figure 4.1, has the advantage of providing a “restoring force” toward its center. If the rod is contacted anywhere but the center of this contact plate, there will be a lateral force toward center; the F/T sensor would sense this force, and, through an active force control scheme, could reposition itself to be centered. One flaw with this design is the lack of information regarding distance from the center, i.e. the lateral force versus the longitudinal force (dependent upon contact angle) is of constant magnitude. If the interior slope were quadratic, for

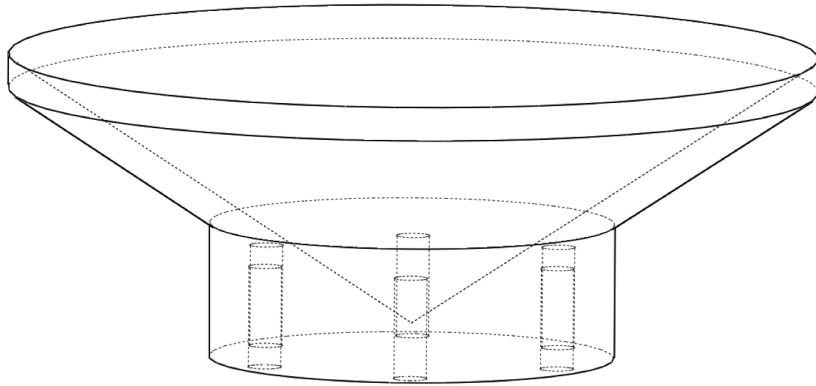


Figure 4.1: Conceptual conical contact plate

instance, then the $\frac{\textit{lateral force}}{\textit{longitudinal force}}$ would be greater near the edges, so the user and/or robot would then know the magnitude of the alignment error in addition to its direction.

The end-effector used in the hard-automated solution in the APPS was used here, as it was designed to support a pneumatic cylinder. Because it was concluded early-on that a linear force would be sufficient, in order to start development and testing as quickly as possible, the pneumatic cylinder is used to apply the linear force.

4.2 Camera Mount

4.2.1 Background

The camera mount design utilizes the camera mount and securing clamps from the APPS. A few modifications were made to the original design to better fit this new application. The thickness of the back face was increased to

allow a through hole for the bolt that would secure the mount to the mounting frame and serve as the pivot point for the angle variability. The recessed bolt holes on the back face were then used to attach a T-strut to the camera mount that was used to lock the camera mount in the desired orientation.

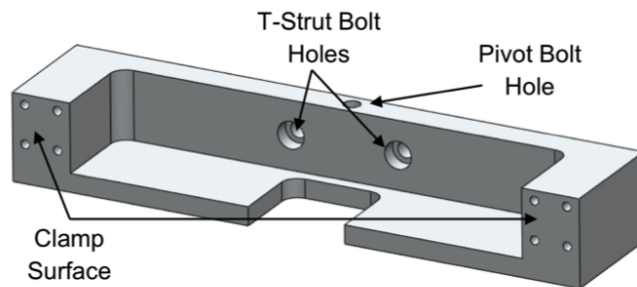


Figure 4.2: Individual camera mount

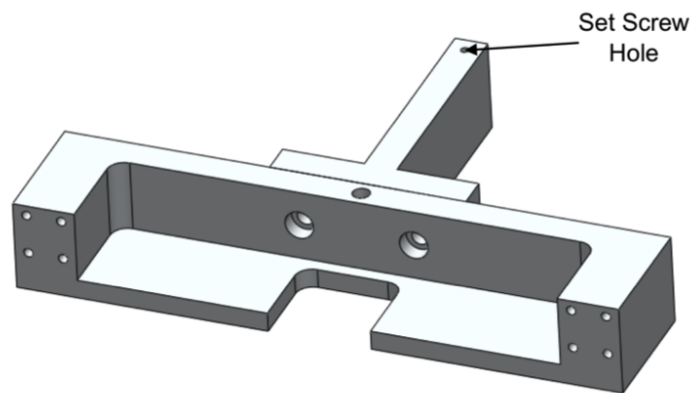


Figure 4.3: Individual camera mount with T-strut

A mounting frame was designed to hold two cameras facing in opposite directions to allow it to still detect the brake rod while moving away from it in case it was unable to detect it while moving towards it due to an obstruction.

The mounting frame was split into four separate pieces to allow for compliance and decrease the overall build time, the top plate, the bottom plate, the rear plate, and the middle support. It was important that the bolt holes of the top and bottom plate align to ensure that the set screw holes would properly align with the set screw hole in the T-strut, so the geometry of each piece was designed such that each piece could lock the position of another.

All the vision mounting pieces were 3-D printed with a precision of $\pm 40\mu\text{m}$. The dimensions of the mounting frame plates had to comply with the build space of the printer.

4.2.2 Angle Testing

In order to determine the optimal orientation of the camera, 2 arrays of set screw holes were designed for the top plate of the mounting frame, one for each camera mount. The arrays consisted of 19 holes in a quadrant around the pivot bolt hole, set in 5° increments spanning from parallel, 0° , to perpendicular, 90° .

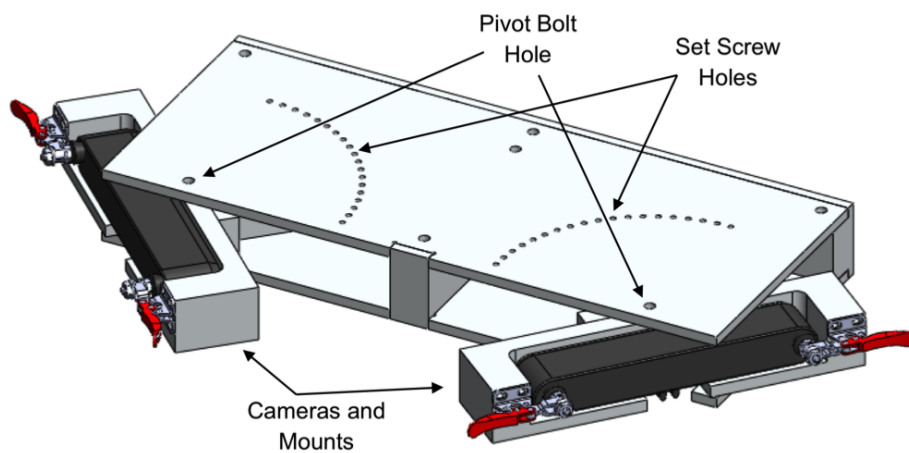


Figure 4.4: Cameras and mounts placed in dual camera mounting frame

Chapter 5

Robot/Vision Hardware Requirements

5.1 Rod Height

There are two primary rod locations on rail cars, and the two heights around which the types are centered must be planned for with the hardware. Height is the only changing factor that needs to be addressed with hardware; the change in position will be accounted for by the mobile platform and its operating scheme.

5.1.1 Robot and Mobile Platform During Actuation

The varying height is critical because it results in a change in moment arm of the actuating force (see Figure 5.1). A higher rod, with the same force applied, produces a greater moment about the robot platform wheels (h_{ef}), and requires both more powerful actuators and a more stable base.

As long as the necessary force is within the specified workspace and maximum payload of the robotic arm, the actuators should not cause a problem. The key, then, is using a robotic arm with a large enough workspace to encompass both heights, and having a platform that can get the arm into a position that allows for proper operation.

The mobile platform itself, however, may need additions to account for a significant rolling torque. Mobile platform recommendations are discussed in Chapter 12, and possible modifications are discussed in Section 12.2.

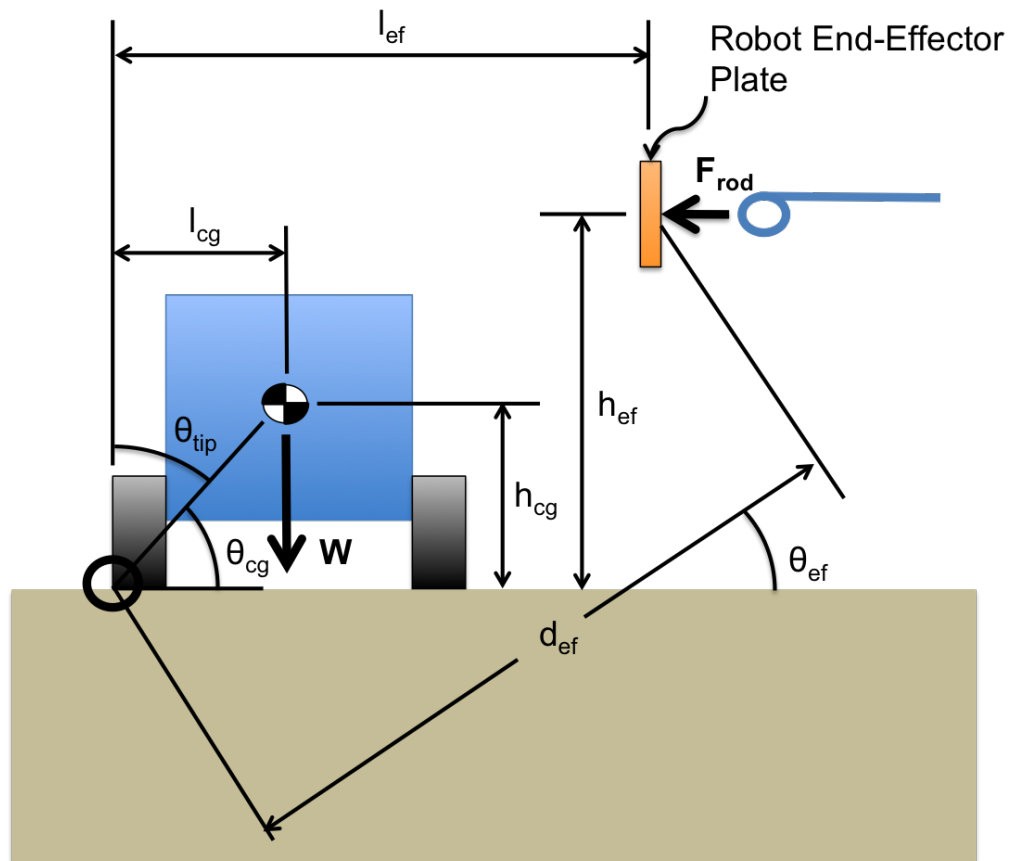


Figure 5.1: Mobile Manipulator platform loads at different rod heights

5.1.2 Mobile Platform Stability Requirements

The following equations, using the variables as defined in Figure 5.1, can be used to determine the viability of a particular mobile platform. The condition in Equation (5.1) must be met for the platform to prevent any tipping at all. l_{tip} , as defined in Equation (5.6), is the actuation distance (horizontal) required, at the end-effector, to fully tip the platform over backward. It can be easily seen from Equation (5.1) that if a platform meets the criteria for the upper rod location, the lower rod location will not cause a problem.

$$l_{cg} > \frac{Fh_{ef}}{W} \quad (5.1)$$

$$\theta_{cg} = \tan^{-1}\left(\frac{h_{cg}}{l_{cg}}\right) \quad (5.2)$$

$$\theta_{tip} = 90 - \theta_{cg} \quad (5.3)$$

$$\theta_{ef} = \tan^{-1}\left(\frac{h_{ef}}{l_{ef}}\right) \quad (5.4)$$

$$d_{ef} = \sqrt{h_{ef}^2 + l_{ef}^2} \quad (5.5)$$

$$l_{tip} = d_{ef} \cos(\theta_{ef} + \theta_{tip}) - l_{ef} \quad (5.6)$$

5.1.3 Vision

Two cameras (one facing forward, the other facing backward) need to be operating at each rod height. The two cameras account for a possible obstruction from either direction and insure that the rod will be seen.

Originally, it was concluded that cameras facing away from each other would be ideal (shown in Figure 6.1), due to possible interference between the two cameras. The main drawback of this configuration is that it results in a “dead zone” in between the cameras, where neither camera can see. However, testing has shown no degradation of rod visibility with cameras looking at it from opposite sides simultaneously. The walls of the car (on which both cameras’ IR arrays are projected) become spotty with both cameras on. However, because each camera projects onto its own respective side of the rod, the rod images are not degraded.

The primary benefit offered by this configuration is a constant view of the rod in all scenarios. If there are no visual obstructions, a coordinate will be received from both cameras. If there is an obstruction from either side, the other camera will still have a view of the rod up until the end-effector moves forward into view.

As is discussed further in Section 8.2.2, detections while the platform is stationary are the most reliable, largely due to the lack of precision in moving on unknown terrain. Maintaining a continuous view of the rod means that there is no “drop-out” zone, where the mobile platform is moving to a previously known position without current information.

Another configuration that was considered uses cameras facing away from each other, but has an additional camera mounted to the end-effector. This additional camera would either be mounted above the contact plate, looking down, or below the contact plate, looking up. Neither of these

mounting positions is feasible, however, because there are some cars with features that would collide with a camera at either position.

5.2 Detection Camera Position

The potential for detection of the bleed rod depends entirely upon the view of the rod that the camera sees. The easily distinguishable side view must be directly in view, and that image must be sufficiently clear. Less than ideal images can potentially still be used for detection, but, once degraded past a certain point, proper detection becomes unlikely. A useful workspace can be defined for camera position, relative to the bleed rod, such that detection is possible anywhere within the workspace. Additionally, a metric for detection potential has been defined to aid in optimizing camera position.

The two parameters that affect the image requirements are viewing angle and viewing depth (see Figure 5.2). These parameters affect the detection potential equally, as either a poor viewing angle or a poor viewing depth can make detection difficult.

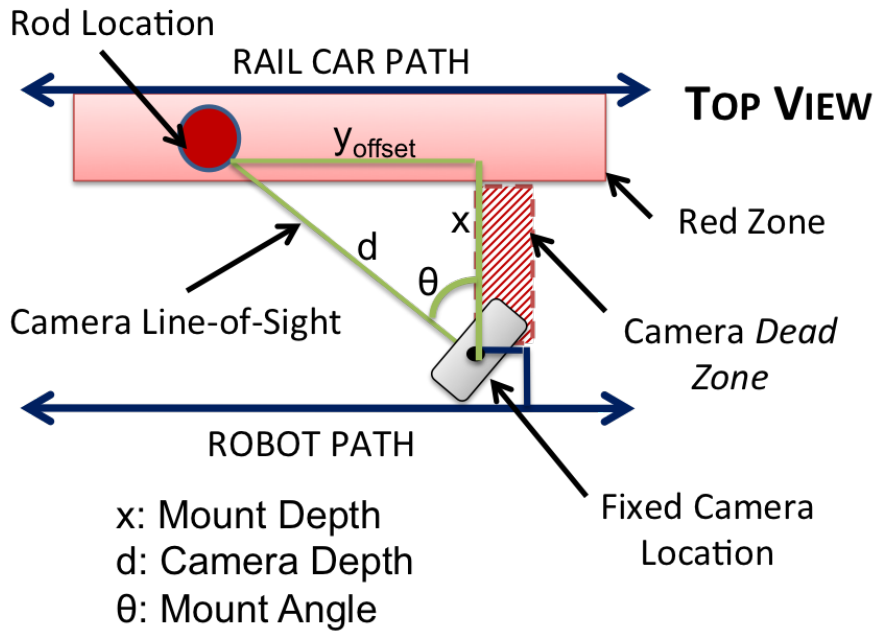


Figure 5.2: Top view schematic of key parameters that affect detection potential

5.2.1 Degrees of Freedom

The camera position has only two degrees of freedom: camera depth (x), and distance along the rail (y); all other positions and orientations are fixed, as of now. The mounting angle of the camera does not change the detection performance for a particular camera position, but rather it affects which viewing angles lie within the camera's field of view (FoV).

5.2.1.1 Camera Depth

The “camera depth” is the distance from the camera to the tip of the rod, perpendicular to the length of the rail. The shorter this distance, the more distance the camera can move along the rail while still maintaining a good viewing angle. However, keeping the camera too close to the side of the cars may result in a collision, if anything is protruding out from a car further than expected.

The addition of a pre-emptive set of cameras to a wayside camera suite could greatly improve the efficiency of the whole break release task. The wayside cameras currently are $\sim 4'$ from the sides of most railcars. This depth is likely too great to allow for successful detection; the process that led to this conclusion will be explained in Section 5.2.5.

5.2.1.2 Distance Along Rail

The “distance along rail” is the distance from the camera to the tip of the rod, parallel to the length of the rail. Ideally the rod should be visible both close up and further away. Close-up vision is needed when determining a location to send the robot’s end-effector. Longer range visibility means the mobile platform will become aware of the rod’s presence sooner, allowing more distance to be travelled when approaching for its final alignment.

In the case of a wayside camera, the cars will be traveling approximately 40 mph during the attempted detection. In order to get as many images of

the rod as possible at this high speed, the range of distances along the rail at which the rod can be detected should be maximized.

5.2.1.3 Mounting Angle

The horizontal viewing angle for the ASUS Xtion Pro camera is 58° . The viewing angle must then be within $\pm 29^\circ$ of the mounting angle in order for the rod to be in view of the camera at all. If the camera is to be mounted statically (as opposed to adding a servo motor to the mount, or something similar), the mounting angle should be chosen to encompass the greatest range of good camera positions possible.

5.2.2 Simulation Layout

For each camera depth, the full range of distances along the rail (as specified by the camera's data sheet) is tested. This simulates the physical process either of a mobile platform traversing a railcar at a constant depth from its wall while searching for the rod, or of a railcar passing a static camera at a wayside sensing station. The depth that has the greatest range of detectable "distances along the rail" should be used, if no other constraints are in place. Note that everything here is done for a left-facing camera, but nothing changes for a right-facing camera except the definitions of positive distance along the rail.

5.2.3 Viewing Angle

The viewing angle is defined (as shown in Figure 5.2) as 0° when looking parallel to the rails, and 90° when looking perpendicular to the rails. A 0° viewing angle, then, provides a full view of the side of the bleed rod. Following a small angle assumption, there is very little change in the view up to about 12° . As the viewing angle approaches 90° , the projection of the rod's side view onto the camera lens becomes more flattened, until it is gone at 90° .

5.2.3.1 Performance Metric

The metric used to describe the detection potential based on the viewing angle is the **ratio of projected rod length to actual rod length**. This value is 1 at 0° (the full length of the rod is visible), and is 0 at 90° . It is taken as a ratio, instead of just the projected length, so that the value is normalized in $[0,1]$, regardless of the actual rod length; normalizing the metric allows for comparison across data sets. The projection ratio is calculated as shown below:

$$\theta_{eff} = \tan^{-1}\left(\frac{x}{y_{offset}}\right) \quad (5.7)$$

$$L_{proj} = L \cos(\theta_{eff}) \quad (5.8)$$

$$projection\ ratio = \frac{L_{proj}}{L} \quad (5.9)$$

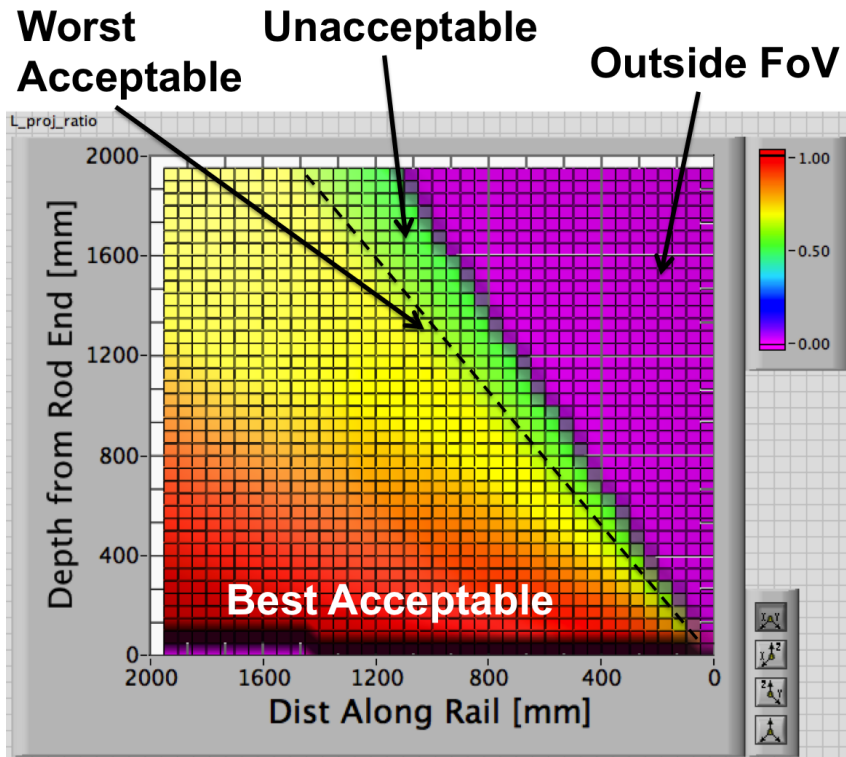


Figure 5.3: Ratio of visible image size to actual size; the camera is at the origin, and locations on the graph represent rod locations

5.2.4 Viewing Distance

The viewing distance is defined as the length of a straight line from the camera lens to the tip of the bleed rod (the rod and camera are assumed to be the same height from the ground). The further from the camera the rod lies, the smaller the rod appears in the image (smaller portion of the FoV is occupied by the rod). When the rod gets small enough, the inherent noise/uncertainty of the camera starts to affect the image of the rod, causing gaps and/or full dropouts in the rod.

5.2.4.1 Performance Metric

Without information from the manufacturer on the camera's interpolation algorithm, the following is an educated guess as to the exact mechanism causing dropouts at great viewing depths. As a finite number of IR points are emitted and tracked, there is likely a linear interpolation of depths between points in the IR array. Each point is assumed to properly represent the region closest to that point. If a point senses no object, a small circle around it will then be assumed to contain no object, within the camera's sensing range. If two of these regions happen to overlap ("Nearly Overlapping Empty Regions" in Figure 5.4), then the camera assumes there is no object within that entire space.

Extrapolated holes like this are what cause the fluctuation in what should be a straight line (like the top and bottom edges of the rod). When holes from the top and bottom edges of the rod overlap, the resulting image shows a gap in the rod. There also appears to be a threshold set, such that when an area becomes too narrow (sensed depth may play a role in a variable threshold) the presence of an object is assumed to be too unreliable, so it completely drops out.

The noise in the camera's images appears to be constant with changing depths. The only variation then is the size of the rod's image in the frame. The relevant metric is the ratio of the rod width to the camera noise (w/ϵ), which is effectively a standard **signal-to-noise ratio**. The sketch in Figure 5.4 depicts a decreasing signal-to-noise ratio (S/N Ratio). This metric is not

inherently normalized, but does decrease to a limit of 0. Below a S/N Ratio of 1, there will be definite flaws in the image. To be safe, the minimum allowable ratio is set to 3.

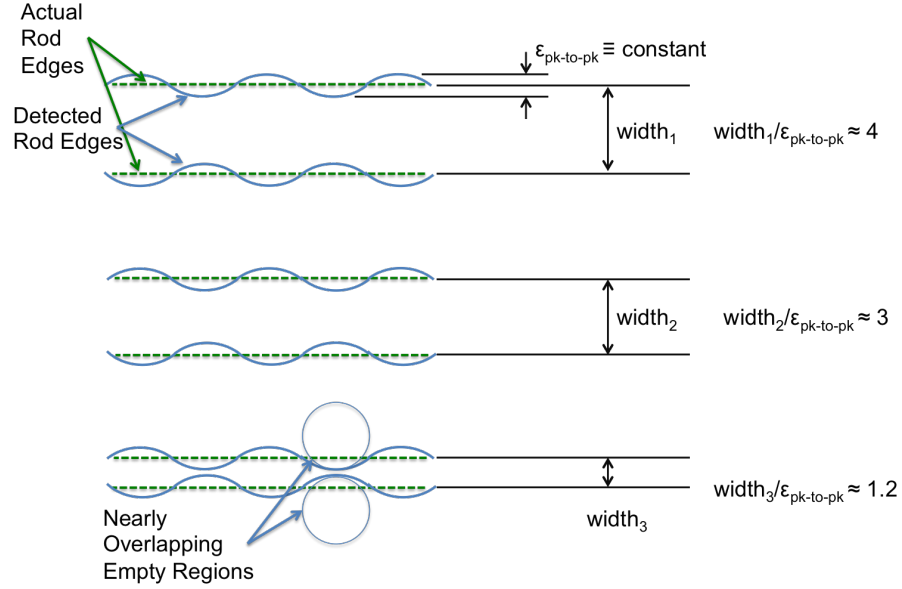


Figure 5.4: Decreasing image signal-to-noise ratio sketch

The S/N Ratio is calculated according to the following:

$$d = \sqrt{x^2 + y_{offset}^2} \quad (5.10)$$

$$w_{eff} = \frac{w}{2d \tan(FOV_v/2)} \quad (5.11)$$

$$w_{pixels} = w_{eff} * YRES \quad (5.12)$$

$$S/N \text{ Ratio} = \frac{w_{pixels}}{error}, \quad (5.13)$$

where x is the camera depth, y_{offset} is the distance along the rail from camera to rod tip, FOV_v is the vertical viewing angle of the camera (58° for the ASUS

Xtion), $YRES$ is the vertical resolution of the camera (480 pixels for the ASUS Xtion), and $error$ is the peak-to-peak noise amplitude of the straight edge for the camera (set to 3 pixels for this computation).

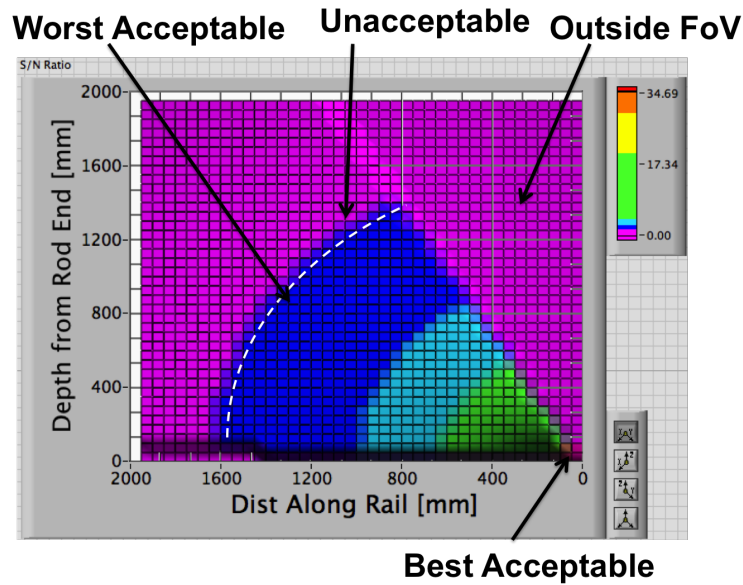


Figure 5.5: Ratio of rod width in the image (in pixels) to edge noise; the camera is at the origin, and locations on the graph represent rod locations

5.2.5 Detection Potential Metric

A single metric is necessary to evaluate the overall “goodness” of a camera position with respect to the rod. There are many common methods that could be used to combine relevant information, some producing a metric with physical meaning, and some producing a metric only for comparison.

5.2.5.1 Fusion Methods

The following methods are a few of many that could be used to “fuse” the two performance metrics:

1. Bayesian condition probabilities; this would require test data to determine the probabilities/rates of detection at all combinations of values
2. Linear weighting scheme; this requires normalizing both metrics to the same scale, and assigning “importance” values to determine the weight given to each metric
3. Non-linear weighting scheme; this also requires normalizing both metrics to the same scale, but “importance” values would be set based on some predetermined function rather than being constant
 - This scheme would be useful if the relative importance of the metrics varied greatly across the workspace, and those importance variations could be accurately modeled with a function. The result would be similar to that of a Bayesian approach, but based on different prior information.

All three of these methods result in metrics that have no physical meaning, therefore should only be used for comparison in optimization. For simplicity’s sake, a linear weighting scheme has been used here, with equal

weight placed on each metric. This should be sufficient, possibly with different weights after more testing is done, but either of the other schemes could provide more accurate information across the workspace; it should be noted that extensive testing is required to successfully implement either of the others, and so should only be done if the additional information is deemed necessary for successful operation.

5.2.5.2 Normalization

In order to add values from different performance metrics with desired weights, they must first be normalized to the same scale so the desired weights are meaningful. Also, 0 and 1 (assuming a $[0,1]$ scale) must have a consistent meaning across metrics (e.g. 0 is always bad, 1 is always good); these requirements for metric/performance map fusion are detailed by Ashok, and Tesar in [2].

Viewing Angle The viewing angle metric (projected length ratio) is inherently normalized to $[0,1]$. A value of 0, meaning the view of the rod projected onto the camera lens has disappeared, is bad, where a 1, a full view of the rod, is good.

Signal-to-Noise Ratio The S/N Ratio at 0, meaning there is effectively no “signal” or rod width, is bad. This matches the meaning of the viewing angle metric.

To normalize the S/N Ratio, two values were set: the maximum ratio at which the vision is known to fail (1.5), and the minimum value for ideal performance without the possibility of degradation (4). Values between 1.5 (S/N_{min}) and 4 (S/N_{max}) are decreased by 1.5 and normalized by dividing them by 2.5 (4-1.5). Anything below 1.5 is set to 0, because it is known to cause failure therefore is unacceptable, and anything above 4 is set to 1, because it is known to be fully reliable. This method provides a continuous range of values, where anything between 0 and 1 is acceptable but lies on a gradient that should be climbed to best avoid failure. The normalized ratio is calculated as

$$S/N_{norm} = \frac{S/N - S/N_{min}}{S/N_{max} - S/N_{min}}. \quad (5.14)$$

The final performance metric, using the linear weighting scheme, is

$$performance = 0.5 S/N_{norm} + 0.5 projection_ratio.$$

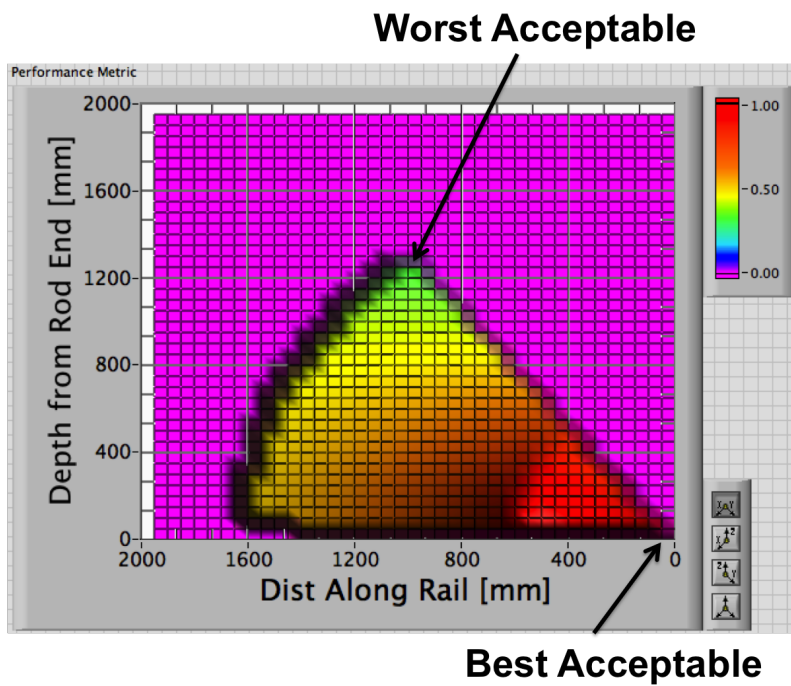


Figure 5.6: Potential for rod detection (scaled 0 - 1); camera is at the origin, and locations on the graph represent rod locations; non-zero values represent feasible rod locations

Chapter 6

Primary Operational Steps

The steps of operation will be the same for every car. These steps are based on assumptions that the conclusions from the San Antonio Yard visit (relatively small set of cars) hold true for the entire population (within reason) of rail cars. The key assumptions are the following:

1. Every *non-broken* bleed rod can be operated by being pushed
2. Every bleed rod-end can be seen clearly from at least one of the two sides (left or right)
3. All bleed rod-ends lie within one of the two measured position distributions (shown in Figure 2.5)
4. All bleed rod-ends lie within the measured distribution of the rail (horizontal distance from closest rail)

6.1 Bleed Rod Detection

A pair of cameras will be mounted at the expected rod height. Each pair of cameras will have one camera pointing forward (along the length of the cars, in the direction the robot is moving) and one camera point along the

length of the cars in the opposite direction (Figure 6.1). From the bleed rods seen at the San Antonio Yard visit (Chapter 2), every bleed rod end is visible from one side or the other (only one of the sides ever has an impeded view, shown in Figure 6.2).

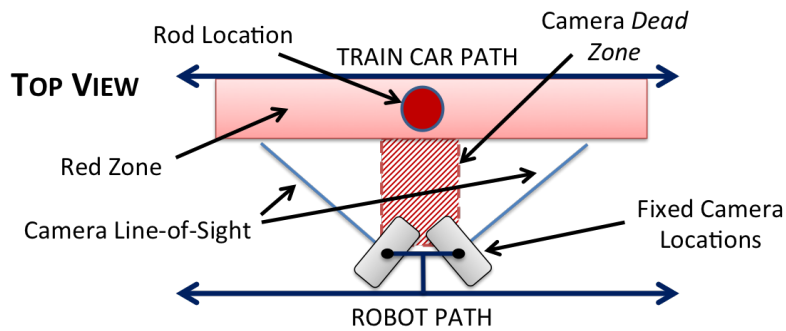


Figure 6.1: Top view of dual camera configuration; cameras facing away from each other

6.1.1 Wayside Sensing

Existing sensor suites already are set up adjacent to rails. Pairs of depth cameras could be added to these suites to serve as a rod locator. Because positions and speed are already known at the time the train passes the wayside sensors, the uncertainty associated with ground vehicle position tracking, discussed in the following sections, is not an issue.

A pair of cameras should be mounted at each of the two expected rod heights. The closer the cameras can be to the rod, the better; expected performance for different relative locations is described in Section 5.2.

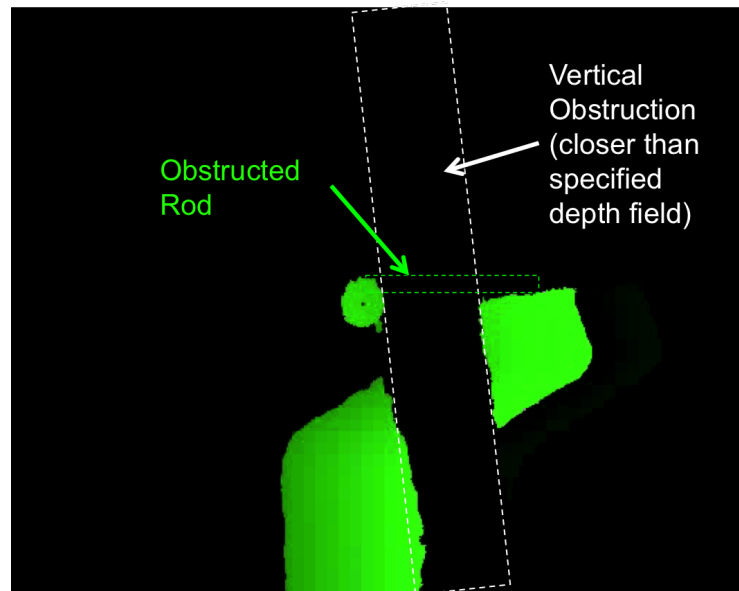


Figure 6.2: ASUS Xtion camera, rod view impeded from right side by a vertical bar

As rods are detected while the train is passing, the location of each rod will be referenced to something on the length of cars, such that a mobile platform can later locate and use that point. Position error, however small it may be, will increase with distance away from the reference point. Each rod's position should be stored both globally and relative to the closest detected rod.

Not all yards have Wayside Sensor Suites installed, and there is no guarantee that the Wayside Sensors will be properly operating at all times. Wayside-mounted depth cameras would provide invaluable information that could improve both reliability and speed of the operation, but the ABRS should not be dependent upon them.

6.1.1.1 Speed Issues

According to UP, trains pass the wayside cameras at an average speed of 40mph. The ASUS Xtion cameras currently used have a frame rate of 30 fps. At these speeds, there will be one frame taken every 1.96 feet along the cars. This should be enough distance to see the rod in at least one frame, but consecutive frames greatly aid detection confidence.

Possibilities for a greater frame rate include using different cameras and using overlapped, out-of-phase cameras. For example, having 3 ASUS Xtion cameras on top of each other, each offset by 11ms, would provide an effective frame rate of 90 fps; work would need to be done to verify the viability of this option.

6.1.1.2 Distance Issues

The Wayside cameras are approximately 4' from the sides of most cars. Getting the cameras closer than this is imperative for reliability; see Section 5.2 for more detail on camera position's effect on detection potential.

6.1.2 Searching Scheme Without Wayside Sensing

6.1.2.1 Blind Searching (No Start-of-Car Reference)

If the robot does not know where it is relative to the front end of the car, then it has no way of knowing whether to be looking for the high region or the low region (shown in Figure 6.3). In this case, two pairs of cameras will be mounted, one pair centered at each of the two possible bleed rod heights,

and both will be checked at all times. This will put a significantly greater processing burden on the on-board computer than if only one pair of cameras were being monitored.

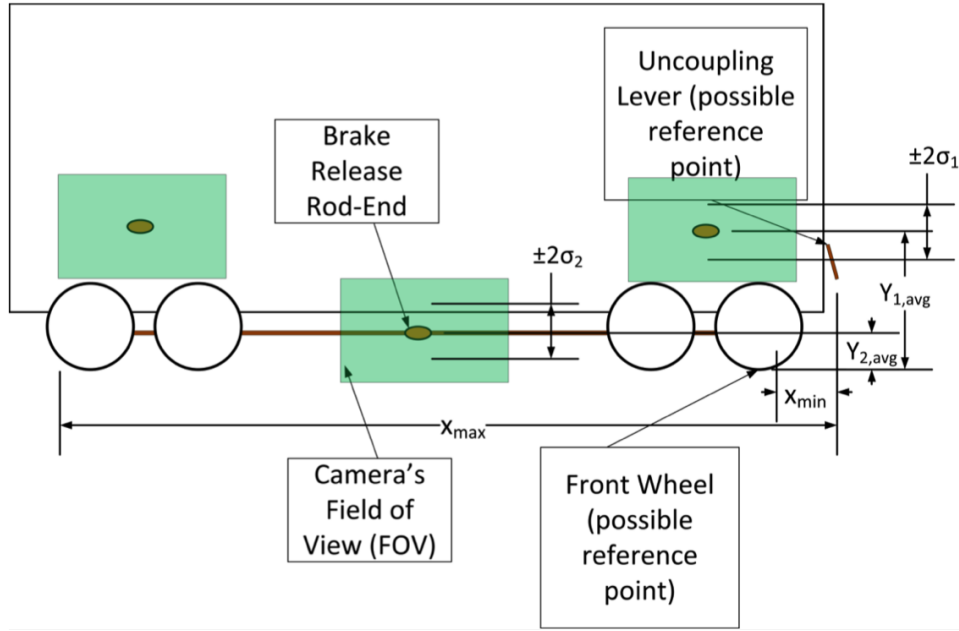


Figure 6.3: Rod location schematic

6.1.2.2 Optional Reference Point Detection

The two primary search regions are known on each car with respect to its leading edge. Using each car's leading edge as a reference point would provide for a more intelligent search scheme. However, it is not completely necessary and would just serve to make the operation more deterministic, and to decrease the processing load on the on-board computer. If AEI tags were fully reliable, then the robot could expect to know the number of axles per car

and the car length, both of which would provide knowledge about a start-of-car reference point. However, it has been widely said (Evan Freilich, Dan Rubin, various yard operators) that the AEI tags neither reliably contain the correct information nor reliably work properly.

Wheel Detection A wheel is always at the leading edge of a railcar. Minimal work was done during the Automated Pin-Pulling System (APPS) development to find wheels using an RGB camera and Hough Circle transforms. While it was not successful, it did show promise and would be a primary avenue to pursue to use as a reference point.

Uncoupling Lever Detection Every car has an Uncoupling Lever at its leading edge as well. In the APPS, reliable lever detection was demonstrated in a controlled environment. It is believed that the detection could be made far more robust by simply collecting a much larger set of training images for the Artificial Neural Network. This would also be a very reliable, consistent start-of-car reference point.

6.1.3 Searching Scheme With Wayside Sensing

A mobile robot would start by locating the reference point used by the wayside sensors as the center of the coordinate system. Uncertainty and error in position increase as a wheeled vehicle moves; the specific sources of error in this case are explained in the following section. The mobile robot would then

move to the expected position of the rod, where its on-board cameras should start locally detecting the rod, enabling fine position adjustment.

Having the rod coordinates referenced to each other allows the robot to use every rod it detects as a new reference point, thus resetting the error it has accrued while moving to the rod.

The three primary benefits provided by wayside sensing are added robot safety, improved reliability, and improved operation speed.

A mobile robot, without wayside sensing information, has to search for rods while traversing the entire car, which requires moving very close to the cars. While there should not be objects protruding to the distance of the cameras, this is not necessarily guaranteed. A mobile robot traveling to a new point, without searching for a rod, can stay further away from the cars, completely out of harm's way.

While a mobile robot is looking for rods, its maximum speed is limited by the camera's frame rate; if that speed is exceeded, a rod may be skipped over in between frames. If the robot can skip a large region and move straight to a known rod location, the ground speed is no longer constrained by the camera's performance, but by a mobile platform's capabilities.

6.2 Bleed Rod Alignment

6.2.1 Coordinate Determination

Once the rod-end has been recognized, its real-world coordinates with respect to the robot will be saved. As the coordinates are being saved, an averaging scheme is used to decide on one point to use for alignment.

6.2.2 Ideal Alignment

6.2.2.1 Robot-Favored Alignment

The ideal position for alignment will likely have the rod-end centered on the robot base (waist axis). This would place the rod-end in the center of the robot's workspace. However, as there is minimal risk involved in this task, it may be beneficial to forgo some dexterity in favor of continued visual information.

6.2.2.2 Camera-Favored Alignment

While a depth camera facing perpendicular to the length of the car cannot pick out a bleed rod from an entire scene, it should be able to pick it out of a very small region that is already loosely known. An additional depth camera could be aligned to the detected real-world coordinate and used to more precisely determine the rod location after the robot has become stationary.

6.3 Approaching the Rod

The robot will position the end-effector to directly face the rod end. Then, the end-effector will slowly move toward the rod end, and stop once contact has been made. This will be the starting positing for rod actuation.

6.3.1 F/T Sensor

The Force/Torque sensor on the robot end-effector will sense a force as soon as contact is made. In the APPS, latency from F/T sensing to stopped robot motion was verified to be within the limits of safety of the robot. This latency is sufficiently small to be relied upon to effectively place the end-effector right at the tip of the rod end.

6.4 Bleed Rod Actuation

Based on the loosely constrained tests performed during the San Antonio Yard visit (see Section 2.4.1), a linear path, without any fine control scheme, will successfully operate all bleed rods. Because some rods have notches worn into them (Figure 6.4), a slight upward angle is necessary to prevent failures in those cases. A slight upward angle was tested on non-flawed rods as well (approximately the same angle), and the operation appeared to perform just as well as with a fully horizontal path. However, if the upward angle does not operate the notched rod successfully (has occurred in the lab, not in the yard, as a result of the lab's notch shape, which may or may not occur

in reality), a second operating scheme is required; this alternate operation is described in “Case 2” of Section 10.2.3.

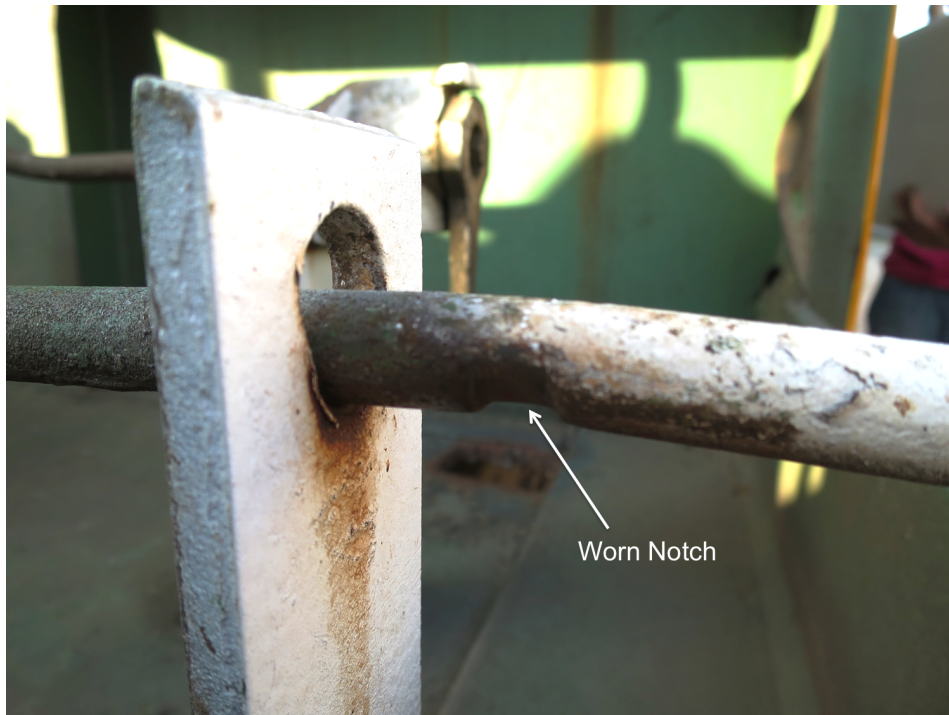


Figure 6.4: Notched Rod

6.5 Valve-Release Confirmation

The only considered method that can viably confirm brake release is Auditory Recognition. All of the other methods can indicate a failure to release the brake, but not a success.

When the valve is opened, there is a clearly defined, audible blow-off sound. This sound is louder than other ambient yard sounds, and is within

close enough proximity to the robot to be reliably heard. The audio processing sequence is the following:

1. Bandpass filter between 5500 Hz and 7000 Hz (3rd order Butterworth filter)
2. Signal amplitude threshold
3. Exceeded amplitude time duration threshold

Chapter 7

Bleed Rod Detection Algorithm

The brake-release task first requires detection of the bleed rod. Significant development went into the Neural-Network-based object detection algorithm used in the APPS (Chapter 6), and, as the detection tasks are similar, this algorithm was modified and used as a starting point for the Bleed Rod detection algorithm.

The primary function of the detection algorithm uses an artificial neural network (ANN) to scan an image and find a learned object. The entire algorithm and its background are explained in detail in Chapter 6 (algorithm-specific text in Section 6.4) of the APPS report [3] .

Definitions of the following key terms as used in this context are necessary to understand the explanation of algorithm modifications:

- Depth Range: defined by the maximum and minimum depths from the camera that are not filtered out; this is where the rod is expected to provide a clear image
- Frame: one image from the camera's video feed

- Window: one region of interest (ROI) from the frame—many windows from each frame are checked for a rod end; e.g. if the frame is 160x120 pixels, a window may be 20x30 pixels
- Hit count: number of windows that the ANN thinks contain a rod within which a certain pixel was contained (see explanation of "hit matrix" in [3] , Section 6.4.1)
- ANN output: the ANN's output is mapped to range from -1 to 1; images that have been trained out using "False" training images will be mapped to -1, and actual rod images should have an ANN output close to 1
- Window step-size: the number of pixels between a corresponding side of successive windows in a particular frame; increasing this effectively decreases frame resolution, but improves processing speed
- Pre-processing: processing that is done on each window before it is run through the ANN; if the pre-processing checks are not met, the ANN won't process the particular window, which greatly decreases false-positive detections and improves speed

7.1 Getting Started with Bleed Rod Detection

The first task, after a rough camera position/orientation was chosen, was to gather training images and train a new ANN. This process, as it was done for the uncoupling lever in the APPS, is detailed in [3], Section 6.3.3 - Artificial Neural Network (ANN).

After adequate bleed rod training images were gathered and used for training, the detection was tested on recorded videos from the San Antonio Yard Visit to establish a performance baseline. From here, additional training images were taken, both True images (bleed rod images) and False images (images of objects frequently giving false-positive detections), and the ANN was retrained to improve performance. Three “True” rod training images are shown in Figure 7.1.

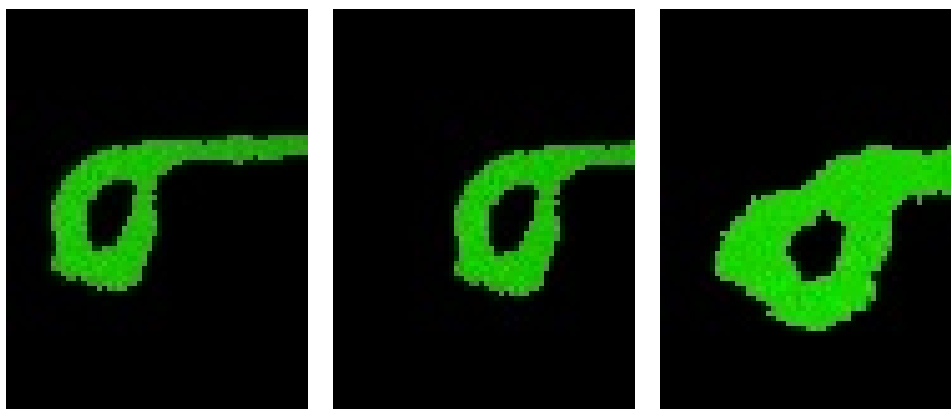


Figure 7.1: “True” rod training images

7.2 Comparing the Vision Problems

Both during and after the ANN retraining process, algorithm parameters were tuned, and pre-processing heuristics were developed.

7.2.1 Parameter Modifications

7.2.1.1 Window Dimensions

The window size for the bleed rod is set to 15x20 pixels. The smaller window size (the APPS uses 20x30 pixel windows) allows the ANN to process each window faster. The size was set with the rod at a depth of 650 mm. The window could, technically, be even smaller, but these dimensions allow space around the rod itself to aid in detection as well (given knowledge of what should/should not be in close proximity).

7.2.1.2 ANN Output Threshold

The ANN Output Threshold should be set high enough to eliminate most windows not containing a rod, but never reject a significant number of “True” windows. The ideal threshold can be set higher the more distinct the separation is between rods and non-rods. The ANN Output Threshold for the bleed rod is set to 0.8.

7.2.1.3 “Hit Count” Threshold

The Hit Count Threshold specifies the minimum number of overlapped windows within a frame that will indicate a “True” recognition; anything below the threshold will be disregarded. The Hit Count Threshold for the bleed red is set to 30. Most rods within the frame have at least 20 windows register as positive detections (20 was the lowest hit threshold that did not induce false-positive detections), which is higher than that used for the pin-pulling lever. This is partially due to the tight heuristic constraints set for the lever and partially due to the small clearance around the lever, as compared to the bleed rod.

Looser heuristic constraints can be set for the bleed rod than for the uncoupling lever without inducing false-positives because the bleed rod is much more unique.

7.2.1.4 Window Step-Sizes

The step size has been set to 1 pixel, to ensure no missed detections (defined as 4 in the code, which corresponds to 1 pixel in the 4-times down-scaled frame). Speed is not as crucial here as in the pin-pulling system, but the algorithm still runs fast within this small step size. Because almost half of each frame in this view is black, far fewer windows need to be processed by the ANN than for the APPS.



Figure 7.2: Window step-size schematic (for clarity, the labeled step is much larger than the actual step size, which gives the spacing between the “true positive detections”)

7.2.2 Heuristic Parameter Modifications

7.2.2.1 Populated Pixel Ratio

The ratio of null pixels (empty space) to populated pixels (space occupied by an object) is consistent between windows containing a bleed rod, but depends on distance to and viewing angle of the object. For the camera placement chosen for here for testing, the ratio range is set to $0.833 > null_pixel_ratio > 0.3$, which corresponds to 90-250 null pixels in a 15x20 window.

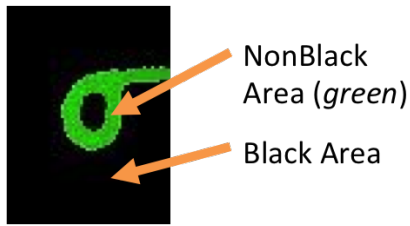


Figure 7.3: Horizontal centroid heuristic

7.2.2.2 Centroid Location

The centroid location within True windows is also consistent, but can take a wider vertical range, as the bleed rod's precise height within the window isn't critical (due to the training images used). However, the trained ANN, in this case, rejects images whose vertical centroids differ significantly from those of True images. As such, setting limits only on the horizontal centroid provided a performance improvement.

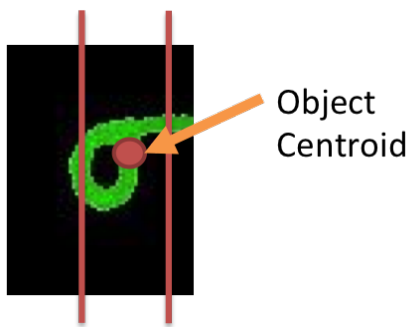


Figure 7.4: Horizontal Centroid heuristic

7.2.3 Additions (Heuristics)

While artificial neural networks learn and generalize very well, they do not as readily learn what humans would establish as hard-fast rules. Using heuristic checks in pre-processing both adds an additional layer of intelligence/oversight and improves program speed by reducing the computational load.

7.2.3.1 Image Height

Instead of using a vertical centroid heuristic, setting an allowable range on the overall height of the rod within the window did improve performance. This range is taken as the difference between the highest non-null row and the lowest non-null row. The allowable range that performs optimally is $18 > \text{row height} > 5$.

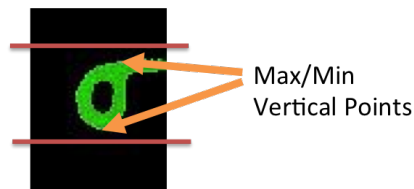


Figure 7.5: Image height heuristic schematic

7.2.3.2 Window Edge Criteria

The bleed rod-end enters every window from the right side and should terminate before reaching the left side of a window. Additionally, no part of the rod should ever be touching either the top or bottom of the window. These conditions are in keeping with the position of the rod-end in training images.

If any of these three ideally untouched sides contains a pixel that is not empty, the ANN is skipped, and its output is assumed to be -1 (corresponding to the state “confidently not the object of interest”).

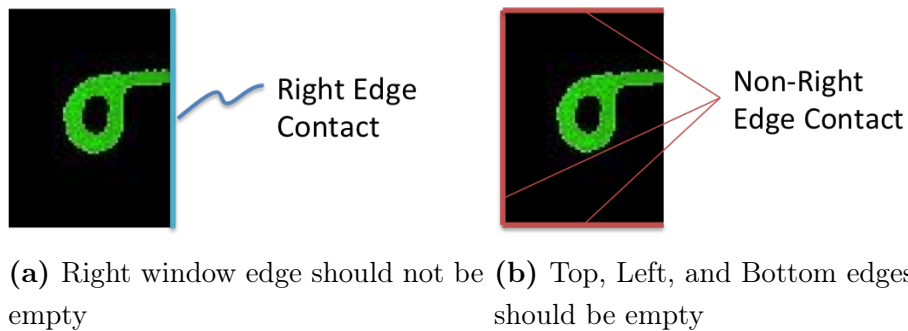


Figure 7.6: Window edge heuristic criteria

7.2.3.3 Rod vs. Loop Width

The loop bent into the end of each rod is just over 3 times the diameter of the stock rod ($\sim 1.6:0.5$); the acceptable ratio range that has been seen in testing is $4.5 > ratio > 1$. This ratio is mostly consistent in images across scaling due to depth changes (at greater distances, thin features become noisier). For the 15-pixel-wide window, the 4 pixel columns just right of center are checked for maximum and minimum values (positions of non-zero pixels),

as are the two right-most columns of the window (see Figure 7.7). These averages account for most of the high-frequency noise in the image.

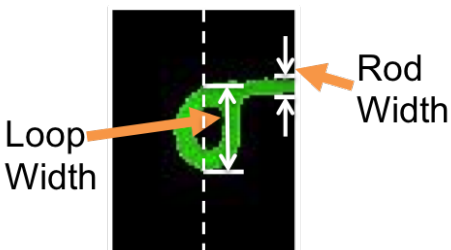


Figure 7.7: Width comparison schematic for pre-ANN heuristic filtering

7.3 Results

The testing stages listed in Table 7.1 represent a progression of algorithm development. The state of the algorithm at each stage of testing was the following:

1. 1053 “True” training images were gathered, used for training, then 1343 “False” images were gathered based on observed false-positives. Stage 1 results include all of these training images.
2. A bug was discovered, and fixed, in the parsing of the training images; only a portion of the image was being used. Stage 2 results are after the bug was fixed, with no other changes.
3. The program used to gather training images (point-and-click application) had a preview box added to it, so the user can see the exact framing of the training image about to be saved. This enables “False” images to

be saved, while ensuring that they aren't images that will already be excluded through heuristic checks. An additional ~ 1300 "False" training images were added to the training set. Stage 3 results use the expanded training set.

4. For Stage 4, the boundary check was increased from 1 to 2 rows/columns.

Table 7.1: ANN Performance Results through 4 stages of development

	Stage 1		Stage 2		Stage 3		Stage 4	
Video #	T	F	T	F	T	F	T	F
2	255	0	3288	0	2686	0	900	0
10	375	145	5425	710	3838	230	1192	13
12	129	73	2097	294	1915	97	781	8
14	207	35	3708	219	2877	80	930	6
15	175	59	4241	200	3794	53	998	3
16	388	93	4452	232	4123	46	1256	2
18	20	2	197	31	126	21	35	0
19	235	0	1411	16	1298	4	417	0
24	67	76	2520	17	2217	3	668	0
25	71	73	1686	305	1361	122	461	5
Totals:	1,922	556	29,025	2,024	24,235	656	7,638	37
Success Rate:	77.56%		93.48%		97.36%		99.52%	

Chapter 8

Robot Alignment/Positioning

Once the bleed rod has been found, its position needs to be determined so the robotic arm can align to it. The specific alignment position (of mobile platform and of robot pose) dictate the reliability of the operation and size of the operation workspace within which the operation can take place.

8.1 Robot Pose and Branching Avoidance

The ease with which a robotic arm's end-effector can change positions largely depends upon its current joint positions. A robotic arm is most reliable far away from geometric singularities (an effective reduction of degrees of freedom of motion), where branching is also likely to occur without control oversight. Positioning the rod directly in front of the center of the robot base makes singularities far more likely; these singularities should be avoided.

8.2 Position Determination

Globally, the rod-end remains fixed as the robot moves; the change in the camera's output coordinate should correspond, oppositely, to the change in the robot's position as it moves. While it is moving into place, it gets many

updated “rod coordinates.” While these are all reasonably accurate with a robot on a precision track, a mobile robot with imprecise movement and/or sensing will have global-coordinate estimations that differ by significantly more.

A simple average will likely work to combine these uncertain data points, but a Kalman filter could account for the uncertainty of each point (where the rod is in the frame, how far away it is, encoder precision at different speeds, etc.). Additionally, expected errors due to ground inconsistencies, wheel slip, or vibrations could be modeled in the robot controller for testing while in a controlled environment with precise position baselines.

This may be a significant issue if high-speed rod locating is desired. However, if the mobile robot can be slowed, or even stopped, once the rod is in view, then this is a far less of a concern.

8.2.1 Fully Static Robot

When the robot platform (or mobile platform) is stationary, the problem is greatly simplified. In this case, the only factors in play are detection accuracy and alignment accuracy.

8.2.1.1 Detection Accuracy

The accuracy of the depth sensor (ASUS Xtion, in this case) varies slightly depending on position in the frame (due to optical distortion). During the development of the Automated Pin-Pulling System, tests were performed

to verify the repeatability of alignment, given a real-time coordinate from the detection algorithm. In these tests, no variations in alignment were apparent upon visual inspection. Because allowance for alignment error was built into the end-effector design, this repeated visual confirmation was sufficient, and no precise measurements were taken. The RRG believes the ASUS Xtion camera can be assumed to be accurate within $\pm 3\text{mm}$.

8.2.1.2 Alignment Accuracy

The accuracy of the alignment to the detected coordinate is dependent solely upon the potential error in the robotic arm's motion, given that the base is sufficiently grounded during the alignment. The accuracy will vary between robotic arms, but any error should be well within the allowable contact offset, described in Section 9.2.1.

In addition to these certainty issues that must be accounted for, additional features necessary for a switchyard-specific mobile platform may have side-effects on position accuracy and sensing errors.

8.2.2 Fully Moving Robot

If the robot is in motion, additional factors must be accounted for. Ground inconsistencies will cause a mobile platform to pitch and roll, possibly so much so that the effects on the cameras and/or other sensors will need to be accounted for. A loose ground surface will also lead to wheel slip for a mobile platform, which would increase uncertainty associated with dead-reckoning.

With motion between camera frames, latency associated with a filter must be accounted for.

8.2.2.1 Ground Inconsistencies

In the lab test-bed, the cameras are always assumed to be horizontal (not angled up or down). However, a mobile platform may end up on uneven terrain, causing it to pitch or roll. Without on-board sensing to account for this, these variations will result in a wider spread of determined “rod coordinates.”

8.2.2.2 Position Error (Slip)

When determining a global coordinate for the rod, large variations in detected lateral position of the rod can be accounted for using the change in robot position (robot is assumed to be moving parallel to the rails). However, a wheeled platform will likely slip considerably, so its wheel rotations can not be assumed to directly correspond to lateral motion. This slip must be sensed and accounted for as well as possible to retain accurate position estimation.

However, minimizing the wheel slip as much as possible is always the best option. This requires real-time slip detection and torque-vectoring to prevent excess slip.

8.2.2.3 Filtering Latency

The object detection algorithm uses a moving average to filter the detected coordinates. There are 33ms between each frame; for only a 5-value moving average, that is a 165ms delay between the oldest saved coordinate and the time the transmission of the filtered coordinate is initiated. Additionally, if the mobile robot has been moving the whole time, those saved values are spanning a range that may or may not follow a normal distribution.

8.3 Rod-Centered Reference Frame

An alternative way to view this problem, once the rod is in view, is to use a rod-centered reference frame, and the goal is to track and control the position of the robot within the rod's fixed frame. This now becomes a vehicle localization problem.

With the robot in the lab test-bed on a precision track, it has a precise measure of its location within a grounded reference frame. A mobile platform would not have this benefit, so relating everything back to a hypothetical reference frame introduces unnecessary error into the control.

All of the previously mentioned uncertainties except position error still apply. However, rather than tracking the mobile robot's position as a ground truth (which becomes increasingly more erroneous with additional slipping), the rod is a ground truth that can be re-referenced with each frame.

Chapter 9

Initial Rod Contact

9.1 Contact Confirmation

9.1.1 F/T Sensor Threshold

The F/T Sensor will feel a force as soon as the end-effector contacts the rod. When contact is made, there will be a force region where the rod will begin to deform prior to the valve beginning to open. The threshold will need to be sufficiently large to avoid sensor noise triggering a false-positive contact, and sufficiently small to be detected well before the valve begins to move.

With approximately ± 1 lbf of noise in the ATI-Omega160 F/T sensor, (± 1500 lbf F_z limit), the noise in a ± 100 lbf F/T sensor could be expected to be ± 0.06 lbf. Given this low level of noise the threshold then could be set anywhere between ~ 0.12 lbf and 20 lbf.

9.1.2 Safety Concerns

Because this operation takes place with both the robot and the railcar stationary, and the rod has compliance in the direction of the approach (as it is also the direction of actuation), the safety concerns during this stage are minimal. Issues that would arise here would be due to an unseen obstacle

(this is very unlikely), or making contact with the wrong object, which is very unlikely given a robust detection algorithm.

9.2 Safe Rod-End Placement

9.2.1 Robot Wrist Torque Limits

Any offset between the rod-end contact point on the end-effector and the axis of the last pitch/yaw joint (Joint T) will induce a moment that the robot will have to account for. This moment arm should be minimized to avoid putting undue loads on the robot. However, if the system is precise enough to sufficiently minimize position error without requiring Active Force Control, then the system will be simpler and more robust.

The Motoman MH80 robotic arm, and most commercially available robotic arms, have a pitch/yaw joint on the wrist (Joint B on the MH80). This joint could be the weakest link in resisting an induced moment from a rod alignment offset, depending on arm orientation; this is shown in the schematic in Figure 9.1. It should be noted that the robot OS will account for the unanticipated force reliably, within the hardware's capability.

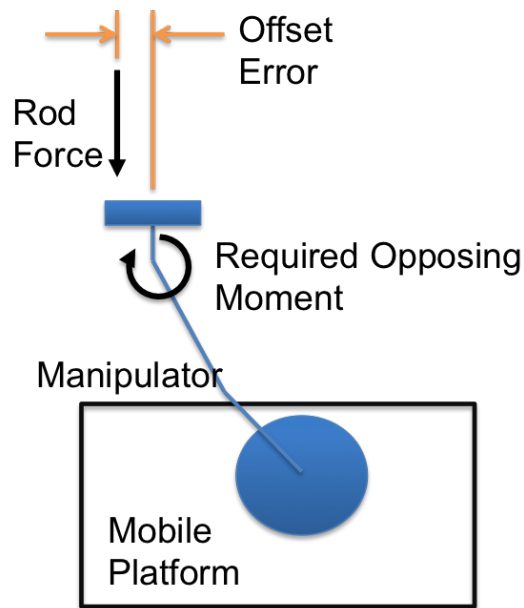


Figure 9.1: Top view schematic of the resisting moment required by an alignment error

9.2.2 Expected Initial Alignment Accuracy

The alignment accuracy is based on robot accuracy and camera accuracy. The accuracy of the Motoman MH80 robot within the workspace is $\pm 0.07\text{mm}$. The Schunk LWA4D arm, discussed in more detail in [8], has an accuracy of $\pm 0.15''$. For this operation, both of these accuracies can be assumed to be effective ground truths, as any other inaccuracies will fall well outside of this limitation. The ASUS Xtion camera's depth sensing has 1mm precision and, in laboratory testing, has been repeatable to $\pm 3\text{mm}$. Again, compared to the inaccuracies of a tire-ground interface, or the play in the rod's connections, the camera's accuracy will not be a limiting factor.

Chapter 10

Rod Actuation

10.1 Actuation Methods Considered

Testing using the lab test-bed has confirmed that a purely linear motion can reliably operate the bleed rod, in the absence of rod abnormalities.

10.1.1 Pneumatic Piston

A pneumatic cylinder was used in the test-bed test to verify linear actuation. A linear actuator of any sort provides a very simple means of performing a linear operation, with just on/off functionality. One significant advantage of a pneumatic (or hydraulic) cylinder is that it is force-based, not position-based. This means that it will not apply more than a specified force, even if it doesn't reach full extension.

Primary disadvantages are weight and the necessity of compressed air. Neither of these are issues in the lab environment, but they will be for a mobile robot.

A mobile robot will have a much smaller arm than the MH80, so its maximum payload will be limited. Taking up a significant portion of the

available payload with a heavy cylinder would severely limit its potential for motion.

Having compressed air available of a mobile platform would require either an on-board compressor or a large air tank. Space on the mobile platform will be limited, so either of these components on-board would not be ideal. Additionally, a compressor would require significant power, which would greatly reduce the running time of the mobile robot.

10.1.2 Robotic Arm Manipulation

Given that a simple, linear motion has proven to be successful, a robotic arm performing the motion would require minimal complexity and/or real-time control, making it very robust and repeatable.

Robot control is position-based, not force-based, so it would require real-time force/torque monitoring using a sensor on the end-effector. Without force monitoring, as the robot is pushing the bleed rod, if the rod were to get stuck, or stop for any reason shorter than the expected travel, then the robot would push increasingly harder until the specified torque limits were exceeded; this could damage the robot; while force/torque monitoring and control is not too complex, it does increase the on-board processing and power requirements of a mobile platform.

An advantage of having a lightweight end-effector that is outfitted with a F/T sensor is the potential for more complex tasks using the same platform.

F/T sensing is needed for active force-control, which is necessary for most complex manipulation.

10.1.3 Conclusion

For a mobile robot solution, performing the rod actuation with the robotic arm doing the manipulation is the better of the two options. This minimizes the required maximum payload of the mobile manipulator, eliminates the need for on-board pneumatic equipment, and, therefore, minimizes the space and power requirements as well.

10.1.3.1 Arm Selection

The requirements of an arm for this task, if it is to be mounted to a mobile platform, are fulfilled by a very narrow range of robots on the market. This is discussed in further detail in [8], Chapter 4.

The maximum payloads of all of the appropriate robotic arms on the market are either right at or just below the requirement. However, maximum payload is specified as the largest mass at the end-effector with which the arm can move at full speed within its entire reachable workspace. The actual limiting factors are the available power and the maximum design loads of the actuators. The power supply for a robot should be matched to its actuator capabilities, so these two parameters should be very similar. The critical measure, then, is current being supplied to the actuator motors. As long as

there is no over-current condition in any of the actuators, the robot is operating safely.

Because the brake-release is a well-defined task that will take place in a small, well-known region of the workspace, there are methods to safely exceed the specified maximum payload within the small range necessary.

Maintenance and applicability to additional tasks are two additional factors that are critical for a long-term hardware selection. In the case of a robotic arm, modular components (actuators and linkages) that allow for reconfiguration provide the greatest possible range of operability, as the same components can be used for greatly varying payload and workspace requirements. This minimal set of necessary components also minimizes up-front cost and down-time required for maintenance.

10.2 Key Factors for Rod Actuation

10.2.1 Travel

Each "Valve Handle" requires the same rotation to initiate release and has the same length tip (to which the rod connects); this means that the end of the rod connected to the valve (see Figure 10.1) always travels the same distance. However, rod configuration and geometry will change the required travel of the rod-end (detected end).



Figure 10.1: Valve handle with rod connection

10.2.1.1 Different Bleed Rod Geometries

Bleed rods are all made custom for the rail car they are to be used on. The configuration of the Valve Handle varies from nearly in-line with the rod guide plate to at most $\sim 2'$ offset from the rod (this measurement was neglected at the San Antonio Yard Visit, but it has been agree upon by those present that this is approximately what was seen; this lack of precision minimally changes any of the following). The required rod-end travel is a minimum when the valve handle and guide plate are aligned and the rod is straight and increases with increasing offset and additional bends in the rod.

Bent rods behave like springs, deforming under load. The greater the moment arm (shown in Figure 10.2a, and as "B" in the schematic in Figure 10.4), the more the rod will deflect under the specified lateral load. Rod-end travel was not measured in the San Antonio yard, but an expected range will be assumed from the "Depth to Nearest Face" measurement that was taken.

The minimum travel is about 0.5", which was measured directly from the valve and assumed zero deflection in the bleed rod. The assumption has

been made that the greatest rod-end travel will not exceed (unless in case of a fault) the maximum Depth to Nearest Face value. This gives a travel range of 0.5" - 7.5".

This range may be used to spec a linear actuator for an end-effector, if that is chosen to be used in the future. However, for use with the robotic arm performing the operation, the range would be used only as a safety check in the software. Moving the end-effector outside of the expected travel range (less than the minimum, or more than the maximum) would indicate a fault of some sort.



(a) Bent rod, with labeled internal bending moment

(b) Straight rod, with no bending moment

Figure 10.2: Rod geometry variations

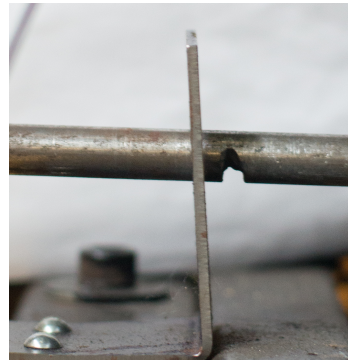
10.2.2 Possible Problems

10.2.2.1 Notched Rod-End

Some rods have a notch worn into them by the guide plate (see Figure 10.3). The notched rod seen in the San Antonio yard has a relatively shallow notch. Two different scenarios can arise while trying to push a notched rod: either the rod lifts enough for the notch to clear and not get stuck (with a slightly inclined force), or the notch gets pressed against the guide plate, and the friction force has to be overcome.



(a) Rod notch in San Antonio Yard



(b) Rod notch in RRG lab

Figure 10.3: Actual and replicated worn rod notches

The notched rod seen in San Antonio never got stuck with the friction force, as long as an upward inclined force was used. The notched rod in the lab has a deeper notch, and about 50% of the time does get stuck.

The full recommended solution to these events is outlined in Section 10.2.3.

10.2.2.2 Insufficient Exposed Rod Length

One of the rail cars tested in the San Antonio Yard had insufficient exposed rod to be actuated one of the two directions (every rod can be either pushed or pulled from both sides of the car). The distance between the back of the looped rod-end and the nearest face (guide plate) was less than the travel required to operate the valve; this caused the rod to bottom-out before the valve was opened. This rod only bottomed-out pushing from one side; it could still either be pulled from the same side, or pushed from the opposite side of the car. Within the scope of this project, this case is assumed to be a mechanical fault, and should be flagged to be fixed, rather than accounted for by the system. If desired, a mobile robot could go to the opposite side of the cut of cars and push the rod from there (however unlikely), or could have a gripper to enable pulling as well.

This fault case could be flagged both by the absence of a confirmed valve-release and possibly by rod travel that is less than the expected minimum.

10.2.2.3 Empty Reservoir Tank

The purpose of opening the service valve is to allow pressurized air (from the Reservoir Tank) to open a second valve, which then stays open and allows the emergency brake piston to retract. The blow-off sound that is heard is the pressurized reservoir air releasing through the valve handle.

If the Reservoir Tank is empty, no sound will be heard, and the second, internal valve will not be opened. The required force to open the valve does not change with air pressure, nor does the travel. This is one of the possible fault conditions that can be detected by the absence of a confirmed valve-release (detailed in Chapter 11).

Table 10.1: Summary of potential fault events during rod operation, and associated sensor responses

Fault #	Fault Description	Audio	Force	Extension
1	Rod-End not contacted	no sound	never increases to expected contact force	extension increases past expected initial contact
2	End-effector contacted rigid, non-rod object	no sound	force spike	extension shorter than expected before bottoming-out
3	Rod-notch stuck on guide plate	no sound	force spike	extension shorter than expected before bottoming-out
4	Rod doesn't have enough travel, loop bottoms out on guide plate	no sound	contact force seen, then force spike	extension may be less than minimum for actuation

10.2.3 Notched Rod Solution

10.2.3.1 Case 1

It was shown in San Antonio that a linear push at a slight upward angle would lift the rod notch off of the guide plate and enable a proper actuation. This has been repeated in the lab.

The relevant forces are shown in Figure 10.4. The rod must be lifted, so its weight must be exceeded by the vertical component of the force. The valve must be operated, so the required operating force must be applied by the horizontal component of the force. Cars are, on average, 10.5 ft wide, so the rod on each side of the valve is approximately 5.25 ft long. The weight of the active half of the rod is then given by $W = LR^2\pi\rho$, where L is rod length in inches, R is rod diameter in inches, ρ is the density of steel, $\frac{L}{2}$ is the moment arm of the center of mass under the assumption of equal weight distribution, and L is the moment arm of the applied vertical force. On average, the weight of the active rod half is 1.75 lbs.

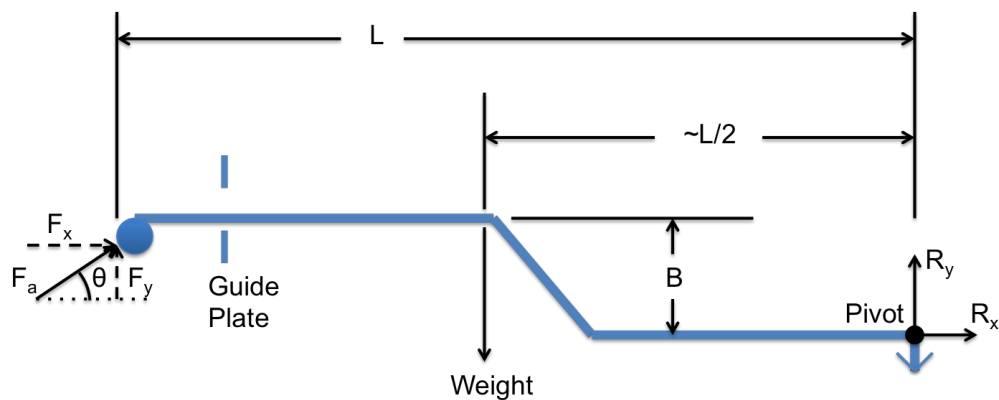


Figure 10.4: Free-body diagram for “Case 1”

The required horizontal actuation force measured between 20 and 30 lbs; the valve begins to move at 24 lbs, appears to be fully released at 30 lbs, and reaches the end of its range of motion around 35 lbs. Here, a 30 lbf requirement is assumed. The required vertical force, F_y , can then be calculated using a net moment equation, with the valve handle as the center of rotation, as $F_y = (W\frac{L}{2} - F_x B)\frac{1}{L - \mu B}$. The moment about the pivot due to the vertical offset, B (as shown in Figure 10.4) may effectively make the rod heavier or lighter, depending upon the direction of the offset. Without knowing the geometry of each rod in advance (all are unique), the necessary vertical force cannot be predicted based on the bending moment. Because B can be in either direction, it is likely best to assume that $B = 0$ (i.e. that the rod is straight), and use active force control to modify the force angle to account for rotation.

Using $F = \sqrt{F_x^2 + F_y^2}$ and $\theta = \tan^{-1}(\frac{F_y}{F_x})$, the applied force should be 30.18 lbs. at a 3.35° incline.

10.2.3.2 Case 2

If the solution for Case 1 fails, then Case 2 has almost definitely occurred. In this case, the inward-facing face of the notch has pressed up against the guide plate, and any horizontal force on the rod now induces a resisting vertical friction force. The minimum required vertical force component now includes the friction, which would likely require over a 45° incline to overcome, given the high friction coefficient between corroded steels.

Alternatively, the robot should remove the horizontal force, and, instead, lift the rod notch off of the guide plate before continuing to push horizontally.

A lip on the bottom of the contact plate (shown in Figure 10.5) protrudes enough to catch the loop; the lip must not extend past the loop, otherwise the lip may bottom-out before full valve travel is reached. The procedure will be very similar to that for the initial rod contact.

1. The robot gets a baseline vertical load from the F/T sensor
2. The end-effector slowly moves upward until a contact force is seen; this should not exceed the weight of the rod
3. From this reference point, the robot continues lifting for another 0.25" (this assumes that a notch will always be less than half the diameter of the bleed rod); alternatively, the robot can lift until the rod contacts the top of the guide plate, then drop until only the weight of the rod is felt by the F/T sensor
4. After the notch has been cleared from the guide plate, the end-effector continues with the normal horizontal actuation procedure.

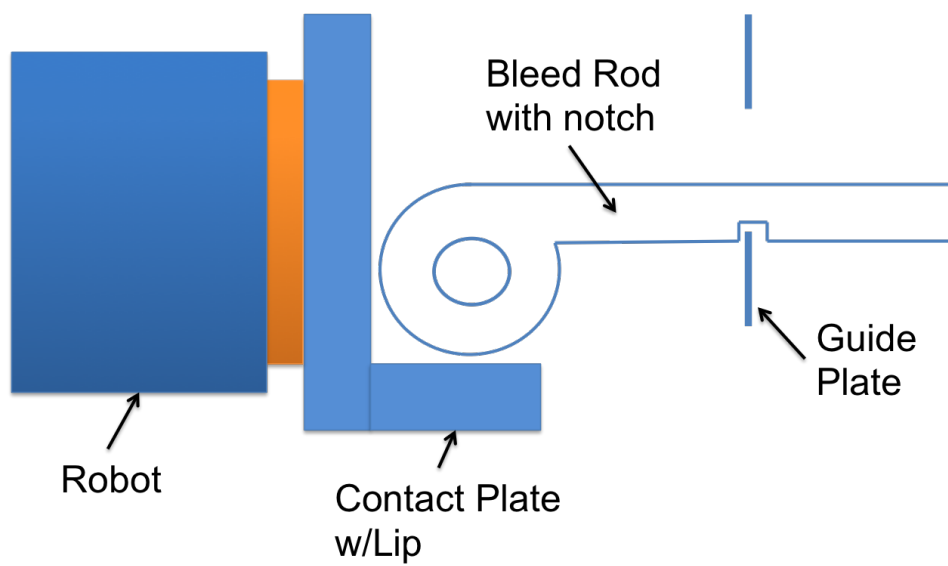


Figure 10.5: Sketch of end-effector contact plate with lower lip to lift the rod in “Case 2”

Chapter 11

Valve-Release Confirmation

The only considered technique that has been deemed a viable brake release confirmation is the Auditory Recognition. However, all of these methods combined can provide a good deal of information and insight for fault diagnosis. The possible fault cases described in Section 11.1 are compiled in Table 10.1, along with information states associated with each.

11.1 Considered Techniques

11.1.1 Auditory (microphone)

There is a clearly defined audible blow-off sound when the valve is opened. Before the San Antonio Yard visit, the only information available regarding this sound was that operators know the valve is released as soon as the sound is heard. This indicator gives no intermediate technical information, making the operation unobservable in the interim; according to operators, this is the only fully reliable method of confirmation.

11.1.2 Force Monitoring

We anticipated a distinct force profile during the valve operation, on which a specific point could be found as a confirmation. However, as the lack or presence of air pressure has no effect on the required operating force, force-monitoring would not be useful as a release confirmation.

Force information is still useful for other observations. The applied force on the end-effector will spike when something has bottomed-out. This bottoming out could be the valve stopping at the end of its range of motion, it could be a rod-notch getting stuck on the guide plate, or it could be the end-effector hitting something rigid other than the rod. In any case, knowledge of a force spike will allow time for the robot to safely stop without damaging an actuator.

11.1.3 Piston-Extension Monitoring

In the lab demonstration, a pneumatic cylinder is being used for actuation. Monitoring the extension of the piston within the cylinder is analogous to monitoring position of the robot end-effector, if the robot were applying the linear motion. Knowing how far the end-effector has travelled when something occurs (sound, force spike, etc.) allows for intelligent interpretation of the situation.

Because there could be cases when the rod can be fully actuated but the valve not released, simply having a full extension does not indicate success. These cases are the following:

- The connection between the rod and the valve handle has broken, possibly due to a missing or broken pin
- The rod-end is too close to the guide plate, and the loop (or “L”) at the end has bottomed-out against the guide plate, rather than the valve handle fully opening
- The compressed air reservoir tank is empty

11.2 Auditory Feature Recognition

The bleed-off sound heard in the yard (that has been recreated in the lab) is significantly louder than other ambient yard sounds and has a distinct sound signature. The key features of any signal are Amplitude, Frequency, and Waveform. Additional, more abstract, features may be extracted for increased robustness[4][6]; the standard methods for using more abstract features for classification are described in Section 11.2.4. All of these features are useful for identifying this particular noise.

LabVIEW has built-in audio processing toolboxes that were used to develop and test the audio Valve-Release Confirmation. The RODE VideoMic Pro microphone that was used for testing is a directional stereo microphone. Directional microphones have a narrow “cone,” outside of which the sounds are greatly attenuated (see Figure 11.2). As shown in the polar plot, attenuation at the frequencies of interest (5 kHz - 7 kHz) begins around 20° off of direct



(a) Front View



(b) Rear View

Figure 11.1: RODE VideoMic Pro directional microphone

alignment. This feature helps to prevent unnecessary noise that could produce false-positive detections.

Polar pattern

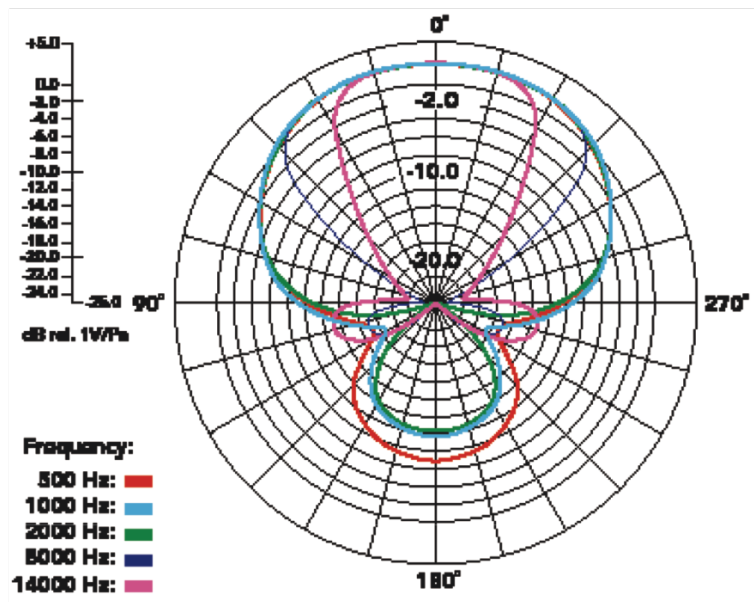


Figure 11.2: “Cone” for the RODE VideoMic Pro

For the brake release detection, a directional microphone is ideal, as the location of the source of the noise is relatively well-known and consistent. However, for other applications requiring sensing in all directions, such as improved awareness for localization or danger avoidance, an omnidirectional microphone would be better suited

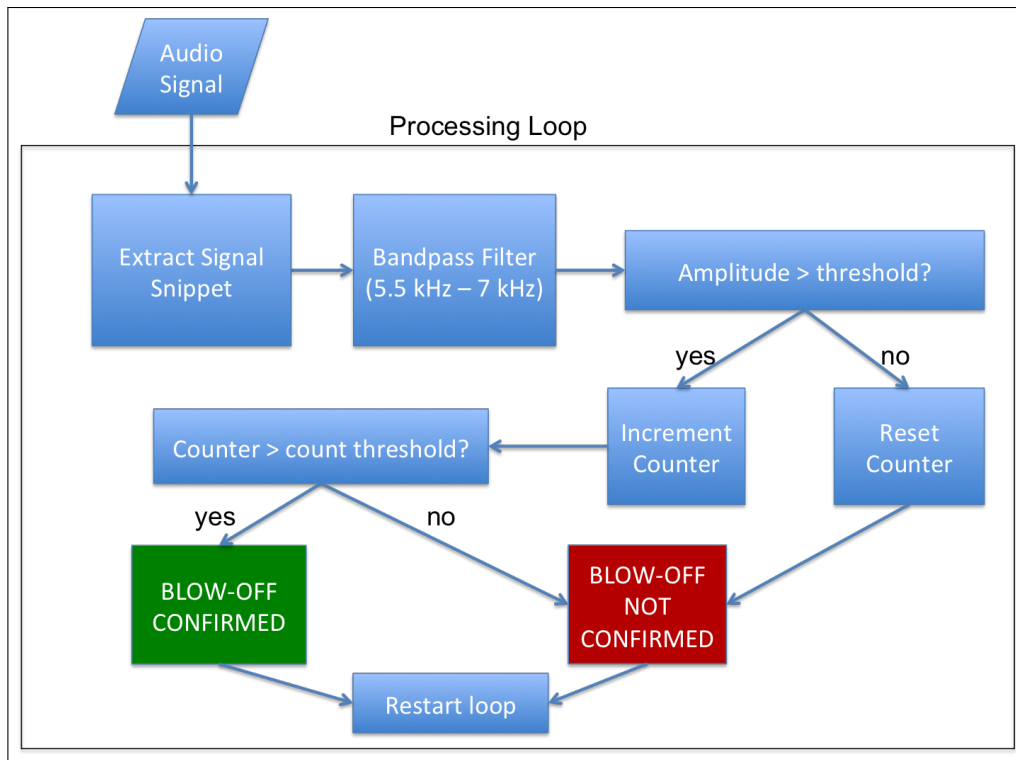


Figure 11.3: On-line audio processing flow chart

11.2.1 Signal Frequency

The first step in the audio filtering is targeting the expected frequency range. This was done initially using a 3rd order Butterworth Bandpass Filter, with cutoff frequencies set to 5500 Hz and 7000 Hz. Additional testing should be done using different types of filters, as something other than Butterworth may be better suited to this application.

The cutoff frequencies were chosen using a “primary frequency extractor” in LabVIEW. First, the primary frequency (frequency present in the signal that has the highest amplitude) from audio recordings was extracted

and recorded. This provided a starting point for a frequency range of interest, then the cutoff frequencies were tuned manually during on-line processing.

11.2.2 Signal Amplitude

Once the signal has run through the bandpass filter, all frequencies outside the desired range are attenuated, so the amplitude of the signal doesn't require further processing. An amplitude threshold is set, so if the amplitude is too low, the signal is rejected as a valve-release confirmation.

The signal amplitude seems to be specific to the computer being used and its sound card. The sound cards in the computers that have been tested have some form of a "high sensitivity" mode for the microphone input; this appears to be necessary to put the signals in a useful amplitude range.

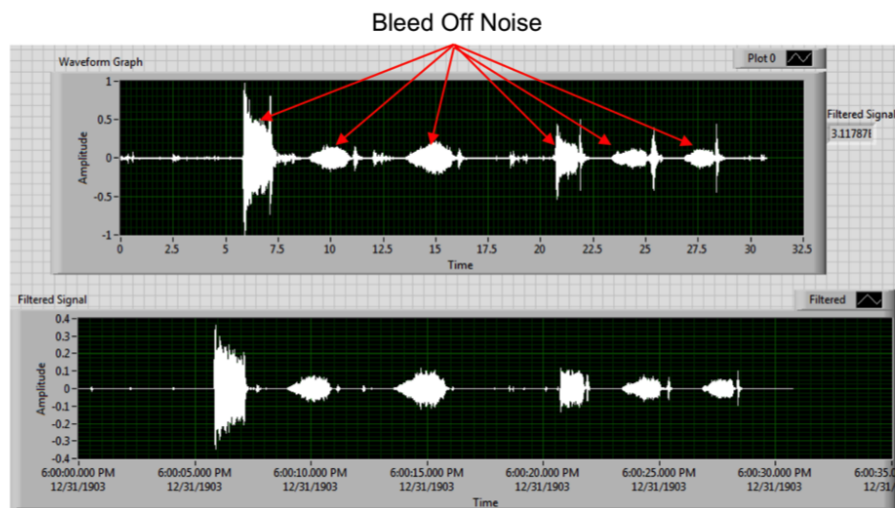


Figure 11.4: Example of audio signal without before (top) and after (bottom) bandpass filter

The signal amplitude threshold largely eliminated background noise from the San Antonio Yard recordings (talking, passing trains, etc), but still there were some false-positives, as well as false-positives that were able to be induced.

11.2.3 Time Threshold

The graphs in Figure 11.5 are plotting detection (a Boolean value, 0 or 1) vs. time. In the “before” graph of Figure 11.5, brief spikes can be seen adjacent to all but the third “True Detection.”

The false-positive spikes from the recording have been attributed to metal-on-metal impacts, for instance, the rod hitting the side of the opening in the guide plate. Similar spikes have been reproduced in the lab by hitting metal together (wrench on steel square tube). Both the amplitude and frequency of these sounds drop outside the detected range quickly. Because the ABRS is going to hold bleed rods in for 3 seconds, there will be no important information lost by requiring a minimum duration measure.

The bandpass filter is using a sample rate of 44,100 samples/second/channel (2 channels, 16 bits per sample per channel) and 4000 samples/channel as each sample set, which yields 0.0907 s (90.7 ms) per sample set. This means that 90.7 ms is the resolution of the detection algorithm. The minimum allowable number of detected sample sets has been set to 4, which is a duration of 362.8 ms. 362.8 ms, then, is the minimum blow-off duration for a reliable detection. The size of the sample set could

likely be decreased from 4000 samples/channel, but must be done carefully, as too great a decrease proved detrimental in testing.

Setting a minimum detection time removed the false-positive spikes, as can be seen in the “after” of Figure 11.5.

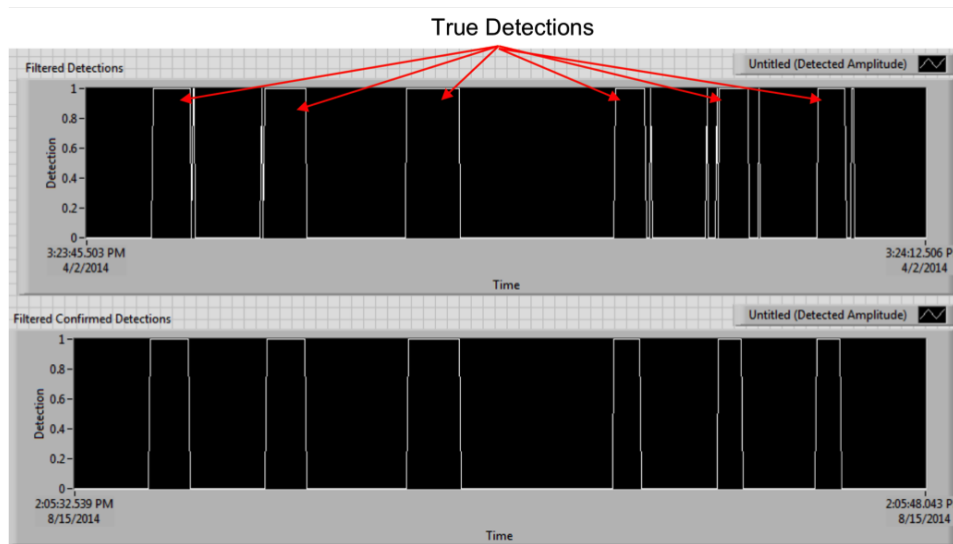


Figure 11.5: Audio detections from amplitude threshold, before (top) and after (bottom) minimum duration threshold

The only sound produced in the lab that has yet to be filtered out as a false-positive is metal-on-metal banging that causes resonance. If the square tube is hit in the right place, the sound resonates loudly enough, and at a close enough frequency, that it is picked up for multiple detection samples. If the square tube is hit repeatedly to create resonance, the vibration sound remains continuous, and the detection program can mistake it as a blow-off sound.

This is a very particular case that, as far as UP knows, is not replicated in the yard by anything within close proximity of the service valve. If either this conclusion is incorrect, or additional robustness is desired, there are options for future development.

11.2.4 Sound Signature Classification

Just as musical instruments all sound different because of distinct sound signatures (wave shapes), the sounds mentioned above are easily distinguishable by a human and, therefore, should have distinguishable numerical sound signatures.

Standard practices for audio classification typically utilize the following scheme[4][6]:

1. Signal conditioning
2. Feature extraction
3. Trained Classifier

11.2.5 Characterized Noise Rejection

A rail yard is a very well understood environment, with well known operations that take place. The possible locations of these operations relative to the location of a confirmation microphone are also known. The sounds associated with any operation within close proximity to the microphone can be recognized, and explicitly filtered out of the “Release Confirmation” audio

signal to prevent false positives. Another benefit of these separate operation recognitions is added spatial awareness. Anything providing a mobile robot additional information that can aid robot platform localization is beneficial

11.3 Microphone Placement

When checking for the valve-release confirmation, to obtain the best results from the confirmation algorithm, the audio signal should be as undistorted as possible. The position of the valve relative to the microphone and the presence of obstructions could interfere with the audio signal, causing the reliability of the audio confirmation algorithm to degrade. In order to maximize the reliability, the microphone should be positioned such that the interference and dampening due to obstructions and/or location of the valve are minimized.

11.3.1 Experiments

A series of experiments were performed to examine the effect microphone positioning has on the robot's ability to gain auditory confirmation of the brake release using the current auditory detection algorithm. The primary characteristics tested were the presence of obstructions, the distance from obstructions, and the total distance to the audio source, the valve. These characteristics were evaluated based on the overall observed change in the detection duration. The duration of detection was measured as the amount of time the detection algorithm continuously identifies the blow off noise from

the brake-release valve, up to the first interruption. The microphone was positioned in 4 different locations, seen in Figure 11.6. The two obstructions used were a thin-walled, hollow rectangular prism and a solid, concrete cylinder. All the microphone positions were in the same horizontal plane as the valve, that is, the vertical displacement between each microphone position and the valve was zero. Vertical displacement changes were not tested for because the effect of a vertical displacement change is the same as a change in the overall distance to the valve, given no obstructions are present. A vertical displacement change that alters the presence of an obstruction and results in a noticeable change in the detection duration is a direct result of the obstruction, and possibly the change in the distance to the valve, and not the change in vertical displacement.

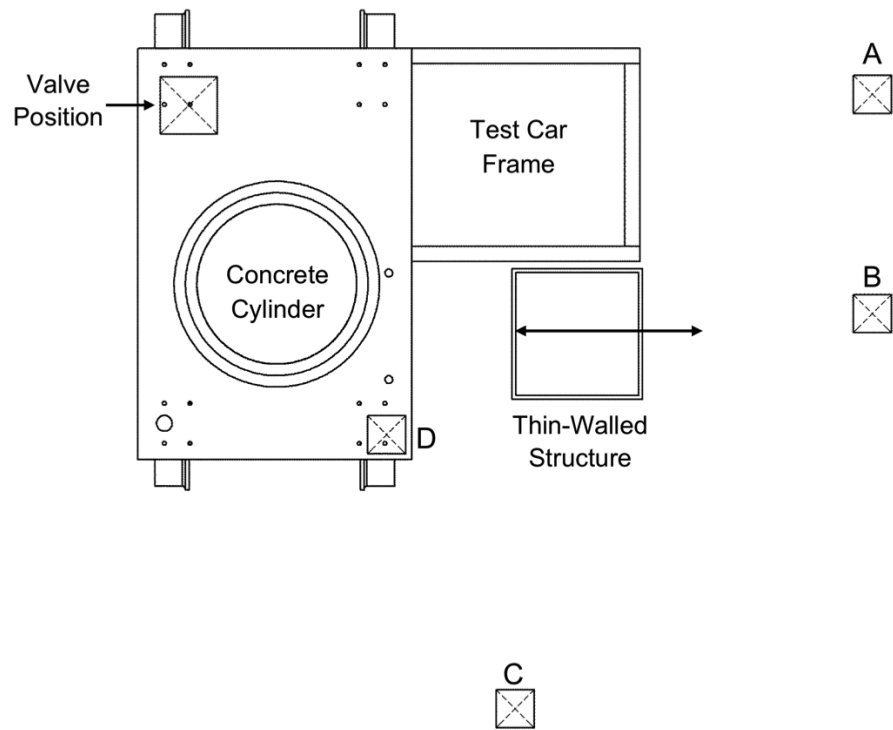


Figure 11.6: Overhead view of experimental setup, showing the concrete cylinder and thin-walled structure obstructions and positions A, B, C, and D

11.3.2 Results

Table 11.1: Tabulated results from the experimental microphone conditions

Test #	Pos.	Obstruction	Type	Dist. from Obstruct. [in.]	Dist. to Valve [in.]	Duration [s]	Press. @ Duration [psi]
1	A	None	N/A	N/A	100	12	47
2	B	None	N/A	N/A	105	12	47
3	B	Yes	Thin Walled	40	105	12	47
4	B	Yes	Thin Walled	25	105	12	47
5	B	Yes	Thin Walled	6	105	6	55
6	C	Yes	Concrete Cylinder	75	105	2	71
7	D	Yes	Concrete Cylinder	15	45	3	65

When the microphone was moved further away from the valve a small amount the duration was not affected, as observed in test 2. When a relatively small object was placed more than 2' from the microphone there is still no effect on the detection duration, as seen in tests 3 and 4. The thin-walled structure only affected the detection duration when placed less than a foot away from the microphone; the detection duration was halved when the thin-walled structure was placed 6" from the microphone. When the larger obstruction, the concrete cylinder, was placed between the microphone and valve the detection duration was reduced by a factor of 6, see test 6. Reducing the distance between the microphone and the valve, while holding the position of the obstruction relative to the valve constant, only increased the detection duration by 1 second.

11.3.3 Practical Application

When objects are placed in front of the microphone the amount of time the detection lasts without interruption decreases based on the object's proximity to the microphone, its geometry, and its composition. For example, a slender rod placed at almost any point between the microphone and the valve will have little to no effect on the detection duration, but a semi-infinite plate of moderate thickness placed at almost any point between the microphone and the valve will significantly decrease, perhaps even nullify, the detection duration. How could these scenarios arise in the practical application? Figure 11.7 demonstrates a situation in which the valve is most likely positioned closer to the ground and the line of sight to the valve is completely obstructed for the

microphone. Figure 3 shows a similar valve position as Figure 11.7, however in Figure 11.8 there are no obstructions in front of the valve. Figure 11.9 demonstrates a situation where the valve is positioned far away from the bleed rod and is highly obstructed.



Figure 11.7: Lower position valve, line-of-sight is completely obstructed



Figure 11.8: Lower position valve with no direct obstructions

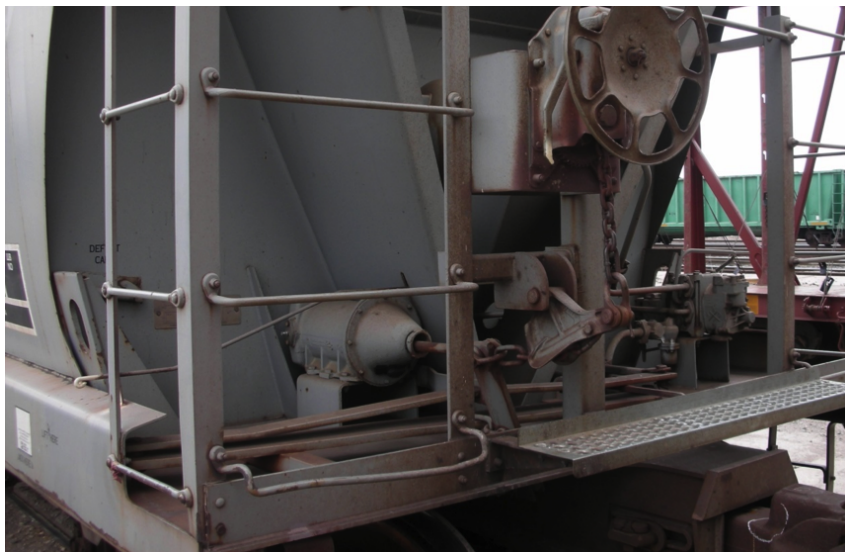


Figure 11.9: Upper position valve with obstructions

The situation depicted in Figure 11.7 would call for a microphone that is positioned lower to the ground to minimize the effects of the obstruction. However, if the microphone is positioned too low to the ground, gaining reliable detections for a situation like the one depicted in Figure 11.9, an upper valve position with moderate obstructions, could prove problematic. It would be beneficial to have the ability to vary the height of the microphone in order to optimally position the microphone for situations like those seen in Figures 11.7 and 11.9. This height variation could facilitate the robot's ability to detect the valve-release by reducing the effect of the obstructions present.

Small objects, like humans, slender rods, etc., should have no effect on the reliability of the audio detection as long as the object is not within a foot of the microphone. Large objects however, significantly hinder the detection. The robot may have a difficult time gaining a positive detection for the situation shown in Figure 11.9. The presence of the large metal plate between the bleed rod and the valve paired with the large distance between the rod and valve could result in an unreliable detection.

Chapter 12

Adaptation for Mobile Manipulator

This chapter will address anticipated issues that apply to the Automated Brake Release System (ABRS) when using a mobile manipulator. Some things discussed here are critical parameters that should drive hardware selection, and problems expected to arise when testing outside of the controlled laboratory environment.

12.1 Manipulator Requirements

12.1.1 Modularity Between Tasks

12.1.1.1 Workspace

A manipulator will have a required range of positions it must reach for any specific task, defining the minimum required workspace. The workspace of a robotic arm is defined by its configuration and the geometry of its linkages. Requisite workspaces should overlap between tasks that are intended to be performed by the same mobile robot.

However, if the arms used on the mobile platforms are all fully modular, reconfiguration can allow the same hardware to fulfill unique workspace

requirements. In this case, if two versions of mobile robots in a yard have different workspace requirements, the same set of parts can still be used.

One of the off-the-shelf robotic arms presented in Section 4.4 in [8] is fully modular.

12.1.1.2 End-Effector Versatility/Interchangeability

For a set of tasks to be performed by one mobile robot, the end-effectors required should have as much functional overlap as possible. Components that can be expected to be necessary in most tasks are a collision sensor and a force/torque sensor (see Sections 8.1.1-8.1.3 in [3] for details about their operation, and Chapter 9 for details about how they apply to the ABRS). Connected to the end of these components will very likely be a unique, custom piece of hardware for each task (gripper, vision sensor, etc).

As long as the sensors meet the requirements for all desired tasks, the custom hardware can be stored on the mobile platform, and swapped automatically when necessary. One common device for this purpose is a pneumatic tool-changer (ATI has a range of models available).

12.1.2 Field-Repairability/Modularity

Multiple levels of possible modular architecture exist for this application:

1. The manipulator can be its own module, with regard to the whole mobile robot system, that can be replaced in full in the case of hardware failure
 - This option requires having multiple robotic arms on standby, ready to be swapped out
 - Replacing the entire mounted robot as one unit is simpler than the alternative, requiring less training for employees doing the maintenance, but has a high up-front and continuing cost

2. The manipulator can be split into its own modular components (links and a small set of actuators)
 - This option requires keeping spare links (arm pieces) and actuators on-site
 - Replacing single components of a manipulator would require more training, and a somewhat more in-depth understanding, for employees who would do the maintenance
 - Individual components, of a minimum set (as few variations as possible) would require less up-front cost

12.2 Static Vehicle Stability

12.2.1 Base Stability Through Workspace

Without any custom modifications, only one of the off-the-shelf (OTS) mobile robots presented in [8] meet the stability criteria outlined in Section 5.1.2 – the Seekur. Table 12.1 shows the relevant geometric parameters for each mobile robot, and the results of their expected stabilities.

Table 12.1: Mobile Robot criteria and static stability evaluations

	Husky	Segway LE	Seekur
W (including 60 lbf robot) [lbf]	170	180	680
l_{cg} [in]	13	16.5	16
h_{cg} [in]	13	21	38
l_{ef} (width + 12") [in]	38	45	51
minimum l_{cg} [in]	13.65 $\not\leq$ 13 FAIL	12.89 < 16.5 BARELY	3.41 \ll 16 OK
l_{tip} [in]	52 \gg 8 OK	45 \gg 8 OK	26 \gg 8 OK
cost	Low	Medium	High

12.2.2 Outriggers

The critical factor in evaluating the static stability for the brake rod application is the moment about the rear pivot of the weight of the mobile manipulator platform. This factor must be maximized, by maximizing MMP weight (W , in Figure 12.1), and the horizontal distance from the rear pivot to the center of gravity (l_{cg} , in Figure 12.1). For power savings and controllability, the weight should be kept to a minimum.

A rear outrigger, depicted in Figure 12.1, serves to move the effective pivot point further back from the center of gravity. All three of the minimum l_{cg} values shown in Table 12.1 could be met using a simple rear outrigger, making all three mobile platforms feasible options, from a static stability perspective.

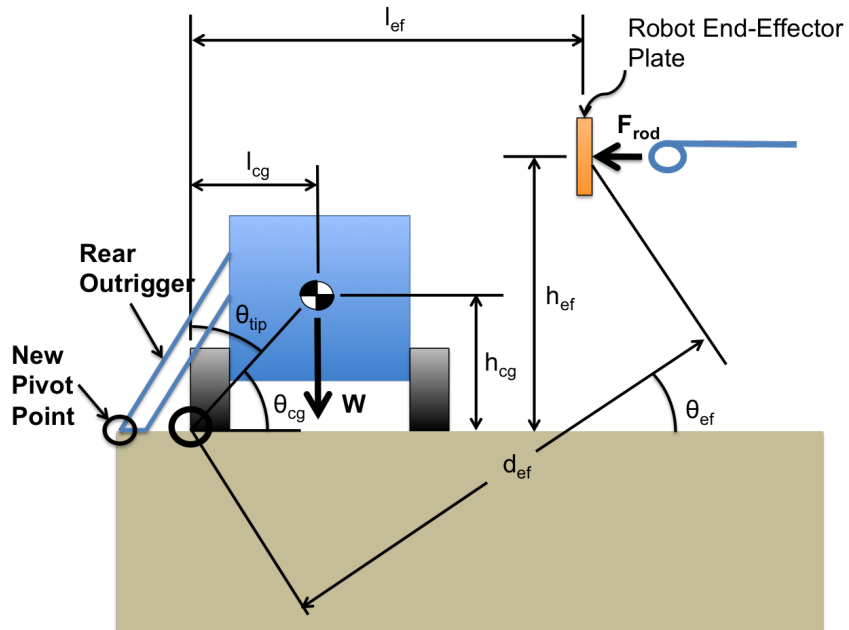


Figure 12.1: Mobile manipulator platform loads with rear outrigger

12.3 Vehicle Maneuverability

12.3.1 Steering

12.3.1.1 Skid-Steer/Differential Steering

Skid-Steer (a.k.a. Differential Steering) is imprecise, inefficient, and prone to failure. Skid-steer is typically used either in educational environments (simple, low-cost) or in applications that very rarely require turning (similar to the justification for un-steered wheels on 18-wheeler bogies).

The principle behind differential steering (skid-steer) is that when wheels on each side of the vehicle (left/right) are driven at different speeds, the vehicle will turn. For example, to turn left without moving forward, the left wheels would be driven in reverse while the right wheels are driven forward, both at the same speed; if forward motion were also desired during the turn, then the right side wheels would need to turn faster than the left side wheels.

This technique causes wear on the wheels every time the vehicle turns. The accuracy of a turn depends on the uniformity of the ground friction, and still would be very error prone even with a uniform surface due to the inherent non-linearity of the friction. The skid-steer approach also requires twice the power of that using turned wheels. During a skid-steer turn, the wheels are intentionally loaded laterally, which unnecessarily stresses the vehicle. Designing a vehicle with this in mind requires a more robust chassis, which makes it heavier without added benefit.

The only justified case for this steering method would be for a mobile platform that intends to move in straight line with limited turning. Otherwise, skid-steered mobile platforms should be avoided at all costs.

12.3.1.2 Ackerman Steering

Ackerman steering uses front-wheel steering (FWS) and is based on the idea of instantaneous center (I) of rotation. In any turn, all wheel axle centerlines should pass through the instantaneous center of rotation. For a left turn, then, the left wheels will be closer to the instant center than the right wheels, which means a smaller turn radius on that side, and a greater angle between front and back wheels. This geometry is shown in Figure 12.2.

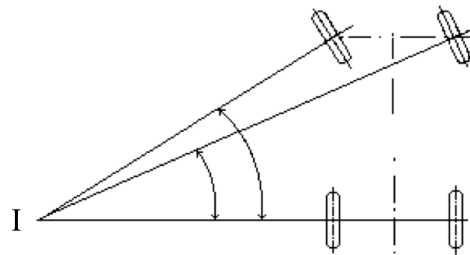


Figure 12.2: Ackerman steering schematic [1]

Ackerman steering works by mechanically linking the front wheels together, typically using a 4-bar mechanism, so that the correct relative angles between the front wheels are always attained. This has the advantage of requiring minimal active control, and only one input (steering angle). Because a linkage is placed between the two wheels, there is a mechanical limit on the maximum steering angle.

The disadvantage to Ackerman Steering is the inherent coupling between forward motion and rotational motion. Because the rear wheels are fixed, the minimum radius of a turn is determined by the length of the vehicle's wheelbase, and can neither turn in place (with no linear motion) nor move laterally without forward motion.

Ackerman steering could be useful for a lower cost, energy efficient mobile platform, meant for tasks that require minimal multi-dimensional adjustment and can have a largely planned out path in advance (that has only small, large-radius turns).

12.3.1.3 Independently Steered Wheels

Independently steered wheels (independent all-wheel steering is what is referred to here) utilize the instantaneous center principle as well, but do not have the same constraints inherent to Ackerman Steering. In AWS vehicles, **any** steering angle can independently be reached for all 4 wheels. This then decouples the vehicle's rotational capabilities from its forward motion. Being de-coupled and unconstrained, it can now turn in place, about its own center, as well as move linearly in any direction.

Because the wheels are all independently steered, a control scheme is required to functionally operate the 4 wheels to obtain the desired motion. This has been fully treated analytically at UTexas over the past 5 years. This is, obviously, more complex than Ackerman Steering, but enables the system to take advantage of 3 fully uncoupled degrees of freedom.

Independent All-Wheel Steering is necessary for any mobile platform that requires a complex, un-planned path in a complex environment. Most tasks to be performed with a mobile robot in a switch yard will fall under this category. Both Independent AWS and Ackerman Steering minimize lateral wheel slip (as it is not inherently necessary) and are, therefore, far more precise than any skid-steer system.

12.3.2 Tires

Tires/wheels (whichever is directly in contact with the ground) must be suited to the terrain and surface requirements. Traction is primarily affected by the tread pattern and material type. Ability to traverse uneven terrain is largely determined by the tread clogs and their angle-of attack.

12.3.2.1 Custom for Rail-Traversing

The wheel angle-of-attack of an obstacle to be climbed increases with increasing obstacle height, and decreases with increasing wheel radius. If tall, unknown obstacles need to be traversed, a large enough wheel size to account for the majority of obstacles is used. If, however, the obstacles are known, they can be specifically planned for with specialized wheels.

In the case of needing to traverse rails in a switch yard, the mobile robot must be able to fit in between a pair of rails; this restricts the maximum allowable wheel size.

12.4 Sensor Preservation

Any on-board sensors that need to collect data while the mobile platform is moving must be resistant to the operating conditions. They must be low cost and easily replaced.

12.4.1 Problems to Anticipate

Vibrations will occur in the presence of any rough contact interface. The following non-continuities can be expected in the switch yard:

1. Grousers (clogs)
2. Standard tire tread
3. Gravel surfaces

Each of these will cause a vibration of varying frequency and amplitude, and should all be examined closely at expected operating speeds to anticipate possibilities of induced sensor error.

12.4.1.1 Camera Vibration

Vibrations will affect a 3D camera in two ways:

1. If each pulse takes place **during** the sensing/imaging operation, the data will deteriorate

2. Successive samples from the camera will show inherent noise over time.

It is possible that this noise could be generically filtered to produce adequate data. However, aliasing (if the sampling rate is not significantly greater than the frequency of the greatest oscillation, the signal may appear to follow a trend it does not actually follow) may be possible, and if the data is needed for use in real-time, then requiring a filter is unacceptable.

Suspended Camera Mount Solution One common solution to vibrational noise is hardware isolation, through a spring-damper suspension. This acts as a mechanical low-pass filter. The inherent problem with this option is a delay in propagation of position changes. If the camera **always** has to be rigidly connected/calibrated to the mobile platform, then this is not a viable option. If, however, the primary goal is always getting a good image or good data, then this may be a good option.

Programmatic Solution Instead of solving the vibration problem mechanically, post-processing the data can also rectify the problem. If the post-processing method is fast enough, real-time use of the data may be possible.

Filters Similar to the mechanical "low-pass filter" provided by a suspended mounting option, a digital filter can be applied to the data to smooth out noise.

Accelerometer Cancellation An alternative to a passive, frequency-dependent filter, on-board sensing could be used to intelligently account for the inconsistencies. On-board accelerometer information could be used to remove platform shocks from the data. This would both be more accurate, and would prevent filtering out artifacts in the data of a similar frequency that may actually be of interest.

12.4.1.2 Microphone Noise

An on-board microphone should be suspended, because slight fluctuations in position will not affect the recorded data, and a rigid mount would transfer all platform noise to the microphone.

In addition to the isolation provided by a suspended mount, active noise canceling should be utilized. A second on-board microphone that only picks up sounds coming from the robot should be monitored. The signal from this second microphone can then be automatically removed from the signal of the primary microphone, leaving the sounds of interest (i.e., those external to the mobile robot platform).

Chapter 13

Phase II Recommendations

13.1 Further Development Needs

13.1.1 Vision Algorithm

The Object Detection algorithm has been shown to perform adequately for deployment, given the proper conditions. The “proper conditions” include positioning the camera within the acceptable workspace (as defined in Section 5.2) acceptable ambient conditions for the camera to be used (lighting, temperatures, dust/precipitation).

13.1.1.1 ANN Retraining

Currently, only the bleed rod with a looped end has been worked with. The other, less prevalent, type has an L-shaped end instead of a loop (not seen in San Antonio, but UP said they should be accounted for). Training images will be needed for these rods. The ANN will then need to be restructured to two outputs, instead of just one, and retrained. The modified ANN would be classifying the images as “looped rod,” “L-shaped rod,” or “not a rod.”

13.1.1.2 3D Camera Choice

With the ASUS Xtion Pro (and Pro Live, and Microsoft® Kinect 2, which all use the same internal 3D sensor), ambient sunlight shining directly either on the IR sensor or on the object of interest will severely degrade the image. The Microsoft® Kinect for Xbox One has recently become available, and shows great potential. This updated camera stays around the same price point of its predecessors (\$149.99, making it still far cheaper than any viable alternatives) and uses Time-of Flight technology in addition to the IR-based structured light techniques utilized by the previous cameras. The updates reportedly (no reliable sources could be found for verification) were focused on improving accuracy and eliminating the camera's sensitivity to lighting conditions, which have been two of the greatest performance concerns. This new camera should be seriously considered before moving to more expensive alternatives.

13.1.2 Multiple On-Board Camera feeds

As explained in Section 2.7.2.2, the PrimeSense 3D sensor (used in the ASUS and Microsoft® Kinect cameras) uses the majority of the available bandwidth of most usb hubs. Currently, only two of these cameras have been run simultaneously with success; even these two required a non-automated work-around. Running many of these 3D cameras simultaneously should be possible, but some development and testing will be required to determine a reliable hardware configuration and integration with the available drivers.

13.1.3 Pulling Rather than Pushing

All Bleed Rods are designed for operation from either side of a railcar. If a Bleed Rod cannot be operated from one of the two sides, it requires human attention. It has been concluded that, in the event that a brake cylinder cannot be bled by pushing the bleed rod, a yard operator should inspect the system, and he/she could do whatever necessary to trigger the valve by hand. Therefore, there is no clear reason to plan to ever pull the rod instead pushing it, as it is a significantly more complex operation. Should this conclusion be proven wrong, through testing or otherwise, an end-effector will need to be either chosen (off-the-shelf) or designed for the task of grasping the rod-end.

13.2 Mobile Manipulator Yard Navigation

Yard Navigation, for the Automated Brake Release task, can be considered within the following steps:

1. Path Planning, to avoid known obstacles and establish a nominal path to follow
 - (a) Requires a well-defined map of the switch yard
2. Reaching a reference point (start of the first car in a consist)
 - (a) Requires localization, to know a mobile robot's location relative to the desired path

- (b) Requires the ability to manage the switch yard terrain, possibly including climbing rails, and locally avoid obstacles
3. Recognizing a reference point
 - (a) Requires reliable recognition methods for confirmation
 4. Traversing the consist of rail cars
 - (a) Requires sufficient mobility to move along a straight line, accounting for ground inconsistencies
 - (b) Requires ability to manage terrain expected to be adjacent to a rail
 - (c) Requires local obstacle avoidance
 - (d) Requires control scheme to identify the side of the rail cars and maintain a constant distance

As outlined above, any task involving mobile platform movement requires a platform that is capable of traversing the known terrain, and is able to avoid unknown/unexpected obstacles.

Additionally, any path that a mobile robot will take should be planned intelligently; that plan must account for as much prior knowledge as possible, such that only unexpected obstacles need to be accounted for locally. In order to plan a path successfully, known obstacles in the switch yard must be fully documented.

13.3 Mobile Robot Recommendations

Numerous switchyard physical tasks are currently performed by yard personnel, which have some performance uncertainty and a concern for safety. This thesis has detailed the development and demonstration of a robotic system for brake release in switchyard operation. Continuing along the lines of automating the physical interface with switchyard operations to reduce uncertainty, improve safety, and increase operator capability, we propose the development of a mobile robot system to remotely perform these tasks under supervised teleoperation using an open architecture (modularity for plug-and-play) to enable trained personnel to operate and maintain (repair on demand) based on a minimum set of spares for a fleet of robot systems

We recommend that such a system be assembled from readily available commercial manipulator or platform systems to accelerate the proposed switchyard mobile platform development. In particular, Schunk and Universal offer useful modular manipulators and Seekur offers an attractive 8 DOF mobile platform with capable internal processor, vision sensors, battery, communication interfaces, etc. infrastructure.

13.3.1 Mobile Platform

The mobile platform must navigate a rail-yard and along a length of cars while carrying the necessary power supply, sensors, computing equipment, and manipulator arm(s). The platform also must provide a stable, rigid

platform where the manipulator arm operates to apply forces and perform the brake release task.

The switchyard is necessarily compact with a complex layout of tracks and facilities. No known robot platform is capable of getting out of the well formed by two parallel 7" high rails. Numerous vehicles, such as railroad cars, locomotives, rail inspection platforms, etc., and yard personnel move continuously in the switchyard. Clearly, this requires a well-trained consort of yard personnel. Unfortunately, yard personnel are not always aware of all train consist movements. This uncertainty leads to safety issues and there is the potential for imperfect task performance which can cause movement failures and significant downtime. Thus, a mobile platform should be used to enhance operators in performing tasks and should be capable of determining sufficient detail of surrounding environment.

13.3.1.1 Recommendation

Given the complexity and diversity of switchyard tasks, we propose to develop a cost-effective Mobile Manipulator Platform (MMP) to eventually reduce manual task performance to less than 5%, to act in a supervisory mode, and to perform tasks in "emergencies". The first issue is supervised teleoperation where the goal is to have a human supervisor to remotely handle up to 5 robot systems. The second issue is task performance based on a library of tools to automatically reconfigure the MMP to carry out a given task on demand. The third issue is that the life cycle cost must be managed by building

the system out of standardized, low-cost, but highly certified, components enabling repair (by plug-and-play) by yard personnel. This would enable a minimum set of components (spares) to fully populate a large number of similar robot platforms. The governing requirement is cost-effective technology maintainable and operated by yard personnel. This requires that the system be modular with standardized interfaces for plug-and-play. The modules must be in a minimum set to populate a large collection of robot platforms (MMP). A minimum set will minimize the total number of modules that need to be designed and maintained while maximizing the overall capability of the family of modules. All platforms would be battery operated with a generic machine vision and sound-sensing system. The CPU would contain a continuously updated and universal operating software system. Unfortunately, the platform must maneuver over and among a very complex set of obstacles, particularly rails, switches, sensor boxes, etc. This requires a 4-cornered platform (each corner has 2 DOF for high dexterity), illustrated in Figure 13.1, with a specialized grouser wheel to climb over 7" rails.

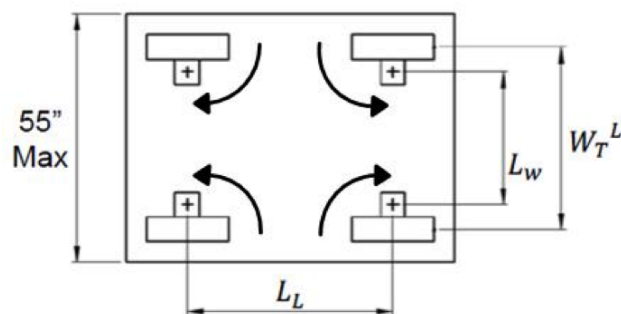


Figure 13.1: Mobile platform offset wheel configuration

13.3.1.2 Others Considerations

Having a grouser every five inches means that there would be 30 grousers per wheel. This means that there would be $4 \times 30 = 120$ “bumps” for every revolution of the four wheels. This might be acceptable if we can isolate the vision system from the impact shock of these bumps. Special coordination between this design detail and the design of the vision system must occur to ensure desired performance.

The height \mathbf{h} and width \mathbf{w} of these “parallel” grousers would govern the magnitude of these shocks. Say you were running at 5 mph or $7.3\text{ft}/\text{sec} = 88\text{in}/\text{sec}$, which would represent $(88/150)(120) = 70.4$ bumps per sec., which is a relatively high frequency. Reducing the tire stiffness (lower tire air pressure) would reduce the sharpness of the shocks (higher harmonic content). Then, if the grousers were shaped to bend a little to enable a better grasp of the rail, they could be of a shorter height \mathbf{h} which would also reduce the shock harmonics. Finally, lowering/raising the pressure in the tires on demand might be considered, but it would add expense and increased maintenance.

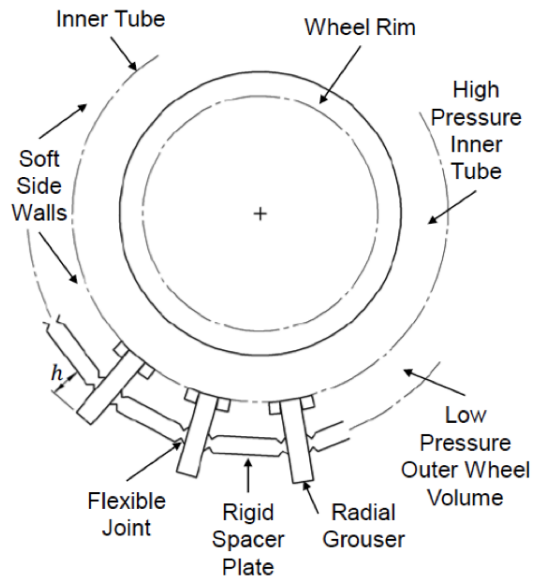


Figure 13.2: Low impulse grouser wheel

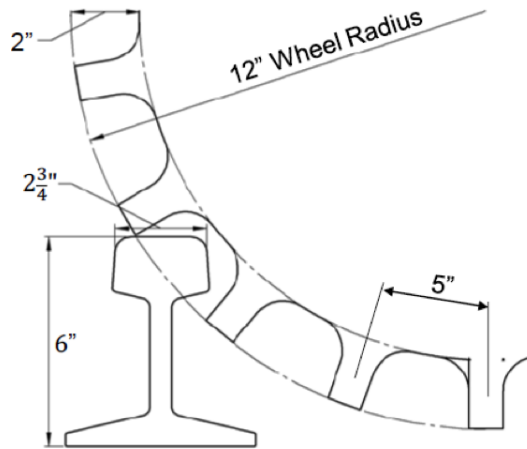


Figure 13.3: 24" grouser wheel climbing a 6" rail

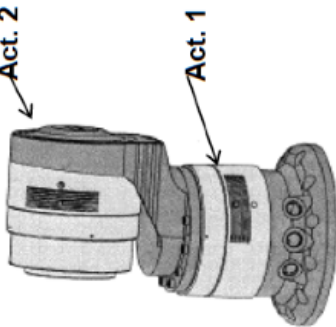
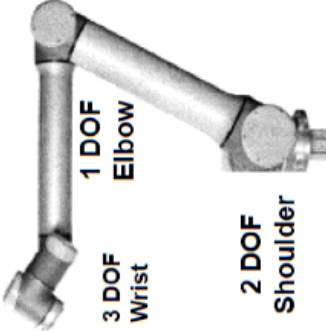

Commercially Available MMP Systems (For Early Switchyard Tasks/Navigation Demonstrations)		
SCHUNK DOF Module		<ul style="list-style-type: none"> • Plug-and-Play • Serial Manipulators • 2 to 8 DOF • Higher Cost
UNIVERSAL 6 DOF Manipulator		<ul style="list-style-type: none"> • Somewhat Modular • High Dexterity • Lower Cost • Medium Loud
SEEKUR 8 DOF Mobile Platform		<p>2 DOF Wheel Corners</p> <ul style="list-style-type: none"> • High Dexterity • Good Sensors • Communication Links • Battery Powered

Figure 13.4: Off-the-shelf mobile robot system components

13.3.1.3 Future Applications

Two important mobile platform tasks in the RR switchyard involve brake release of a cut of cars before moving through the car decoupling station and the decoupling/re-coupling (hose lacing) operations of the pneumatic air brake hose line between cars. It is recommended that one MMP be equipped to perform these functions on either side of the automated pin-pulling station.

Air Hose Lacing To complete the process of taking a cut of cars apart and rearranging them into a new cut of cars for the next movement requires the automated decoupling (before the pin-pulling step), and the automated recoupling of the air hoses between cars (after the pin-pulling step). This is called hose lacing. Unfortunately, the air hose is physically close to the RR car coupler, which is bulky and hard to maneuver around. To decouple the hoses may require a simple end-effector which grasps the air hose coupling and allows the two ends to rotate relative to each other, while a manipulator raises the end-effector (without hitting the car coupler). To re-lace the air hoses is much more complex.

Hose Lacing/Brake Release MMP A reasonably dexterous MMP is needed to perform both the brake release and hose lacing operations in the switchyard. It is assumed that the mobile platform and its structured navigation exist and that the machine vision can be augmented. Here, we want to concentrate on the primary hardware for this BH (Brake release/Hose

lacing) MMP. An important requirement is that it uses as much standard, low-cost robot technology as possible. Unfortunately, the following two constraints exist:

1. The brake release rod can be low next to the ground or high at the platform height.
2. The air hose near the coupler may be 5 ft. from the nearest access area for the mobile platform, which can require a long and ungainly manipulator reach.

Figure 13.5 shows a mobile platform with the following three principal active subsystems:

1. An outrigger to stabilize the platform from overturning forces/movements due to the actions of active forces generated at the end-effector of the manipulator and gravity forces from the extended manipulator.
2. A 2 DOF vertical mast/horizontal boom to raise/lower and extend the position of a lightweight dexterous manipulator held to the end of the boom.
3. A standard, low-cost 6 DOF manipulator at the end of the boom that can perform dexterous motions to achieve the hose lacing operation.

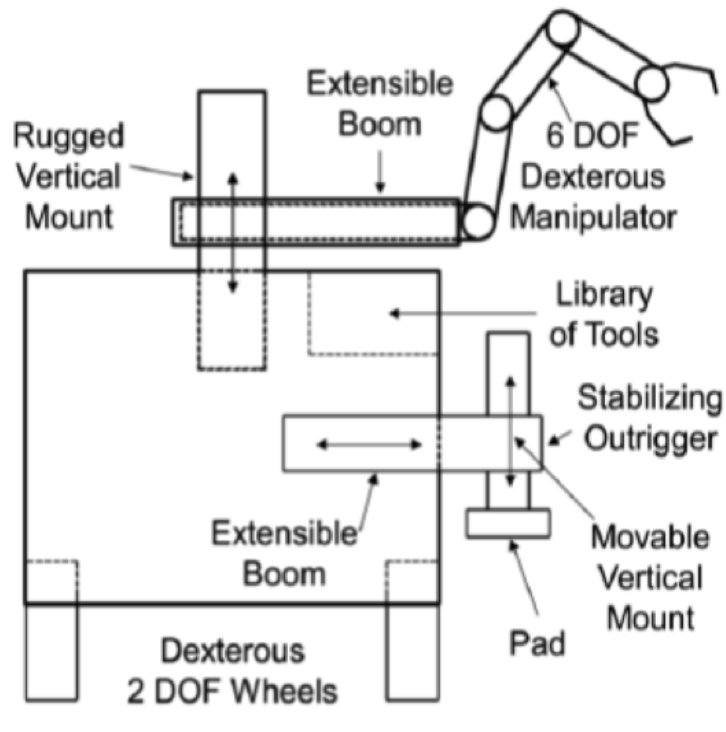


Figure 13.5: Versatile MMP for brake release & air hose lacing

None of this technology is unusual (it uses standard components for rapid repair and updating), and it could be operated by available software with specialized commands for the unique tasks.

In summary, this MMP should be able to carry out the brake release and the hose lacing operations. The brake release and hose uncoupling before the pin-pulling station could be performed by a simpler version of the MMP. The hose lacing (coupling) after the pin-pulling station to form a cut of cars would require a fully equipped MMP with special machine vision, a library

of end-effectors, two manipulators (one more complex than the other), a rigid outrigger, and a useful structured navigation operating system.

Other Considerations Sensors on the mobile platform must be used for navigation and task performance. The following is a list of viable sensor types for this sort of robot design.

1. **Lidar:** Lasers are used to survey an area by measuring reflections of light and generating a 3D profile to map out surface characteristics and detect objects
2. **Vision:** Vision could be used for obstacle avoidance, motion detection, object detection, or any number of other tasks. Two cameras placed adjacent to each other can provide depth information using its stereo vision.
3. **Time-of-Flight:** A light source is synchronized with an image sensor to calculate distance based on the time between the pulse of light and the reflected light back onto the sensor.
4. **Ultra-Sonic Sensor:** Ultrasonic sensors are used to detect obstacles meters away from the platform as an aid to obstacle avoidance.
5. **Touch/Collision Sensor:** A touch sensor is used to detect collisions with the external environment; these would be necessary for safety and protection of the robot while operating in an uncertain environment.

6. **GPS (DGPS):** Global Positioning Satellite (system) and Differential Global Positioning System. GPS is used to locate the latitude and longitude of the system with resolution from meters to centimeters. GPS can be used to set way-points within a yard to help navigation to and from charging/maintenance/action stations.
7. **RFID:** Radio-frequency identification (RFID) is the wireless use of electromagnetic fields to transfer data, for the purposes of automatically identifying and tracking tags attached to objects.
8. **Sound Microphones:** A lot of distinct noises exist in the switchyard to assist in the motion planning of the MMP (and manipulator) augmenting other more precise task oriented sensors to maximize use information in the structured navigation and task performance of the robot system. Directional audio detection is also possible.
9. **Magnetic Compass:** Similar to a handheld magnetic compass, Digital Magnetic compass provides directional measurements using the earth's magnetic field which guides your robot in the right direction to reach its destination. These sensors are cheap compared to GPS modules.
10. **IMU:** Inertial Measurement Units combine properties of two or more sensors such as Accelerometer, Gyro, Magnetometer, etc, to measure orientation, velocity and gravitational forces. IMUs are capable of providing feedback by detecting changes in an objects orientation (pitch, roll and yaw), velocity and gravitational forces.

11. **Infrared Distance:** IR circuits are designed on triangulation principle for distance measurement. A transmitter sends a pulse of IR signals which is detected by the receiver if there is an obstacle and based on the angle the signal is received, distance is calculated. IR cameras can be effective vision systems.
12. **Hall Effect:** A Hall effect sensor is a transducer that varies its output voltage in response to a magnetic field. Hall effect sensors are used for proximity switching, positioning, speed detection, and current sensing applications. Hall effect sensors are used to detect cars for traffic light control.
13. **Laser Range-Finding:** Laser(s) used to determine accurate relative distance by targeting a known object (reference), possibly detected by a vision system.

13.3.2 Manipulator

13.3.2.1 Recommended Development

The manipulator arm serves to position the end-effector and to apply a force (or torque if desired) in a prescribed manner. To accomplish these functions, the robot must have joints that provide forces or torques and links that connect the joints, the platform, and the end-effector. The number, nature, and placement of joints, or DOF, and the geometry of the links determines the motion possibilities of the manipulator workspace). The joint

torques or forces determine the forces the manipulator is able to provide at the end-effector. The workspace requirements and the end-effector force requirements form the primary design of the manipulator arm. Based upon the variation of the brake release rod location, an appropriate workspace can be generated. Figure 5 shows the expected operational layout and the required workspace of the robot arm. If the mobile platform sets the base of the arm close to the midpoint of the workspace range, then the arm needs to have a reach slightly larger than half of the 50". A reach of 31.5" provides a factor of safety near 1.25 for the required reach. Each joint must supply a torque or force to move the manipulator arm and apply a force to the release rod. It is anticipated that pushing the release rod will primarily drive the actuator load requirements, as opposed to the acceleration, pose, or weight due to gravity of the end-effector. A measured load of 25 lbs. activation force required for the lab installed release valve was found during testing. This value is expected to be typical to activate the release valve. A factor of safety of 1.25 brings the required load capacity to 33 lbs. With the workspace and load requirements known, a specification sheet can be made. The manipulator arm spec sheet in Table 1 contains the primary requirements and secondary desirable attributes, which includes low weight, low power consumption (to increase operational time), and high DOF (higher dexterity).

Table 13.1: Manipulator Requirements [8]

Requirement	Desired Value	Importance (Max: 5 Min: 1)
Large Payload	15 kg	4
Large Reach / Useful Workspace	800 mm	4
Low Power Consumption	n/a	3
Light Weight	n/a	3
High DOF	≥ 5	2

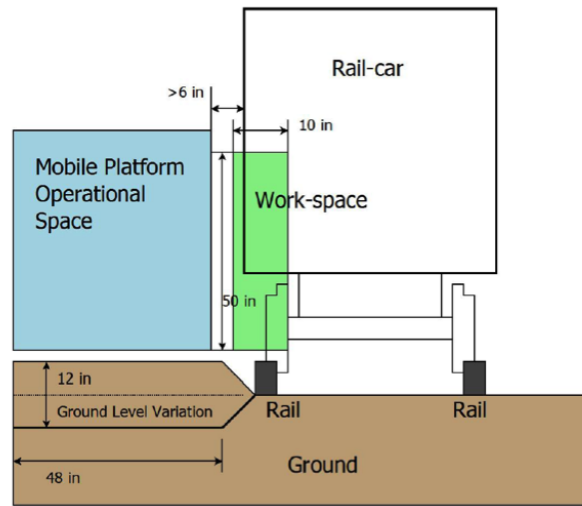


Figure 13.6: Required manipulator workspace [8]

A fundamental specification of robot manipulators is the degrees of freedom (DOF) and the motion (rotary or linear) of each joint. Rotary joints are stiff, low weight, and can simplify the architecture of the manipulator arm; therefore, they are preferred over prismatic (linear) joints wherever possible. The number of DOF is determined by the number of motions required to achieve a given task. The release rod moves nominally along a straight line, so a robot to follow such a path, 5 DOF minimum is required. If an end-effector allows motion between the arm and the release rod, then 3 DOF is possible.

However such a design would be specific to this task, so a minimum of 5 DOF and most likely a 6 DOF robot will be assumed for this design. Ultimately the manipulators must be modular, lightweight, and very low cost (from 3 to 6 DOF). A library of tools for inspection and physical tasks would be on each robot platform and automatically interchanged for each succeeding task. Much of this technology has been formulated over the past two decades at UTexas, so its implementation should proceed with nominal uncertainty.

13.3.2.2 Recommended Actuators

Schunk This German-based company has been making standard actuator modules for two decades. The module architecture is self-contained (wiring, connections) using D.C. prime-movers and harmonic gear drives. These units remain expensive, have poor shock resistance, have high lost-motion, and are relatively inefficient. Schunk offers four standard actuator modules with a collection of standard connecting links to enable the user to assemble any 2 to 8 DOF manipulator on demand. The Schunk module in Figure 13.7 has two actuators, connected by a standard link (using a circle of bolts).

The LWA (Light-Weight Arm) 4D is a 7 DOF arm that has a good payload-to-weight ration (1:2). The LWA-4D has internal actuator power supplies, and the robot is powered using a 24V power supply. This robot is attractive due to its light weight (18 kg), adequate reach, modularity, low power requirement, commercial support with ROS (Robot Operating System), and the brakes within the actuators.



Figure 13.7: Schunk LWA 4D robotic arm

Universal Robots (UR) This is also a European company which produces a low cost (and somewhat modular) 6 DOF manipulator. Last year, they sold 2,000 and will sell 7,000 units in 2014. The UR10 (named for its 10 kg payload; shown in Figure 13.8) has a maximum reach of 51", weighs 63.7 lb., lifts 22 lb., has a resolution of 0.004 in., built of aluminum/plastic, and has an expected life of 35,000 hours. The volume of its control box is 2 ft³ and uses a standard 12 in. touch screen to enable user input. This is a rather low powered system, so it would not be able to rapidly move loads. It is, however, very slim and should provide exceptional dexterity throughout its work space.

Universal Robots makes a 6 DOF arm which is fairly light with a 10 kg payload. This robot is attractive due to its ROS compatibility, adequate reach, simplicity, and low cost. More needs to be known about the actuators because it would be helpful if the actuators could brake during operation of the release rod.



Figure 13.8: UR10 & UR5

Most available robots are industrial sized and packaged which means they are heavy and meant to be run off of AC mains power supply and bolted to the floor. A few robots such as the Schunk LWD-4D and the UR10 arm are easily fitted onto a mobile platform and used with ROS. Unfortunately this lists contains robots mostly capable of only 10 kg payloads, when the requirement is closer to 15 kg. It would be possible either to brake the actuators on a 10 kg payload robot and have a pneumatic cylinder at the end-effector actuate the rod, or to work in a small portion of the workspace that has greater capacities (this is explained in a bit more detail in Section 10.1.3). In such a case the Schunk LWA 4D or UR10 robot arm should be adequate for this application as indicated in Table 13.2. For additional detail on commercial solutions, see Section 4 in [8].

Table 13.2: Manipulator Comparison [8]

Requirement	Schunk LWA4D	UR10	SAI10D	Ilwa R820
Payload	10 kg	10 kg	10 kg	14 kg
Reach	1113 mm	1300 mm	1203 mm	1291 mm
Power Supply	24 V 15 A	100-240 VAC	240 VAC	???
Weight	18 kg	28.9 kg	60 kg	29.5 kg
Has Separate Control Box	No	Yes	Yes	Yes
Estimated Cost	\$\$\$	\$	\$\$	\$\$\$\$
ROS Compatible	Yes	Yes	Yes	No

Chapter 14

Conclusions & Results

14.1 Introduction

Union Pacific railroad (UP) is interested in performing repetitive physical tasks on stationary cars in switch yards using mobile robots. This thesis introduces a standard framework for evaluating and planning for desired autonomous (or semi-autonomous) operations; then applies, in detail, the framework to the task of automated brake release before car decoupling.

A significant hurdle to be accounted for is the lack of standardization of much of the hardware of interest. Non-standardized rail car components must be formally structured as fully as possible to improve the reliability of the robotic automation.

The second application of mobile robotics UP is focusing on is releasing the emergency brakes on cars, which must take place prior to uncoupling (which was demonstrated as their first task of interest). This task requires either pushing or pulling a “bleed rod” that protrudes from the side of each rail car. The requirements for each step of the process will be laid out in this conclusion, and hardware recommendations are given based on these requirements.

Thousands of rail cars go through hump yards and/or switch yards each day, and each and every car must be operated on at least once. Because the cars are stationary during the brake release tasks (suggested to be done by a mobile robot), safety is not as great a factor as for other dynamic operations (e.g. pin-pulling), where the cars are necessarily moving during the process. The three major factors motivating a move away from human operators in this case are safety, reliability, and efficiency, which include the possibility of human operators supervising multiple robots for increased capability.

These tasks are monotonous and repetitive, and are burdens for human beings, who are capable of higher-level thinking and decision making, especially given the proper information.

An easily adaptable object-detection algorithm was developed for the Automated Pin-Pulling System (APPS); the algorithm utilizes a depth camera (3D camera) and an Artificial Neural Network (ANN). The most taxing modification that is needed for any new hardware is collecting images of the object of interest for retraining the ANN. Other parameters in the algorithm should be tuned for the other expected images.

The next, more open-ended task is developing heuristic requirements. The heuristics are checked during pre-processing (before running the computationally expensive ANN); if the conditions aren't met, further processing is not required. Heuristics serve to speed up the object search process, as well as increase the algorithm's robustness to false-positive detections.

14.2 San Antonio Yard Visit - Understanding the Problem

The primary goal of visiting this rail yard in San Antonio was to gather information about the brake release rods (bleed rods). The information needed to establish the constraints of this task includes rod positions, mounts, clearances, travels, forces, and any other still unknown variations or constants.

Each valve housing is very similar, with a few different valve models, but all can be assumed to be equivalent for this purpose. The housings seem to all be standardized, and have the same mounting configurations (bolt-hole placement, etc.). Each housing has a release valve, a release reservoir input, a bleed output, and a brake piston pressure output. When the valve is opened, the release reservoir (which contains air pressurized to at least 70 psi) is opened to the brake piston release valve. This pressurized air opens the brake piston release valve, which locks open and bleeds the brake cylinder, allowing the brake piston to retract over a period of about 30 seconds.

The rods themselves are made from standard 0.5" steel rod stock. Each rod is custom-bent for the particular car, because valve-mounting positions are not consistent enough for standardization. While the valve connections always requires the same amount of travel, there may be much additional travel at the rod-end due to compliance along the rod (bending due to large force/geometry moment arms).

There is little to no documentation or specifications regarding the bleed rod position, shape, or size. In order to automate a search process

as robustly as possible, the problem needs to be standardized and constrained wherever applicable. Measurements were taken with respect to all standard, documented, recognizable baselines. From this point, the data can be analyzed to discover any patterns or groupings that can be utilized.

Table 14.1: Features of brake release valve

Topics	Feature Descriptions	Impact On Automation
1. Locations	Bleed rods are located in two primary positions, avg. 56" from the ground and 3.9' from the pin-pulling lever (hoppers), or avg. 27" from the ground and 28' from the lever. See Figure 2.5.	A pair of cameras is needed at each average height, as the vertical separation is too great for a single frame at close distance. The center and each end of the cars are all that need to be searched carefully when finding the rod. Avg. rod depth from rail is 29.6" for all positions
2. Forces	Horizontal force required to fully displace the valve handle, thereby opening the valve. All service valves have the same specifications with regard to bleed rod operation.	24 lbf consistently begins to operate the valve handle; the valve appears to be adequately displaced by 30 lbf, and displaced to its fullest extent by 34 lbf.
3. Internal Bending Moment	Non-straight bleed rods bend when pressed. The amount of bending depends on this built-in bending moment. Bending can cause an additional travel requirement, and can cause the rod to move up or down under load.	The bending moment arms were not measured, but they are believed to be less than $\pm 2'$, requiring as much as a 17° upward or downward incline to resist rotation and eliminate any scraping on the guide plate opening. See Section 10.2.1.1.

Continued on next page

Table 14.1 – *Continued from previous page*

Topics	Feature Descriptions	Impact On Automation
4. Motion Range	Total distance the rod end must travel for the valve to be fully operated	The required travel ranges from 0.5" to 7.5". 0.5" is the travel at the valve handle, assuming zero deformation in the rod or connection. If improperly mounted, the guide plate can impede full travel. See Section 10.2.1.
6. Worn Notch	Most likely complication; in transit, a notch can be worn into the rod, where it rests on the guide plate. The notch can get stuck during rod operation.	A stuck-notch occurrence should be detectable (Table 10.1), and will require a different operating scheme (see Section 10.2.3).
7. Reservoir Pressure	The air reservoir that releases air upon valve operation loses pressure over time. Without pressure, even if the valve is opened, the brake won't be released.	The reservoir is at ~85 psi when fully pressurized. Testing is needed to determine the lowest functional pressure.

14.3 Laboratory Test-Bed Modifications

The valve acquired from the train yard in San Antonio is the standard service valve used in most train cars. The bleed rod is attached to the service valve via the release valve handle. On the train car the service valve is attached

to all the other components via the pipe bracket. When the bleed rod is pressed, air is released from a pressure reservoir and a loud bleed off sound can be heard. This sound will later be used to verify that the rod has, indeed, been pressed and the brake release accomplished. The various rod designs are needed to test the range of conditions observed in the San Antonio yard. However the desired sound, which emulates the sound heard in the train yard, is a loud, extended, bleed-off noise at about the volume of the short initial burst that was obtained.

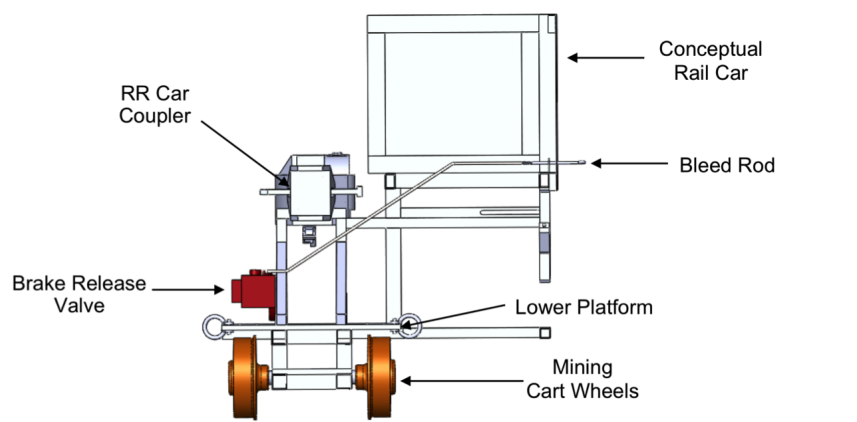


Figure 14.1: Left View (SolidWorks) of bleed rod mockup

14.4 Robot Modifications

The bleed rod requires only a linear force (30 lbf is taken to be the nominal required force), in most cases, applied in a direction directly perpendicular to the rails. The simplest contact piece that can be used is a flat plate. In the field-implemented version, a Force/Torque sensor should

be used, possibly in conjunction with a more complex contact plate geometry (shown in Figure 4.1 and discussed in Section 4.1).

The camera mount design utilizes the camera mount and securing clamps from the APPS. A few modifications were made to the original design (original design consists only of the piece into which the camera directly fits, labeled “mount” in Figure 14.2 to better fit this new application. The thickness of the back face was increased to allow a through hole for the bolt that would secure the mount to the mounting frame (also shown in Figure 14.2) and serve as the pivot point for camera the angle variability. The recessed bolt holes on the back face were then used to attach a T-strut to the camera mount that was used to lock the camera mount in the desired orientation. The mounting frame was designed to hold two cameras facing in opposite directions to allow it to still detect the brake rod while moving away from it in case it was unable to detect it while moving towards it due to an obstruction. The mount can set the cameras in 5° increments spanning from parallel, 0° , to perpendicular, 90° .

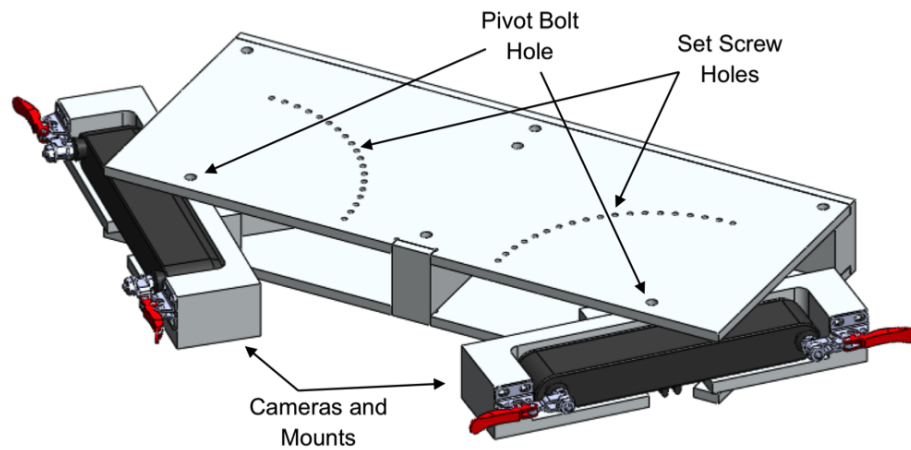


Figure 14.2: Cameras and mounts placed in dual camera mounting frame

14.5 Robot/Vision Hardware Requirements

14.5.1 Mobile Platform

There are two primary rod locations on rail cars, and the two heights around which the types are centered must be planned for with the hardware. Height is the only changing factor that needs to be addressed with hardware; the change in vertical position (27" - 56") will be accounted for by the mobile platform and its operating scheme.

The varying height is critical because it results in a change in moment arm of the actuating force (see Figure 14.3). A higher rod, with the same force applied, produces a greater moment about the robot platform wheels (h_{ef}), and requires both more powerful actuators and a more stable base. The further a point force is applied from a robot's center, the greater the torque required ($F \times d = \tau$).

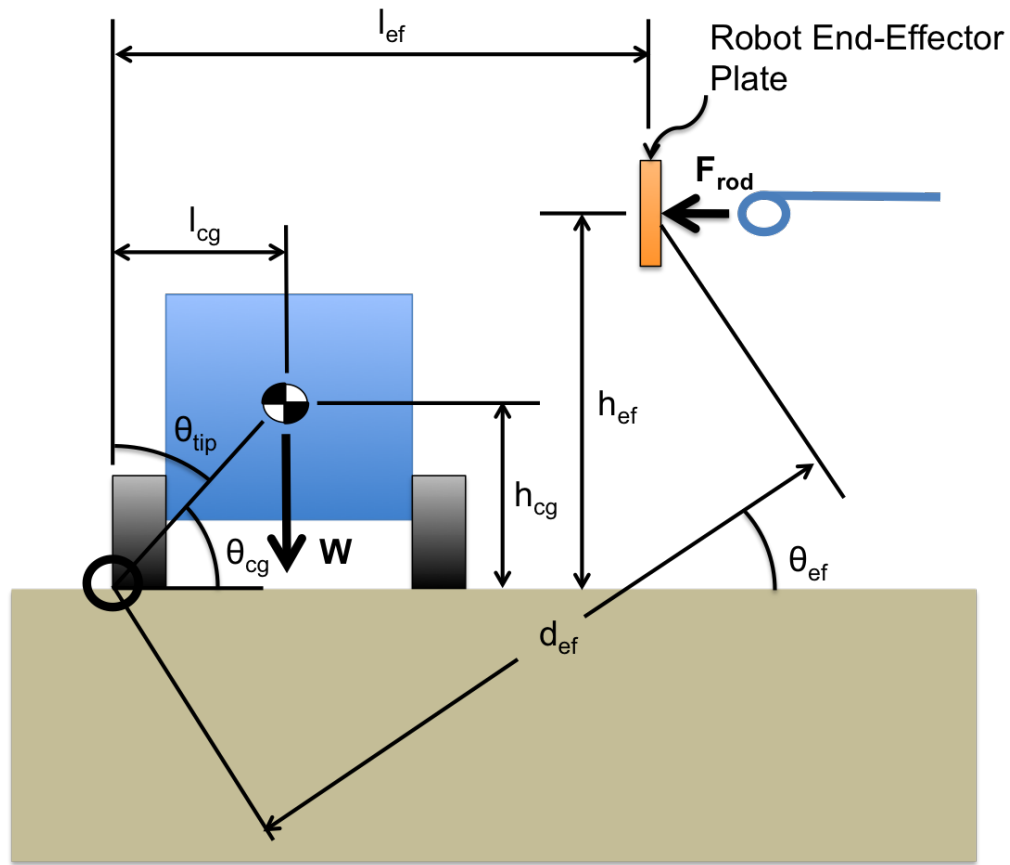


Figure 14.3: Mobile manipulator platform loads at different rod heights

14.5.2 Vision Hardware

Originally, it was concluded that cameras facing away from each other would be ideal, due to possible interference between the two cameras. However, testing has shown no degradation of rod visibility with cameras looking at it from opposite sides simultaneously. The primary benefit offered by this configuration is a constant view of the rod in all scenarios. If there are no visual obstructions, a coordinate will be received from both cameras. If there

is an obstruction from either side, the other camera will still have a view of the rod up until the end-effector moves forward into view.

Another configuration that was considered uses cameras facing away from each other, but has an additional camera mounted to the end-effector. This additional camera would either be mounted above the contact plate, looking down, or below the contact plate, looking up. Neither of these mounting positions is feasible, however, because there are some cars with features that would collide with a camera at either of these positions.

The potential for detection of the bleed rod depends entirely upon the view of the rod that the camera sees. The easily distinguishable side view must be directly in view, and that image must be sufficiently clear. Less than ideal images can potentially still be used for detection, but, once degraded past a certain point, proper detection becomes impossible. A useful workspace can be defined for camera position, relative to the bleed rod, such that detection is possible anywhere within the workspace. Additionally, a metric for detection potential has been defined to aid in optimizing camera position.

The metric takes into account detection potential due to both viewing angle and viewing distance. It uses a simple linear weighting, where $performance = 0.5 S/N_ratio + 0.5 projection_ratio$ (see Figure 14.4, and Section 5.2 for further explanation).

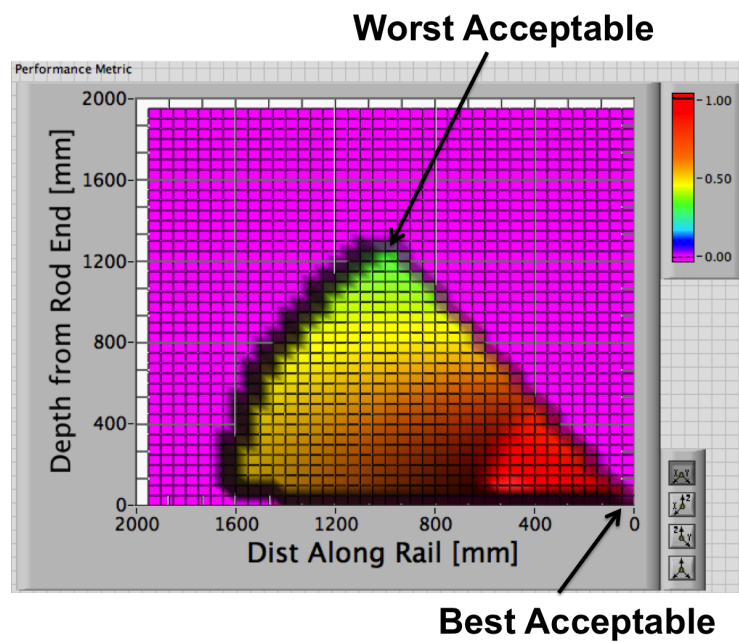


Figure 14.4: Potential for rod detection (scaled 0 - 1); camera is at the origin, and locations on the graph represent rod locations; non-zero values represent feasible rod locations

Table 14.2: Required Vision System Features

Topics	Feature Descriptions	Impact On Automation
1. Depth Camera	3D camera, creates a point cloud that can be analyzed in a variety of ways to interpret the surroundings	ASUS Xtion (based on the Xbox Kinect) has been used for development, but does not work in direct sunlight; new “Kinect for Xbox One” fuses structured IR light techniques with Time-of-Flight technology to improve precision and robustness to ambient IR conditions
2. Camera Framerate	Rate at which the camera captures images (frames), in frames/second	The standard frame rate of 30 fps is sufficient for an MMP; A wayside camera, searching rail cars at ~40mph, needs higher frame rates. Alternatively, more expensive cameras, can be used, or multiple inexpensive cameras could likely be synchronized and fused together.

Continued on next page

Table 14.2 – *Continued from previous page*

Topics	Feature Descriptions	Impact On Automation
3. Effective Resolution	Minimum width of object that can be detected; lower than actual pixel resolution, likely due to internal interpolation methods	Determines the maximum depth from which an object (especially a narrow object) can be detected. ASUS Xtion allows detection from up to appx. 5.25' (1600 mm) away. See Section 5.2.4.
4. On-Board Processing Power	Amount of RAM, processor speed, and capability of graphics card of MMP's on-board computer	Performing on-line object detection in real-time is highly memory and processing intensive; an MMP will need 2, possibly 4, cameras running simultaneously (Section 6.1).
5. Camera Dist From Rail	Distance away from the rail that a camera is mounted; must be close enough to nearly be perpendicular to the rod, but outside the red zone, where it would be in danger of collision with a car.	Wayside camera suites are ~3.5 ft from the sides of cars; Figure 14.4 shows that this depth isn't feasible. An MMP can change depth from rail, and should do so according to this detection potential map (details in Section 5.2). Recommended depth from the side of the cars is no more than 2'.

Continued on next page

Table 14.2 – Continued from previous page

Topics	Feature Descriptions	Impact On Automation
6. Camera Mount Angle	Angle at which the camera is mounted, relative to the rail, either at a wayside suite or on an MMP	Ideal mounting angle for the ASUS Xtion camera is 59° from perpendicular to the rail (θ in Figure 5.2). The mounting angle should be chosen to maximize the range of distances along the rail where the detection is possible, within the horizontal field of view (58°)
7. Current Performance	Detection algorithm performance is given as $\frac{\text{true positive detections}}{\text{all positive detections}}$, and has been evaluated using videos taken in the San Antonio yard, walking along the cars, holding the camera.	Current performance is 99.95% . Each frame will have up to 30 or more detections overlapped on a rod image, so, even 0.05% <i>false</i> positives is not enough to trigger an incorrect rod recognition for the frame (see Chapter 7).

14.6 Primary Operational Steps

The steps of operation will be the same for every car (see Chapter 6). They are the following:

1. Bleed Rod Detection
2. Bleed Rod Alignment
3. Approaching the Rod
4. Bleed Rod Actuation
5. Valve-Release Confirmation

These steps are based on assumptions that the conclusions from the San Antonio Yard visit (relatively small set of cars) hold true for the entire population (within reason) of rail cars. The key assumptions are the following:

1. Every *non-broken* bleed rod can be operated by being pushed
2. Every bleed rod-end can be seen clearly from at least one of the two sides
3. All bleed rod-ends lie within one of the two measured position distributions (shown in Figure 2.5)
4. All bleed rod-ends lie within the measured distribution of the rail (horizontal distance from closest rail)

A pair of cameras will be mounted at the expected rod height. Each pair of cameras will have one camera pointing forward (along the length of the cars, in the direction the robot is moving) and one camera point along the length of the cars the opposite direction (Figure 6.1). From the bleed rods seen at the San Antonio Yard visit (yard visit described in Chapter 2, bleed rod images shown in Figure 7.1), every bleed rod end is visible from one side or the other (only one side ever has an impeded view, shown in Figure 6.2).

Existing wayside sensor suites already are set up adjacent to rails and can be outfitted with pairs of depth cameras to serve as a rod locator. Because positions and speed are already known at the time the train passes the wayside sensors, the uncertainty associated with ground vehicle position tracking is not an issue.

Not all yards have Wayside Sensor Suites (WSS) installed, and there is no guarantee that the wayside sensors will be properly operating at all times. Wayside-mounted depth cameras would provide invaluable information that could improve both reliability and speed of the operation, but the ABRS should not be dependent upon them (see Section 6.1.3 for more detail).

In a blind search (no start-of-car reference like wayside sensing), the robot does not know where it is relative to the front end of the car. The robot also has no way of knowing whether it should be looking for the high region or the low region (shown in Figure 14.5). In this case, two pairs of cameras will be mounted, one pair centered at each of the two possible bleed rod heights, and both will be checked at all times. This will put a significantly greater

processing burden on the on-board computer than if only one pair of cameras was being monitored.

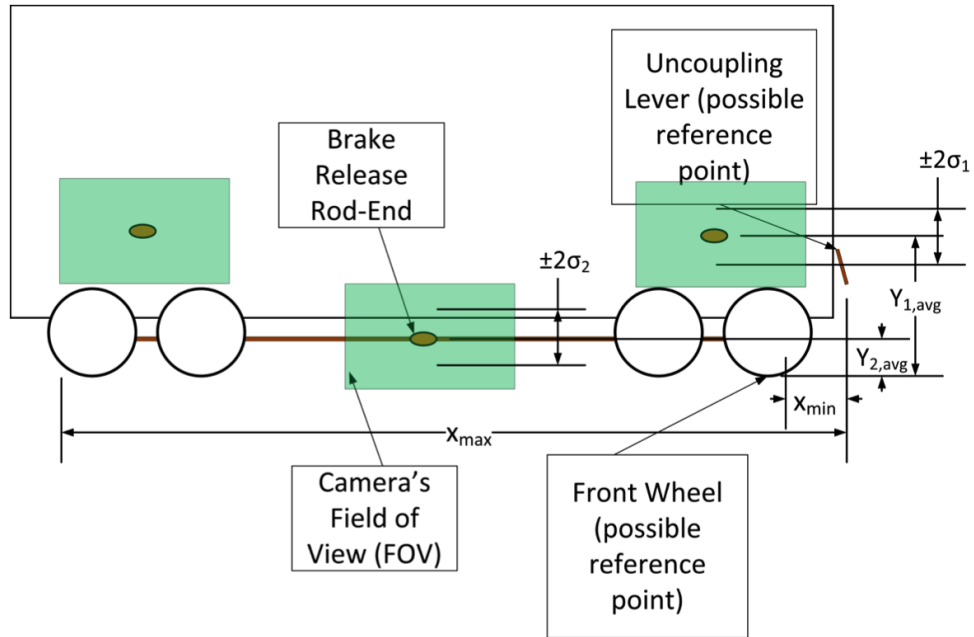


Figure 14.5: Rod location schematic

The two primary search regions are known on each car with respect to its leading edge. Using each car's leading edge as a reference point would provide for a more intelligent search scheme. However, it is not completely necessary and would just serve to make the operation more deterministic, and to decrease the processing load on the on-board computer. If AEI tags were fully reliable, then the robot could expect to know the number of axles

per car and the car length, both of which would provide knowledge about a start-of-car reference point. However, there is a possibility that the AEI tags neither reliably contain the correct information nor reliably work properly.

Other reference methods such as wheel detection and uncoupling lever detection should be considered. Both have been developed and implemented in the APPS [3].

The primary procedure includes: coordinate determination relative to robot, robot approach, robot alignment, and bleed rod actuation. Based on the loosely constrained tests performed during the San Antonio yard visit (see Section 2.4.1), a linear path, without any fine control scheme, will successfully operate all bleed rods. Because some rods have notches worn into them (Figure 14.6), a slight upward angle for the end-effector motion is necessary to prevent failures in those cases. A slight upward angle was tested on non-flawed rods as well (approximately the same angle), and the operation appeared to perform just as well as with a horizontal path. However, if the upward angle does not operate the notched rod successfully (has occurred in the lab, not in the yard, as a result of the lab's notch shape, which may or may not occur in reality), a second operating scheme is required; this alternate operation is described in "Case 2" of Section 10.2.3.

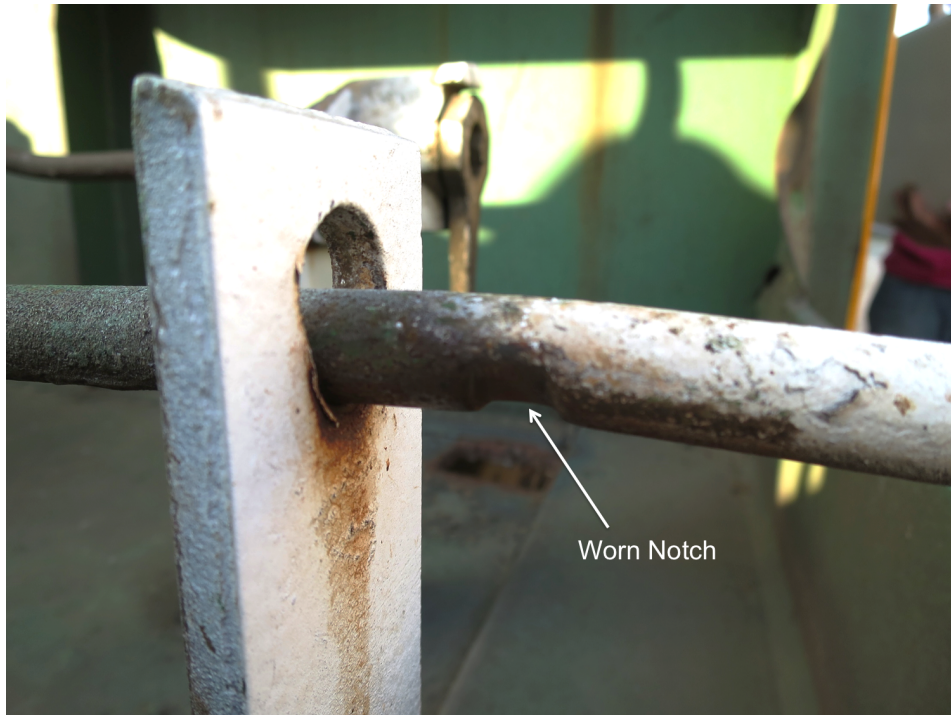


Figure 14.6: Notched rod

14.7 Bleed Rod Detection Algorithm

The brake-release task first requires detection of the bleed rod. Significant development went into the Neural-Network-based object detection algorithm used in the APPS (Chapter 6), and, since the detection tasks are similar, this algorithm was modified and used as a starting point for the Bleed Rod detection algorithm. The primary function of the detection algorithm uses an Artificial Neural Network (ANN) to scan an image and find a learned object. The entire algorithm and its background are explained in detail in Chapter 6.

The first task, after a rough camera position/orientation was chosen, was to gather training images and train a new ANN. This process, as it was done for the uncoupling lever in the APPS, is detailed in [3], Section 6.3.3 - Artificial Neural Network (ANN). After adequate bleed rod training images were gathered and used for training, the detection was tested on recorded videos from the San Antonio yard visit to establish a performance baseline. From here, additional training images were taken, both True images (bleed rod images) and False images (images of objects frequently giving false-positive detections), and the ANN was retrained to improve performance. At the point of the final algorithm performance test, 99.65% of detections in the videos taken at the San Antonio yard visit were true positives. Detections refer to a single window position within a frame (a rod may have 20 overlapped detections in a single frame). See Chapter 7 for detail about the algorithm and its adaptation for the brake rod task.

14.8 Robot Alignment/Positioning

Once the bleed rod has been found, its position needs to be determined so the robotic arm can align itself for actuation. The specific alignment position (of mobile platform and of robot pose) dictates the reliability of the operation and size of the operation workspace within which the operation can take place.

When the robot platform (or mobile platform) is stationary, the problem is greatly simplified. In this case, the only factors in play are

camera detection accuracy and robot alignment accuracy. The accuracy of the alignment to the detected coordinates is dependent solely upon the potential error in the robotic arm's motion, given that the base is sufficiently grounded during the alignment. The accuracy will vary between robotic arms, but any error should be well within the allowable contact offset (depends largely upon the end-effector and manipulator; expect value is $\sim \pm 2$ "), described in Section 9.2.1.

If the robot is in motion, additional factors must be accounted for. Ground inconsistencies will cause a mobile platform to pitch and roll, possibly so much so that the effects on the cameras and/or other sensors will need to be accounted for. A loose ground surface will also lead to wheel slip for a mobile platform, which would increase uncertainty associated with dead-reckoning. With motion between camera frames, latency associated with a filter must be accounted for (a 5-point moving average is used to smooth any noise in the detections).

In the lab test-bed, the cameras are always assumed to be horizontal (not angled up or down). However, a mobile platform may end up on uneven terrain, causing it to pitch or roll. Without on-board sensing to account for this, these variations will result in a wider spread of determined "rod coordinates." When determining a global coordinate for the rod, large variations in detected lateral position of the rod can be accounted for using the change in robot position (robot is assumed to be moving parallel to the rails). However, a wheeled platform will likely slip considerably, so its wheel rotations

can not be assumed to directly correspond to lateral motion. This slip must be sensed (using on-board accelerometers, encoders, etc.) and accounted for in an internal model as well as possible to retain accurate position estimation.

However, minimizing the wheel slip as much as possible is always the best option. This requires real-time slip detection and torque-vectoring to prevent excess slip. These methods and the required hardware to make them possible are described in detail in Appendix B in Chapters 7 and 8.

The object detection algorithm uses a moving average to filter the detected coordinates. There are 33ms between each frame; for only a 5-point moving average, that is a 165ms delay between the oldest saved coordinate and the time the transmission of the filtered coordinate is initiated. Additionally, if the mobile robot has been moving the whole time, those saved values are spanning a range that may or may not follow a normal distribution. Impact shocks from a moving platform could cause significant distortion in the averages, forcing the platform to move at lower speeds for more accurate performance. Shock effects are addressed in Chapter 13.

14.9 Initial Rod Contact & Rod Actuation

The force/torque sensor will feel a force as soon as the end-effector contacts the rod. When contact is made, there will be a force level where the rod will begin to deform prior to the valve beginning to open. The threshold level will need to be sufficiently large to avoid sensor noise triggering

a false-positive contact, and sufficiently small to be detected well before the valve begins to move. The threshold could fall anywhere in the range $\sim 0.12 - 20$ lbf, but should be as low as possible without allowing noise-induced triggerings.

Because this operation takes place with both the robot and the railcar stationary, and the rod has compliance in the direction of the approach (as it is also the direction of actuation), the safety concerns during this stage are minimal. Issues that would arise here would be due to an unseen obstacle (this is very unlikely), or making contact with the wrong object, which is very unlikely given a robust detection algorithm.

Testing using the lab test-bed has confirmed that a purely linear motion can reliably operate the bleed rod, in the absence of rod abnormalities. A pneumatic cylinder was used in the test-bed test to verify linear actuation. A linear actuator of any sort provides a very simple means of performing a linear operation, with just on/off functionality. One significant advantage of a pneumatic (or hydraulic) cylinder is that it is force-based, not position-based. This means that it will not apply more than a specified force, even if it doesn't reach full extension. Primary disadvantages are weight and the necessity of compressed air. Neither of these are issues in the lab environment, but they will be for a mobile robot.

Given that a simple, linear motion has proven to be successful, a robotic arm performing the motion would require minimal complexity and/or real-time control, making it very robust and repeatable. For a mobile robot solution, performing the rod actuation with the robotic arm doing the manipulation

may be the better of the two options. It minimizes the required maximum payload of the mobile manipulator, eliminates the need for on-board pneumatic equipment, and, therefore, minimizes the space and power requirements as well.

14.10 Valve-Release Confirmation

The only considered technique that has been deemed a viable brake release confirmation is the Auditory Recognition. However, all of the considered methods (sound, force monitoring, position monitoring) combined can provide a good deal of information and insight for fault diagnosis. The possible fault cases described in Section 11.1 are compiled in Table 10.1, along with information states associated with each.

There is a clearly defined audible blow-off sound when the valve is opened. Before the San Antonio Yard visit, the only information available regarding this sound was that operators know the valve is released as soon as the sound is heard. This indicator gives no intermediate technical information, making the operation unobservable in the interim; according to operators, this the only fully reliable method of confirmation.

We anticipated a distinct force profile during the valve operation, on which a specific point could be found as a confirmation. However, as the lack or presence of air pressure has no effect on the required operating force, force-monitoring would not be useful as a release confirmation.

The bleed-off sound heard in the yard (that has been recreated in the lab) is significantly louder than other ambient yard sounds and has a distinct sound signature. The key features of any signal are amplitude, frequency, and waveform. Additional, more abstract, sound features may be extracted for increased robustness; the standard method for using more abstract features for audio classification are briefly described in Section 11.2.4. All of these features are useful for identifying this particular noise.

For the brake release detection, a directional microphone is ideal, as the location of the source of the noise is relatively well-known and consistent. However, for other applications requiring sensing in all directions, such as improved awareness for localization or accident danger avoidance, an omnidirectional microphone would be better suited

The first step in the audio filtering (spectrum analysis) is targeting the expected frequency range. This was done initially using a 3rd order Butterworth bandpass filter, with cutoff frequencies set to 5500 Hz and 7000 Hz. Additional testing should be done using different types of signal recognition filters, as something other than Butterworth may be better suited to this application.

Once the signal has run through the bandpass filter, all frequencies outside the desired range are attenuated, so the amplitude of the signal doesn't require further processing. An amplitude threshold is set; so if the amplitude is too low, the signal is rejected as a valve-release confirmation.

The false-positive spikes from the recording have been attributed to metal-on-metal impacts, for instance, the rod hitting the side of the opening in the guide plate. Similar spikes have been reproduced in the lab by hitting metal together (wrench on steel square tube). Both the amplitude and frequency of these sounds drop outside the detected range quickly. Because the ABRs is going to hold bleed rods in for 3 seconds, there will be no important information lost by requiring a minimum duration measure.

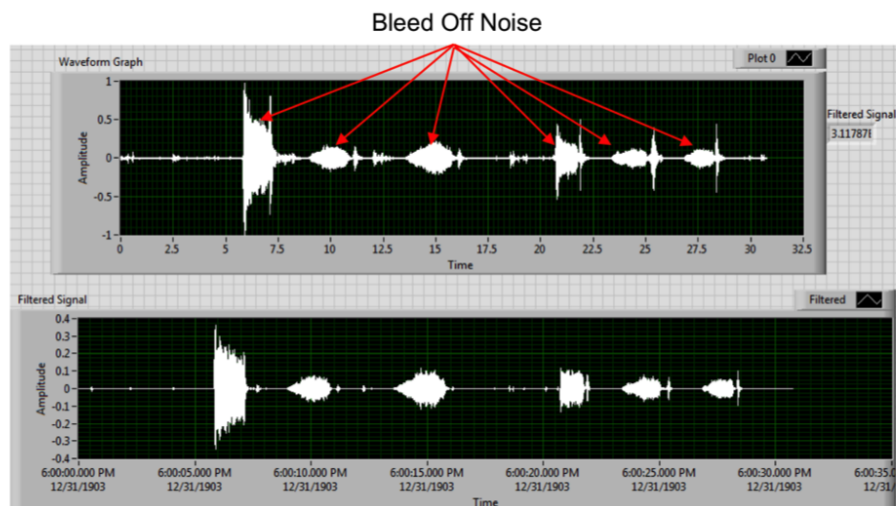


Figure 14.7: Audio signal before (top) and after (bottom) bandpass filter

The bandpass filter uses a sample rate of 44,100 Samples/second/channel (2 channels, 16 bits per sample per channel) and 4000 samples/channel for each sample set, which yields 0.0907 s (90.7 ms) per sample set. This means that 90.7 ms is the resolution of the detection algorithm. The minimum allowable number of detected sample sets has been set to 4, which is a duration of 362.8 ms. The duration of 362.8 ms, then,

is the minimum blow-off duration to reliably omit false-positives. Setting a minimum detection time removed the false-positive spikes, as can be seen in the “after” of Figure 11.5. The size of the sample set could likely be decreased from 4000 samples/channel, but must be done carefully, as too great a decrease proved detrimental in lab testing.

The only sound produced in the lab that has yet to be filtered out as a false-positive is metal-on-metal banging that causes resonance. If the square tube is hit in the right place, the sound resonates loudly enough, and at a close enough frequency, that it is picked up for multiple detection samples. If the square tube is hit repeatedly to create resonance, the vibration sound remains continuous, and the detection program may see it as a blow-off sound.

This is a very particular case that, as far as UP knows, is not replicated in the yard by anything within close proximity of the service valve. If either this is incorrect, or if additional robustness is desired, there are options for future development.

When checking for the valve-release confirmation, to obtain the best results from the confirmation algorithm, the audio signal should be as undistorted as possible. The position of the valve relative to the microphone and the presence of obstructions could interfere with the audio signal, causing the reliability of the audio confirmation algorithm to degrade. In order to maximize the reliability, the microphone should be positioned such that the interference and dampening due to obstructions and/or location of the valve are minimized.

A series of experiments were performed to examine the effect microphone positioning has on the robot's ability to gain auditory confirmation of the brake release using the current auditory detection algorithm.

When the microphone was moved a small amount from the valve, the duration was not affected. When a relatively small object was placed more than 2' from the microphone there is still no effect on the detection duration. The thin-walled structure only affected the detection duration when placed less than a foot away from the microphone; the detection duration was halved when the thin-walled structure was placed 6" from the microphone. When the larger obstruction, the concrete cylinder, was placed between the microphone and valve the detection duration was reduced by a factor of 6, see test 6. Reducing the distance between the microphone and the valve, while holding the position of the obstruction relative to the valve constant, only increased the detection duration by 1 second.

Table 14.3: Sound Recognition Features

Topics	Feature Descriptions	Impact On Automation
1. Microphone Line-of-Sight	Straght line between the on-board microphone and the valve	A clear path enables the developed detection method to sense the distinct sound down to 47 psi reservoir pressure (nominal is 85 psi). An impeded line-of-sight can make the sounded currently undetectable as high as 71 psi (see Section 11.3 for application examples)
2. Primary Frequencies	Every sound/signal has many frequencies, each of a different magnitude. The primary frequencies can be extracted (spectral analysis).	The blow-off sound, at the reservoir's high pressures, reliably has its primary frequencies between 5.5 kHz and 7 kHz. A bandpass filter can attenuate frequencies outside of this range, leaving only the relevant frequencies at a high amplitude (see Section 11.2.1).

Continued on next page

Table 14.3 – *Continued from previous page*

Topics	Feature Descriptions	Impact On Automation
3. Primary Frequency Duration	Amount of time that primary frequencies in the expected range last; frequency and amplitude change with decreasing reservoir pressure	Metal-on-metal contact can produce the expected frequency range at high amplitude, but only for a brief time; checking for a minimum duration (~360 ms with current algorithm) helps eliminate false-positive detections (see Section 11.2.3)
4. Sound Signature Classification	Features can be extracted from a signal (concrete and abstract), and used as inputs to a trained classifier (Neural Network, Support Vector Machine, etc.)	Standard techniques can be used (have not yet been tested) to robustly classify sounds, in a more descriptive, detailed way, similar to human learning. (Section 11.2.4)
5. Noise Rejection	Detection of sounds that are expected, but not specifically of interest, enables rejection as possible sounds of interest.	A Rail yard has many sounds (clanging, pneumatic blow-offs, wheels scraping, braking, etc.), most of which are expected and have repeatable sound signatures. Detecting these sounds would improve valve blow-off detection reliability by diminishing false-positives (Section 11.2.5)

Continued on next page

Table 14.3 – *Continued from previous page*

Topics	Feature Descriptions	Impact On Automation
6. Directional Microphone	Attenuates sound as its source moves further from the microphone's designed direction (Figure 11.2)	Directional microphones can be used, mounted to point toward the valve (behind the rod-end), to decrease the possibility of non-relevant noise interference (Section 11.2)

14.11 Adaptation for Mobile Manipulator

There are a few anticipated issues will be addressed that apply to the Automated Brake Release System (ABRS) when using a mobile manipulator. Some things discussed here are critical parameters that should drive hardware selection, and problems expected to arise when testing outside of the controlled laboratory environment.

A manipulator will have a required range of positions it must reach for any specific task, defining the minimum required workspace. The workspace of a robotic arm is defined by its configuration and the geometry of its links. Requisite workspaces should overlap between tasks that are intended to be performed by the same mobile robot.

However, if the arms used on the mobile platforms are all fully modular, reconfiguration can allow the same hardware to fulfill unique workspace requirements. In this case, if two versions of mobile robots in a yard have

different workspace requirements, the same set of parts (system module) can still be used.

For a set of tasks to be performed by one mobile robot, the end-effectors required should have as much functional overlap as possible. Components that can be expected to be necessary in most tasks are a collision sensor and a force/torque sensor (see Sections 8.1.1-8.1.3 in [3] for details about their operation, and Chapter 9 for details about how they apply to the ABRS). Interfaces at the end of these components should give study towards standardization for hardware for each task (gripper, vision sensor, etc). Without any custom modifications, only one of the commercially off-the-shelf (COTS) mobile robots presented in [8] partially meets the stability criteria outlined in Section 5.1.2 – the Seekur. Table 12.1 shows the relevant geometric parameters for each mobile robot, and the results of their expected stabilities.

The attributes of a switch yard that directly affect mobile platform maneuverability (wheel DOF, wheel height, wheel clogs, body clearance, etc.) are explained in detail in Section V of Appendix B. Steering is one of the important aspects of mobile platforms. Here, a comparison of common steering methods: skid steering, Ackerman steering and independent wheel steering is discussed.

Skid-steer (a.k.a. differential steering) is imprecise, inefficient, and prone to failure. Skid-steer is typically used either in educational environments (simple, low-cost) or in applications that very rarely require turning (similar to the justification for un-steered wheels on 18-wheeler bogies). The only justified

case for this steering method would be for a mobile platform that intends to move in a straight line with limited turning. Otherwise, skid-steered mobile platforms should be avoided.

Ackerman steering uses Front-Wheel Steering (FWS) and is based on the idea of instantaneous center of rotation. In any turn, all wheel axle centerlines must pass through the instantaneous center of rotation. For a left turn, then, the left wheels will be closer to the instant center than the right wheels, which means a smaller turn radius on that side, and a greater angle between front and back wheels.

Ackerman steering could be useful for a lower cost, energy efficient mobile platform, meant for tasks that require minimal multi-dimensional adjustment and can have a largely planned out path in advance (that has only small, large-radius turns).

Independent Steered Wheels (independent all-wheel steering is what is referred to here) utilize the instantaneous center principle as well, but do not have the same constraints inherent to Ackerman Steering. In IWS vehicles, any steering angle can independently be reached for all 4 wheels. This then decouples the vehicle's rotational capabilities from its forward or sideways motion. Being de-coupled and unconstrained, it can now turn in place, about its own center, as well as move linearly in any direction.

Independent Wheel Steering (IWS) is necessary for any mobile platform that requires a complex, un-planned path in a complex environment. Most

tasks to be performed with a mobile robot in a switch yard will fall under this category. Both independent IWS and Ackerman Steering minimize lateral wheel slip (as it is not inherently necessary) and are, therefore, far more precise than any skid-steer system.

Any on-board sensors that need to collect data while the mobile platform is moving must be resistant to the operating conditions. As mentioned earlier, impact shocks can significantly affect the performance of the vision system. Solutions must be low cost and easily replaced.

Vibrations will occur in the presence of any rough contact interface. Grousers (clogs), standard tire treads, and gravel surfaces represent non-continuities that can be expected. Each of these will cause a vibration of varying frequency and amplitude, and should all be examined closely at expected operating speeds to anticipate possibilities of induced sensor error.

If each grouser pulse takes place during the sensing/imaging operation, the data will deteriorate. Successive samples from the camera will show inherent noise over time. It is possible that this noise could be generically filtered to produce adequate data. However, aliasing (if the sampling rate is not significantly greater than the frequency of the greatest oscillation, the signal may appear to follow a trend it does not actually follow) may be possible, and if the data is needed for use in real-time, then requiring a filter is unacceptable.

One common solution to vibrational noise is hardware isolation, through a spring-damper suspension. This acts as a mechanical low-pass filter.

The inherent problem with this option is a delay in propagation of position changes. If the camera always has to be rigidly connected/calibrated to the mobile platform, then this is not a viable option. If, however, the primary goal is always getting a good image or good data, then this may be a good option. Instead of solving the vibration problem mechanically, post-processing the data can also rectify the problem. If the post-processing method is fast enough, real-time use of the data may be possible. Using accelerometers to programmatically cancel out undesired motion is also feasible.

Table 14.4: Mobile Manipulator Features

Topics	Feature Descriptions	Impact On Automation
1. Maximum Payload	Defined as the greatest load mass at the end-effector with which a manipulator can move about its entire workspace with the rated maximum actuator torques and velocities	Should be high enough to account for end-effector weight plus the required operating force. Max payload can be exceeded, as long as servo drive current limits are not exceeded (Section 10.1.3.1)
2. Workspace Size	Workspace is defined as the space containing all points the end-effector can reach	Must be capable of reaching both rod heights, and applying full extension at both heights (see Section 12.1.1.1, and Section 4.2 of [8])
3. Degrees of Freedom (DOF)	Number of independently controllable joints (either rotational or linear) of a manipulator	6 DOF allows the end-effector to reach any position in any orientation. Even if only 4 or 5 DOF is needed, it is often more expensive as a non-standard. more than 6 DOF allows multiple joint configurations to provide the same end-effector position/orientation, which offers multiple real-time options

Continued on next page

Table 14.4 – *Continued from previous page*

Topics	Feature Descriptions	Impact On Automation
4. Power Requirement	Amount of energy required for the desired task, and required means of getting power	A minimum power requirement allows more functional use between recharges, as well as smaller on-board power requirement; DC power source compatibility is ideal, so the drives can be run directly from on-board batteries.
5. Servo Drive Configuration	Drives to power the manipulator joints can be external, in various sizes, or internal.	External drives are typically too large to be on-board a mobile platform, and would require tethering; some external drive boxes could fit on an MMP, but use much of the available real-estate. Manipulators with internal drives (e.g. Schunk LWA 4D) are ideal for a fully mobile solution.

Continued on next page

Table 14.4 – *Continued from previous page*

Topics	Feature Descriptions	Impact On Automation
6. Modular Design	Fully interchangeable actuators and linkages, with standardized interfaces allow on-demand reconfiguration and repair with a minimum set of parts	Different MMPs may have slightly varying manipulator requirements, all of which could be met with the a single set of parts. Maintenance becomes replacing a component, rather than requiring the cost and time of replacing the entire robot.

Bibliography

- [1] 2014 (accessed December 23, 2014). URL: <http://www.auto-ware.com/setup/ackrman2.gif>.
- [2] Pradeepkumar Ashok and Delbert Tesar. “A Visualization Framework for Real Time Decision Making in a Multi-Input Multi-Output System”. In: *IEEE Systems Journal* 2.1 (2008).
- [3] P. Ashok et al. *Pin-Pulling Robot Development for Union Pacific Application*. 2014.
- [4] Michael Cowling and Renate Sitte. “Comparison of techniques for environmental sound recognition”. In: *Pattern Recognition Letters* (2003), pp. 2895–2907.
- [5] *DB-60 Control Valve Instruction Manual MU-21*. New York Air Brake Corporation. 746 Starbuck Ave. Watertown, New York 13601, Feb. 2000.
- [6] Richard S. Goldhor. “Recognition of Environmental Sounds”. In: *IEEE* 1 (1993), pp. 149–152.
- [7] Dr. Delbert Tesar. “Mobile Platform / Manipulator Development / Switchyard Operations (Emphasis on Navigation/Task Performance)”. Sept. 2014.
- [8] Timothy Woodard, Delbert Tesar, and Christopher Hammel. “Union Pacific Brake Release Mobile Platform and Manipulator Arm Specification”. Aug. 2014.
- [9] Timothy Woodard and Dr. Delbert Tesar. “Union Pacific Mobile Manipulator Platform Development”. Nov. 2014.