

**ULTRA HIGH FREQUENCY (UHF) RADIO-FREQUENCY IDENTIFICATION
(RFID) FOR ROBOT PERCEPTION AND MOBILE MANIPULATION**

A Thesis
Presented to
The Academic Faculty

by

Travis Deyle

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Electrical and Computer Engineering

Georgia Institute of Technology
December 2011

**ULTRA HIGH FREQUENCY (UHF) RADIO-FREQUENCY IDENTIFICATION
(RFID) FOR ROBOT PERCEPTION AND MOBILE MANIPULATION**

Approved by:

Dr. Charles C. Kemp, Advisor
Biomedical Engineering
Georgia Institute of Technology

Dr. Ayanna Howard
Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Mary Ann Ingram
Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Gregory Abowd
College of Computing
Georgia Institute of Technology

Dr. Emmanouil Tentzeris
Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Matt Reynolds
Electrical and Computer Engineering
Duke University

Date Approved: 6 October 2011

Dedicated to my loving wife, Fei.

ACKNOWLEDGEMENTS

I must begin by thanking my wife, Fei. It's rare to find someone as supportive, loving, and patient as my wife – she is my soul mate. She shouldered all the burdens of a driven, grad-student husband without (hardly) a complaint. I would not have completed my PhD without her support. I would also like to thank my parents. Mom and Dad, thank you for always encouraging my quirky curiosity. To my brothers, Casey and Toby, I'd like to thank you for always challenging me. You have kept me fresh and on my toes all these years; I hope I've been a good role model. I would also like to thank my grandfather. That man could build anything with his bare hands, and we had a blast building some of our early robots. He taught me many valuable lessons; may he rest in peace.

After family, my thanks must go to my advisor, Charlie Kemp. Charlie is a brilliant roboticist who always seems to intuit the “right” way to do things (sometimes to my chagrin). Charlie has molded me as a researcher and provided me with unparalleled access to robot hardware – heck, our lab had (almost) one large robot per grad student! I'm still perplexed by his innate ability to tease out salient details in any conversation. Perhaps most importantly, I want to thank Charlie for assembling the current (and former) students of the Healthcare Robotics Lab. He created one of the most synergistic and capable robotics labs in the world, and I'll forever be indebted.

I would also like to thank my mentor (and committee member), Matt Reynolds. Matt and I share an uncanny overlap in interests and thought processes. I look forward to working with him during my forthcoming postdoc at Duke University. Many thanks to my other committee members. In particular, I'd like to thank Gregory Abowd for taking time very early in my academic career to get me thinking about what I wanted to do with my future. Unfortunately, I still don't have a clue (beyond “do cool stuff and make a difference”), but at least he got the ball rolling.

During my tenure at the Healthcare Robotics Lab, I had the pleasure of working with a number of talented and thoughtful individuals: Advait Jain, Marc Killpack, Hai Nguyen, Cressel Anderson, Aaron King, Tiffany Chen, Kelsey Hawkins, Tapo Bhattacharjee, Phillip Grice, Young Sang Choi,

Alex Trevor, Zhe (Joseph) Xu, Sahbi Bellamine, and Zhengqin Fan. You guys are more than lab-mates. I consider each of you friends. In particular, my special thanks goes to: Cressel Anderson and Zhe Xu for their early help on EL-E's RFID systems; Hai Nguyen for his leading-role on PPS-Tags and supportive-role on RF Vision; Advait Jain for all the infrastructure code on EL-E; and Stewart Thomas for your RF measurements on the in-hand RFID reader – your contributions were instrumental to this thesis.

Georgia Tech is a great school, but it's really the people who made it special for me. I'd like to thank: Thad Starner for taking me under his wing when I first arrived; the Human Factors and Aging Lab (led by Wendy Rogers) for all the older-adult user-study support; Shwetak Patel for the insightful discussions; and Brian Jones for running the Aware Home. I'd like to give a shout-out to the “noon-thirty” lunch crew (Nicolas Dudebout, Francesco Barale, and Romain Pelard) for dragging me away from the computer each and every day. You guys taught an American like me to appreciate quality espresso.

My work was made possible through the generous support of the National Science Foundation (NSF) Graduate Research Fellowship Program (GRFP), NSF grant CBET-0932592, NSF grant CBET-0931924, and NSF grant IIS-0705130.

I also benefited greatly from Willow Garage, who supplied our lab with a PR2 as part of their Beta Program. Plus, in just a few short years, Willow Garage's Robot Operating System (ROS) has fostered an open-source community that is changing the landscape of robotics. This thesis would not have been possible without the many talented people with whom I anonymously collaborated via open source software. To pay it forward, all of the code and data from this thesis (in particular on the PR2) will be released as open source in Georgia Tech's ROS package, *gt-ros-pkg*.

Finally, I would like to thank the community of science fiction authors, such as: Arthur C. Clarke, Vernor Vinge, Neal Stephenson, Robert Heinlein, Isaac Asimov, etc. You have provided me with countless hours of entertainment and just as many creative ideas. Hopefully we can turn some of your musing into reality while avoiding the less palatable outcomes.

There are many other people who were vital to my maturation, as a grad student and as a person – you all have my sincere gratitude.

TABLE OF CONTENTS

DEDICATION	iii
ACKNOWLEDGEMENTS	iv
LIST OF TABLES	x
LIST OF FIGURES	xii
SUMMARY	xvi
I INTRODUCTION	1
1.1 Motivation	1
1.2 Contributions	2
1.3 Related Work	4
1.3.1 Pressing Challenges in Healthcare	4
1.3.2 Robots for Assistive Manipulation	7
1.3.3 RFID Sensing for Robotics	9
1.3.4 Competing non-RF Technologies for Robotics	13
1.4 Our Robots	16
1.5 Organization of This Work	18
II UHF RFID AND THE FRIIS RADAR EQUATION	21
2.1 The Friis Radar Equation(s)	21
2.1.1 The Forward-Link Friis Equation	21
2.1.2 The Forward-Backward-Link Friis Equation	23
2.1.3 Antenna Selection	25
2.1.4 Remarks	29
2.2 Tag Detection Model	31
2.2.1 Tag Detection Regions	33
2.2.2 Data-Driven Evaluation	34
2.3 RSSI Model	36
2.3.1 RSSI Model Fitting and Evaluation	39
2.4 Simulated RFID Sensor Model	43
2.5 Model Deviations	45

2.5.1	Environmental Conditions	45
2.5.2	Robot Properties	46
2.5.3	Tagging Objects	46
2.5.4	Tags' On-Board Energy Storage	47
2.5.5	Data-driven Examples	48
2.6	Conclusions	49
III	PROBABILISTIC APPROACHES TO LOCALIZE UHF RFID TAGS	52
3.1	Bayesian Filter Formulation	53
3.2	Related Work	53
3.3	UHF RFID Sensor Models	55
3.3.1	Data-Driven Tag Detection Regions	55
3.3.2	Tag Detection Regions Derived from the Friis Model	56
3.3.3	Data-Driven RSSI Sensor Models	63
3.3.4	Other UHF RFID Sensor Models	64
3.4	Discussion	64
3.4.1	Probabilistic Models From Data Captures Under Relatively Ideal Settings	65
3.4.2	Antenna-By-Antenna Variation	65
3.4.3	Variation Due to Tag Orientation	67
3.4.4	Variation Due to Tagged Object Material Properties	69
3.5	Conclusions	70
IV	OPTIMIZATION-BASED APPROACHES TO LOCATE AND APPROACH UHF RFID TAGS	75
4.1	Bearing Estimation (Azimuth Only)	78
4.1.1	Non-Planar Bearing Estimation	78
4.1.2	Evaluation	80
4.1.3	Extensions to Bearing Estimation	85
4.1.4	Bearing Estimation for Sub-360° Scans	86
4.2	RFID Servoing	92
4.2.1	Bearing-Only RFID Servoing	92
4.2.2	RFID Servoing	94
4.2.3	RFID Servoing in Simulation	96

4.2.4	RFID Servoing on the PR2	97
4.2.5	Discussion	102
4.3	Bearing Estimation – Azimuth (Θ) and Elevation (Φ)	103
4.3.1	Similarity to Azimuth-Only Bearing Estimation	103
4.3.2	RSSI Images	103
4.3.3	RSSI Images in Practice	105
4.4	Global RFID Search	112
4.4.1	Related Work	112
4.4.2	Exhaustive Brute-Force Search	115
4.4.3	Practical Constraints	115
4.4.4	On-Robot Implementation	117
4.5	Hybrid Global-Local RFID Search	119
4.5.1	Evaluation	119
4.5.2	Comparison to Bayesian Estimation	143
4.6	Conclusions	145
V	SHORT-RANGE UHF RFID	148
5.1	Related Work	150
5.2	Design of a Wrist-Mounted Reader	151
5.3	Design of an In-Hand Reader	154
5.4	Characterizing In-Hand Reader Detection Performance	155
5.5	Short-Range UHF RFID Robot Behaviors	161
5.5.1	Locating a Tagged Object	161
5.5.2	Short-Range RFID Servoing	164
5.5.3	Raster-Scan RSSI Images	166
5.6	Additional Robot Behavior: Verifying a Grasped Object’s Identity	167
5.7	Conclusions	171
VI	UHF RFID SENSING FOR MOBILE MANIPULATION	173
6.1	Physical, Perceptual, and Semantic (PPS) Tags	175
6.1.1	Related Work	177
6.1.2	Three Examples of PPS-Tags	178
6.1.3	Interacting with PPS-Tags	184

6.1.4	Evaluating EL-E's Interaction with PPS-tags	192
6.1.5	Extending Clickable World Using PPS-Tags	194
6.1.6	Discussion and Conclusions	200
6.2	Multi-Sensor Fusion and Mobile Manipulation Using RSSI Images	202
6.2.1	Related Work	203
6.2.2	Sensor Fusion	205
6.2.3	Sensor Fusion Evaluation	214
6.2.4	Mobile Manipulation System	216
6.2.5	Mobile Manipulation Evaluation	216
6.2.6	Discussion and Conclusions	218
6.3	UHF RFID For Medication Delivery and Adherence	220
6.3.1	UHF RFID Medication Delivery on the EL-E Platform	221
6.3.2	UHF RFID Medication Delivery on the PR2	222
6.3.3	Discussion and Conclusions	226
6.4	Conclusions	226
VII	CONCLUSIONS	228
REFERENCES	233

LIST OF TABLES

1	Particle Filter Tag Localization Results	61
2	Bearing Estimation Statistics For the Head-Mounted Antenna	82
3	Bearing Estimation Statistics For the Shoulder-Mounted Antennas	86
4	RFID search results (summary) for tagged objects at location #0.	134
5	Hybrid RFID search results (detailed) for tagged objects at location #0.	134
6	RFID search results (summary) for tagged objects at location #1.	135
7	Hybrid RFID search results (detailed) for tagged objects at location #1.	135
8	RFID search results (summary) for tagged objects at location #2.	136
9	Hybrid RFID search results (detailed) for tagged objects at location #2.	136
10	RFID search results (summary) for tagged objects at location #3.	137
11	Hybrid RFID search results (detailed) for tagged objects at location #3.	137
12	RFID search results (summary) for tagged objects at location #4.	138
13	Hybrid RFID search results (detailed) for tagged objects at location #4.	138
14	RFID search results (summary) for tagged objects at location #5.	139
15	Hybrid RFID search results (detailed) for tagged objects at location #5.	139
16	RFID search results (summary) for tagged objects at location #6.	140
17	Hybrid RFID search results (detailed) for tagged objects at location #6.	140
18	RFID search results (summary) for tagged objects at location #7.	141
19	Hybrid RFID search results (detailed) for tagged objects at location #7.	141
20	RFID search results (summary) for tagged objects at location #8.	142
21	Hybrid RFID search results (detailed) for tagged objects at location #8.	142
22	RFID Search Using Hybrid Optimization-Based Approach Versus Bayesian Localization	145
23	Tagged Object Detection for Grasped Objects – Overall Results	158
24	Tagged Object Detection for Grasped Objects (Under Various Orientations) Using In-Hand RFID Reader (1 of 2)	159
25	Tagged Object Detection for Grasped Objects (Under Various Orientations) Using In-Hand RFID Reader (2 of 2)	160
26	Experimental Evaluation Results of EL-E Operating PPS-tags	192
27	PPS-Tag Experimental Results	199

28	Sensor Fusion Evaluation Results	214
----	--	-----

LIST OF FIGURES

1	Overview: Sensor Models Vary When Tagging Objects	3
2	Overview: RFID Search in a Real Home Environment	4
3	Overview: Using UHF RFID Sensing for Robot Applications	5
4	Environmental Augmentations by Individuals with Motor Impairments	7
5	Examples of passive RFID tags	10
6	A Collection of Passive UHF RFID Tags	11
7	Commercially Available ThingMagic RFID Readers	12
8	UHF RFID Antennas	12
9	Backscatter Modulation	14
10	Visual Fiducial Technologies	15
11	Combining RFID with Visual Fiducials	16
12	Robots Used in Thesis	17
13	Friis Forward-Link Equation	22
14	Six Degree-of-Freedom (6-DoF) Relationship Between Reader Antenna and Tag	23
15	Friis Forward-Backward-Link Equation	24
16	Theoretical Dipole Radiation Pattern	27
17	Alien Technologies “Squiggle” Tag Radiation Pattern	27
18	Theoretical Patch Radiation Pattern	28
19	Patch Antennas on EL-E Robot	30
20	Tag Detection In Monostatic Mode	32
21	Tag Detection Probabilities Given Power Incident at Tag	33
22	Theoretical Tag Detection Regions	34
23	Recording Ground Truth Tag Pose	35
24	Ground-Truth Data Captures Using the Head-Mounted Antenna	37
25	Ground-Truth Data Captures Using the Shoulder-Mounted Antennas	38
26	Data-Driven Tag Detection Model Using the Friis Equation	39
27	RSSI Versus Attenuated Power for M5e Reader	40
28	Fitting to the RSSI Model	41
29	Understanding the RSSI Model	43

30	Data Capture Under Realistic Circumstances	48
31	Tag Detection Model Under Realistic Circumstances	50
32	Variation in Tag Detection	51
33	Bayesian Inference Using Data-Driven Tag Detection Regions	56
34	Friis Model Tag Detection Sensor Model	58
35	Friis Model Tag Detection Rig	59
36	Friis Model Tag Detection Results	62
37	Data-Driven RSSI Sensor Model	64
38	2D Tag Detection and RSSI Statistics (Comprehensive)	66
39	2D Tag Detection and RSSI Statistics (Differences Between Antennas)	68
40	2D Tag Detection and RSSI Statistics (Differences Between Tags)	72
41	2D Tag Detection and RSSI Statistics (Cluttered Objects 1 of 2)	73
42	2D Tag Detection and RSSI Statistics (Cluttered Objects 2 of 2)	74
43	Planar Bearing Estimation Algorithm	79
44	Bearing Estimation Using Noisy Measurements	80
45	Performing data captures using a head-mounted RFID antenna	81
46	Bearing Estimation Algorithm Individual Successes (On-Robot)	83
47	Bearing Estimation Algorithm Individual Failures (On-Robot)	84
48	Bearing Estimation Error Versus RSSI and Distance from Tag	85
49	Bearing Estimation Algorithm Individual Sub-360° Scans	87
50	Bearing Estimation Error Versus RSSI and Distance for Sub-360° Scans	88
51	Bearing Estimation by Fitting Measured RSSI to Reader Antenna Radiation Patterns	89
52	Using a Fixed Antenna, The Robot's Mobile Base Provides Actuation	90
53	Signals for RFID Servoing	95
54	Simulated RFID Servoing on the PR2	97
55	The PR2 Robot Serves to an RFID Tag	98
56	Servoing Trajectories to Approach an Isolated Tag (On-Robot)	99
57	Servoing Trajectories to Approach a Tagged Medication Bottle (On-Robot)	100
58	EL-E Servoing Results	101
59	RSSI Images of a Visible Tagged Object	104
60	RSSI Images of a Concealed Tagged Object	105

61	EL-E Robot Performing a RSSI Image Scan	106
62	RSSI Image Bayesian Sensor Model	108
63	RSSI Images Under Various Environmental Conditions (1 of 3)	109
64	RSSI Images Under Various Environmental Conditions (2 of 3)	110
65	RSSI Images Under Various Environmental Conditions (3 of 3)	111
66	RFID Search as a Function of Distance from a Tag	116
67	Limitations of Brute-Force RFID Search	117
68	Photos of Georgia Tech’s Aware Home	120
69	Tagged Objects and Locations During RFID Search Experiments	121
70	The Robot Pose After Performing Hybrid RFID Search	123
71	Distance Error Measure for RFID Search	124
72	RFID Search Results at Location 0	125
73	RFID Search Results at Location 1	126
74	RFID Search Results at Location 2	127
75	RFID Search Results at Location 3	128
76	RFID Search Results at Location 4	129
77	RFID Search Results at Location 5	130
78	RFID Search Results at Location 6	131
79	RFID Search Results at Location 7	132
80	RFID Search Results at Location 8	133
81	Probabilistic Estimation of Robot Location and Orientation using RFID Search Data	144
82	Short-Range UHF RFID Systems	149
83	EL-E Robot with a Short-Range Wrist-Mounted Antenna	151
84	EL-E Robot with a Short-Range Wrist-Mounted Antenna	153
85	EL-E Robot With In-Hand RFID Reader	155
86	In-Hand Reader Hardware Design	156
87	Short-Range Wrist-Mounted Antenna Performance	163
88	Behavior to Verify Grasped Object Identification	165
89	Short-Range RFID Servoing	167
90	Short-Range “Raster-Scan” RSSI Images	168
91	Grasped Object Verification Behavior Using In-Hand Reader	170

92	Physical, Perceptual, Semantic (PPS) Tags Being Manipulated by EL-E	176
93	Examples of Physically Assistive Devices	179
94	EL-E Robot Used For PPS-Tag Experiments	181
95	PPS-Tag Semantic Database Entry	182
96	Contextually-Aware Graphical User Interface	185
97	PPS-tag Perceptual Processing Algorithm	188
98	EL-E Detecting the Change In Lighting To Determine PPS-tag Success or Failure .	190
99	Overview of Clickable World	194
100	EL-E For PPS Tag Experiments	195
101	PPS-Tags Used for Clickable World	197
102	Perceptual Visualizations for the Trashcan PPS-Tag	198
103	3D Point Serves as a Mobile Manipulation Abstraction Layer	202
104	EL-E Used for Sensor Fusion and Mobile Manipulation	203
105	Generating a Range Feature Image From a 3D Point Cloud	207
106	Color Image Tag Detection Distributions	209
107	RSSI Image Tag Detection Distributions	210
108	Sensor Fusion Diagram	212
109	Multi-Sensor Fusion Using RSSI Images	213
110	Sensor Fusion Evaluation Objects and Locations	215
111	Bearing estimation for Sensor Fusion	217
112	RFID User Interfaces	218
113	EL-E Delivering Medication	221
114	PR2 Medication Delivery Procedure	224
115	PR2 Delivering Medication to an Older Adult	225

SUMMARY

Personal robots with autonomy, mobility, and manipulation capabilities have the potential to dramatically improve quality of life for various user populations, such as older adults and individuals with motor impairments. Unfortunately, unstructured environments present many challenges that hinder robot deployment in ordinary homes. This thesis seeks to address some of these challenges through a new robotic sensing modality that leverages a small amount of environmental augmentation in the form of Ultra High Frequency (UHF) Radio-Frequency Identification (RFID) tags. Previous research has demonstrated the utility of infrastructure tags (affixed to walls) for robot localization; in this thesis, we specifically focus on tagging *objects*. Owing to their low-cost and passive (battery-free) operation, users can apply UHF RFID tags to hundreds of objects throughout their homes. The tags provide two valuable properties for robots: a unique identifier and receive signal strength indicator (RSSI, the strength of a tag’s response). This thesis explores robot behaviors and radio frequency perception techniques using robot-mounted UHF RFID readers that enable a robot to efficiently discover, locate, and interact with UHF RFID tags applied to objects and people of interest. The behaviors and algorithms explicitly rely on the robot’s mobility and manipulation capabilities to provide multiple opportunistic views of the complex electromagnetic landscape inside a home environment.

The electromagnetic properties of RFID tags change when applied to common household objects. Objects can have varied material properties, can be placed in diverse orientations, and be relocated to completely new environments. We present a new class of optimization-based techniques for RFID sensing that are robust to the variation in tag performance caused by these complexities. We discuss a hybrid global-local search algorithm where a robot employing long-range directional antennas searches for tagged objects by maximizing expected RSSI measurements; that is, the robot attempts to position itself (1) near a desired tagged object and (2) oriented towards it. The robot first performs a sparse, global RFID search to locate a pose in the neighborhood of the tagged object,

followed by a series of local search behaviors (bearing estimation and RFID servoing) to refine the robot's state within the local basin of attraction. We report on RFID search experiments performed in Georgia Tech's Aware Home (a real home). Our optimization-based approach yields superior performance compared to state of the art tag localization algorithms, does not require RF sensor models, is easy to implement, and generalizes to other short-range RFID sensor systems embedded in a robot's end effector. We demonstrate proof of concept applications, such as medication delivery and multi-sensor fusion, using these techniques. Through our experimental results, we show that UHF RFID is a complementary sensing modality that can assist robots in unstructured human environments.

CHAPTER I

INTRODUCTION

Personal robots with mobility, autonomy, and manipulation capabilities have the potential to dramatically improve quality of life for various user populations, such as older adults and people with motor impairments. Unfortunately, unstructured human environments pose myriad challenges that hinder home robot deployment. This thesis seeks to address some of these challenges through a new robotic sensing modality that leverages a small amount of environmental augmentation in the form of Ultra High Frequency (UHF) Radio-Frequency Identification (RFID) tags. Furthermore, unlike UHF RFID sensing techniques commonly employed in factory settings where readers are statically fixed in the environment, the behaviors and algorithms explored in this thesis explicitly rely on the robot's ability to change its configuration or that of the world through mobility or manipulation, thereby providing multiple opportunistic views of the RF landscape, improving robot perception, and ultimately facilitating advanced robotic capabilities.

1.1 Motivation

The overarching motivation behind this research is to hasten the deployment of personal robots to *your* home. Key insights of our work are that UHF RFID “smart labels” are low cost, passive (battery-free), can be read from across a room, and can be ubiquitously attached to important objects throughout the home. In this thesis we show that autonomous mobile manipulators can leverage this modest form of environmental augmentation using short-range and long-range UHF RFID perception to enable advanced robot capabilities. For example, consider this scenario:

During a trip to a department store, you purchase a robot and a box of standard labels: “dish”, “dish washer”, “clothing”, “washing machine”, “toy”, “litter box”, “scoop”, and “storage bin”. You return home, apply the labels as directed, unbox the robot, and turn it on. The robot connects to the web, downloads the appropriate behaviors, and is instantly able to operate in the labeled world: loading the dishwasher with labeled

dishes, putting away labeled toys, cleaning the cats' litter box, and washing labeled clothing. Improved or new functionality, such as delivering medicine, is just a few labels and an internet connection away.

Of course, this is just an interim vision; ultimately we would like robots that are capable of robustly operating in unaugmented home environments. Unfortunately, the economics of large (human-scale) personal robots are untenable. The robots in this thesis cost between \$150,000 to \$400,000 USD in small (research-level) quantities; plus, no single, suitably-compelling end-user application yet exists to justify the exorbitant costs. However, there is one area where burgeoning costs and the need for automation may make an expensive personal robot a compelling reality: *healthcare*. As we will show later, RFID is well-matched to several pressing healthcare applications (eg. medication delivery and medication adherence, discussed in Section 6.3). Perhaps the robot capabilities afforded by UHF RFID can help bootstrap a new market for personal robots and pave the way for general, in-home robots operating in unaugmented environments.

1.2 Contributions

- In Chapter 3, we examine existing approaches to UHF RFID perception that employ two degree-of-freedom (2-DoF), probabilistic, data-driven models to localize tags. These techniques have shown promise when control can be exercised over tag placement (orientation and nearby material properties), as with tags statically attached to a building's infrastructure. However, we show that tagging *objects* introduces a slew of challenges: varied object material properties, changing positions and orientations, and changing environmental RF properties. We show that tagging *objects* can cause drastic deviations to classic 2-DoF UHF RFID sensor models; some of these deviations are pictured in Figure 1.
- The electromagnetic properties of RFID tags change when applied to common household objects. Objects can have varied material properties, can be placed in diverse orientations, and be relocated to completely new environments. In Chapter 4, **we present a new class of optimization-based techniques for RFID sensing that are robust to the variation in tag performance caused by these complexities**. We discuss a hybrid global-local search algorithm where a robot employing long-range directional antennas searches for tagged objects

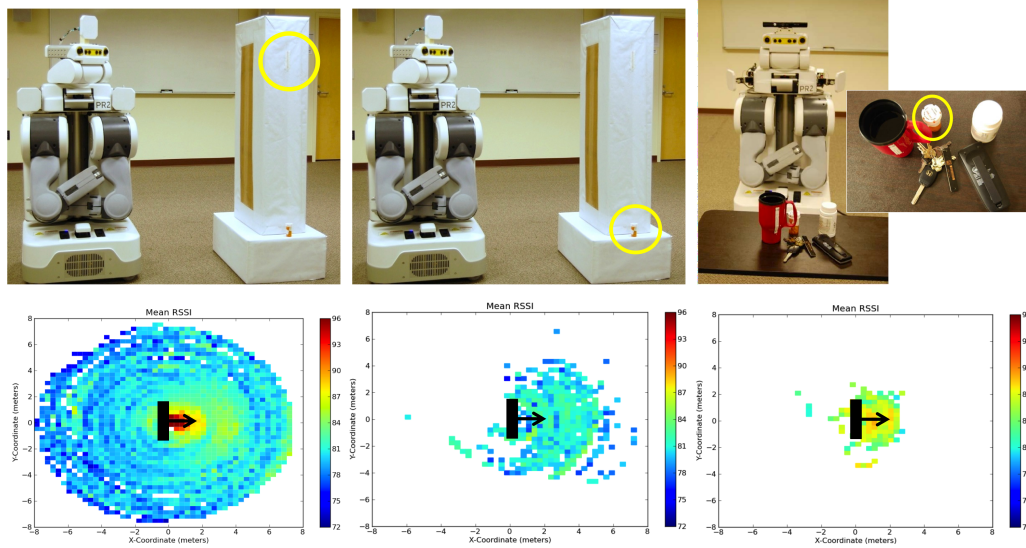


Figure 1: The RFID sensor models (bottom) can vary drastically between different scenarios (top). For example, the model for a tag in relatively ideal, controlled conditions (left) is drastically different from a tagged medication bottle (middle and right). The model for the medication bottle changes depending on its pose in the environment and in relation to nearby objects, ie. in isolation (middle) and in clutter (right). We explore this complexity in Chapter 3.

by maximizing expected RSSI measurements; that is, the robot attempts to position itself (1) near a desired tagged object and (2) oriented towards it. The robot first performs a sparse, global RFID search to locate a pose in the neighborhood of the tagged object, followed by a series of local search behaviors (bearing estimation and RFID servoing) to refine the robot’s state within the local basin of attraction. We report on RFID search experiments performed in Georgia Tech’s Aware Home (a real home). **Our optimization-based approach yields superior performance compared to state of the art tag localization algorithms, does not require RF sensor models, is easy to implement, and easy to generalize.** Some of the final robot positions obtained after performing hybrid global-local search are shown in Figure 2.

- In Chapter 5, we **develop novel UHF RFID antennas mounted-near or embedded-in a robot’s end effector**; these antennas interact with the *same* UHF RFID tags at short-range (an atypical mode of operation), and we generalize the optimization-based behaviors to this regime.
- Finally, in Chapter 6, we **develop a series of mobile manipulation capabilities** that utilize tags’ unique ID coupled with UHF RFID perception and optimization-based behaviors to:



Figure 2: We tested our hybrid global-local RFID search algorithm in a real home environment (Georgia Tech’s Aware Home, top left). As shown in the photographs, the robot obtains positions near the desired tagged object(s) and oriented toward them for various locations inside the home. Details are available in Chapter 4.

perform multi-sensor fusion, build context-aware user interfaces, lookup semantic information (eg. perceptual or kinematic) from a tag-indexed database, and deliver medication. Some of these applications are illustrated in Figure 3.

1.3 Related Work

1.3.1 Pressing Challenges in Healthcare

Healthcare is one example of a motivating application for autonomous robots. For millions of people, motor impairments diminish quality of life, reduce independence, and increase healthcare costs. Assistive mobile robots that manipulate objects within everyday settings have the potential to improve quality of life by augmenting people’s abilities with a cooperative robot. People with motor impairments would benefit from robotic assistance in a broad spectrum of daily living activities,




Figure 3: We develop a new type of tag (dubbed a PPS-tag) that combines a UHF RFID tag with additional forms of augmentation (eg. compliant materials with visually-distinct properties) to provide physical, perceptual, and semantic assistance to robots (Chapter 6). EL-E interacts with PPS-tags to turn on-and-off lights and open-and-close drawers (left). We apply RFID perception and optimization-based RFID behaviors to medication delivery (right – details also in Chapter 6).

including object fetching, object carrying, personal hygiene, food preparation, and the operation of drawers, cabinets, and doors [140]. With over 250,000 people with spinal cord injuries and 3,000,000 stroke survivors in the United States alone, the impact of affordable, robust assistive manipulation would be profound [2, 6]. Moreover, the elderly population worldwide is increasing substantially as a percentage of overall population, with over 16,000,000 people currently over the age of 75 in the United States [120]. This aging population creates a real need for affordable, robust robotic assistance. In the United States, 20% of people between 75 and 79 years of age have been shown to require assistance in activities of everyday living, and this percentage increases precipitously with age, with 50% of people over 85 years of age requiring assistance.

People with motor impairments have consistently placed a high priority on robotic retrieval of objects from the floor and shelves [140], and studies have confirmed this preference for motor-impaired users with amyotrophic lateral sclerosis (ALS, or Lou Gehrig’s disease) working with an autonomous mobile robot [28, 29]. The object retrieval task is also a prerequisite for a variety of other assistive tasks related to food preparation, eating, and cleaning. Current autonomous mobile robots are not yet robust enough to perform these tasks reliably in a real home. We seek to address these challenges using UHF RFID tags affixed to objects, locations, and people of interest. This thesis will examine several foundational robot capabilities for locating, approaching, and manipulating tagged objects and people.

Advances in robotics technology are already helping physically impaired patients through prosthetics, rehabilitation machines, powered orthoses, robotic exoskeletons, and wheelchair-mounted robot arms [133]. Another alternative is an autonomous mobile robot with manipulation capabilities, or autonomous mobile manipulator (AMM). An AMM can perform tasks on its own, move through an environment independently, and manipulate objects. In contrast to other approaches to robotic assistance, an AMM performs tasks independently from the user, does not require donning and doffing, and does not directly encumber the user. AMM's have the potential to be a generally useful assistive technology that assists a wide range of users. For example, unlike a wheelchair-mounted robot arm, a user can benefit from an AMM without being in a wheelchair. Due to the general utility of this type of technology, there is the potential for broad adoption, commoditization, and economies of scale in the design and manufacture of AMMs.

Currently, assistance to people with severe motor impairments is most often provided by a human caregiver, such as a spouse or nurse, which reduces patient privacy and independence while either placing a heavy burden on a family caregiver or entailing costly help. For some patients, highly trained animals, such as service dogs or helper monkeys, provide physical assistance by performing tasks such as fetching objects, opening drawers and doors, wiping a person's face, scratching an itch, or flipping pages in a book [7]. Yet service animals come with a host of other complications, including high training cost (up to \$17,000 for dogs and \$35,000 for monkeys), long wait lists for placement (2-7 years), extensive user training, imperfect reliability, and their own care requirements [1, 75]. Autonomous assistive robots have the potential to offer an important subset of these services at a lower cost, with less maintenance, better reliability, and better sanitation.

In spite of this promising opportunity, autonomous mobile robots are not yet robust enough for daily operation in real homes [77]. This thesis directly addresses the challenges of achieving this competency through novel, innovative research that combines UHF RFID sensing with autonomous mobile manipulators. Of course, there is a trade-off: UHF RFID tags are a form of environmental augmentation. Users will be required to modify their home for the explicit benefit of the robot. Thankfully, many environmental affordances are made by humans for machines; barcodes are a canonical example. Barcodes are of no direct use to human perception. They are included exclusively for machines. RFID tags are nowhere near as mainstream as barcodes, but

they are gaining traction. In recent years, Walmart stores have initiated a field trials of RFID-tagged clothing for inventory control [27]; US passports come embedded with tags [164]; and an alternative short-range RFID-like technology called Near-Field Communication (NFC) is seeing a massive surge, with readers being embedded in smartphones. In addition, the US Food and Drug Administration’s E-Pedigree program, which seeks to track prescription medications from manufacture to end-consumer, is strongly considering RFID as a candidate solution. As with barcodes, robots with RFID capabilities may benefit from these augmented environments.

Furthermore, we have shown that people with motor impairments, who could potentially benefit from robotic assistance, not only accept environmental augmentation but actively precipitate it to assist in mobility and manipulation [29]. In Figure 4, we show scenarios where individuals with motor impairments have modified their own homes to ease mobility and manipulation [29]. Similarly, people with service dogs augment metallic objects (eg. door handles) with towels to facilitate canine manipulation [109]. These populations are already amenable to environmental augmentation; if tagging objects can precipitate more capable mobile manipulators and hasten in-home deployment, it could have significant impact on this population’s quality of life.



Figure 4: Examples of environmental augmentations by people with motor impairments. *Left:* Wheelchair ramp to aid in mobility. *Middle:* Pull-tab on a microwave oven to aid in manipulation. *Right:* Service dog uses a towel on the doorknob to assist in manipulation.

1.3.2 Robots for Assistive Manipulation

Research into assistive robots with manipulation capabilities has a long history, beginning in the 1960s [56, 34]. The following is a brief sampling of the vast related work.

Researchers have developed stationary robotic workstations for assistance with office and factory tasks [33, 147, 148] as well as eating and drinking [145]. Wheelchair mounted robot arms

have provided mobile assistance via joystick control and are beginning to exhibit more autonomy [88, 59, 146]. Until recently, autonomous mobile manipulators had not been investigated widely in an assistive context. The MoVAR and MOVAID projects represent early work in this area [147, 35]. Recent research has sought to develop assistive mobile manipulators with autonomy using methods such as model-based planning and visual servoing [52, 124]. These approaches have not yet yielded assistive robots capable of robustly providing daily service in the home.

There has been a recent surge of interest in autonomous mobile manipulation in human environments [77, 166, 21, 154, 65, 112, 60, 85]. Most research systems have made use of complex pre-defined models of the environment and objects, which the robot registers to its sensory data [139, 11]. In contrast to work that uses complex planning based on detailed geometric models of the environment, other approaches use sparse, low-dimensional, task-relevant features coupled with task-specific robot behaviors [25, 30]. Since no models are required in advance, and no detailed geometric models need be built upon encountering something new, the robot can immediately begin to perform useful actions within new, never before encountered environments, and can manipulate new objects. Through the use of task-specific behaviors, the robot's design is modular and its actions are predictable. This approach is in the tradition of behavior-based robotics.

Similarly, autonomous grasping has been an area of research since the dawn of computer-controlled robots [45]. The great majority of these efforts have focused on grasping modeled objects and the use of grasp planning. Methods capable of grasping unmodeled objects are less common, but early examples do exist [72]. Recent research has demonstrated that autonomous grasping of novel objects is feasible via approaches that take advantage of compliant grippers, tactile sensing, and machine learning [43, 104, 129, 130]. The approach employed by the Healthcare Robotics Lab builds on this work, having shown successful grasp attempts from 85% of the top 25 objects for robotic retrieval (as ranked by patients with ALS) when the object is in relative isolation [67]. In this thesis, we will explore robotic algorithms and robot behaviors that specifically leverage the information provided by UHF RFID to enable mobile manipulation capabilities.

1.3.3 RFID Sensing for Robotics

RFID tags may be a boon for robotic systems. Owing to their unique identifiers and the RF perception techniques developed in this thesis, RFID tags may help bootstrap advanced applications by providing high-fidelity object recognition, coarse robot and tagged object localization, and semantic understanding via databases indexed by tags' unique IDs. We will revisit all of these topics in later chapters; for now, we will briefly examine an overview of RFID technologies and competing non-RF technologies.

1.3.3.1 RFID Technologies

Radio-Frequency Identification (RFID) is a generic term that refers to a wide variety of underlying technologies, ranging from centimeter-range read-only tags used for access control badges, to battery-powered active transponders hosting sensors and long range transceivers. For example, recent work in active tagging demonstrates mobile-robot guidance to a battery powered target tag in a cluttered environment [82]. However, battery powered tags have three major disadvantages compared to passive tags that make them ill-suited for in-home use: limited battery life, larger size and weight, and higher cost. Because of these drawbacks, we prefer passive (battery-free) technologies.

Examples of passive tag technologies are shown in Figure 5. The different technologies are distinguished by their operating frequency, which also dictates how the tag is powered and its read range (the maximum distance between the reader antenna and tag for a successful tag detection). Low frequency (LF) RFID operates at 125 kHz, and high frequency (HF) RFID operates at 13.56 MHz. Both LF and HF RFID tags are powered magnetostatically, a coupling mechanism that limits read ranges to mere centimeters. However, ultra high frequency (UHF) RFID tags operate at 915 MHz and are powered by electromagnetic coupling, allowing tags to be read at distances exceeding 6 meters. [47]

The long-range sensing capabilities of UHF RFID tags makes them particularly appealing, as they can be sensed from across a room. This thesis focuses on long-range, battery-free passive UHF (902-928 MHz) tags and readers based on the Generation 2 (Gen2) standard [44]. Gen2 UHF RFID tags come in many varieties with varied antenna designs, material properties, and underlying protocol implementations; thirty-two commercially-available tags are shown in Figure 6. In this




Figure 5: Examples of passive RFID tags, not to scale. *Left:* 125kHz low frequency (LF) RFID tag. *Middle:* 13.56MHz high frequency (HF) RFID tag. *Right:* 915MHz ultra-high frequency (UHF) RFID tag.

work, we predominantly employ the Alien Technologies ALN-9640 Squiggle tag. This tag is also shown (approximately to scale) in Figure 6. This tag is commonly available and retails for less than \$0.25 USD per tag in bulk.

UHF RFID readers are comprised of two main components: a reader module and one (or more) transmit / receive antennas. The reader modules are comprised of integrated circuits (and printed circuit boards) that implement PC communications, the Gen2 protocol, and RF signal input / output. The RF signals are transmitted and received via attached antennas. UHF RFID reader modules are readily available from multiple vendors, from \$300-\$1000 USD. Owing to our collaboration with Dr. Matt Reynolds (Duke University, and co-founder of ThingMagic), we employ several variants of ThingMagic readers: Mercury 4e (M4e), Mercury 5e (M5e), and Mercury 5e-Compact (M5e-C). Two of the more modern (M5e-variety) reader modules are shown in Figure 7. The particular choice of reader for each experiment depends on the date of the research; our research evolved lock-step with the evolution of ThingMagic’s readers.

The ThingMagic UHF RFID readers can sense tags beyond six meters under ideal conditions (depending on the antenna); however, they do *not* provide precise tag location information. Rather, they provide a binary indication of tag presence or absence (a detection); in the event of a positive tag detection they also provide a receive signal strength indicator (RSSI) measurement, a scalar value that indicates the strength of the tag’s response.

UHF antennas are available from numerous manufacturers. Figure 8 shows two antennas used





Figure 6: *Top:* Thirty two different varieties of passive UHF RFID tags. The RFID reader modules used in this thesis are capable of communicating with any of these tags. *Bottom:* A Alien ALN-9640 Squiggle tag, approximately actual size.




Figure 7: We used the ThingMagic Mercury 5e (M5e, left) and Mercury 5e Compact (M5e-C, right) reader modules in this thesis. These two modern readers provide both tag detection and RSSI measurements.

through the course of our research. Each of these antennas has different properties (eg. size, directivity, etc) that make them well-suited for particular tasks. We will explore these properties later in the context of system operation, as well as during experimental evaluation.



Figure 8: UHF RFID Antennas. From left-to-right: Laird Technologies S9028PC12NF long-range patch antenna, a quarter (US \$0.25) for scale, Laird Technologies S9025P long-range patch antenna, and Johanson Technology (part number 0920AT50A080E) short-range ceramic microstrip antennas.

At this point, it is useful to understand the basic principle behind UHF RFID tag operation. RF signals are generated from the RFID reader module; these signals are modulated (changing over time) using phase modulation to communicate with nearby tags. The RF signals are radiated by the transmit antenna, propagate through the medium (eg. air), and are electromagnetically coupled to an RFID tag's antenna. At the tag, three things occur. First, some of the incident energy is harvested by the tag and used to power-up the tag's onboard circuitry; the rest of the incident energy is reflected. Second, the tag senses the phase-modulated signal from the reader. And thirdly, if the tag should respond (per the Gen2 protocol), the tag will electronically alter its radar cross section to amplitude-modulate the amount of incident energy reflected back to the reader. This reflected signal is received and decoded at the reader. This process is called backscatter modulation; it is described in detail in the literature and depicted in Figure 9 [48].

For readers unfamiliar with backscatter modulation, consider this scenario for emergency (S.O.S.) signaling: A rescuer riding in a helicopter scans a spotlight (the transmitter) around to look for a stranded traveler (the tag). Our stranded traveler is holding a signalling mirror; when the spotlight's beam illuminates the mirror, the traveler very slightly adjusts the mirror's angle to reflect a small amount of light back to the rescuer. This return signal is detected (eg. electronically with a photodiode), and can be used to signal (eg. Morse code) back to the rescuer. To complete the analogy, the mirror is not actually being held by the traveler. Our traveler is quite clever; the mirror is actually a reflective solar cell hooked to a motor. Some of the spotlight's incident light is captured by the solar cell, which powers a motor to spin the solar cell. The rest of the spotlight's incident light is reflected back to the rescuer. The solar cell simultaneously draws power from the incident beam while signaling a small amount back to the rescuer. This is the basic premise behind UHF RFID tag operation, except that we're using longer-wavelength radio frequencies instead of visible light (just a different wavelength of the electromagnetic spectrum). This relationship is shown in Figure 9.

1.3.4 Competing non-RF Technologies for Robotics

There are many other non-RFID forms of environmental augmentation used for the benefit of robots and other machines. Visual fiducials are probably the most ubiquitous form of augmentation; we




Figure 9: *Top:* UHF RFID tags operate on the principle of backscatter modulation. The tag harvests some of the incident power (unmodulated transmission) in order to power up. If the tag should respond (per the Gen2 protocol), it reflects some of the incident radiation by modifying its radar cross section to send an (amplitude-modulated) signal back to the reader. *Bottom:* Backscatter modulation is analogous to a S.O.S. emergency signaling, where a stranded traveler signals to their would-be rescuers.

show several examples in Figure 10. Barcodes are a canonical example; their marginal cost is less than one cent each, and they are commonplace on most mass-market commercial products. They provide strong identity information, and are used by many automated systems [138]. Other visual fiducials, such as AR Toolkit tags and QR codes, also provide identification information and have

been used in robotic systems for object recognition. When sensed with calibrated cameras, AR Toolkit tags and checkerboard patterns can yield 3D position and orientation information, which also aids in navigation, localization, object recognition, and manipulation [46]. These fiducials also find extensive use in computer vision and augmented reality. Distinctive colors or color patterns can also be used as a visual fiducial. Researchers from the Healthcare Robotics Lab employed bright red towels affixed to doors and drawers; the bright color simplified perception (detection via cameras), and the compliant nature of the towel aided in robotic manipulation [109]. Most high performing systems in RoboCup (robot soccer) competitions depend on accurate color segmentation algorithms and unique colored patterns to identify and track their own position and that of competitors [26, 155].




Figure 10: Examples of visual fiducial technologies. From left-to-right: barcodes, AR Toolkit tags, QR codes, and checkerboards.

There are also plenty of non-visible forms of environmental augmentations to assist robots too. The Northstar system projects a static infrared pattern onto the ceiling; robots can use a camera to detect the pattern and determine their position in a room [161]. Laser “tags” that respond when struck by a laser rangefinder can help a robot locate tags [165]. Ultraviolet-sensitive paint can define navigation paths for cleaning robots [100]. Underground wires can be used to demarcate boundaries for lawn-mowing robots or define paths for factory robots [4]. Motion capture devices, such as Optitrack, provide precise (millimeter-scale) accuracies for manipulation [113].

Despite all of these competing forms of augmentation, RFID tags are still finding useful applications in robotics. For example, Aethon’s TUG robot uses RFID tags to identify hospital equipment, which is then relocated to maximize resource utilization [20]. UHF RFID has a few unique characteristics, which we believe are profoundly enabling for autonomous mobile manipulators: they can be read from a long distance (greater than 6 meters under ideal conditions, and unlike visual

fiducials, UHF RFID tags do not require line of sight and are less sensitive to orientation.

UHF RFID tags are not a panacea; they cannot single-handedly solve all of robotics' challenges. The UHF RFID sensing techniques developed in this thesis are complementary in nature to other sensing modalities. In Chapters 6, we will demonstrate ways to combine sensor data from UHF RFID, cameras, and laser rangefinders to perform useful tasks. In fact, we also combine UHF RFID tags with other forms of augmentation (visual fiducials) to provide additional robustness; indeed, several varieties of UHF RFID tags already possess visual fiducials for precisely this reason. For example, The Sontek Stick Tag (an on-metal UHF RFID tag) shown in Figure 11 possesses a UHF RFID tag, a barcode, and a QR code.



Figure 11: The Sontek Stick tag (an on-metal UHF RFID tag) combines UHF RFID with a barcode and QR code.

1.4 Our Robots

We used three different robot platforms during the course of our research (shown in Figure 12): EL-E, PR2, and a mobile multi-antenna test rig. All of these robots are shared resources at the Healthcare Robotics Lab. The robot hardware and software evolved over the duration of the research, incorporating several different RFID configurations. In each chapter and section we will examine the pertinent robots and their configurations.

EL-E, pronounced “Ellie,” was designed and constructed in the Healthcare Robotics Lab. EL-E is comprised of several off-the-shelf components: a differential-drive mobile base, a linear actuator, a 6 degree-of-freedom (6-DoF) position-controlled arm, numerous UHF RFID antennas, and many other sensors (cameras, laser rangefinders, etc). EL-E is a semi-custom design whose bill of materials was (roughly) \$150,000.

The PR2 robot is a commercial platform from Willow Garage, a startup company located in Palo



Figure 12: Over the course of performing our research, we used several different robots. The robots were in a continuous state of evolution. Here are some snapshots of the robots: EL-E (top), PR2 (middle), test rigs (bottom).

Alto, CA. The PR2 retails for \$400,000 before open-source discounts. The PR2 comes stock with a holonomic powered-caster mobile base, a telescoping linear actuator “spine”, two 7-DoF compliant arms, a pan-tilt head, and numerous cameras and laser rangefinders. We modified the PR2 by adding different sensors: pan-tilt UHF RFID antennas on the two shoulders, and either a Microsoft Kinect depth camera or additional UHF RFID antenna on the head.

The test rigs were custom creations used at the onset of our research. They have six UHF RFID antennas equally spaced in a circular configuration, a laser rangefinder, and a magnetometer.

1.5 Organization of This Work

UHF RFID systems are essentially short-range radars that track nearby tags rather than (for example) planes in the sky. Chapter 2 examines the Friis radar equation, which is commonly used by the RF community to understand radar design considerations under ideal, free-space conditions. It is well-known that the Friis equation is a crude approximation to real-world radar system performance [136], especially for tags operating in extreme clutter (compared to planes in the open sky). We use insights from the Friis equation throughout the remainder of this thesis.

Chapter 3 discusses state-of-the-art probabilistic (Bayesian filter) techniques that, as a robot moves around, integrate multiple RFID measurements to localize tags relative to the robot (or on a map). These techniques require a motion model (eg. odometry) and a RFID sensor model. Predominantly, researchers (including us) have developed sensor models that describe the likelihood of tag detection and RSSI measurements given the 2 degree-of-freedom (2-DoF, planar) relationship between tag and reader antenna. These methods have primarily been applied to infrastructure tags (eg. affixed to walls or floors), where the tag pose and nearby material properties are controlled. Placing a tag on an object introduces a slew of challenges that may alter RF signal propagation: object material properties (eg. metal components), the specifics of the tag mounting (eg. wrapped around the object), and object configuration (eg. an empty cup versus a full cup). Furthermore, tagged objects can be relocated throughout the environment, which changes the tag position and orientation (in 6-DoF) and nearby environmental (RF) properties. Through extensive ground-truth data captures, we show illustrative examples where 2-DoF sensor models degrade when dealing with tagged objects.

In Chapter 4, we describe a new class of optimization-based techniques allowing robots to discover, locate, and approach UHF RFID tags affixed to objects of interest in unstructured home environments. These techniques use insights from the early radar literature. Explicitly, the robot uses directional transmit and receive antennas to maximize measurements of received signal strength (RSSI), a scalar value that indicates the strength of a tag’s response, to obtain positions and orientations that are both (1) near the desired tagged object and (2) oriented toward it. We segment the optimization-based RFID search techniques into two categories: global and local. In global RFID search, the robot uses its mobility and sparse sampling to locate a pose in the neighborhood of the tagged object. Subsequently, the robot performs a series of local optimization-based search behaviors that allow the robot to estimate the direction toward a tag (bearing estimation) and/or approach the tag (RFID servoing). The local search techniques refine the robot’s state within the local basin of attraction. We combine these two search techniques into a hybrid global-local optimization-based search algorithm that is easy to implement, generalizes across situations, doesn’t require training data to build sensor models, and yields comparable position estimates and superior angular estimates compared to state-of-the-art probabilistic methods during in-home experiments with tagged objects.

In Chapter 5 we examine novel antennas mounted-near or embedded-in a robot’s manipulator that interact with the *same* UHF RFID tags at short-range – an atypical mode of operation. We generalize our optimization-based behaviors to this regime. We characterize the antennas’ ability to detect tagged objects being held in a robot’s manipulator and develop robot behaviors that use the short-range antennas to distinguish between visually-identical objects and to determine the identity of a grasped object.

In Chapter 6, we develop a series of mobile manipulation capabilities that utilize both long-range and short-range optimization-based UHF RFID behaviors:

1. We develop a new type of tag (dubbed a PPS-tag) that combines a UHF RFID tag with additional forms of augmentation (eg. compliant materials with visually-distinct properties) to provide physical, perceptual, and semantic assistance to robots. We present five exemplar PPS-tags along with a set of robotic behaviors that utilize UHF RFID perception and

optimization-based RFID behaviors (eg. bearing estimation and RFID servoing) to accomplish various tasks, such as: turning on and off light switches, opening and closing drawers, operating lamps, and disposing of trash in a waste bin.

2. We develop a framework for multi-sensor fusion using UHF RFID sensing. We discuss how to use the tag's unique identifier as an index into a semantic database, where we can store and retrieve information about the tagged object's appearance in other sensing modalities (eg. cameras and laser rangefinders that produce 3D point clouds). Combined with the optimization-based RFID behaviors, we describe a mobile manipulation system that is capable of locating, approaching, perceiving (through multi-sensor fusion), and grasping tagged objects in relative isolation on the floor.
3. We apply RFID perception and optimization-based RFID behaviors to medication delivery. We describe a system wherein an autonomous mobile manipulator confirms the identity of a grasped (tagged) medication bottle, uses optimization-based RFID behaviors to locate and approach the intended (tagged and seated) recipient, and then hands off the medication. This system is currently undergoing testing in a realistic home environment with older adults.

Finally, in Chapter 7, we conclude with a discussion about the future of UHF RFID and robotics.

CHAPTER II

UHF RFID AND THE FRIIS RADAR EQUATION

UHF RFID systems are essentially short-range radars that track nearby tags rather than (for example) airplanes in the sky. In this chapter, we examine the Friis radar equation, which is commonly used by the RF community to understand radar design considerations under ideal, free-space conditions. Throughout the remainder of this thesis, we use insights gleaned from the Friis equation: to describe the pertinent relationships between reader and tag (eg. relative pose and antenna properties), to discuss probabilistic UHF RFID sensor models used by robotics researchers for Bayesian tag localization, and to explain the intuition behind our optimization-based techniques that rely on numerous (relative) RSSI measurements.

2.1 *The Friis Radar Equation(s)*

Passive UHF RFID tags harvest all of their operating power from nearby, interrogating readers. Tracing the power from reader-to-tag (and back) can provide insights into both tag detection and received signal strength (RSSI, a scalar value indicating the strength of a tag’s response) measurements. The Friis equation is a simplified, line-of-sight formulation based on the physics of RF propagation that is commonly used to trace RF power in ideal radar systems [136]. It makes no attempt to model non-ideal conditions such as multipath, shadowing, diffraction, material properties, or atmospheric conditions; thus, it is a crude approximation to real-world performance. However, it can give RFID system designers insights into the *physics* trade-offs inherent with radio frequency operation.

2.1.1 **The Forward-Link Friis Equation**

The *forward-link Friis equation* computes the power incident at the tag (P_{tag}^{inc}) by tracing the RF power through the system, as shown in Figure 13,

$$P_{tag}^{inc} = P_{rdr} \cdot G_{rdr} \cdot PL \cdot G_{tag} \quad (1)$$




Figure 13: The Friis forward-link equation predicts the power incident at the tag (P_{tag}^{inc}) by tracing the flow of RF power from the reader (P_{rdr}), through the reader antenna (G_{rdr}), propagating through free-space (falling off as $1/r^2$), through the tag's antenna (G_{tag}).

It starts by considering the RF power emitted by the reader (P_{rdr}). This power is radiated according to the gain of the reader transmit antenna (G_{rdr}), and is then dispersed as the RF power propagates through free-space causing path loss (PL), which falls off as the inverse-squared of the distance ($1/r^2$). At the tag, the radiating power is multiplied by the tag antenna's gain (G_{tag}), yielding the power incident at the tag (P_{tag}^{inc}). We specify all quantities for this analysis in linear units (Watts); occasionally we will specify them in logarithmic units (dB).

The reader transmit power (P_{rdr}) is directly controlled by the system designer; for the Thing-Magic M5e readers used in this work, the reader power can be programmatically adjusted from 3.2 mW to 1000.0 mW (5 to 30 dB milli-Watts or dBm).

The free-space path loss (PL) is a wavelength-dependent quantity that accounts for radiated power falling off as the inverse-squared of radius (r) between antennas. As a function of wavelength (λ , in meters) or frequency (f , in MHz) it can be expressed as follows,

$$PL = \left(\frac{\lambda}{4\pi \cdot r} \right)^2 = \left(\frac{3 \cdot 10^2}{4\pi \cdot r \cdot f} \right)^2. \quad (2)$$

The remaining two parameters are the antenna gains for the reader transmit antenna (G_{rdr}) and tag antenna (G_{tag}). In general, antenna gain is a function of transmit or receive direction. When

we know the 6-DoF pose of the tag relative to the reader antenna (eg. in spherical coordinates as shown in Figure 14), the gain is expressed as a two-dimensional function of azimuth angle (θ) and declination angle (ϕ) with respect to each antenna's coordinate frame. This two-dimensional function is often referred to as a *radiation pattern*, and is determined by the physical properties of the antenna. The *directivity* of an antenna's gain function is a measure of how much radiation is emitted in a particular direction relative to the gain of an isotropically-radiating antenna. We examine several common antenna gain patterns in Section 2.1.3.




Figure 14: The six degree-of-freedom (6-DoF) relationship between the reader antenna (aligned with frame F_{rdr}) and the tag antenna (aligned with frame F_{tag}) can be described (in part) by the spherical coordinates shown.

For now, we make the gain functions' dependent variables explicit in the forward-link Friis equation,

$$P_{tag}^{inc} = P_{rdr} \cdot G_{rdr}(\theta_{rdr}, \phi_{rdr}) \cdot \left(\frac{\lambda}{4\pi \cdot r} \right)^2 \cdot G_{tag}(\theta_{tag}, \phi_{tag}). \quad (3)$$

2.1.2 The Forward-Backward-Link Friis Equation

UHF RFID tags do not actually transmit power. Rather, they reflect *some* of their instantaneous incident power back to the RFID reader's receive antenna by modulating their radar cross-section in a process called backscatter modulation. We can use the *forward-backward-link Friis equation* to compute the power incident at the reader's receive antenna (P_{rdr}^{inc}) by again tracing the power from

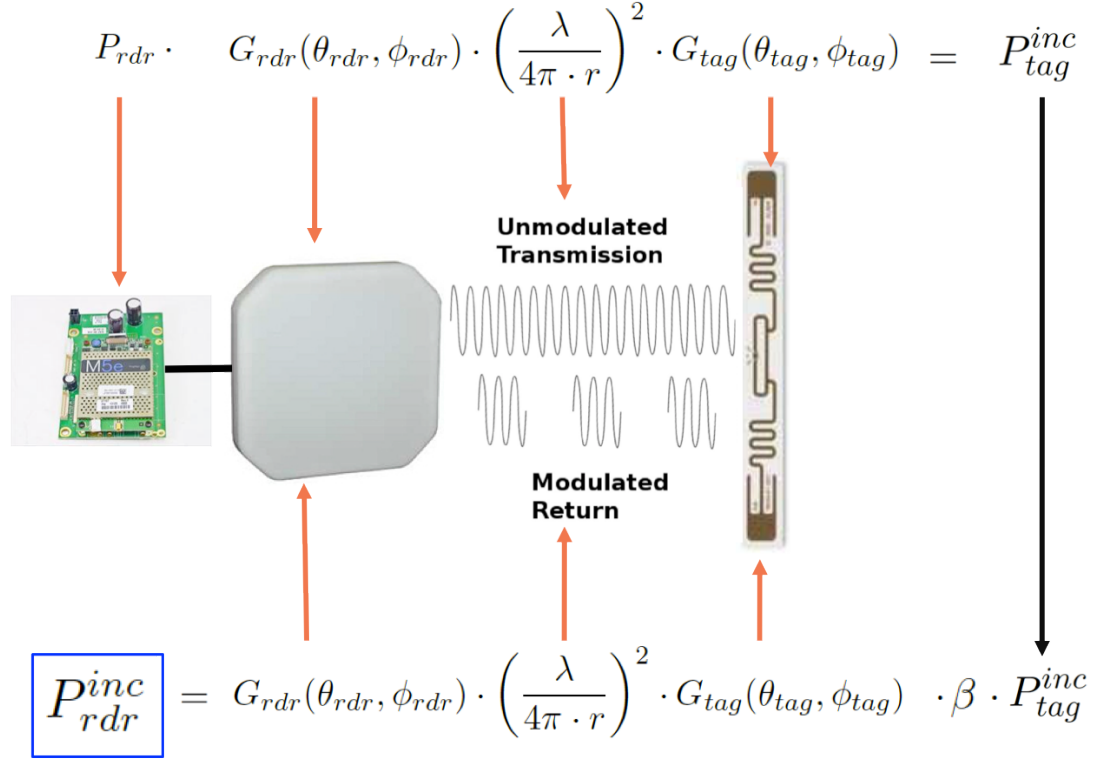


Figure 15: The Friis forward-backward-link equation predicts the power incident at the reader (P_{rdr}^{inc}) by tracing the flow of RF power reflected back from the tag (a fraction, β , of the power incident at the tag, P_{tag}^{inc}), through the tag antenna (G_{tag}), propagating through free-space (falling off as $1/r^2$), through the reader's receive antenna (G_{rdr}).

the tag back to the reader, as shown in Figure 15.

$$P_{rdr}^{inc} = P_{tag}^{inc} \cdot \beta \cdot G_{tag}(\theta_{tag}, \phi_{tag}) \cdot \left(\frac{\lambda}{4\pi \cdot r} \right)^2 \cdot G_{rdr}^{rec}(\theta_{rdr}, \phi_{rdr}). \quad (4)$$

This time, the power being transmitted is actually the power reflected by the tag – the incident power at the tag (P_{tag}^{inc}) multiplied by the backscatter efficiency (β). This power is radiated according to the tag antenna's radiation pattern (G_{tag}), propagated through free space (PL), and modified by the RFID reader's *receive* antenna gain (G_{rdr}^{rec}) before being sensed and processed by the RFID reader's internal circuitry.

The same physical antenna need *not* be used for both transmission and reception, in which case the gain functions (G_{rdr} and G_{rdr}^{rec}) could take very different forms and would operate with respect to their own independent coordinate frames. For the ThingMagic Mercury5e readers, this is referred to as “bistatic mode.” For all of the work presented in this thesis, we operate exclusively in the “monostatic mode” where we use the same physical antenna for both transmission and reception,

so that $G_{rdr}^{rec} = G_{rdr}$. Thus, we can further simplify the forward-backward-link Friis expression in Equation 4 by substituting the power incident at the tag (P_{tag}^{inc}) calculated in Equation 3, yielding

$$P_{rdr}^{inc} = P_{rdr} \cdot \beta \cdot \left\{ G_{tag}(\theta_{tag}, \phi_{tag}) \cdot \left(\frac{\lambda}{4\pi \cdot r} \right)^2 \cdot G_{rdr}(\theta_{rdr}, \phi_{rdr}) \right\}^2. \quad (5)$$

Several of the Friis equation arguments are directly controlled by the system designer: the reader transmit power (P_{rdr}), the choice of reader antenna with its associated gain (G_{rdr}), and the RF operating frequency / wavelength (λ). Others arguments may be less controlled. The choice of tag dictates its antenna gain (G_{tag}) and backscatter efficiency (β). The remaining arguments (θ_{tag} , ϕ_{tag} , θ_{rdr} , ϕ_{rdr} , and r), whether controlled or not, are determined by the relative position and orientation (pose) between the reader antenna and tag antenna. When we know the full six degree-of-freedom (6-DoF) pose of the tag relative to the reader, as shown in Figure 14, we can ascertain these arguments' values.

2.1.3 Antenna Selection

Antenna selection is one of the most crucial UHF RFID system design considerations, and there are many of types to choose from: dipole, monopole, helical, patch, Yagi, rabbit ears, horn, isotropic, etc. For each type, there is a rich body of literature about antenna design – how the physical properties of the antenna manifests into the varied radiation patterns [63]. Fortunately, two types of antennas dominate the literature for UHF RFID sensing in robotics: dipoles and patches.

Dipoles are the most commonly used tag antennas in the UHF RFID sensing literature related to robotics, owing mainly to their simplicity, size, low cost (sub-\$0.25 USD), and commercial availability. We employ several dipole tag antenna variants in this research, but predominantly focus on the Alien Technologies ALN-9640 ‘‘Squiggle’’ tag – which is frequently used in the RFID literature and has been extensively characterized: on humans [3], near liquids [9], near metals [10], on cellphones [36], and under various GEN2 protocol configurations [110]. Meanwhile, patch antennas are the most commonly used reader transmit and receive antennas. This owes largely to their commercial availability and high *directivity* (high gain in one direction). Also, patch antennas in the UHF spectrum tend to be of a more manageable size compared to other antenna variants, which makes them more amenable to mobile robot integration and construction. To fully appreciate

the forward-link and forward-backward-link Friis equations, it is important to understand the gain functions for these two types of antennas.

Gain functions, or radiation patterns, can be determined in a number of ways: from the antenna manufacturer, measured empirically using anechoic chambers, using closed-form “ideal” expressions, or via RF simulation software tools such as Ansys’ HFSS. Now, we look at dipole and patch antennas in more detail.

2.1.3.1 Dipole Tag Antennas

Dipole antennas have a well-known closed-form expression for their radiation pattern when the length of their two radiating elements is exactly one-quarter wavelength each ($\approx 16.4\text{cm}$ total length for 915MHz) [14], which is given by

$$G_{tag}(\theta_{tag}, \phi_{tag}) = 1.5 \cdot \sin^2(\phi_{tag}). \quad (6)$$

Dipoles have a maximum gain of 1.5 (1.76 dB) for $\phi = 90^\circ$ and a minimum gain of zero at $\phi = 0^\circ$ and $\phi = 180^\circ$. In practice, gains near zero are not physically realizable; to compensate antenna specifications commonly specify a lower-bound on the antenna gain, sometimes referred to as a front-back-ratio (*FBR*) and is specified with respect to the antenna’s maximum gain (G_{max}). Accounting for *FBR*, the gain function becomes

$$G_{tag}(\theta_{tag}, \phi_{tag}) = \max \left(1.5 \cdot \sin^2(\phi_{tag}), G_{max} \cdot FBR \right). \quad (7)$$

Since dipoles’ radiation patterns do not depend of azimuth angle (θ), the radiation patterns have a toroidal shape, as shown in Figure 16.

As previously mentioned, in this thesis we (and most researchers using UHF RFID in robotics) use the Alien Technologies ALN-9640 “Squiggle” tag. Though not a pure quarter-wave dipole antenna, it is reported to be a close approximation with a front-back-ratio around 0.16 (-8 dB) under common conditions [36]. A photograph of this tag (approximately to scale) is shown in Figure 6. The tag’s coordinate frame and ideal radiation pattern are shown in Figure 16.

Often, *full* radiation patterns for all directions (θ and ϕ) will not be provided by antenna manufacturers; instead, radiation patterns will be described by one or more *slices* along dominant planes.

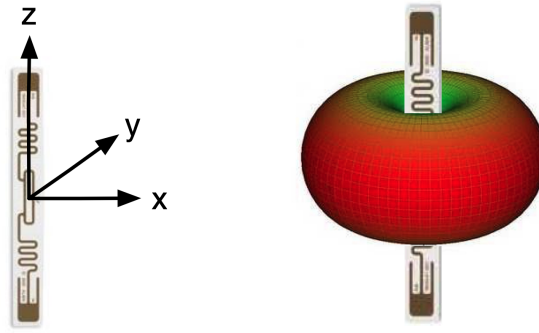


Figure 16: *Left:* Alien Technologies “Squiggle” tag. *Right:* Theoretical dipole radiation pattern from Equation 6.

Frequently, this corresponds to the dominant excitation planes: the electric field plane (E-plane) and magnetic field plane (H-plane). Also, these slices will often display *directivity* (with a maximum value of 1.0 in the direction of maximum gain) rather than the gain itself. In Figure 17 we show gain slices for the E-Plane (xz-plane, or $\theta = 0^\circ$, $\phi = [0^\circ, 180^\circ]$) and H-Plane (xy-plane, or $\phi = 90^\circ$, $\theta = [-180^\circ, 180^\circ]$) for an ideal Alien Squiggle (dipole) tag with $FBR = -8dB$.

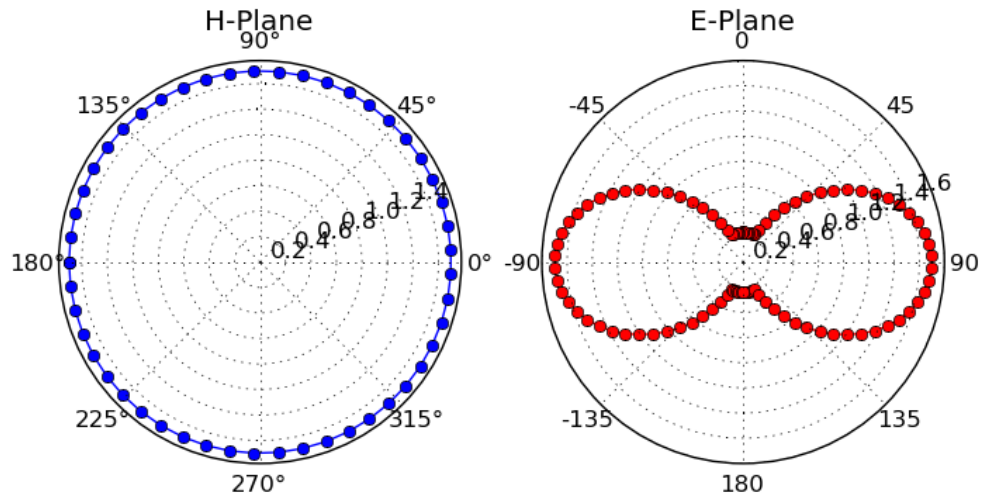


Figure 17: H-plane (left) and E-plane (right) cross sections of a an ideal dipole’s radiation pattern (Equation 7) with $FBR = -8dB$; this corresponds to the gain and FBR of an Alien Technologies “Squiggle” tag.

2.1.3.2 Patch Antennas

Patch antennas are commonly used as the RFID reader transmit and receive antennas. Compared to dipoles, patch antennas have much greater directivity, and larger gains. They have unimodal, semi-conical radiation patterns with a global maximum at $\theta_{rdr} = 0^\circ$, $\phi_{rdr} = 90^\circ$ – see Figure 18. Further, for any fixed value of ϕ_{rdr} (ie. a cross section of the radiation pattern), the patch antenna still has a local maximum at $\theta_{rdr} = 0^\circ$. These properties will prove useful for the methods developed in later chapters.

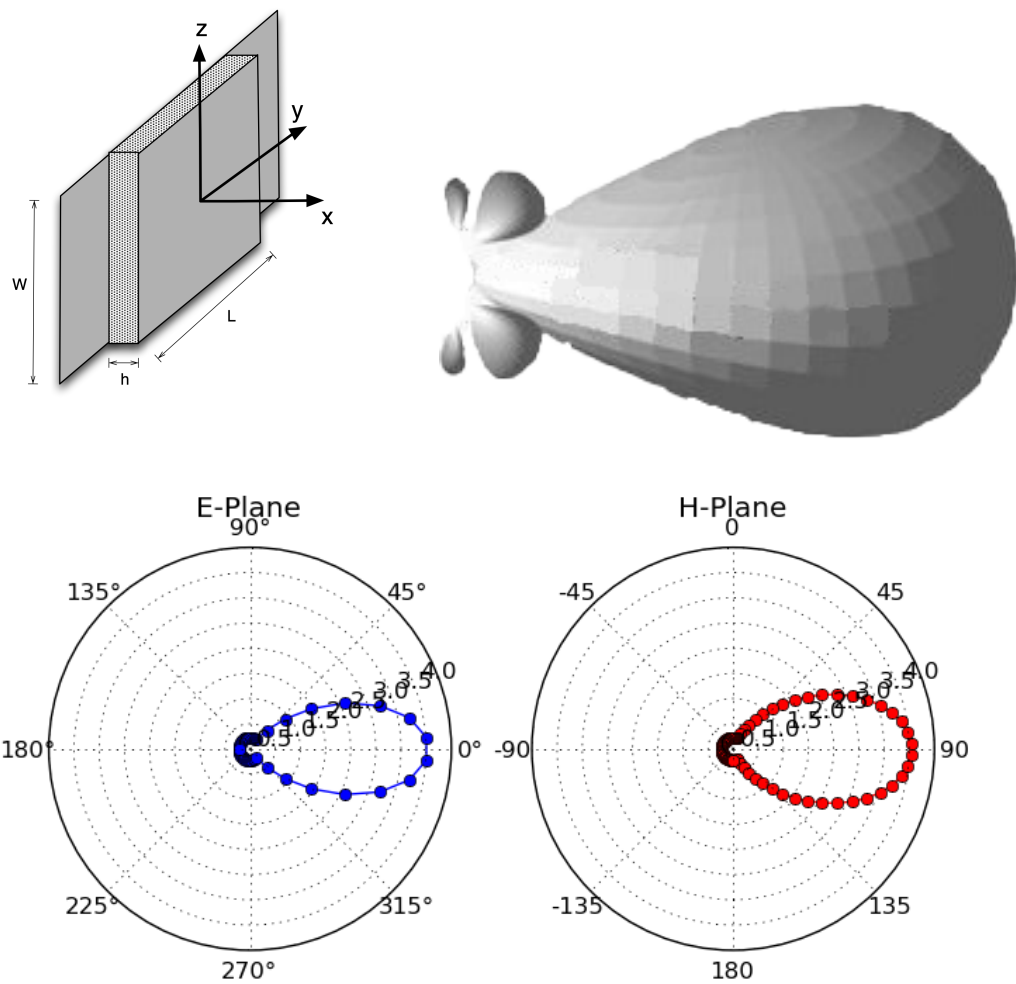


Figure 18: *Top Left:* Patch antenna drawing. *Top Right:* Patch antenna radiation pattern. *Bottom:* E-Plane (left) and H-Plane (right) cross sections of an ideal patch antenna’s radiation pattern (Equation 8) using the physical properties of the Laird Technologies S9025P depicted in Figure 8.

While not as straight forward as a dipole, there are closed-form theoretical expressions for the

gain of a patch antenna [14]. Assuming $k_0 \cdot h \ll 1$, the gain function is given by

$$G_{rdr}(\theta_{rdr}, \phi_{rdr}) \simeq \gamma \cdot \left\{ \sin(\phi_{rdr}) \frac{\sin\left(\frac{k_0 W}{2} \cdot \cos(\phi_{rdr})\right)}{\cos(\phi_{rdr})} \cdot \cos\left(\frac{k_0 L_e}{2} \cdot \sin(\phi_{rdr}) \sin(\theta_{rdr})\right) \right\}^2. \quad (8)$$

This equation has several parameters related to the physical properties of the patch antenna: effective length (L_e), width (W), height (h), and permeability / permittivity constant (k_0). The gain has a maximum at $\phi = 90^\circ, \theta = 0^\circ$. The multiplier (γ) can be determined from the maximum gain (G_{max}) supplied by the manufacturer. For example, if given G_{max} , then for some constant κ related to the physical antenna parameters

$$G_{max} = G_{rdr}(\phi_{rdr} = 90^\circ, \theta_{rdr} = 0^\circ) \simeq \gamma \cdot \kappa, \text{ and thus} \quad (9)$$

$$\gamma \simeq G_{max} / \kappa. \quad (10)$$

The ideal patch antenna radiation pattern is depicted in Figure 18. In reality, fringing electric / magnetic fields will cause deviations from the ideal, particularly where the gain is at or near zero – as on the backside of the antenna.

We use two different patch antennas in this thesis; they were shown in Figure 8 in Chapter 1. The first is the Laird Technologies *S9028PC12NF*, a 25×25 cm air-dielectric circularly polarized patch antenna with 65° 3dB-beamwidth (half-power beamwidth), 7.5 dB gain, and 20 dB *FBR* that retails for around \$215. The second is the Laird Technologies *S9025P*, a 13.5×13.5 cm circularly polarized patch antenna with 100° 3dB-beamwidth, 5.5 dB gain, and 8 dB *FBR* that retails for around \$100. In Figure 19, we can see the drastic difference in antenna directivity. The *S9028PC12NF* antenna offers a narrower beamwidth (region of high gain), which is useful for applications such as bearing estimation and RSSI Images (Chapter 4); however, the *S9025P* is easier to integrate into a robot due to smaller overall size.

2.1.4 Remarks

The Friis equation(s) gives us some indication about the amount of power incident at the tag and the amount reflected back to the reader. If we knew the 6-DoF pose (x) between the reader antenna and tag, and we knew the radiation patterns of the antennas, then we could use the Friis equation(s) to

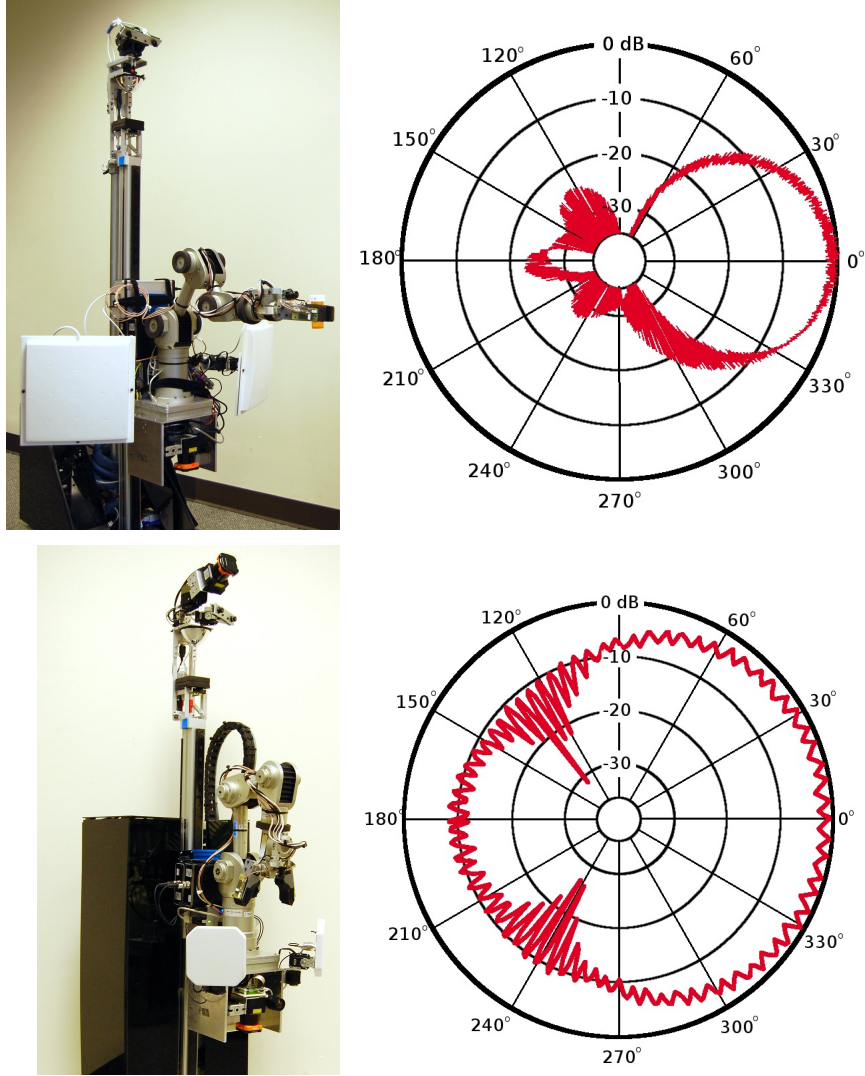


Figure 19: *Top:* Two Laird Technologies *S9028PC12NF* patch antennas attached to EL-E (top left) have a highly-directional manufacturer-supplied radiation pattern (top right). *Bottom:* Two Laird Technologies *S9025P* patch antennas attached to EL-E (bottom left) are much smaller, making them easier to mount on the robot. However, their radiation patterns, supplied by the manufacturer (bottom right), are much less directional.

estimate the power incident at the tag, $\hat{P}_{tag}^{inc}(x)$, and the power incident at the reader $\hat{P}_{rdr}^{inc}(x)$. In the following sections, we will see how these values relate to the probability of a tag detection and the probability of obtaining specific RSSI measurements – in effect, defining a Friis-based RFID sensor model.

2.2 Tag Detection Model

Ignoring on-board energy storage, a tag is *capable* of a positive response when the instantaneous incident power (P_{tag}^{inc}) exceeds a threshold – sufficient power to run the tag’s on-board circuitry. The tag power-up threshold depends on properties of the small silicon integrated circuit (IC) that comprises the “brains” of the tag. Advances in IC fabrication processes, power harvesting circuitry, low-power oscillators, and low-power state machines all seek to reduce the power-up threshold and facilitate longer read ranges. In 2006, the power-up threshold for Alien Squiggle tags was reported in the literature at $50\mu\text{W}$ (-13 dBm) [37]. In 2009, after subsequent improvements, the power-up threshold for the Squiggle tags was reported at $15\mu\text{W}$ (-18 dBm) [31]. Thus, we expect the tag to power-up and send a signal back to the reader when the Friis-predicted power incident at the tag (Equation 3) exceeds the tag power-up threshold,

$$P_{tag}^{inc} \geq 15\mu\text{W} (-18\text{dBm}) \quad (11)$$

$$P_{rdr} \cdot G_{rdr}(\theta_{rdr}, \phi_{rdr}) \cdot \left(\frac{\lambda}{4\pi \cdot r}\right)^2 \cdot G_{tag}(\theta_{tag}, \phi_{tag}) \geq 15\mu\text{W} (-18\text{dBm}). \quad (12)$$

In general, the tag power-up threshold alone is insufficient to ensure a successful tag detection. The RFID reader must also possess sufficient sensitivity to detect the tag’s backscattered signal. For the ThingMagic M5e, the reader’s receive sensitivity is guaranteed to be better than 10 pW (-80 dBm) for Miller-encoded tag responses with 4.256 kHz backscatter link rate[125]. These are the default M5e settings, so a positive tag detection also requires that

$$P_{rdr}^{inc} \geq 10\text{ pW} (-80\text{dBm}) \quad (13)$$

$$P_{rdr} \cdot \beta \cdot P_{tag}^{inc} \geq 10\text{ pW} (-80\text{dBm}). \quad (14)$$

Recall that the Friis-predicted power incident at the tag (P_{tag}^{inc} , Equation 3) and the Friis-predicted power incident at the reader (P_{rdr}^{inc} , Equation 5) have very similar forms,

$$P_{tag}^{inc} = P_{rdr} \cdot \left\{ G_{tag}(\theta_{tag}, \phi_{tag}) \cdot \left(\frac{\lambda}{4\pi \cdot r}\right)^2 \cdot G_{rdr}(\theta_{rdr}, \phi_{rdr}) \right\} \quad (15)$$

$$P_{rdr}^{inc} = P_{rdr} \cdot \beta \cdot \left\{ G_{tag}(\theta_{tag}, \phi_{tag}) \cdot \left(\frac{\lambda}{4\pi \cdot r}\right)^2 \cdot G_{rdr}(\theta_{rdr}, \phi_{rdr}) \right\}^2 \quad (16)$$

We assign an intermediate variable α ,

$$\alpha = \left\{ G_{tag}(\theta_{tag}, \phi_{tag}) \cdot \left(\frac{\lambda}{4\pi \cdot r} \right)^2 \cdot G_{rdr}(\theta_{rdr}, \phi_{rdr}) \right\}, \quad (17)$$

and assume that the backscatter modulation efficiency $\beta = 1.0$ (though any value of $\beta > -44dB$ is sufficient). Substituting α into Equations 11 and 13 and expressing them in logarithmic units yields a pair of inequalities,

$$P_{tag}^{inc}(dB) = P_{rdr} + \alpha \geq -18dBm, \text{ and} \quad (18)$$

$$P_{rdr}^{inc}(dB) = P_{rdr} + 2 \cdot \alpha \geq -80dBm. \quad (19)$$

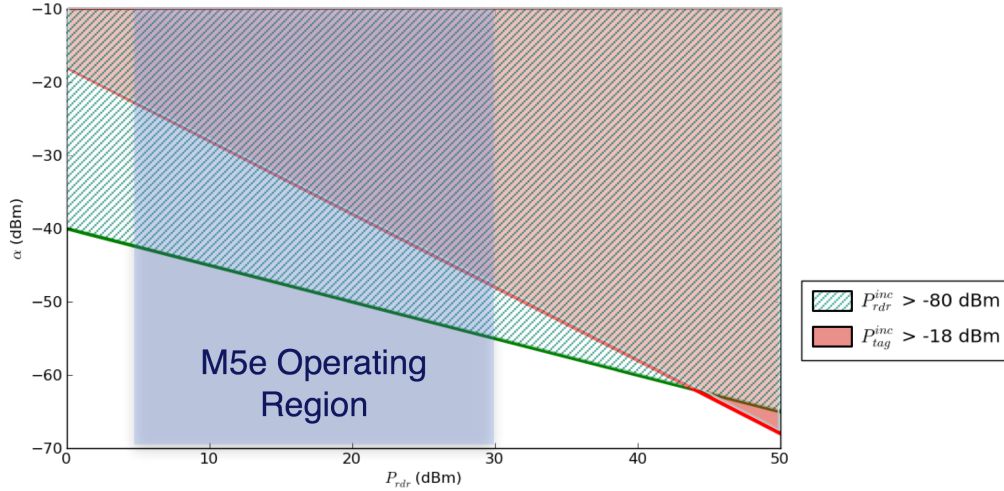


Figure 20: For any supported ThingMagic M5e reader output power (5 to 30 dBm), if the power incident at the tag (P_{tag}^{inc}) exceeds its power-up threshold, then the power incident at the reader (P_{rdr}^{inc}) also will exceed the necessary reader sensitivity.

This system of inequalities is plotted in Figure 20, and shows that for any supported ThingMagic M5e reader transmit power (5 to 30 dBm), if the power incident at the tag exceeds its power-up threshold ($P_{tag}^{inc} \geq -18dBm$), then the power incident at the reader will also exceed the necessary reader sensitivity ($P_{rdr}^{inc} \geq -80dBm$) and the reader should accurately detect the tag’s positive response. This condition is known as “forward-link limited;” the forward-link power-up threshold is the limiting condition for a positive tag detection,

$$P_{tag}^{inc} \geq 15\mu W (-18dBm). \quad (20)$$

Thus, the probability of a positive tag detection (a binary random variable \mathbf{d}) given the estimated power incident at the tag ($\hat{P}_{tag}^{inc}(x)$), calculated from the 6-DoF pose, x) would be given by,

$$p(\mathbf{d}|\hat{P}_{tag}^{inc}(x), x) = \begin{cases} 1.0 & \text{if } \hat{P}_{tag}^{inc}(x) \geq 15\mu W (-18dBm) \\ 0.0 & \text{otherwise.} \end{cases} \quad (21)$$

This distribution is shown in Figure 21 and is a rudimentary RFID sensor model.

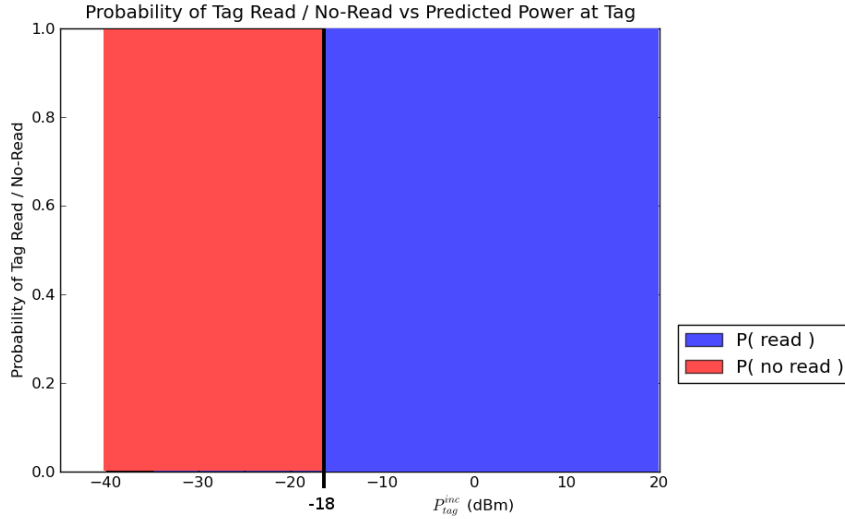


Figure 21: If the predicted power incident at the tag is greater than -18 dBm, then ideally the probability of a tag detection should be 1.0 (certainty). If the predicted power incident at the tag is less than -18 dBm, then the tag detection probability should be 0.0 (impossible).

2.2.1 Tag Detection Regions

In effect, Equation 21 defines a “tag detection region” in 6-DoF space (of relative pose between the reader and tag antenna). For example, consider a tag with an ideal dipole antenna sitting in the xy-plane of the reader’s antenna ($\phi_{rdr} = 90^\circ$), where the tag is oriented such that it’s gain is constant (ie. $\phi_{tag} = 90^\circ$). Under these conditions, $G_{tag} = 1.5$ is constant and the likelihood of a tag detection depends only on the tag’s $\langle x, y \rangle$ position in the reader antenna’s xy-plane. We can make the tag detection region bigger or smaller by adjusting the reader transmit power (P_{rdr}). In Figure 22 we show the tag detection regions for an ideal (patch) reader antenna with 5.5 dB gain and 8 dB *FBR* (ie. the *S9025P* specification), and an ideal dipole with 1.76 dB gain and 8 dB *FBR* (ie. the the Alien “Squiggle” tag antenna).

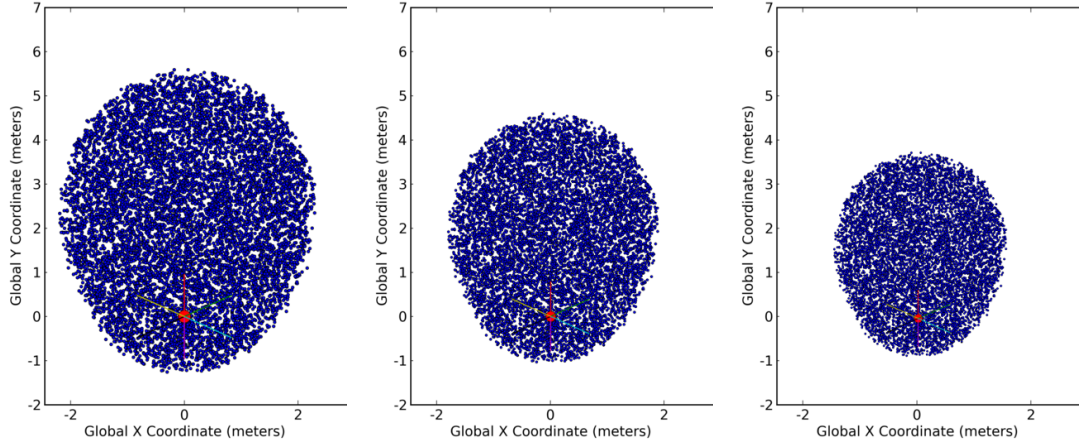


Figure 22: Detection regions predicted by the Friis equation with $\phi_{rdr} = \phi_{tag} = 90^\circ$, a *S9025P* reader antenna, and a Alien “Squiggle” tag antenna, and P_{rdr} equaling **Left:** 30 dBm, **Middle:** 28 dBm, and **Right:** 25 dBm. Note that the tag detection regions for the various powers are larger for the higher power, and smaller for the lower powers.

During our early work with UHF RFID, commercial readers did not provide RSSI measurements. As a proxy to RSSI, we pragmatically adjusted the reader transmit power (adjusted the tag detection region) to determine the “minimum required reader power” (MRRP) that still resulted in a positive tag detection [37]. Similar to RSSI, MRRP provided us with a scalar value that (approximately) indicated the strength of a tag’s response. However, MRRP and RSSI were inversely related; a low MRRP indicates that the tag is very near the reader, which would yield a strong RSSI on modern UHF RFID readers. When describing our earlier work (in Chapters 3 and 5), we will occasionally refer to MRRP.

2.2.2 Data-Driven Evaluation

In practice, tag detection likelihoods are as deterministic as pictured in Figure 21. To evaluate the tag detection sensor model, we performed semi-autonomous ground-truth data captures using the PR2 robot (as shown in Figures 23, 24, and 25). The PR2 robot had a total of three reader antennas: one Laird Technologies S9025P on each shoulder (articulated with Robotis servos) and one on the PR2’s pan-tilt head. We placed two Alien Squiggle tags, each at different orientations (“vertical” and “horizontal,” as shown in Figure 23), on a large cardboard box in a 12×8 meter large, empty room (relatively ideal, free-space conditions). We positioned the tags at the same height as the robot’s RFID reader antennas.

We used the ROS (Robot Operating System) *gmapping* software package to perform Simultaneous Localization and Mapping (SLAM); this allows us to track the position and orientation of the robot on a map of the room to within 10cm in position and 5° in orientation. Before performing the data captures, we manually positioned the robot’s end effector onto the tags to obtain their 6-DoF poses (in the map’s frame) using the *gmapping*’s robot pose estimates and the robot’s forward kinematics – as shown in Figure 23.



Figure 23: Using simultaneous localization and mapping (SLAM), we track the robot’s pose on a map (right). Using the forward kinematics of the robot’s arms, we can determine the pose of the tag in the map frame (left and middle).

To perform the data captures, the PR2 autonomously navigated to numerous positions and orientations on a 1-meter grid within the room. At each position and orientation, the PR2 slowly articulated (panned) one of its robot-mounted reader antennas and captured RFID sensor readings targeting one specific tag using a reader transmit power of $P_{rdr} = 30dBm$. The same antenna then repeated the process for the second tag. The grid locations and orientations for each capture are shown in Figures 24 and 25.

During the captures, we obtained the 6-DoF position and orientation of the robot’s reader antennas in the map frame (by way of the SLAM robot pose estimates and proprioception regarding the antennas’ mounting). Thus, for each and every RFID measurement, we used ROS’s *tf* software package to compute the 6-DoF relationship between the relevant reader antenna and tag. Knowing the 6-DoF pose parameters, the type of tag antenna, and type of reader antenna, we used the forward-link Friis equation (Equation 3) to calculate an estimate of the power incident at the tag,

$\hat{P}_{tag}^{inc}(x)$ (where we make the equation explicitly a function of the relative 6-DoF pose between reader antenna and tag, x).

Using this data, we can discretize $\hat{P}_{tag}^{inc}(x)$ into bins (5 dBm wide) and compute the ratio of positive tag detections to total detection attempts for each bin to yield a data-driven tag detection likelihood, $p(\mathbf{d}|\hat{P}_{tag}^{inc}(x), x)$. The tag detection likelihood from all 300,000 tag measurements is shown in Figure 26. Note that this data-driven distribution deviates from the deterministic condition suggested by the forward-link Friis equation (Equation 21). In fact, it is well-known that the Friis equation is a crude approximation to real-world radar system performance [136], even under relatively ideal scenarios (like this). We know from the literature that the Friis equation further degrades when tagging objects or in the presence of clutter [48]. In Section 2.5.5, we show that tagging objects causes further deviations from the ideal, Friis-predicted tag detection model in Equation 21.

2.3 RSSI Model

When a tag is successfully detected, modern UHF RFID readers provide a scalar value called “received signal strength indicator” (RSSI) that indicates the strength of a tag’s response. From communications with engineers at ThingMagic [142], we know that the relationship between the power incident at the reader (P_{rdr}^{inc}) and RSSI is linear when expressed in logarithmic (dB) units,

$$RSSI \propto P_{rdr}^{inc} \text{ (dB)} \quad (22)$$

However, the RSSI measurement is dimensionless; it is *not* calibrated to a particular physical unit (eg. dBW or dBm). To ascertain the true nature of this relationship, we turn to data involving actual tags. We connected a M5e reader in monostatic mode (same antenna for transmission and reception) at 30 dBm transmit power, through a variable attenuator, to a S9028PC12NF patch transmit antenna. We read an Omron Gen2 tag at a fixed position at a distance of 20cm while varying *only* the attenuation. This scenario and its results are depicted in Figure 27. The measured RSSI has some interesting properties. First, no positive reads possess RSSI readings below 71. Second, the RSSI saturates in the low-to-mid 100’s, giving the sensor a dynamic range of approximately 33. Third, the relationship between these extremes is highly linear with respect to the attenuation (A),

$$RSSI = 1.51 \cdot A + 111.1. \quad (23)$$

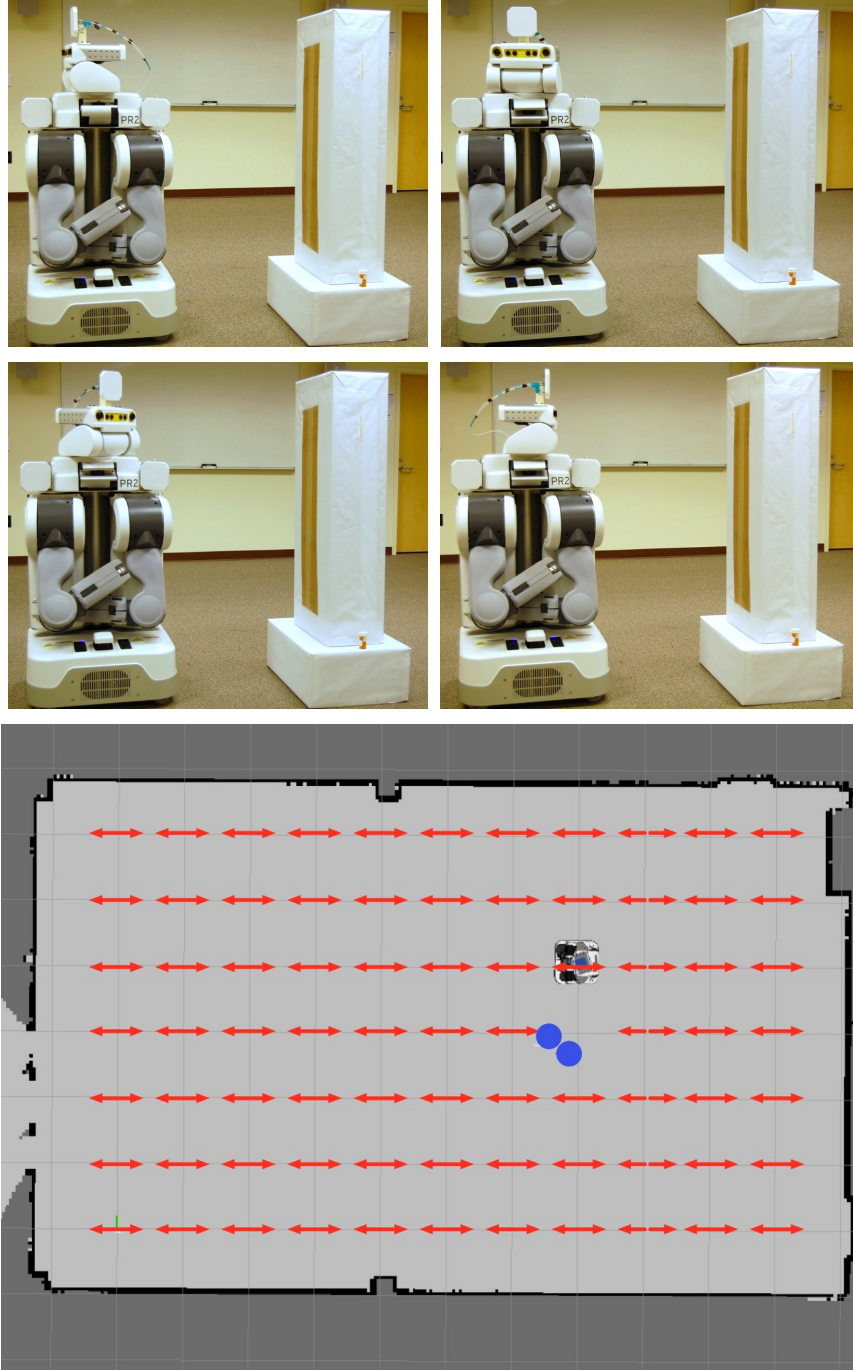


Figure 24: The robot pans a fixed, head-mounted antenna through 360° rotations over a 30-second duration, performing RFID queries of two specific tags (top). Pan captures are made at each of the (red) positions and orientation in a room. The robot knows the poses of the tags in the room (blue dots), as indicated (bottom). The robot recorded a total of 50,000 RFID sensor measurements for each tag, resulting in a total of 100,000 measurements.


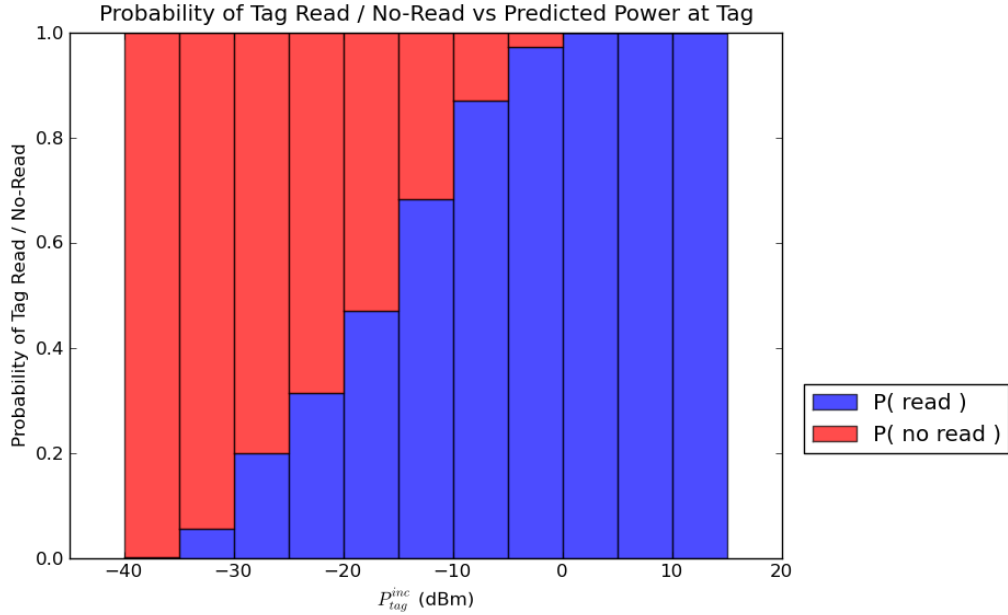


Figure 25: *Top Row:* The robot pans each of two articulated, shoulder-mounted antennas through 130° over a 26-second duration, performing RFID queries of two specific tags. *Bottom Row:* Pan captures are made at each of the (red) positions and orientation in a room. The robot knows the poses of the tags in the room, as indicated (blue dots). The robot recorded a total of 50,000 RFID sensor measurements per tag per antenna, resulting in a total of 200,000 measurements.

However, note that the attenuation is applied *twice* in the signal path for this measurement: once during transmission and once during reception. This means that the *backscattered* power at the reader (ie P_{rdr}^{inc}) computed by the forward-backward-link Friis equation in Equation 5 decreases 2 dB for every 1 dB of additional attenuation. Also, the attenuation is dimensionless, so the proportional relationship between attenuation and actual power (in dBm) is related (through some unknown



n

Figure 26: Discretizing $\hat{P}_{tag}^{inc}(x)$ into bins (5 dBm wide) and computing the ratio of positive tag detections to total detection attempts for each bin, we can develop a data-driven tag detection likelihood, $p(\mathbf{d}|\hat{P}_{tag}^{inc}(x), x)$, based on the Friis equation.

offset, κ_1) by

$$P_{rdr}^{inc}(dBm) = 2 \cdot A + \kappa_1, \text{ so} \quad (24)$$

$$A = 0.5 \cdot (P_{rdr}^{inc}(dBm) - \kappa_1) \quad (25)$$

In essence, the slope measured in Figure 27 is exactly *twice* the value of the slope between P_{rdr}^{inc} (in logarithmic units) and RSSI, so substituting Equation 25 into Equation 23 yields

$$RSSI = 0.75 \cdot P_{rdr}^{inc}(dBm) + \kappa_2. \quad (26)$$

We still do not know the intercept (κ_2) for this linear relationship; next, we will show how to use our data from Section 2.2.2 to determine its value.

2.3.1 RSSI Model Fitting and Evaluation

For the ground truth data captures from Section 2.2.2, we obtained RFID measurements (both tag detection and RSSI) for various 6-DoF poses between the reader antenna and tag (x). Assuming an ideal reader antenna radiation pattern (S9025P patch antenna) and an ideal tag antenna radiation pattern (dipole Squiggle tag), we can compute the Friis-predicted power incident at the reader,





Figure 27: We hold a RFID tag in a fixed position and orientation 20cm from the reader antenna. We make RFID sensor measurements where we vary *only* the attenuation, A (top), to get an idea of the true relationship between power incident at the reader and RSSI (bottom). Note that the attenuator is applied to the RF signal twice.

$\hat{P}_{rd}^{inc}(x)$, using the forward-backward-link Friis equation (Equation 5). In Figure 28, we plot the measured RSSI versus our calculated power incident at the reader antenna (\hat{P}_{rd}^{inc} in dBm), as well as the probability of successful tag detection versus $\hat{P}_{rd}^{inc}(dBm)$ (assuming the tag backscatter efficiency, $\beta = 0dB$).

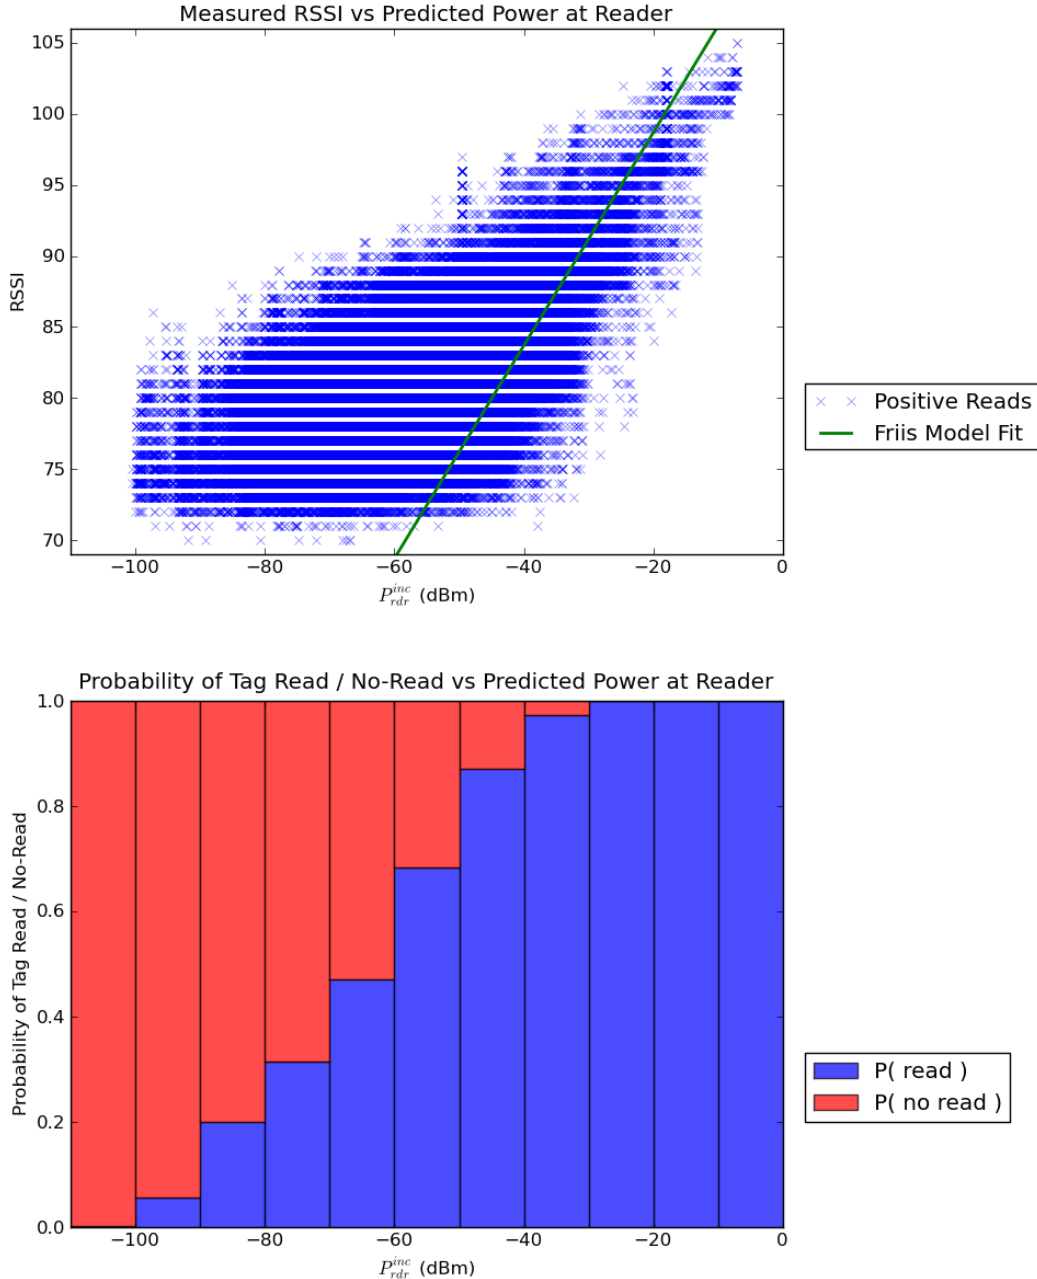


Figure 28: *Top:* A plot of measured RSSI versus the Friis-predicted power incident at the reader (P_{rd}^{inc}) for all 300,000 RFID reads across three antennas and two tags. *Bottom:* The probability of successfully detecting a tag versus the Friis-predicted power incident at the reader.

From Equation 26, we already determined that slope between RSSI and $\hat{P}_{rdr}^{inc}(dBm)$ should be 0.75, but that the intercept (κ_2) remains unknown. We can try to fit this value using the graph for RSSI vs $\hat{P}_{rdr}^{inc}(dBm)$. We could use the entirety of the data to perform a least-squares fit. However, many of the positive detections actually occur when the tag is located outside the 3-dB beamwidth of the antenna. Comparing the ideal patch antenna gain pattern (Figure 18) to the manufacturer-supplied gain pattern (Figure 19), we see that these configurations have large discrepancies in gain (by as much as 25+ dB). Thus, in these configurations, our calculated $\hat{P}_{rdr}^{inc}(dBm)$ may be significantly different than the true $P_{rdr}^{inc}(dBm)$. Instead, we limit ourselves to data with $RSSI \in [95, 100]$, where $\hat{P}_{rdr}^{inc} \in [-35, -10]$ dBm. We do this for three reasons. First, most sensor readings in this region had positive tag detections (the probability of a tag detection was high in this region according to Figure 28). Second, most of the tag detections in this regime occur when the tag was located in the 3-dB beamwidth of the antenna, so that the discrepancy between real and estimated P_{rdr}^{inc} is minimized. Third, $RSSI \in [95, 100]$ falls in the linear region of operation according to our data in Figure 27. Fitting the intercept (κ_2) via least-squares minimization using data from this region yields a final Friis RSSI model given by,

$$RSSI = 0.75 \cdot P_{rdr}^{inc}(dBm) + 113.7. \quad (27)$$

The Friis RSSI model is shown (green line) in Figure 28. To understand the various properties of RSSI versus the Friis RSSI model, consider the regions demarcated in Figure 29.

In Region I, the measured RSSI is less than the Friis calculation as the measured RSSI values begin to saturate (as was the case in our data in Figure 27). In Region II, the measured RSSI closely approximates that predicted by the Friis forward-backward-link calculations. In Region III, there are no readings. Either the tag power-up threshold was unmet, or the RFID reader lacked sufficient sensitivity / dynamic range to sense the tag response. Eighty-five percent (85%) of the readings in Region IV correspond to cases where the tag was physically located outside the 3dB-beamwidth of the interrogating patch antenna *or* the reader was located outside the 3dB-beamwidth of the tag antenna. Again, under these configurations there can be large discrepancies in the antennas' gain. For example, in the ideal patch antenna gain equations (from Section 2.1.3.2) the gain behind the patch antenna is 0. However, the manufacturer-specified gains (Figure 18) show that the gain behind

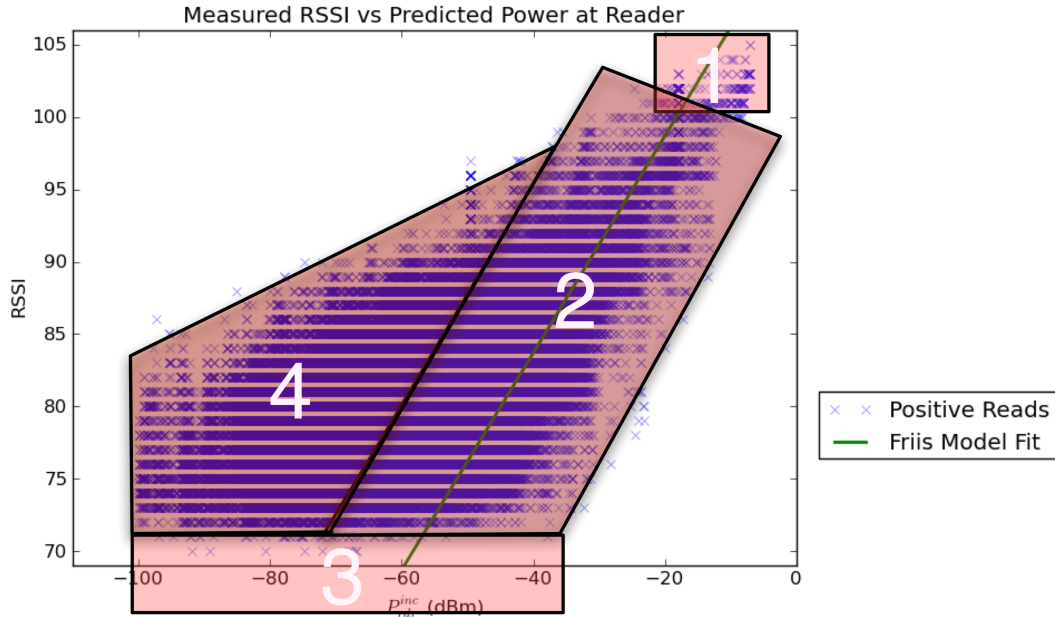


Figure 29: These four regions describe various conditions related to the Friis RSSI model.

the antenna is not only non-zero, but can also vary by as much as 20-30dB from the nominally-specified FBR . This discrepancy is due to the electromagnetic boundary conditions associated with real-world antenna designs (ie. the metal plane behind the radiating elements are of finite physical extent). Furthermore, the simple act of attaching the antenna to a large, unmodeled (and metal) robot likely exacerbates the issue; the antenna radiation patterns behind the antenna will likely experience even greater discrepancies.

2.4 Simulated RFID Sensor Model

The Robot Operating System (ROS) developed at Willow Garage provides extensive robot simulation tools, which can provide a complete PR2 robot in a Gazebo (physics-based) simulation environment [62]. The simulated PR2 can execute the same code and algorithms as a real-life PR2, allowing a roboticist to evaluate their code in a high-fidelity simulation environment *before* deploying to the real world. In this section, we show how to incorporate our Friis-based models (tag detection and RSSI) into this framework – this simulated sensor will be used to develop some of our later RFID behaviors in simulation before migrating the behaviors over to the real robot for subsequent testing.

To employ our models for tag detection and RSSI, we need to know four things for each simulated measurement: the 6-DoF relationship between the reader antennas and tag (x), the reader antenna gain pattern, the tag antenna gain pattern, and the reader transmit power. The system designer directly provides the antenna gain patterns (ie. an ideal patch antenna using the properties of the S9025P, and an ideal dipole Alien squiggle tag). The system designer specifies the 6-DoF pose between each antenna and one of the robot’s built-in coordinate frames; if the antennas are articulated, their 6-DoF relationship to these frames should be updated. The system designer also specifies the 6-DoF pose of the tag in some global (map) frame. When the ROS simulation environment is executed, the robot’s pose is tracked using SLAM; using ROS’s *tf* package, we can compute the 6-DoF pose between the reader antenna and tag. Finally, the system designer selects a reader transmit power. Using these quantities, we can compute the Friis-estimated power incident at the tag (\hat{P}_{tag}^{inc}) and reader (\hat{P}_{rdr}^{inc}). From the former, we can immediately employ the Friis-criterion for tag detection (Equation 20),

$$\hat{P}_{tag}^{inc} \geq 15\mu W(-18dBm). \quad (28)$$

If the tag was detected, we can use the Friis-based RSSI estimate using the relationship determined from our (relatively) ideal conditions (Equation 27),

$$RSSI = 0.75 \cdot \hat{P}_{rdr}^{inc}(dBm) + 113.7. \quad (29)$$

During model evaluation (Sections 2.2.2 and 2.3.1), we noted that the sensor measurements were highly stochastic. To create a more realistic simulation, we can draw from the distributions given by our data captures under relatively ideal conditions. For example, to first determine tag detectability, we can draw from the distribution $p(\mathbf{d} \mid \hat{P}_{tag}^{inc}(x), x)$ shown in Figure 26. If the tag was detected, then the RSSI value can be drawn from a Gaussian distribution fit to the data (ie. the sample mean and standard deviation),

$$P(RSSI \mid \hat{P}_{rdr}^{inc}(x), x) = \mathcal{N}(\mu, \sigma). \quad (30)$$

Later, we will use this stochastic sensor model when developing robot behaviors to acquire some insights into the algorithms (and design trade-offs) before deploying the system on the real robot. We released all of the code necessary to simulate the RFID antennas on a PR2 (or any other) robot

under open-source (BSD) licenses in the *gt-ros-pkg* repository so that other people can evaluate our RFID algorithms.

2.5 *Model Deviations*

In developing our Friis-based RFID sensor model, we considered relatively ideal (free-space) circumstances. Aside from the natural room features (walls, ceiling, floor), we eliminated all objects from our data capture environment that could cause reflections (multipath) or obstructions (shadowing). During our analysis, we assumed ideal (or perfectly specified) antennas operating without interference from nearby structures, and we assumed that tags respond only when the instantaneous incident power exceeds a fixed threshold. Deviations from these assumptions can have significant impact on all aspects of the model, whether tag detection or RSSI measurements. In fact, looking at the real-world data in this relatively ideal case, we *already* observed deviations from the ideal Friis-based model.

In this section, we enumerate several factors that can cause deviations in the Friis models.

2.5.1 **Environmental Conditions**

In radar systems, it is commonly held that the Friis equations are crude approximations to system performance that can be off by an order of magnitude or more – and this is for systems operating in the open air (eg. to track airplanes) [136]. The discrepancies originate from the difference between free-space assumptions (a vacuum) and real-world conditions, where material composition and RF properties (ie. permeability and permittivity) are highly non-uniform and complex. These conditions cause a slew of non-ideal RF propagation effects that are unaccounted for by the Friis models: multipath (reflections), fading (attenuation), shadowing, scattering, and diffraction [80].

In RFID systems, where the reader and tags are being placed in cluttered environments (for example, a home instead of airplanes flying in the open sky), non-ideal propagation effects can become quite significant. Even the simple act of relocating a tagged object to a new position or orientation can drastically alter the unmodeled RF propagation aspects of the Friis model.

If the environment's precise material composition and RF properties were known, then high-fidelity RF simulation software packages may be able to estimate the necessary path loss parameter in the Friis equation through finite element simulations of Maxwell's equations. Unfortunately,

sensing the composition and RF properties of the environment, particularly the non-visible composition, is a difficult challenge and may be intractable for robots operating in unstructured human environments. A large body of the RFID literature seeks to study these effects, but by-and-large it remains an unsolved challenge [92].

2.5.2 Robot Properties

The environment's macro-scale composition affects RF propagation, but its micro-scale composition nearby an antenna can dramatically affect the antenna's gain. The closed-form ideal radiation patterns (Section 2.1.3) are derived from a careful analysis of Maxwell's equations given the boundary conditions imposed by the antenna's (ideal) material properties. When placed into its operating environment, the antenna's nearby material properties are no longer free space and will alter the antenna's radiation pattern. For some materials (eg. the protective plastic coatings, called radomes) the alterations may be minimal. However, other material properties (particularly metal, which introduce significant electromagnetic boundary conditions) can cause dramatic changes – to the point of nullifying the gain or causing significant reflections. These unmodeled effects can create severe deviations from the ideal and manufacturer-supplied radiations patterns.

System designers should give careful consideration given to robot material properties with mounting the RFID antennas. As we will see in Chapter 3, there are systemic differences in the tag detection probabilities and RSSI measurements between the left and right shoulder-mounted antennas. The disparity persisted even when replacing out the antennas, switching antenna mount points, and changing out the RFID reader – the right side of the robot always had a greater tag detectability likelihood. We attribute these observations to complex interactions between the reader antenna and unseen robot material properties (there were moving cables and wires behind the right antenna *inside* the robot). In general, system designers need to exercise caution and be cognizant of the robot's material properties; particularly, designers should try to minimize the amount of nearby metal (especially in front of the patch antenna).

2.5.3 Tagging Objects

When tagging an arbitrary object – whether walls (infrastructure tags), people, medication, TV remotes, etc – we are purposefully placing the tag antenna near unmodeled materials. Particularly

when tagging objects in home environments, the material composition near the tag is likely to be highly non-uniform and many objects also contain metal (eg. a printed circuit boards in consumer electronics). Exacerbating the issue, there are many possible ways to tag objects – the tagging may not be consistent between user applications. Further still, the *curvature* of objects often necessitates significant antenna deformation, where the physical properties of the antenna itself are changed. For example, wrapping a tag around an object may make the tag behave more like a helical antenna than a dipole. Again, under some conditions (such as placing a tag on metal or self-overlapping) these new tag or material properties may nullify the tag gain and render the tag completely unreadable.

To address *some* of these concerns, tag manufacturers have designed a broad spectrum of tags with different properties: size, radiation pattern, etc. Some tags are even designed to be affixed to metal. We confirmed that the ThingMagic RFID reader could interact with 32 different UHF RFID tags, shown in Figure 6. Often, the tag manufacturer does not provide a detailed specification for the tag antenna. The remaining option is to rely on simulation (ie. Ansys HFSS) or anechoic chamber measurements to get a good measure of the tags’ radiation patterns.

Perhaps most poignantly, placing the tagged objects in new environments may completely alter the tag antennas’ properties and nearby RF propagation, resulting in unmodeled behavior. One contribution of this thesis is the design of robot behaviors that can function under these conditions.

2.5.4 Tags’ On-Board Energy Storage

In the Friis formulation for tag detection (Equation 20), we assumed that the tag’s ability to respond was solely dependent upon the *instantaneous* power available at the tag. Most tags have a small amount of on-board energy storage capacity that can accumulate energy while other tags are being read or between successive sensor readings. Some tags have relatively large power requirements since they also possess general purpose computation and/or sensors [117]. For these tags, the on-board energy storage is absolutely crucial to ensure tag power-up, and instantaneous power is a poorer predictor of a tag response (it could take several seconds for a tag to respond). Rather, the instantaneous power will correspond to the *rate* of tag detection events. This property of RFID tags has even been used as a RFID sensor model as a proxy for RSSI – where the *frequency* of tag detections for a particular tag ID corresponds to RSSI [61].

2.5.5 Data-driven Examples

To illustrate the variability experienced when tagging *objects*, we turn to data. Consider the scenario depicted in Figure 30. We tagged five different objects: an orange medication bottle, a white vitamin bottle, a red cup 75% full of water, a set of keys, and a TV remote. The keys and TV remote both employed tags designed to operate on or near metal, as standard Alien Squiggle tags were undetectable when affixed to these objects. We tagged the other objects with Alien Squiggle tags. For the medication and vitamin bottles, it was necessary to wrap the tags around the bottles to conform to their surfaces.




Figure 30: *Top Row:* The PR2 robot performs data captures using its two shoulder antennas for each of the five tagged objects. *Bottom Row:* Pan captures are made at each of the (red) positions / orientation in a room. The robot knows the ground-truth poses of the tags in the room (blue dots).

We captured 25,000 RFID sensor readings from each of the shoulder antennas for each tag, for a total of 250,000 RFID sensor readings. The robot performed its panning captures from the positions and orientations indicated (with red arrows) on the map in Figure 30. The tag detection probabilities deviate significantly from the relatively ideal conditions explored previously, as shown in Figure 31.

For yet another example, consider the scenario depicted in Figure 32. We tagged *thirty seven* different objects located on a bookshelf in the Healthcare Robotics Laboratory. We positioned EL-E at 36 different locations within the room and queried for all tags that could be read in that location while panning EL-E's antennas. In Figure 32 we see that the number of tagged objects detected is highly dependent both on EL-E's location within the room as well as the particular tagged object. Some objects, like the plastic bottles, plastic containers, and toothbrush were detected at nearly every location. Other objects, like the orange medication bottle (with tag wrapped around itself in an overlapping fashion), the metallic toothpaste bottle, and the TV remote went virtually undetected.

2.6 Conclusions

In this chapter, we examined a UHF RFID sensor model based on the Friis transmission equation. We used the Friis equation to build a deterministic model that *should* predict UHF RFID sensor measurements. However, as is well-known in the radar community, the Friis model is a crude approximation to real-world performance. Performing over 600,000 total ground-truth RFID sensor measurements under both ideal and non-ideal conditions, we confirmed that (1) the RFID reads are stochastic in nature, (2) tagging objects can dramatically affect the tag detection probability and the resultant RSSI values. Later in Chapter 3, we will show how researchers typically use these data captures to build data-driven, 2-DoF (planar) probabilistic models of tag detectability and RSSI, and how these models degrade when tagging objects. Finally, as in the radar communities, we will use insights gleaned from the Friis equations in this chapter to explain the operating principles behind our novel optimization-based robot behaviors used to search out, locate, and approach RFID tagged objects in Chapter 4.

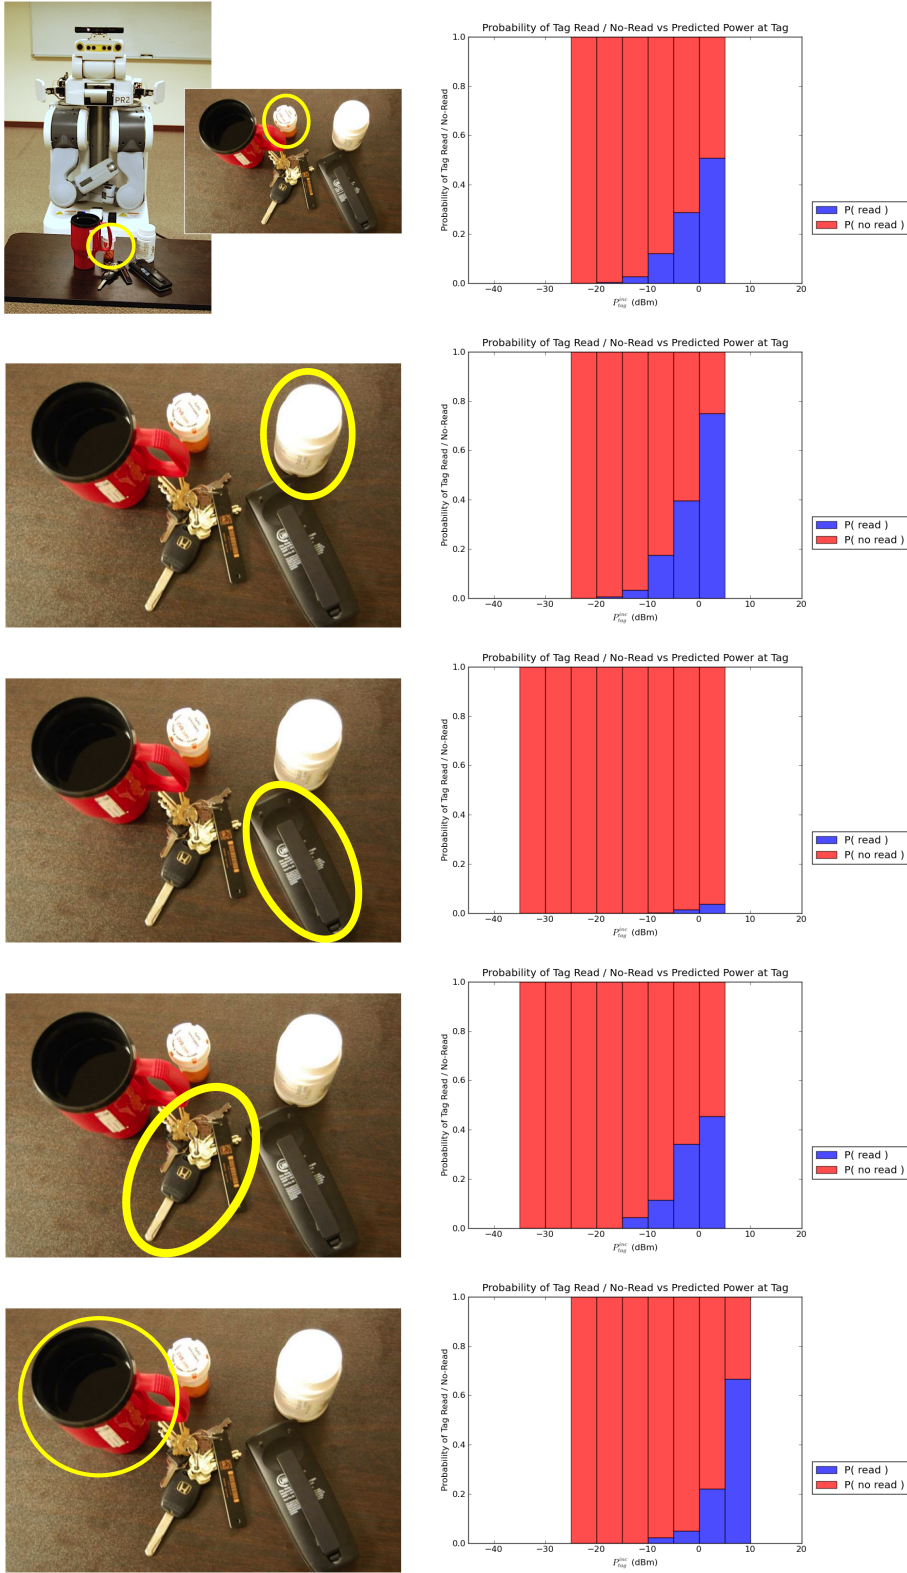


Figure 31: The tag detection probability given the Friis-predicted power incident at the tag, $p(\mathbf{d}|\hat{P}_{tag}^{inc}(x), x)$ (right), substantially deviates from the ideal scenario depicted in Figure 21 when tagging objects (left) such as:

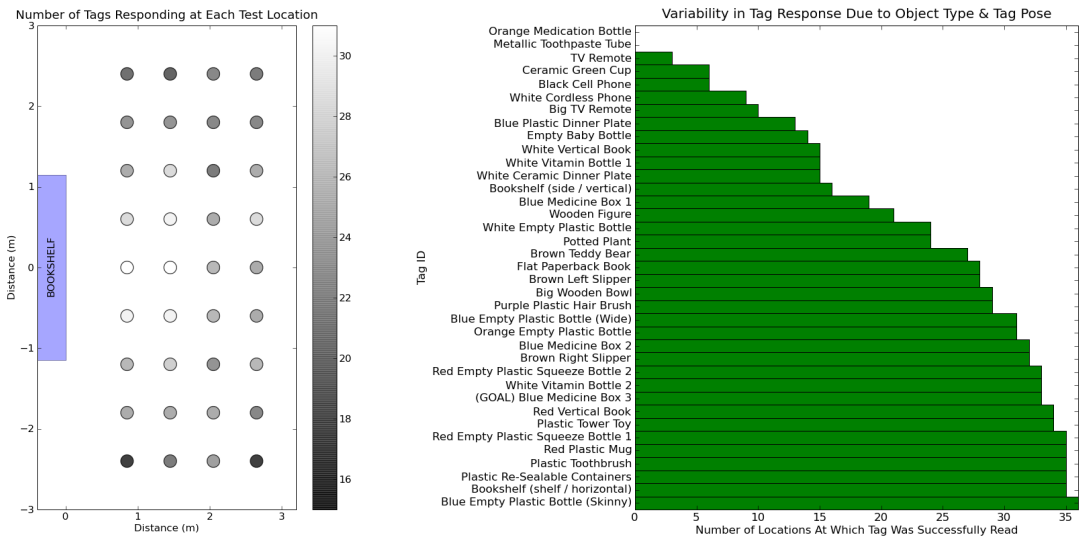


Figure 32: Thirty-seven tagged objects are sitting on a bookshelf (top left). EL-E is moves to various positions and pans her RFID antennas (top right). The number of tags detected at each location is shown (bottom left), and a tag-by-tag breakdown shows that there is a lot of variability depending on the particular tagged object (bottom right).

CHAPTER III

PROBABILISTIC APPROACHES TO LOCALIZE UHF RFID TAGS

Much of the UHF RFID sensing work to date has focused on simultaneous localization and mapping (SLAM) techniques to localize (track) the robot's pose on a map (using tags at known positions) and/or to localize tags relative to the robot. The probabilistic techniques employed belong to the class of Bayesian filters [143], which include Gaussian (Kalman) filters and particle filters. These algorithms incorporate evidence (measurements) over time as the robot moves around the environment to estimate the value of a state variable, such as the position of the robot or the position of a tag. The algorithms require both a control (robot motion) model and a measurement (RFID sensor) model. The latter is the main subject of this chapter.

In this chapter, we make the following contributions:

Bayesian Formulation and Related Work

We examine a common particle filter formulation for UHF RFID tag localization. As is common in the literature, the sensor model is a planar, two degree-of-freedom (2-DoF) model that predicts the likelihood of obtaining an RFID sensor measurement given the $x - y$ or $d - \theta$ position of the UHF RFID tag relative to the reader transmit / receive antenna. We briefly examine several 2-DoF sensor models, both from the literature as well as our own. Our sensor model is based on the Friis radar equation from Chapter 2. We summarize the results of the various sensor models, which we tested under nearly-ideal, controlled conditions: tags on RF-friendly materials, in isolation, and under controlled orientations. These conditions may be characteristic of infrastructure tags (eg mounted on a wall) used for localization, where the system designer can exercise control.

Discussion of Assumptions and Possible Limitations

We discuss some of the limitations of the aforementioned sensor models. In particular, we show that the 2-DoF planar sensor models suffer from severe deterioration when tags are affixed to objects with non-ideal RF material properties and when the objects are placed

in new environments with uncontrolled orientations. Later in Chapter 4, we show that the model deviations incurred when tagging objects cause poor performance relative to our new optimization-based approach to discover, locate, and approach RFID tagged objects.

3.1 Bayesian Filter Formulation

Bayesian filters are a general class of techniques for recursively computing a belief (probability distribution) about the state (x_t) of a mobile robot at time t by taking into account all previous sensor readings $z_{1:t}$ from time $[1, t]$, all previous robot motions $u_{1:t}$ from time $[1, t]$, and all previous state estimates $x_{1:t-1}$. This belief distribution is commonly known as a *posterior* and is given by $p(x_t|z_{1:t}, u_{1:t}, x_{1:t-1})$ [143]. Making a first-order Markov assumption, which assumes that past and future data are independent when we know the previous state x_{t-1} , we can apply Bayes' rule to reduce the complex posterior probability distribution into the more manageable form

$$p(x_t|z_{1:t}, u_{1:t}, x_{1:t-1}) = \eta \cdot \underbrace{p(x_t|u_t, x_{t-1})}_{\text{motion}} \cdot \underbrace{p(z_t|x_t)}_{\text{sensor}}, \quad (31)$$

comprised of a motion (or control) model, a sensor (or measurement) model, and a normalization factor η . The motion model describes how the system's state updates given a control (u_t) and the previous state (x_{t-1}). The sensor model describes the likelihood of obtaining our most recent sensor measurements given the current state; the sensor model is the biggest distinguishing factor among the various RFID sensing techniques and is the subject of this chapter. The normalization factor η is required to turn our belief into a proper probability distribution [143].

3.2 Related Work

There are many ways to realize a Bayesian filter depending on the type of probability distribution (eg. Gaussian or non-parametric) and the underlying system models. For example, if the distributions are multivariate normals (Gaussians), then we are dealing with a Gaussian filter. Perhaps the best known Gaussian filter is the Kalman Filter (KF), which is a special case where the motion and sensor models are linear (with Gaussian noise) and the initial belief (x_0) is Gaussian [143]. Kalman filters have been used to great effect; for example, NASA's early Apollo and Polaris space programs championed Kalman filters shortly after their discovery for high-fidelity attitude estimation [64]. There are many extensions to the basic Kalman Filter that relax these assumptions or provide other

desirable properties; examples include: Extended Kalman Filters (EKF), Unscented Kalman Filters (UKF), and Information Filters. If the distributions are non-Gaussian, non-parametric distributions can be used to approximate the posterior to implement a Bayesian filter. For example, Histogram Filters use a histogram to represent the distribution, whereas particle filters use a weighted set of particles drawn from the posterior to represent the distribution [143].

The underlying system models (motion model, measurement model, and state variable x_t) also define the Bayesian filter. For mobile robots, the motion model is often a form of dead reckoning using feedback from wheel encoders; this may be either a differential drive or holonomic motion model depending on the type of mobile base being employed. The measurement model is often the source of differentiation between RFID methods; we will discuss several UHF RFID measurement models shortly. For RFID sensing (LF, HF, or UHF), the state variable (x_t) of interest is often the robot's pose on a map or the location of a tag (either relative to the robot or on a map). Often, these state variables are analogous; for example, if you are estimating the robot's pose on a map of fixed tags (of known location), then you could use the corresponding coordinate transformations to compute likely tag poses relative to the robot.

All of these types of Bayesian filters have been used in the RFID robotics literature.

- An Extended Kalman Filter is used to estimate the 3-DoF pose (planar position and orientation) of a robot operating on a carpet embedded with LF or HF RFID tags. The measurement model is geometric, a histogram based on data, or model based. [132]
- A Kalman Filter is used to track the position of a HF RFID tag relative to a nearby robot. The position of the tag is recorded to build a map of the environment for navigation, subsequent localization, etc. [22]
- Using multiple UHF RFID tag detection events to define “fingerprints” for robot pose estimation [153].
- Data-driven histograms, closed-form physics models, and tag read rate models are used to estimate a tag's pose relative to a mobile robot. We discuss these particular implementations below.

3.3 UHF RFID Sensor Models

Now, we are going to examine three different UHF RFID sensor models, $p(z|x)$, that have been used to localize a tag relative to a robot using particle filters. This includes one of our own design. The state variable x is the 2-DoF planar pose of the tag relative to the RFID reader transmit / receive antenna (monostatic mode). Using the coordinate conventions in Figure 14, the 2-DoF state variable can be expressed in polar coordinates as,

$$x = \begin{bmatrix} d \\ \theta \end{bmatrix}, \quad (32)$$

where d is the planar distance to the tag (ie. $r \cdot \sin(\phi_{rdr})$) and θ is the bearing toward the tag (θ_{rdr}); these quantities describe the 2-DoF pose of the tag relative to the reader antenna's coordinate frame. To estimate the position of the tag relative to the reader antenna, we make RFID sensor measurements targeting a single tag ID and update our estimate according to the recursive relationship given by Equation 31.

3.3.1 Data-Driven Tag Detection Regions

Early research on RFID mapping and localization used a data-driven approach to build tag detection regions as a sensor model. The measurement z was a binary random variable d_g that expressed the probability of detecting the specific tag g [54]. To construct the sensor model, $p(z|x) = p(d_g|x)$, researchers captured more than 12,000 RFID sensor measurements under various reader-tag conditions (values of x). The data capture occurred under relatively ideal conditions: a tag affixed to a cardboard box standing in free space. Researchers then discretized relative reader-tag poses into grid cells (d,θ) and counted the number of positive ($n_{d,\theta}^+$) and negative ($n_{d,\theta}^-$) detection events (d_g^+ and d_g^-) for tag g within the grid; the resulting sensor model was given by

$$p(z|x) = n_{d,\theta}^+ / (n_{d,\theta}^+ + n_{d,\theta}^-). \quad (33)$$

The researchers did not use this histogram model directly; rather, they used it to hand-designated regions of (approximately) equal tag detection probabilities, as shown in Figure 33. Their final sensor model had three regions: a circular high-likelihood detection region near the antenna, a

medium-likelihood detection region comprised of a 95° arc in front of the antenna, and a low-likelihood “background” detection region everywhere else.




Figure 33: Seminal work on UHF RFID tag localization by Hähnel *et al* used a mobile robot (left) to perform extensive data captures (middle) to determine regions of equal tag detection likelihood (right) [54]. [Figures reproduced with authors’ permission.]

We would like to point out several features about this seminal work:

- This was the first work to use a particle filter to localize UHF RFID tags relative the robot. No quantitative results were reported for this aspect of the work. However, the authors sought to perform subsequent global robot localization using the mapped-out tags and quantitative results *were* reported for this capability. Global *robot* localization error using odometry and RFID sensing was reported to be 50-125cm. We presume that individual tag localization relative to the robot would be around the same magnitude.
- They developed a 2-DoF sensor model based on extensive data capture; however, they ultimately relied on hand-designated regions that “conservatively approximated the histogram” [54].
- The tags used to generate the sensor model and those used for localization were placed in a controlled manner: consistent orientations on nearby walls (consistent material properties) with at least one line-of-sight path from the robot to the tag.

3.3.2 Tag Detection Regions Derived from the Friis Model

In 2008, we devised our own UHF RFID sensor model, $p(z|x)$, based on the Friis model of RF signal propagation rather than an empirically-generated histogram model. Our technique obviated the need for extensive empirical training. We use the Friis free-space RF propagation model, including

the RFID reader's transmission power and antenna characteristics, to predict far-field tag detection regions. Recall from Section 2.2: If we assume that the system is forward-link-limited, then according to the Friis forward-link equation (Equation 11) the tag will be capable of responding whenever the incident power exceeds a the tag power-up threshold, P_{tag}^{th} . Restating the equation:

$$P_{tag}^{inc} \geq P_{tag}^{th} \quad (34)$$

$$P_{rdr} \cdot G_{rdr}(\theta_{rdr}, \phi_{rdr}) \cdot \left(\frac{\lambda}{4\pi \cdot r} \right)^2 \cdot G_{tag}(\theta_{tag}, \phi_{tag}) \geq 15\mu W. \quad (35)$$

This formulation establishes a condition that predicts tag readability when given the relationship between the 6-DoF relative position between the tag and reader antenna. If we can exercise control over tag selection, tag orientation, tag gain, nearby material properties (free space), and antenna positioning, we can ensure that the only pertinent factors in this formulation are the relative 2-DoF *planar* position between the tag and reader antenna.

For example, if the tag has an omnidirectional antenna radiation pattern (constant gain) and the tag is located in the reader antenna's *xy*-plane (ie. at the same *height*), then the only two variable model parameters are the tag-reader distance ($d = r$) and the bearing toward the tag, $\theta_{rdr} = \theta$. This is also the case for an Alien Squiggle tag (dipole) oriented with its *z*-axis parallel to the reader's *z*-axis and located in the reader's *xy*-plane (ie. $\phi_{rdr} = \phi_{tag} = 90^\circ$). This reduces Equation 34 to

$$P_{rdr} \cdot G_{rdr}(\theta_{rdr}) \cdot \left(\frac{\lambda}{4\pi \cdot d} \right)^2 \cdot G_{tag} \geq 15\mu W. \quad (36)$$

This equation defines a 2-DoF planar region whose extent can be programmatically adjusted by altering the reader transmit power, as shown in the sample tag detection regions of Figure 34. We could use this condition as a deterministic sensor model. If the RFID sensor measurement is a positive detection ($z = d_g^+$), then the tag must be within the region; if it is not detected ($z = d_g^-$), then it must be outside, ie

$$p(z = d_g^+ | x) = \begin{cases} 1.0 & \text{if } P_{tag}^{inc} \geq 15\mu W \\ 0.0 & \text{otherwise.} \end{cases} \quad \text{and} \quad p(z = d_g^- | x) = \begin{cases} 0.0 & \text{if } P_{tag}^{inc} \geq 15\mu W \\ 1.0 & \text{otherwise.} \end{cases} \quad (37)$$

Prescribing a zero probability can cause particle filter degradation (starving). Instead, we assign a nominal background probability (similar to [54]) to allow for greater noise immunity. The final




Figure 34: Scaling the read power while making RSSI measurements alters the RFID detection regions. These diagrams show the RFID sensor model for an Alien Squiggle tag with $\phi_{rdr} = \phi_{tag} = 90^\circ$. Here we show the likelihood regions for positive tag detections (d_g^-) for $P_{rdr} = 30, 25, 20$ and dBm. The negative detection regions are simply the complement.

sensor model is

$$p(z = d_g^+ | x) = \eta \cdot \begin{cases} 1.0 & \text{if } P_{tag}^{inc} \geq 15\mu W \\ 0.6 & \text{otherwise.} \end{cases} \quad \text{and} \quad p(z = d_g^- | x) = \eta \cdot \begin{cases} 0.6 & \text{if } P_{tag}^{inc} \geq 15\mu W \\ 1.0 & \text{otherwise.} \end{cases} \quad (38)$$

Where $\eta = 1/1.6$ is a normalization factor.

When performing particle filtering, it is typically ill-advised to integrate additional sensor measurements when the robot is standing still (or hardly moving) [143]. From an RFID perspective, the intuitive explanation is that (all measurements being equal) too many measurements from one vantage biases our estimate toward the posterior sensed at that one location. However, our ability to programmatically alter the reader transmit power to extend the tag detection region offers us the ability to meaningfully integrate several RFID measurements from a single position and (in essence) reduce our uncertainty about the read region (ie. determine information about the read range). If we progress through n distinct reader power levels (P_{rdr}) and treat the RFID measurements as (conditionally) independent, then we could define a new sensor model $p^*(z|x)$ as the multiplicative combination of the reads,

$$p^*(z|x) = \prod_{i=1}^n p_i(z_i|x) \quad (39)$$

For example, if a single antenna were to perform measurements at $P_{rdr} = 30, 25, 20$, and 15

dBm and receive positive detections for the two highest powers, the combined sensor model would yield a resultant posterior probability distribution similar to the one obtained for RSSI measurements in Section 3.3.3. In this manner, the scaling the read power (at least anecdotally) serves as a proxy to later RSSI sensor models. We were the first to employ such a method for UHF RFID tag localization.

3.3.2.1 Evaluation

To evaluate our sensor model, we developed an antenna array with *six* MaxRad MP8906PTNF patch antennas ($G_{max} = 10dB$, $FBR = 15dB$, 90° 3-dB beamwidth) arranged with overlapping beamwidths, as shown in Figure 35.

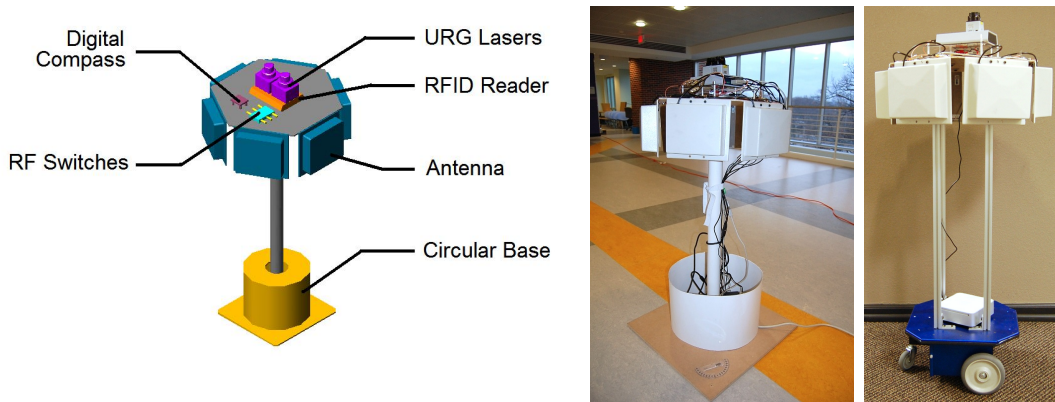


Figure 35: We employed an array of *six* reader antennas to use during the power scaling localization. The early prototype rig, with cad drawing (left) and photograph (middle), were early approximations to the ultimate desired robot (right). Unfortunately, the robot version was never realized / utilized.

We developed a particle filter implementation using the power-scaling tag detection sensor model. We performed all of the tests in a $10m \times 12m$ open room with ceiling height of $\approx 2.85m$. This environment was mostly free space, with the ceilings and reinforced concrete floors as the primary sources of multipath. We relocated the test rig to numerous locations (to simulate a robot path) throughout the room in $\approx 20cm$ increments. At each location, we recorded sensor measurements for *nine* different power levels evenly spaced between 30 dBm and 14 dBm. We obtained ground truth antenna and tag position (in a map frame) using a SICK LMS291 laser range-finder and a circular Hough transform to find the centroids of circular cardboard targets surrounding both the test rig and the tag. We measured the orientation of the robot (in the map frame) using digital compass and cross-validated with manual angle measurements relative to room “north”.

We transform the robot pose changes (from the map frame) into the robot’s local frame and use them to simulate odometric control updates (with Gaussian error profiles); we fed these control updates into the particle filter as the motion model. The particle filter’s initial state consisted of 1000 particles uniformly spaced over a $20m \times 20m$ grid, centered about the robot. We performed Control updates and RFID read events from all 6 antennas over 9 read power levels for every 20 cm of robot movement. The measurement update rate is dominated by tag read times (1-3 ms) and antenna switching times (6ms), which yielded a full 6-antenna update time of 50-100ms.

At each position, we performed an update of the particle filter (Equation 31) according to the motion model and our measurement model (Equation 39), which yielded a posterior probability distribution over tag pose relative to the robot. We estimated the 2-DoF tag pose as the mean of this distribution,

$$\hat{x} = \begin{bmatrix} \hat{d} \\ \hat{\theta} \end{bmatrix} = \text{E} (x_t \mid z_t, u_t, x_{t-1}). \quad (40)$$

We show several particle filter updates (including the robot’s movement, the posterior distribution, and the estimated tag pose) in Figure 36. For each location, we used ground truth tag pose to compute the true distance (d^*) and bearing (θ^*) to the tag. We define the following error measures,

$$d_{|err|} = |\hat{d} - d^*|, \quad (41)$$

$$d_{relative} = \frac{d_{|err|}}{d^*}, \text{ and} \quad (42)$$

$$\theta_{|err|} = |\hat{\theta} - \theta^*|, \quad (43)$$

for the distance error magnitude ($d_{|err|}$), relative distance error ($d_{relative}$), and angular error magnitude (θ_{err}). We calculated the estimation error values after each successive update of the particle filter; we report the results for the entire data set ($\forall d^*$) in Table 1. We also segmented the results into two categories¹ – error for robot-to-tag distances less than four meters ($d^* \leq 4m$) and error for robot-to-tag distances greater than four meters ($d^* > 4m$).

For robot-to-tag distances less than four meters ($d^* \leq 4m$), the RF power received by the reader was predominantly due to the line-of-sight path between the robot and tag; multipath was insignificant, and thus we acquired stronger tag readings. At these close ranges, the mean distance

¹The actual model we developed also accounted for multipath. We did this by assuming all multipath was constructive and summing the power incident from signal paths that included reflecting surfaces [38].

error magnitude ($d_{|err|}$) was 40cm ($\sigma = 20cm$), and the mean angular error magnitude ($\theta_{|err|}$) was 5.1° ($\sigma = 3.6^\circ$). When the distance exceeded four meters ($d^* > 4m$), the mean distance error magnitude ($d_{|err|}$) and angular error magnitude ($\theta_{|err|}$) were worse: 105cm ($\sigma = 39cm$) and 7.2° ($\sigma = 4.5^\circ$), respectively. The degradation in estimation performance at longer reader-to-tag distances is a trend we will see in some of our later behaviors in Chapter 4. The error values reported for our sensor model are consistent with those reported by other data-driven models [54].

Table 1: Particle Filter Tag Localization Results

Distance	Error Measure	Line-Of-Sight Only mean (std dev)	With Ground & Ceiling Bounce mean (std dev)
All d^*	$d_{ err }$	0.71 m (0.45 m)	0.69 m (0.42 m)
	$d_{relative}$	16.2% (6.8%)	15.9% (6.6%)
	$\theta_{ err }$	6.11° (4.19 $^\circ$)	6.11° (4.19 $^\circ$)
$d^* \leq 4m$	$d_{ err }$	0.40 m (0.20 m)	0.41 m (0.21 m)
	$d_{relative}$	13.3% (6.0%)	13.6% (6.3%)
	$\theta_{ err }$	5.12° (3.62 $^\circ$)	5.08° (3.72 $^\circ$)
$d^* > 4m$	$d_{ err }$	1.05 m (0.39 m)	1.00 m (0.38 m)
	$d_{relative}$	19.3% (6.1%)	18.4% (6.1%)
	$\theta_{ err }$	7.17° (4.50 $^\circ$)	7.23° (4.38 $^\circ$)

3.3.2.2 Remarks about Our UHF RFID Sensor Model

We would like to point out several distinctive characteristics about our implementation:

- We derived a RFID sensor model from the Friis radar equation for UHF RFID tag localization using a particle filter. The closed-form nature of the model obviated the need for extensive data captures employed in the prior art.
- Our power-scaling technique used *nine* reads from a single location to provide a posterior distribution that is more discriminative than that provided by a single tag detection alone. The power scaling served as an proxy to RSSI. This alluded to (forthcoming, at the time) RFID hardware that could provide direct RSSI measurements, which could potentially yield implementations with faster operation, faster convergence, and (perhaps) superior localization capabilities. New hardware did yield such improvements, as we will examine shortly.

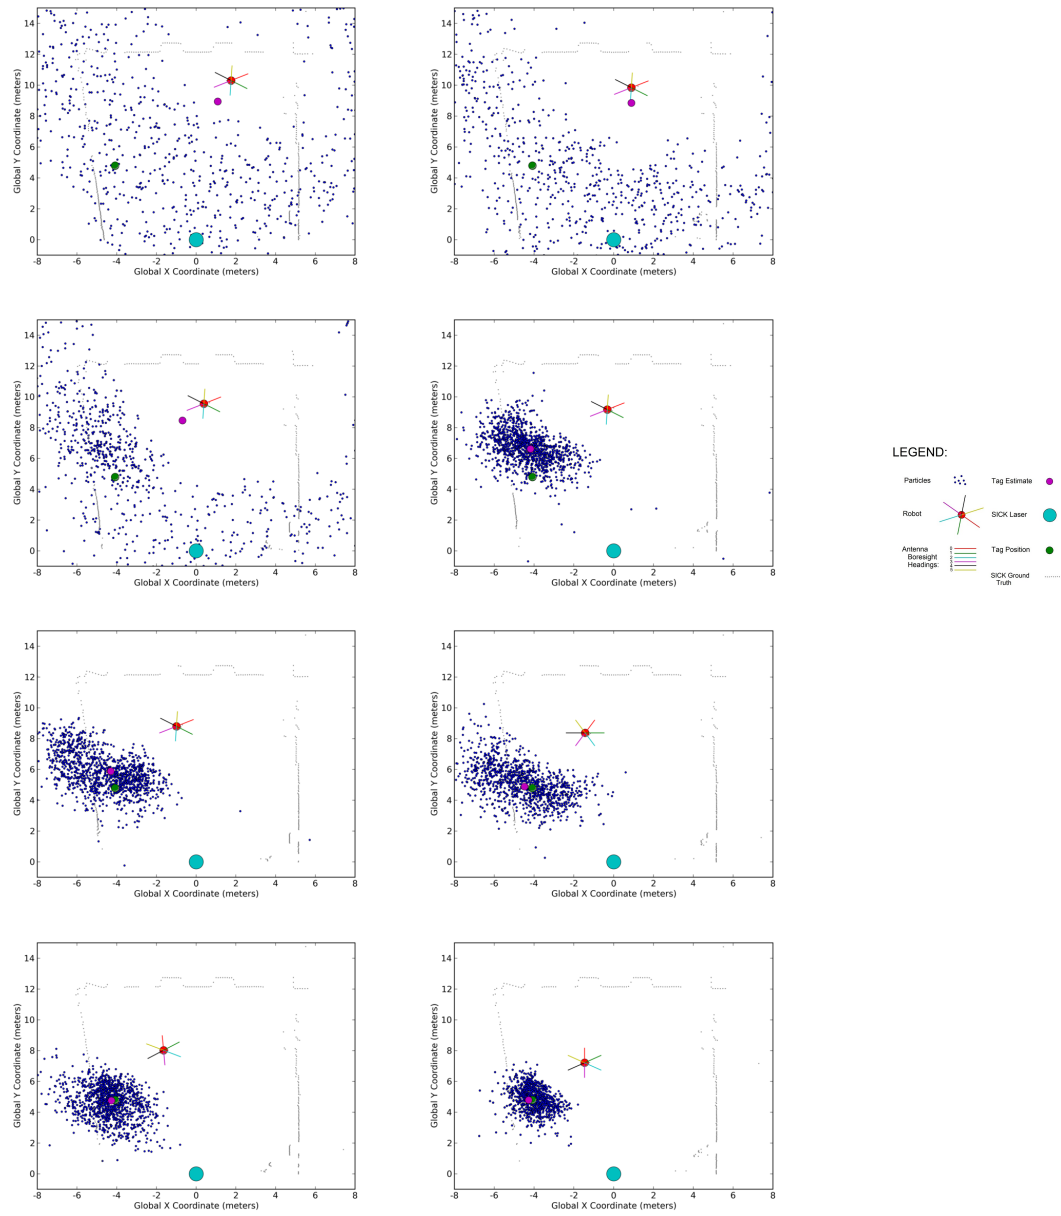


Figure 36: The particle filter converges to the true tag position as the RFID rig moves through the environment while capturing scaled-power RFID reads.

- Unfortunately, we performed limited evaluation of our system. The lack of (modern) Robot Operating System (ROS) navigation software, the tedium of hand-measured odometry, and the delayed on-robot integration conspired to prevent further measurements.
- More recently, researchers have developed techniques to bootstrap data-driven, planar UHF RFID sensor models using an initial “guessed” pattern [152]. Conceivably, the Friis model we have developed could serve as a useful starting point for such algorithms.

3.3.3 Data-Driven RSSI Sensor Models

After the widespread availability of commercial RFID readers supporting direct RSSI measurements, researchers developed revised data-driven RFID sensor models to harness the additional data [69]. The revised sensor yields observations (measurements) comprised two pieces of data, $z = (d_g, s_{rssi})$ where d_g is the binary random variable representing tag-detection, and s_{rssi} is a continuous random variable representing the RSSI measurement associated with a positive detection. The researchers define the combined sensor model,

$$p(z|x) = p(s_{rssi}|x) \cdot p(d_g|x). \quad (44)$$

Researchers discretize the state x in several grid cells. The tag detection probability distribution, $p(d_g|x)$, is computed using the ratio of positive tag detections ($n_{d,\theta}^+$) to total tag detections,

$$p(d_g|x) = n_{d,\theta}^+ / (n_{d,\theta}^+ + n_{d,\theta}^-). \quad (45)$$

This distribution is represented as a histogram. Researchers assume that the RSSI probability distribution, $p(s_{rssi}|x)$, is normally distributed so that

$$p(s_{rssi}|x) = \mathcal{N}(\mu(x), \sigma(x)), \quad (46)$$

which can be represented through two real-valued 2D arrays that store the mean and standard deviation (respectively) of the continuously-valued random variable for each grid cell (d, θ) . An example of this models is shown in Figure 37.

We would like to point out several distinctive features this seminal work by Joho *et al* that used more modern UHF RFID readers, along with a tag detection model *and* a RSSI model, to perform tag localization [69]:




Figure 37: Joho *et al* performed seminal work on RFID tag localization combining both tag detection *and* receive signal strength (RSSI) measurements; they used extensive ground-truth data captures to build histogram-based sensor models of tag detection likelihood (left) and RSSI likelihood (represented by a Gaussian with mean (middle) and standard deviation (right)) [69]. [Figures reproduced with authors’ permission.]

- This work was some of the first to employ a dual tag-detection plus RSSI model for UHF RFID tag localization using a particle filter. Reportedly, this combined model yielded tag localization results with (mean) error on the order of 30cm, which is within a factor of two of previous tag-detection only models.
- The researchers used the RSSI model to perform both tag localization relative to the robot and robot localization on a map of demarcated tags [69].
- Like the other RFID sensor models, the system was evaluated under relatively controlled conditions: consistent tag orientation on nearby walls / shelves on consistent nearby materials. Plus, the evaluation used the same tag locations to bootstrap the sensor model as to evaluate localization; thus, the sensor model is inherently specific to this particular distributions of tags.

3.3.4 Other UHF RFID Sensor Models

These three sensor models we explored are just a subset of the UHF RFID sensor models discussed or examined in the literature. Other sensor models, including UHF RFID tag detection fingerprints [150] and tag detection rates [151], have also been used to perform 2-DoF tag localization and report similar localization errors. The former, RFID fingerprinting methods have also been extended to incorporate RSSI measurements and report similar performance [153].

3.4 Discussion

All of the UHF RFID sensor models used for Bayesian inference in Section 3.3 use the 2-DoF relative position (ie d and θ) between the reader antenna and tag antenna. We know from the Friis

radar equation formulation in Chapter 2 that the physics of RF propagation are inherently 6-DoF, and that they are highly subject to environmental factors such as object material properties, tag orientation, and environmental properties (eg. multipath). This invites the question: What happens to these 2-DoF sensor models when we tag objects and cannot exercise control (ie. under non-ideal conditions)?

To address this question, we use extensive captured data (from Chapter 2) to compare the resultant sensor models during relatively ideal situations (as with all the aforementioned probabilistic models) to the models obtained when tagging objects. We find that the models deviate substantially.

3.4.1 Probabilistic Models From Data Captures Under Relatively Ideal Settings

To understand the impact of tagging *objects*, we first investigated tag detection and RSSI models using our RFID setup under relatively ideal settings. Using the ground truth data captures from Chapter 2 (from both the head antenna and two shoulder antennas) we built a data-driven probabilistic model of tag detection likelihood and RSSI likelihood for each cell in a discretized 2-DoF RFID sensor model (ie. a histogram, as in Section 3.3.3). Our data included measurements of the tag at multiple 6-DoF orientations (ie. both vertical and horizontal orientations, with the tag mounted on a large cardboard box) in a large, empty room. We created the data-driven RFID model using the same method as Section 3.3.3. We computed the tag detection probability for each cell according to Equation 45; we represented the RSSI model by a Gaussian distribution for each cell, where the mean and standard deviation was computed using all the (positive) RFID readings obtained for tag poses within the corresponding cell (ie. the sample mean and sample standard deviation). Figure 38 shows the resulting model.

3.4.2 Antenna-By-Antenna Variation

All of the systems described in Section 3.3 employ two or more antennas, and assume that all antennas share a common sensor model (the same distributions). We can test this assumption by computing the individual antenna distributions using the free-space, cumulative data. For brevity, we will only look at the tag detection distribution and the mean RSSI. The distributions are shown in Figure 39. As we can see, the tag detection and RSSI distributions can vary drastically between different antennas. The head-mounted antenna's tag detection region is noticeably smaller compared




Figure 38: Tag detection probability $p(d_g|x)$ (top), mean RSSI $\mu(x)$ in $p(s_{rssi}|x)$ (middle), and RSSI standard deviation $\sigma(x)$ in $p(s_{rssi}|x)$ (bottom) created from the relatively ideal (free space) data captures using the three PR2-mounted antennas (head, left shoulder, and right shoulder) and both tag orientations (horizontal and vertical) obtained in Figures 24 and 25.

to the two shoulder-mounted antennas. Further, despite almost identical physical arrangements on the robot, there are major differences between left and right shoulder-mounted antennas. The right shoulder-mounted antenna has a much greater tag detection probability compared to the left, but the left shoulder-mounted antenna produces stronger RSSI readings when the tag is detected.

The effect of the antenna-to-antenna model differences on tag localization fidelity remains an open question. However, these results seem to suggest that a common sensor model for all the antennas doesn't necessarily capture all of the system dynamics. In fact, early in evolution of the PR2 UHF RFID system, the extreme difference between left and right shoulder-mounted antenna distributions gave us strong suspicions that something was amiss. We spent considerable time isolating various design variables (antennas, cabling, reader ports, actuators, etc) to locate the discrepancy. Ultimately, performing a *complete* swap of left-right antenna systems yielded the same results. This suggests that there are systemic differences between the left and right mountings, which we attributed to metallic structures *internal* to the robot. This shows that (1) care must be taken to account for antenna placement in relation to other robot components, and (2) we cannot necessarily assume that all antennas will yield the same sensor model distributions, even when using the same make and model of antenna.

3.4.3 Variation Due to Tag Orientation

We can also evaluate the effect of 6-DoF tag orientation on the 2-DoF, data-driven tag detection and RSSI models. For example, during the cumulative data captures, we acquired data with the tag in both a vertical and horizontal orientation (in the world frame). In Figure 40, we show the difference in the resulting distributions for: just the vertical tag orientation, just the horizontal orientation, and the cumulative data (from both orientations).

We observe that the vertical tag orientation has a much greater tag detection likelihood and stronger mean RSSI compared to the horizontal tag orientation. In fact, we could predict this result from the Friis model. Namely, a dipole tag oriented vertically (world frame) will have a isotropic in-plane gain, but a dipole tag oriented horizontally (world frame) will have in-plane gain that varies from a maximum (of 1.5) to a minimum (theoretically, 0.0) depending on the robot's location. When considering RFID measurements over all tag-robot relative positions (*marginalizing* over tag-robot

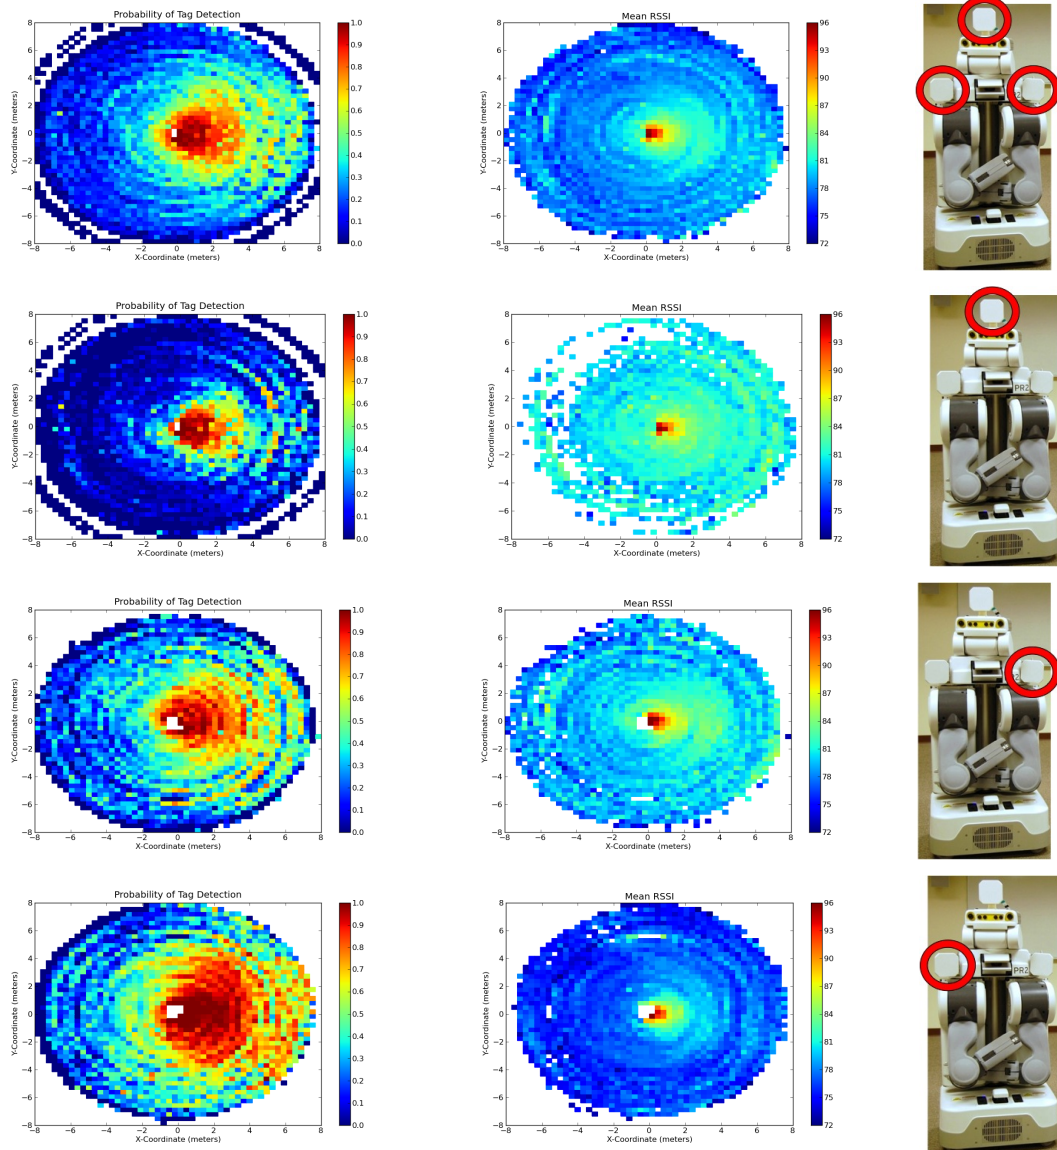


Figure 39: The 2-DoF model statistics vary antenna-by-antenna. Tag detection probability (left column) and mean RSSI (right column) are shown for all the antennas combined (1st row), the head-mounted antenna (2nd row), the left shoulder-mounted antenna (3rd row), and the right shoulder-mounted antenna (4th row).

positions), this has the cumulative effect of reducing the detection region (and mean RSSI) for the horizontally mounted tag. Similarly, the cumulative data (both orientations) shown in Figure 40 represents the “mixed” distribution obtained by combining data over two tag orientations. Since the underlying relationship between the sensor model is 6-DoF (as we examined in Chapter 2), we would need to acquire extensive data from many different tag-robot positions *and* orientations to build a model that accurately reflects the 6-DoF relationship in this 2-DoF model. Statistically, we would need extensive data captures from many different 6-DoF orientations to create a 2-DoF *marginalized* representation that accounts for the prior likelihoods of various orientations – and we would need to know the ground-truth tag pose for each and every data capture.

3.4.4 Variation Due to Tagged Object Material Properties

So far, all of the UHF RFID models assume that the tag is in relatively ideal conditions (free-space), with controlled nearby material properties. This may be a reasonable assumption for infrastructure tags, which are affixed to walls. However, a mobile manipulator will likely work with tagged *objects* of various composition, which may (or may not) require tag deformation, and will likely be located in a diverse set of 6-DoF orientations relative to the robot. We compute the tag detection and RSSI histograms using data from Section 2.5.5 on an object-by-object basis for each of six objects: an orange medication bottle in isolation, a orange medication bottle in clutter, a red mug in clutter, keys in clutter, vitamin bottle in clutter, and a TV remote (with on-metal tag) in clutter.

The resulting sensor model distributions for each of these objects is shown in Figures 41 and 42, and includes the cumulative (free-space) model from Figure 38 for comparison. The deviation from the free-space sensor model is *drastic*.

The UHF RFID tag localization literature is notably devoid of tagged *object* localization results. We believe that the free-space model distributions would be unable to localize a tagged object without severe degradation; similarly, using the tagged object distributions to localize a free-space (or infrastructure tag) would likely result in severe degradation as well. It is certainly possible to construct distributions that marginalize over various free-space distributions, tagged-object distributions, and 6-DoF pose. However, we believe that such a model would (1) require *significant* effort to account for all eventualities / configurations, (2) be highly dependent on the specific scenarios

examined, and (3) would be unlikely to localize tags and tagged objects with the same fidelity as reported in the literature.

Of course, the RFID sensor provides unique identity; thus, we could imagine storing sensor models on a tag-by-tag (object-by-object) basis. However, the practical challenges of ground truth data captures for real-world deployments makes this an unattractive option. In the remainder of this thesis, we examine alternative methods that use robot *behaviors* to either estimate certain parameters of interest (eg. bearing to the tag) or to just approach the tag (or tagged object) without maintaining an explicit estimate of the tag's pose / position.

In the next chapter, we develop a series of optimization-based approaches to discover, locate, and approach tagged *objects*. We show that these new behaviors do not require training data, and when compared to state-of-the-art probabilistic techniques for object localization, they achieve comparable performance in positioning the robot near the tagged object and superior performance in orienting towards the tagged object.

3.5 Conclusions

In this chapter, we presented a probabilistic Bayesian filter technique that incorporates RFID sensor measurements over time as the robot moves about the environment to localize (track) UHF RFID tags relative to a robot. Among the contributions explored in this chapter:

- A literature review summarizing the various methods of RFID tag localization. All of the methods discussed use some form of Bayesian filter: Kalman Filter, Histogram Filter, or Particle Filter.
- We examined three specific UHF RFID sensor models based on 2-DoF relative pose between the RFID reader antenna and tag. The first model used data captures to define tag detection regions. The second model, of our own devise, used the Friis transmission equation to define tag detection regions that could be programmatically adjusted by altering the reader transmit power. Scaling the power yielded tag detection models analogous to RSSI readings. The third model, leveraging advances in RFID hardware, used data-driven tag detection *and* RSSI distributions to quickly localize tags. All three of these models were tested under relatively ideal (free-space) conditions and yielded similar tag localization fidelity.

- We showed that the 2-DoF probabilistic sensor models were subject to extreme deviations on an antenna-by-antenna basis, depending on tag orientations, and depending on the object to which the tag is affixed. We believe the deviations from the free-space models would likely result in tag localization degradation.

In the next chapter, we develop a series of optimization-based approaches to discover, locate, and approach tagged *objects*. We show that these new behaviors do not require training data, and when compared to state-of-the-art probabilistic techniques for object localization, they achieve comparable performance in positioning the robot near the tagged object and superior performance in orienting towards the tagged object.

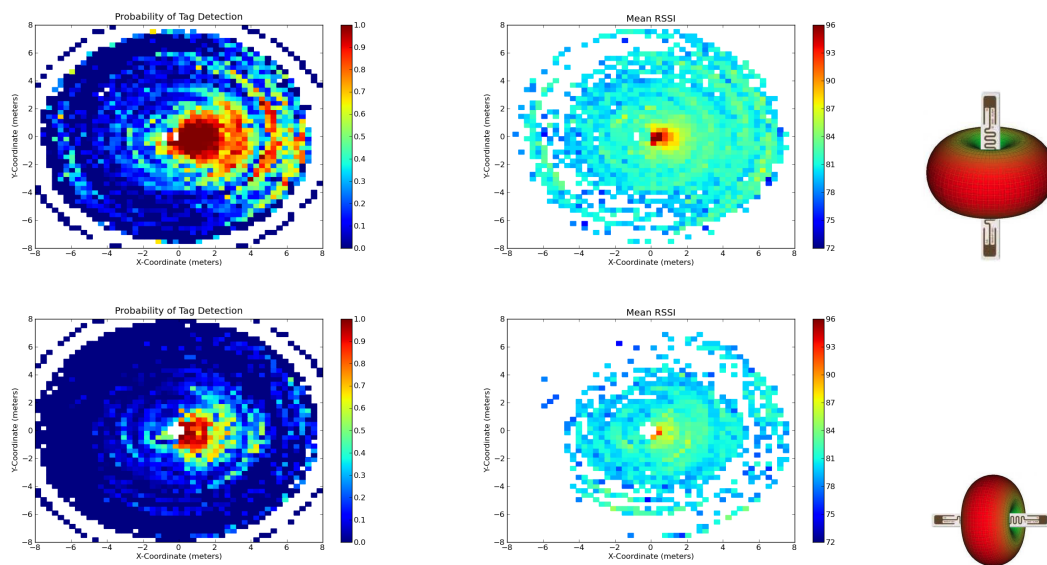


Figure 40: The 2-DoF model statistics vary tag-by-tag. Captured using just the head-mounted antenna, tag detection probability (left column) and mean RSSI (right column) vary between the vertical tag orientation (1st row), the horizontal tag orientation (2nd row). As you can see, the tag’s radiation pattern (ie. the full 6-DoF pose between reader antenna and tag) plays a significant role when the tag is end-on.

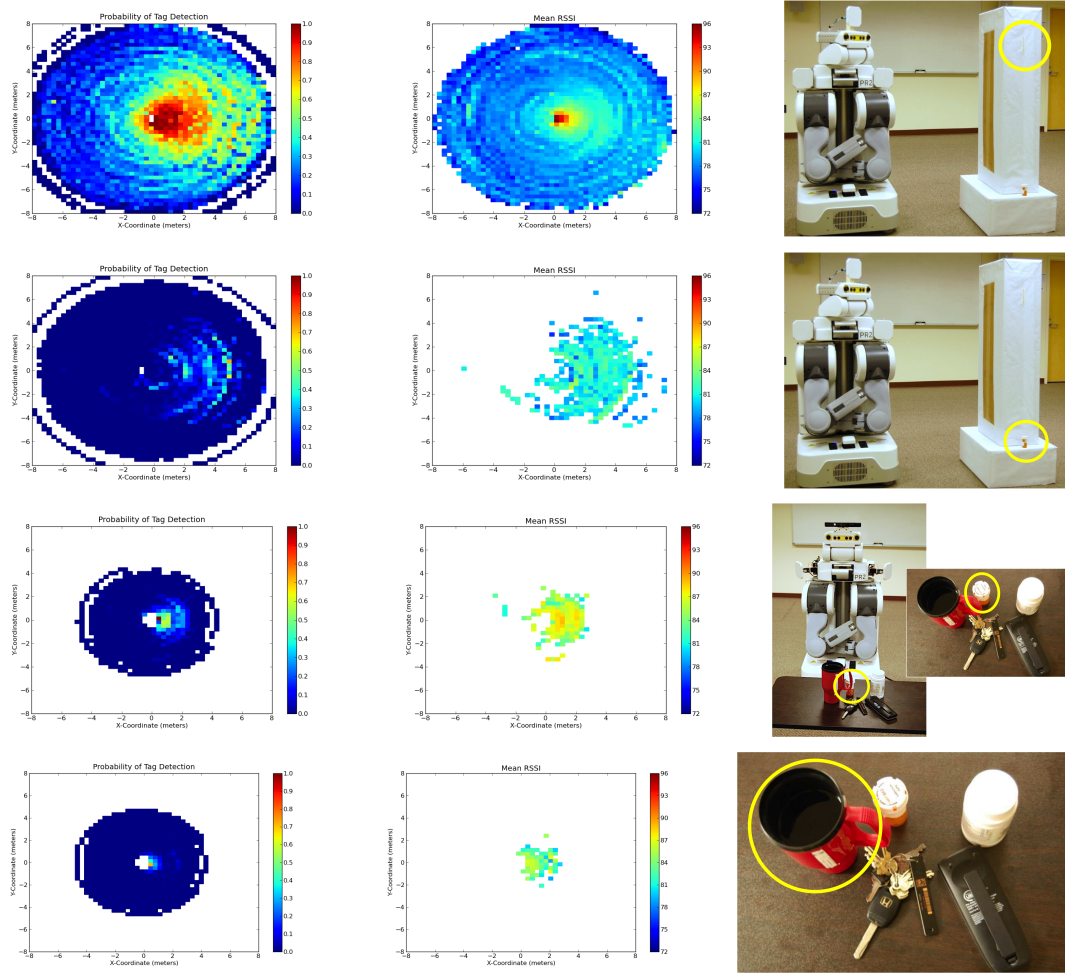


Figure 41: The 2-DoF model statistics vary tag-by-tag. These plots show the tag detection probability (left) and mean RSSI (right) distributions for the tag in free space (1st row), the orange medication bottle in isolation (2nd row), the orange medication bottle in clutter (3rd row), and the 75-percent full red mug (4th row).

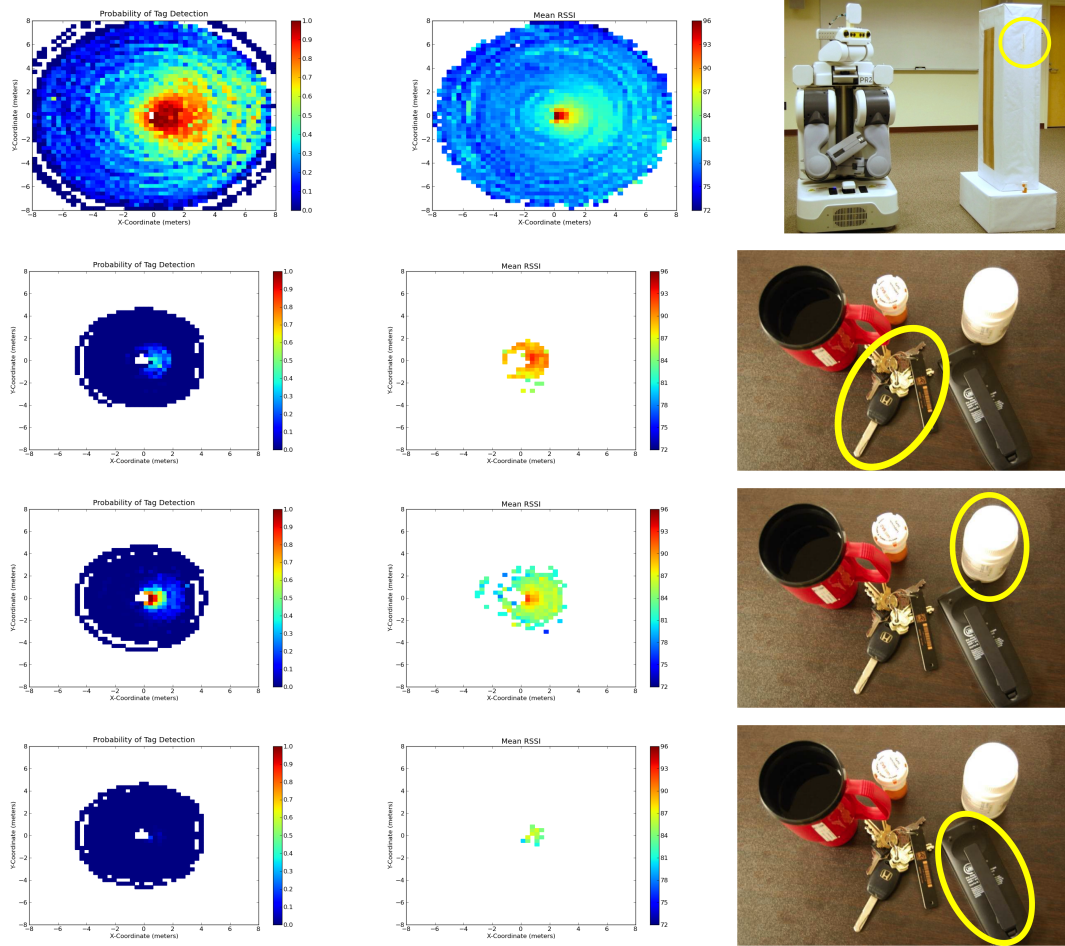


Figure 42: The 2-DoF model statistics vary tag-by-tag. These plots show the tag detection probability (left) and mean RSSI (right) distributions for the tag in free-space (1st row), the keys with “on-metal” tag (2nd row), the white medication bottle (3rd row), and the TV remote with “on-metal” tag (4th row).

CHAPTER IV

OPTIMIZATION-BASED APPROACHES TO LOCATE AND APPROACH UHF RFID TAGS

In this chapter, we take a markedly different approach to UHF RFID perception. Rather than trying to estimate the pose of a tag relative to the robot using Bayesian inference (as in Chapter 3), we present a new, general approach for perceiving UHF RFID tags affixed to locations, people, and objects of interest that is easy to implement, easy to generalize, and does not require training data or sensor models. We formulate our approach as an optimization problem, where the robot actuates its directional (patch) reader antennas and/or uses mobility to provide opportunistic views of the RF landscape and estimate a desired quantity of interest (eg. bearing toward the tag) by maximizing the expected RSSI measurements associated with a desired tag ID, over all possible robot states (x). The following equation represents the general approach:

$$\hat{x} = \operatorname{argmax}_x E (\mathbf{RSSI} \mid \mathbf{X} = x).$$

We develop a series of *local* optimization-based behaviors that allow the robot to estimate the direction toward a tag (bearing estimation) or to approach a tag (RFID servoing) given some initial position in the (local) neighborhood of the tag. These methods are expressed as follows:

- **Bearing Estimation – Azimuth (Θ) Only:**

$$\hat{\Theta} = \operatorname{argmax}_{\theta} E (\mathbf{RSSI} \mid \boldsymbol{\theta} = \theta)$$

Panning the RFID reader's directional (patch) antenna about azimuth angles only (θ), we can estimate the direction (bearing) toward a tag ($\hat{\Theta}$). Since the patch antenna has a unimodal radiation pattern with unique maximum, the pan angle pointing toward the tag will tend to have the largest expected RSSI readings.

- **Bearing Estimation – Azimuth (Θ) and Elevation (Φ):**

$$\hat{\Theta}, \hat{\Phi} = \operatorname{argmax}_{\theta, \phi} E (\mathbf{RSSI} \mid \boldsymbol{\theta} = \theta, \boldsymbol{\phi} = \phi)$$

We can extend bearing estimation to two dimensions by both panning (in azimuth, θ) and tilting (in elevation, ϕ) the RFID reader antenna to estimate the bearing toward the tag ($\hat{\Theta}$ and $\hat{\Phi}$). We develop a method of displaying the RSSI values as images, where the value at each pixel corresponds to the expected RSSI measurement at the associated pan and tilt angles. We dub these visualizations “RSSI images.” They provide a natural and intuitive way to visualize the otherwise invisible RFID sensor values.

- **RFID Servoing – Angular Velocity ($\dot{\Theta}$):**

$$\dot{\Theta} = \kappa \cdot E (\Delta RSSI \mid \theta = \theta)$$

Taking successive RFID readings from two patch antennas, slightly offset in angle from one another, we can determine how the robot should adjust its heading (angular velocity, $\dot{\Theta}$) so that it can face a RFID tag. Intuitively, a stronger reading from the *left* antenna dictates that we turn left; a stronger reading from the *right* antenna dictates that we turn right. We show that this amounts to an online version of bearing estimation, where the difference between left and right RSSI measurements approximates the gradient of a function of RSSI vs azimuth angle. Using gradient ascent, we can rotate the robot to maximize the expected RSSI, thus orienting the robot toward the desired tag. Such processes, which use feedback to reduce system error, are called “servoing.” Adding a constant forward motion while servoing the robot’s heading, we can approach a tagged object; we use this process to approach tagged objects in real time and generally refer to this process as “RFID servoing.”

From an optimization perspective, these techniques perform local search (eg. gradient ascent [19]) to refine the robot’s state (position and/or orientation) within a local basin of attraction, thereby yielding (progressively) greater expected RSSI measurements.

We also develop a global RFID search behavior, where the robot determines the best position and orientation (X, Y, Θ) to observe the desired tag by moving throughout the entire environment and selecting the pose that yields the maximum expected RSSI. This technique is akin to sampling-based global search techniques in the optimization and planning literature [91, 149]; in the preceding framework, global RFID search is expressed as:

- **Global RFID Search:**

$$\hat{X}, \hat{Y}, \hat{\Theta} = \underset{x, y, \theta}{\operatorname{argmax}} E (\mathbf{RSSI} \mid \mathbf{X} = x, \mathbf{Y} = y, \boldsymbol{\theta} = \theta)$$

Mobile robots can relocate themselves completely new locations, $\langle x, y, z \rangle$, to get advantageous views of the RF landscape. By brute-force sampling over all robot (or reader antenna) poses, the robot can return to the configuration $\langle \hat{X}, \hat{Y}, \hat{Z} \rangle$, that yielded the maximum expected RSSI; this configuration is likely to be (1) near the tag, and (2) oriented toward the tag.

It is well-known in the optimization literature that hybrid global-local search can increase the effectiveness and/or efficiency of a search algorithm [51, 91, 15]. Therefore, we develop a hybrid global-local RFID search algorithm that first performs a sparse global RFID search to determine an acceptable initial robot state (position and orientation), followed by a series of local RFID search behaviors (bearing estimation and RFID servoing) to refine the robot’s state within the local basin of attraction and yield a “good” final robot state that is both (1) near the desired tagged object, and (2) oriented towards it. We evaluate hybrid RFID search using nine tagged objects in nine locations in a real home environment for a total of 81 trials.

We compare our results to those obtained when using state-of-the-art Bayesian tag localization techniques (Chapter 3). We show that our hybrid global-local RFID search technique achieves superior performance for several reasonable sensor models. Since our optimization-based approach is easy to implement, easy to generalize, does not require training data, and yields superior results, we believe that it is the superior approach to discover, locate, and approach tagged *objects* in human environments.

The final robot poses obtained through hybrid RFID search are well-matched to multi-sensor fusion and object search: the tagged object is likely nearby and within the field of view of other sensors (eg. cameras and depth sensors), and the tag’s unique identifier can be used to obtain information about the object’s appearance in other sensor modalities or to seed object recognition algorithms. We believe that UHF RFID could be a valuable, complementary sensing modality when used in existing object search algorithms that rely on false-positive-prone object recognition or line-of-sight sensor information [58, 89, 121, 97]. We re-visit these topics in Chapter 6.

4.1 Bearing Estimation (Azimuth Only)

Our bearing estimation technique estimates the direction toward a stationary tag; in the context of local optimization techniques, the direction towards the tag is the gradient (direction) the robot should move to decrease the distance from the tag (and thus increase the expected RSSI). Bearing estimation draws inspiration from early azimuth-only acquisition radars that were used to locate and track aircraft [136]. To explain the intuition behind bearing estimation, we start by considering just the planar case shown in Figure 43, where the desired tag is located in the reader antenna's xy -plane¹. The goal of bearing estimation is to find the bearing toward the tag, also called azimuth angle ($\hat{\Theta}$), relative to the antenna's initial frame of reference (F_{init}).

We assume that the reader antenna is a patch antenna with a unimodal radiation pattern and the tag remains stationary. We pan the reader antenna about its z -axis; for every angle, θ , we perform an RFID query specific to the desired tag ID. When the tag is positively detected, we record the RSSI, yielding a one-dimensional function of RSSI versus rotation angle, $RSSI = g(\theta)$, from which we estimate the bearing toward the tag through the following optimization:

$$\hat{\Theta} = \underset{\theta}{\operatorname{argmax}} g(\theta). \quad (47)$$

However, real-world measurements of RSSI are noisy, as shown in Figure 44. If we assume that RSSI measurements are comprised of a noise-free component, $g(\theta)$, and a zero-mean noise component represented by a Gaussian, $\mathcal{N}(0, \sigma)$, then the distribution of the continuous random variable **RSSI** given rotation θ can be expressed as:

$$P(\mathbf{RSSI} \mid \theta = \theta) = g(\theta) + \mathcal{N}(0, \sigma), \quad (48)$$

so that our optimization now seeks the angle, θ , that maximizes the *expected* RSSI,

$$\hat{\Theta} = \underset{\theta}{\operatorname{argmax}} E(\mathbf{RSSI} \mid \theta = \theta). \quad (49)$$

4.1.1 Non-Planar Bearing Estimation

In general, the tag will not be located in the reader's xy -plane, as shown in Figure 14. However, panning the RFID reader antenna about its z -axis only alters a single 6-DoF pose parameter: the

¹These are sketches to illustrate the overall concept, not real data.

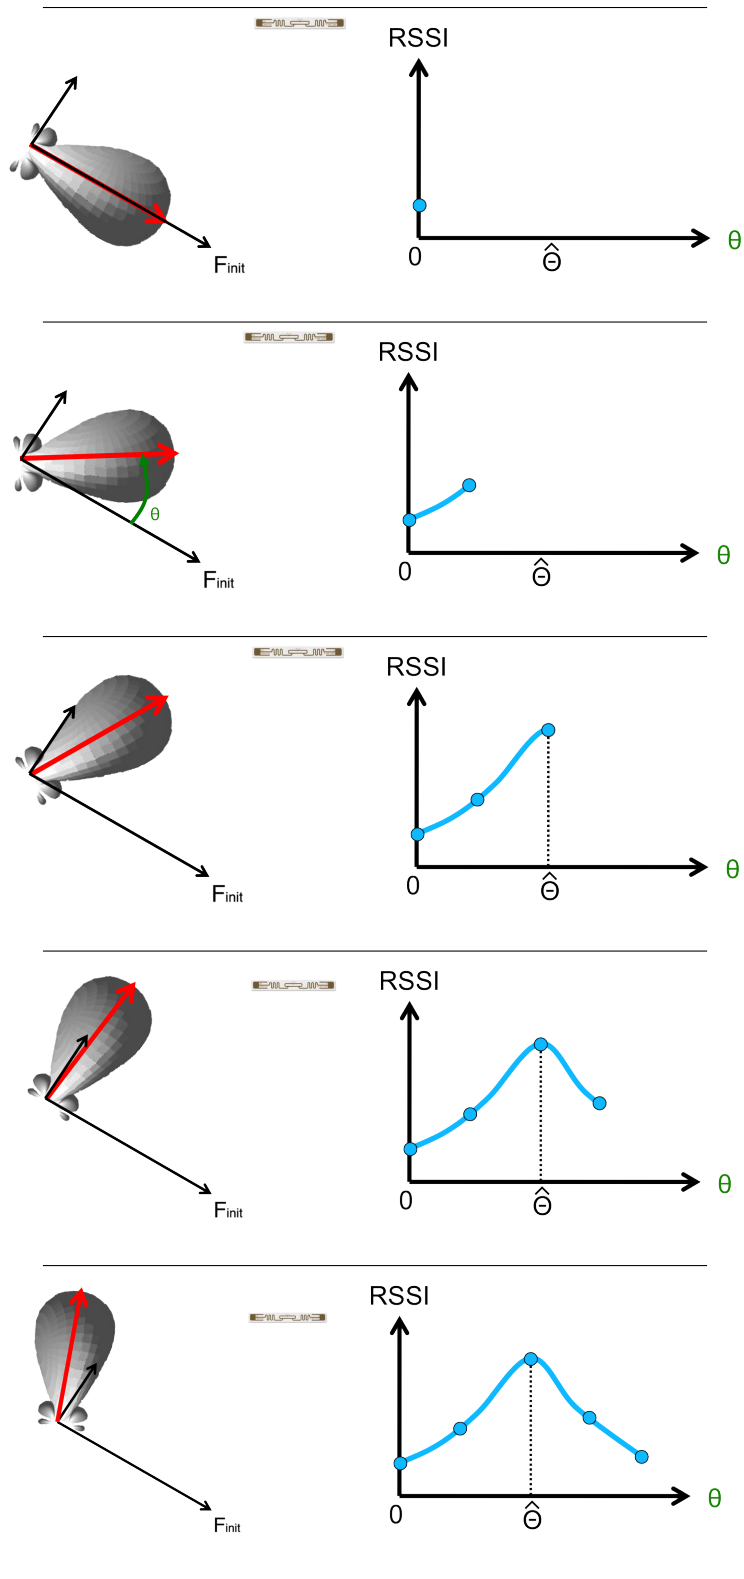


Figure 43: We pan a directional patch antenna by about its z -axis to estimate the bearing, $\hat{\theta}$, toward a stationary tag. Since the patch antenna has a unimodal gain pattern, we expect maximum RSSI measurements when the antenna is pointed directly at the tag ($\theta = \hat{\theta}$)




Figure 44: Actual RSSI measurements as a function of θ are noisy, so we treat $RSSI$ as a random variable and instead optimize for the maximum *expected* RSSI. For this particular bearing estimation attempt, $\hat{\Theta} = 10^\circ$ and the true $\Theta^* = 6^\circ$.

azimuth angle (θ_{rdr}). Thus, under ideal circumstances, the forward-backward Friis radar equation predicts that RSSI ($g(\theta)$) will be proportional to a cross section of the reader antenna’s radiation pattern (in logarithmic units),

$$g(\theta) \propto P_{rdr}^{inc}(dB) \propto G_{rdr}(\theta_{rdr} = \hat{\Theta} - \theta)(dB). \quad (50)$$

Recall from Section 2.1.3.2 that ideal patch antennas, for any fixed value of ϕ_{rdr} , the cross section of the antenna’s gain will have a global maximum at $\theta_{rdr} = 0^\circ$. Thus, bearing estimation yields the azimuth angle pointing directly toward the tag.

4.1.2 Evaluation

To evaluate bearing estimation, we employ a Laird Technologies S9025P patch antenna and Thing-Magic Mercury 5e UHF RFID rigidly affixed to the pan-tilt head of a PR2 robot from Willow

Garage, as shown in Figure 45. We placed three tags in a large (8×12 meter) clutter-free environment: a vertically-oriented Alien ALN-9640 Squiggle tag in the plane of the antenna affixed to a large (white) cardboard box, a horizontally-oriented Squiggle tag in the plane of the antenna affixed to a large (white) cardboard box, and a Squiggle tag wrapped about a medication bottle located 60cm below the xy -plane of the antenna. We recorded the ground-truth pose of the tags on a map and used FastSLAM to localize the robot on the same map (within 10cm and 5° , respectively) [99], yielding ground-truth estimates of the 6-DoF pose between tag and reader antenna (including distance and bearing estimates).

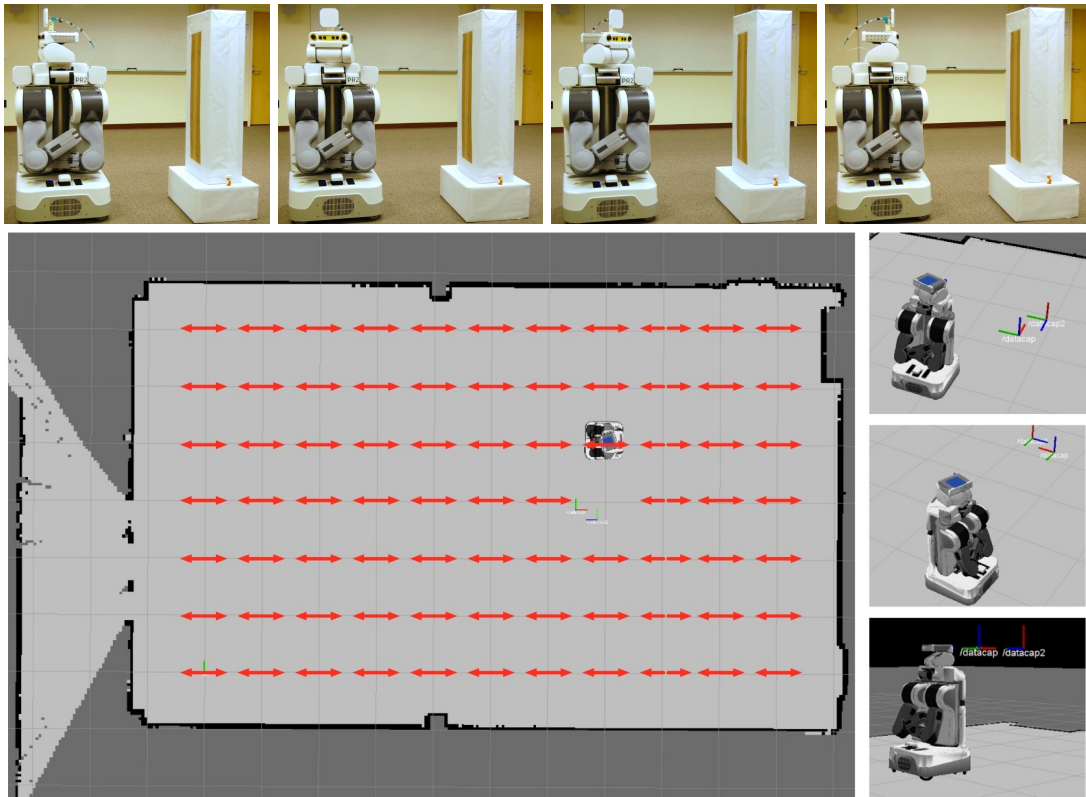


Figure 45: *Top Row:* The robot pans a fixed, head-mounted antenna through 300° rotations to perform bearing estimation. *Bottom Row:* Bearing estimation attempts are made for each indicated positions / orientation in a room (red arrows). The ground-truth pose of the tags are also recorded / known.

We autonomously relocated the PR2 to numerous initial positions and orientations, as shown in Figure 45. For each location, the robot performed bearing estimation for each of the three tags. The robot slowly panned its head-mounted antenna through $\pm 160^\circ$ rotations, made RSSI measurements of the desired tag, and computed the expected RSSI for each rotation angle by discretizing

θ into 10° -wide bins and computing the sample mean (expected RSSI) of all RSSI measurements contained within that bin. According to the optimization for bearing estimation (Equation 49), the robot estimated the bearing toward the tag as the bin-center with maximum expected RSSI. Figure 44 shows a single bearing estimation attempt; additional bearing estimation attempts with good and poor estimates are shown in Figures 46 and Figure 47, respectively.

We define the *angular error magnitude*, $|err|$ between the ground truth bearing (Θ_{gt}) and our estimate ($\hat{\Theta}$) as

$$|err| = | \Theta_{gt} - \hat{\Theta} |, \quad (51)$$

so that we can compile statistics about the fidelity of the robot’s bearing estimation. Table 2 contains the mean (μ) and standard deviation (σ) of this error value on a tag-by-tag basis, as well as for the combination of all three tags. Due to physical limitations, the PR2’s head-mounted RFID reader antenna can only scan through $\pm 160^\circ$, so in some of the data captures (10 of 152) the tag was outside the scan range; these scans were omitted for the statistics calculations.

Table 2: Bearing Estimation Statistics For the Head-Mounted Antenna

Tag	Argmax Estimate	Static Estimate
Vertically Oriented Tag	$\mu_{ err } = 36.5^\circ$ $\sigma_{ err } = 39.4^\circ$	$\mu_{ err } = 82.2^\circ$ $\sigma_{ err } = 56.9^\circ$
Horizontally Oriented Tag	$\mu_{ err } = 57.2^\circ$ $\sigma_{ err } = 52.3^\circ$	$\mu_{ err } = 81.5^\circ$ $\sigma_{ err } = 57.2^\circ$
Tagged Medication Bottle	$\mu_{ err } = 57.7^\circ$ $\sigma_{ err } = 54.8^\circ$	$\mu_{ err } = 82.0^\circ$ $\sigma_{ err } = 57.4^\circ$
All Three Tags Combined	$\mu_{ err } = 50.4^\circ$ $\sigma_{ err } = 50.3^\circ$	$\mu_{ err } = 81.9^\circ$ $\sigma_{ err } = 57.1^\circ$

As a baseline for comparison, we use a “Static Estimator” that always estimates the same constant angle (Θ_{static}) that minimizes the expected angular error magnitude on average, across the data set. We determine this value for each data set through brute-force search. Table 2 shows that optimization-based bearing estimation outperforms the naive static estimate, but that the magnitude of the angular error is still quite large (about half of the antenna’s 3-dB beamwidth of 100°). However, Figure 48 shows that the magnitude of the angular error tends to decrease as the maximum

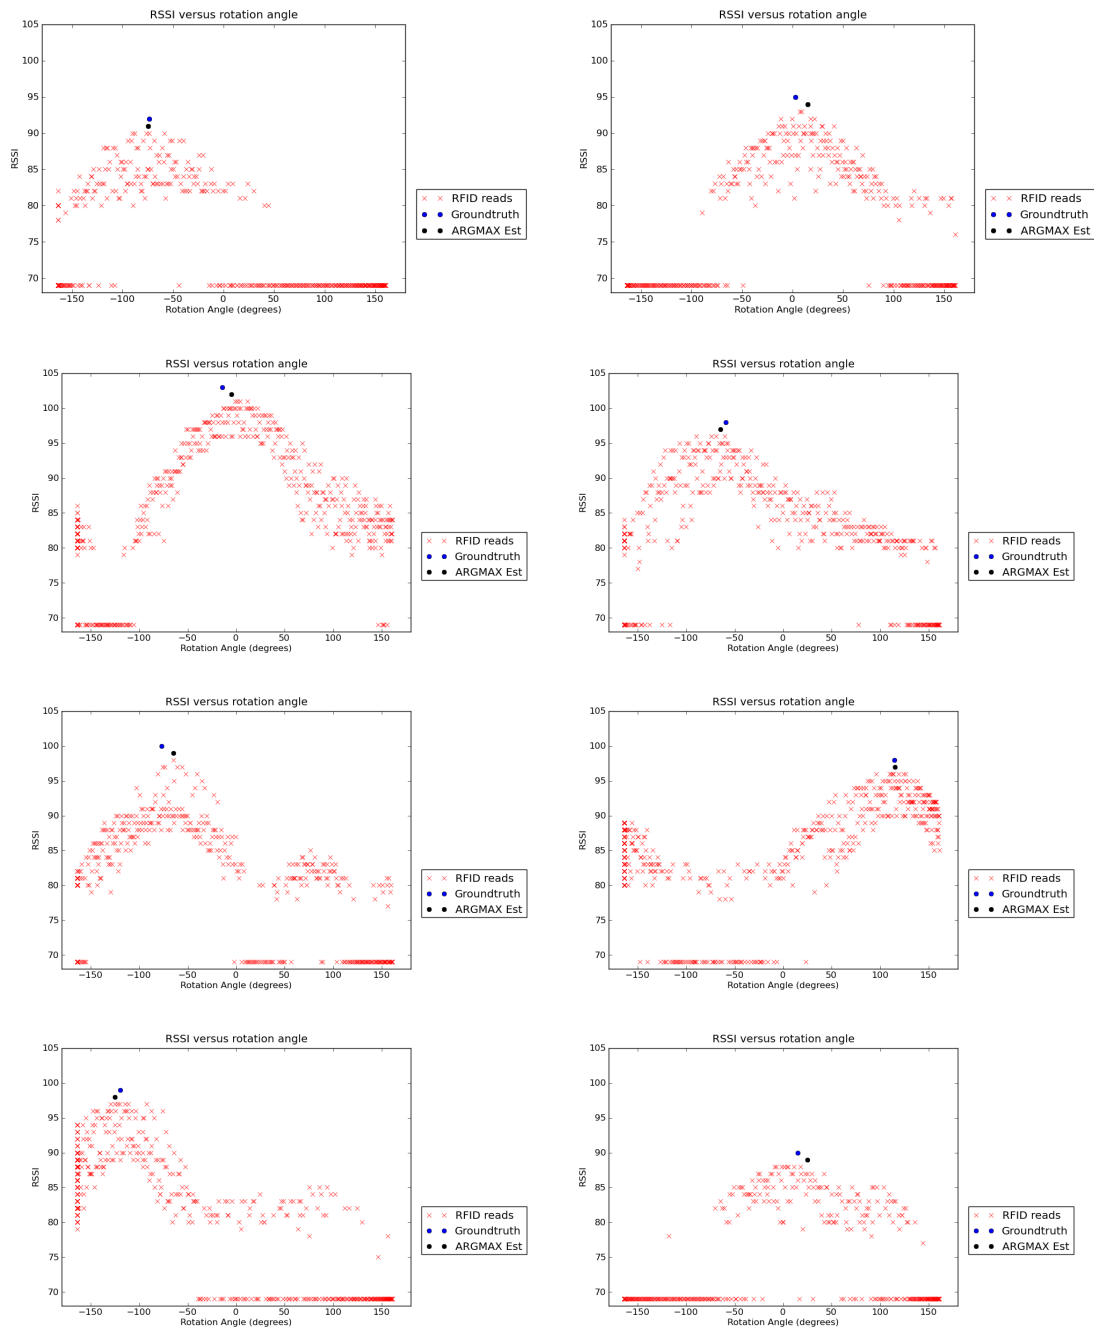


Figure 46: Bearing estimation captures show RFID sensor readings (RSSI) versus rotation angle. RSSI readings of 69 correspond to failed tag detections. The ground truth bearing to the tag, as well as the value estimated through the argmax optimization framework are indicated.

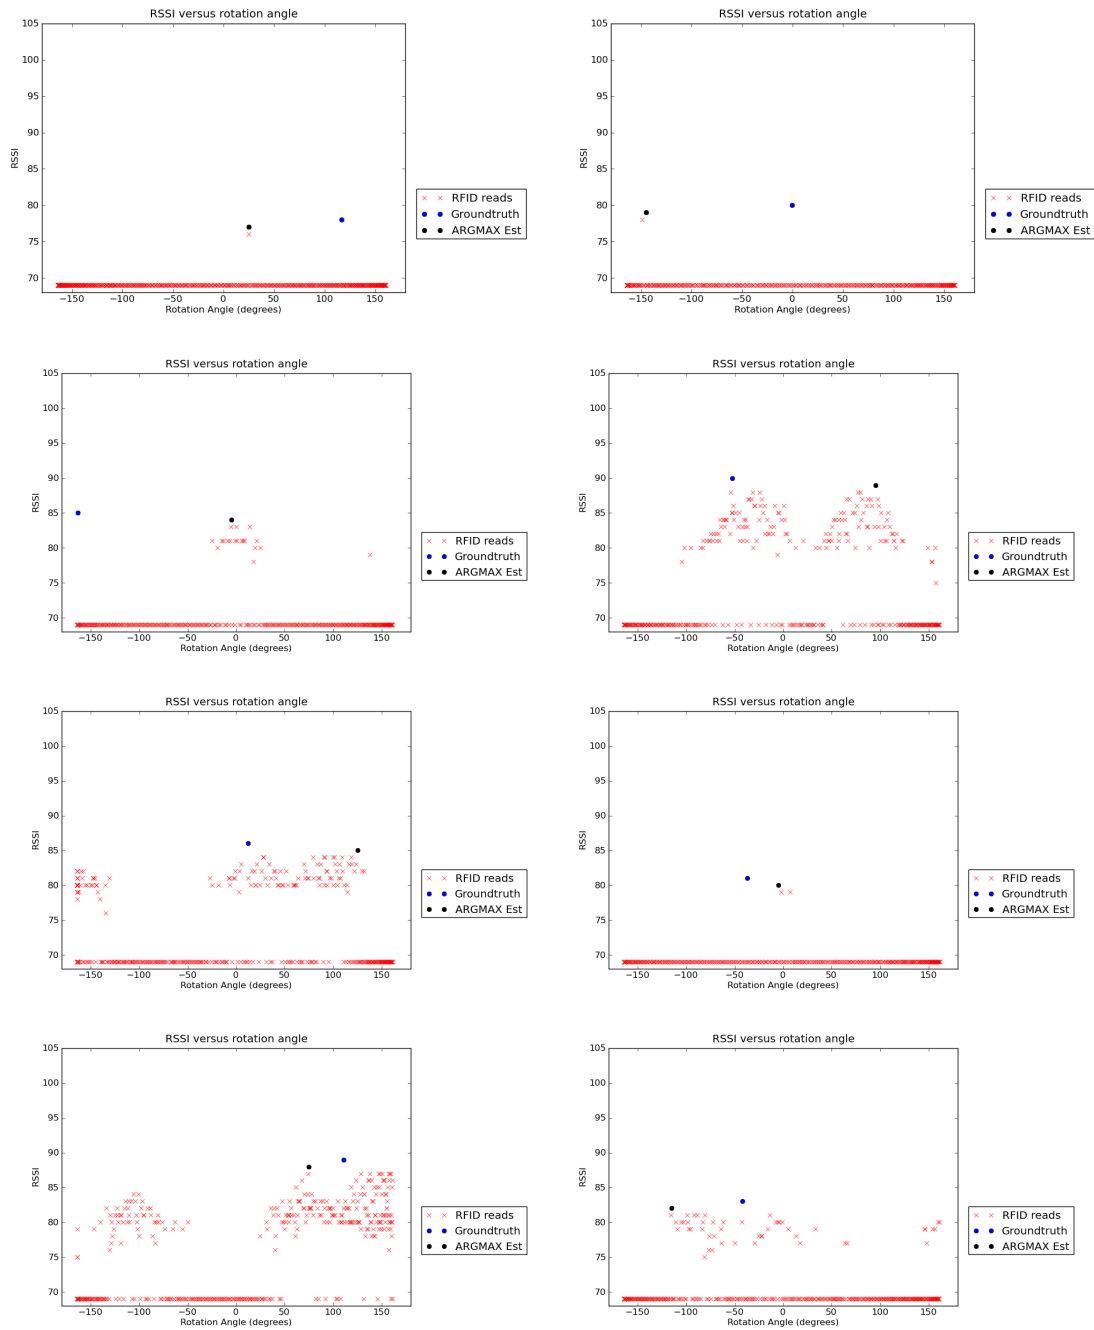


Figure 47: Multipath, spurious or sparse readings, ill-defined peaks, and relatively flat RSSI responses can hamper the bearing estimation algorithm.

expected RSSI increases (to less than 20° when $RSSI > 90$), and that the magnitude of angular error decreases with decreasing distance. This means that bearing estimation (local gradient estimation) performs better gets closer and closer to the tag.




Figure 48: Plotting the bearing estimation error versus the maximum expected RSSI in the scan (left) and the distance from the tag (right), we notice two trends: (1) the error decreases with increasing maximum expected RSSI and (2) the error decreases as the distance to the tag decreases. Thus, bearing estimation improves as the robot gets closer to the tag and obtains larger RSSI measurements.

4.1.3 Extensions to Bearing Estimation

Through the course of this research, we performed several basic extensions to bearing estimation. We explore them briefly here. Furthermore, we discuss other (advanced) techniques from the radar

literature (eg. direction of arrival estimation using phased arrays) that may become relevant with future reader hardware incarnations.

4.1.4 Bearing Estimation for Sub-360° Scans

In Chapter 2 (Figure 25), we examined an alternative data set that utilized the PR2’s shoulder-mounted antennas (Figure 25). However, due to physical constraints, these antennas could only be panned through $\pm 70^\circ$ (a 140° range). This poses an interesting design question: Given our shoulder-mounted scan data and assuming the tag is within the new $\pm 70^\circ$ scan range, how well do we expect this reduced-scan-range bearing estimation to perform compared to the static estimator?

We expect the static estimator to have (roughly) uniform estimation error over $[-70^\circ, +70^\circ]$, so the magnitude of the error will be uniformly distributed over $[0^\circ, +70^\circ]$ with $\mu_{|err|} = 35.0^\circ$. This is problematic – our overall bearing estimation results from the head-mounted antenna in the previous section resulted in $\mu_{|err|} = 50.4^\circ$, which suggests that the static estimator for the reduced scan range (ie. picking the “middle” angle of the scan) ought to perform better than bearing estimation. This scenario is reflected in the actual shoulder-mounted antenna scans (examples of which are shown in Figure 49); the angular error magnitude using bearing estimation *is worse* than the static estimator (the means are approximately equal, but the bearing estimation’s error has larger variance) – see Table 3.

Table 3: Bearing Estimation Statistics For the Shoulder-Mounted Antennas

Antenna	Argmax Estimate	Static Estimate
Left Shoulder Antenna	$\mu_{ err } = 31.5^\circ$ $\sigma_{ err } = 29.3^\circ$	$\mu_{ err } = 35.1^\circ$ $\sigma_{ err } = 19.9^\circ$
Right Shoulder Antenna	$\mu_{ err } = 38.4^\circ$ $\sigma_{ err } = 32.7^\circ$	$\mu_{ err } = 35.3^\circ$ $\sigma_{ err } = 19.9^\circ$

Recall that the full-scan statistics did not tell the entire story. Bearing estimates with larger expected RSSI values yielded improved performance (as did scans near the tag). The same applies for the sub-360° scans. Scans with a maximum RSSI greater than 82 and scans less than two meters from the tag are likely to provide meaningful bearing estimates compared to the static estimator, as shown by the plots in Figure 50.

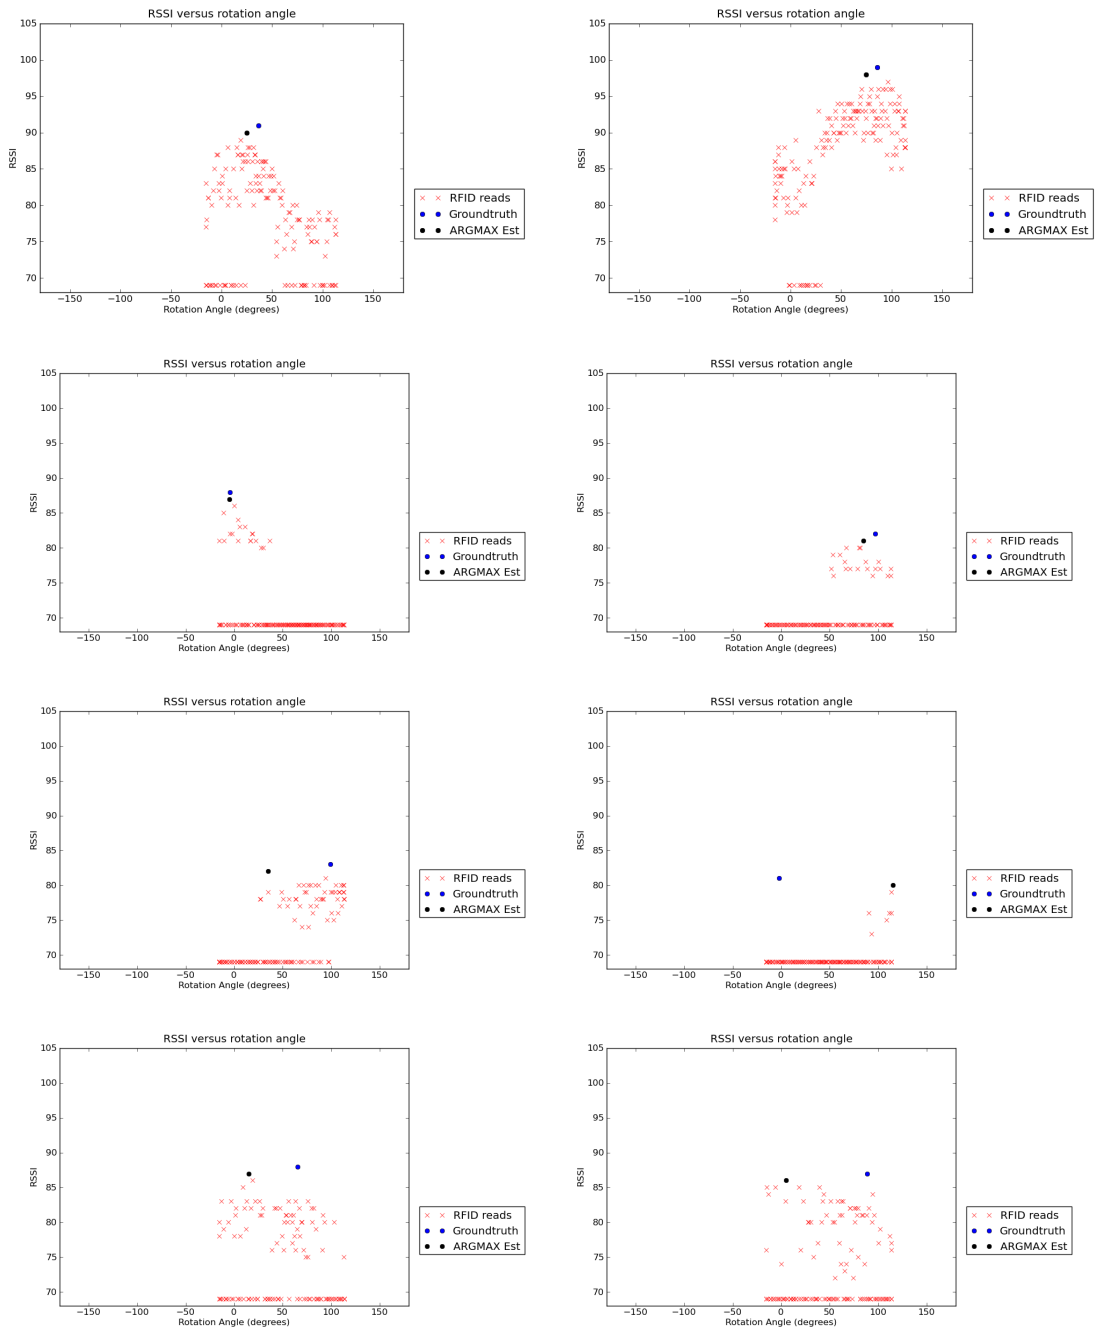


Figure 49: Various bearing estimation attempts using the left shoulder-mounted antenna making sub- 360° scans, where the true bearing to the tag lies within the scanned range.

Even when the bearing estimates for sub- 360° scans provide poor estimates, *comparing* the scans (or more precisely, the maximum RSSI) between the left and right shoulder antennas can be useful. Specifically, the difference between left-and-right can help the robot determine if the tag is more likely to the robot’s left or right. This principle is a key insight for RFID servoing, which we will examine later in this chapter.




Figure 50: We plot the bearing estimation error versus the distance from the tag (left) and the maximum expected RSSI in the scan (right) for the PR2’s left shoulder-mounted antenna (top row) and right shoulder-mounted antenna (bottom row). When the maximum expected RSSI is greater than 82 and when the robot is less than two meters from the tag, our bearing estimation for sub- 360° scans yields better measurements than the static estimator.

4.1.4.1 Leveraging the Shape of Radiation Patterns

In Figure 48, we saw that scans with larger maximum expected RSSI yielded better bearing estimates. In fact, if we look at Figure 46, we note that the measurements from these scans roughly resemble the shape of the antenna radiation pattern. This is not a coincidence; our earlier analysis using the Friis equation predicts that RSSI versus rotation angle (θ) should be proportional to a cross section of our reader antenna’s radiation pattern (Equation 50). This suggests that using the *shape* of the bearing estimation scan could dramatically improve our results.

In some of our earlier work (described in Section 6.2), we briefly used this insight for bearing

estimation. To perform the bearing estimation process using the shape of the reads, we assumed that our patch antenna possesses a 2nd-order exponential (Gaussian-like) profile, which has unimodal shape with global maximum. We considered the positive detections' RSSI readings as a function of the antenna direction (θ), so that we could represent RSSI as a function of theta by

$$g(\theta) = \alpha \cdot e^{-(\theta-\mu)^2/\sigma} + \beta. \quad (52)$$

Performing a non-linear least squares fit to the data yields estimates for these four parameters ($\alpha, \mu, \sigma, \beta$), from which we can obtain a bearing estimation (Θ) that uses all of readings along with their structure,

$$\Theta = \operatorname{argmax}_{\theta} g(\theta) = \mu \quad (53)$$

Later in Chapter 6, we use this technique with the higher-directivity Laird Technologies *S9028PC12NF* patch antenna to estimate the bearing toward a tag [39]. An example of this sort of estimation is shown in Figure 51. This same technique (fitting the shape of the curve) could conceivably be applied to any (directional) antenna radiation pattern, not just unimodal patch antennas. We performed very limited evaluation using this approach; future exploration is probably warranted, as it may yield superior angular estimates.

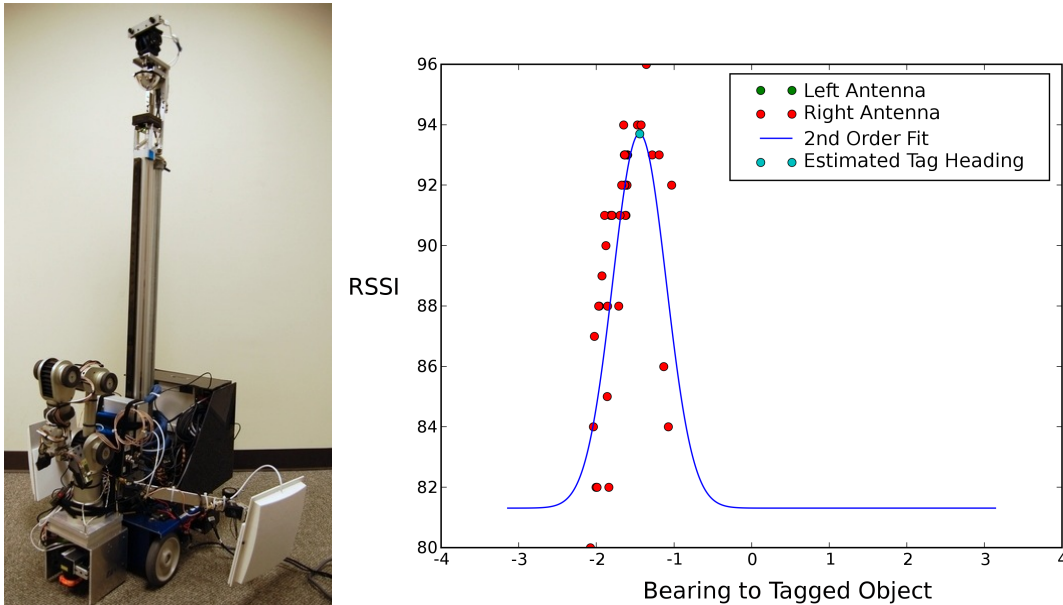


Figure 51: RSSI versus rotation angle should resemble a cross-section of the reader antenna's radiation pattern. Fitting the measured RSSI to the radiation pattern may yield superior angle estimates.

4.1.4.2 Using Robot Mobility to Actuate the Antenna(s)

So far, we have only been considering actuated antennas. An equally-viable solution is to use the robot’s own mobile base as the source of actuation, while keeping an un-actuated antenna rigidly affixed to the robot. We applied this technique in some of our early work using EL-E [37], as shown in Figure 52. When taking this approach, it is important to consider the ramifications of offsetting the antenna from it’s center of rotation – large offsets may have adverse effects on bearing estimation, where we assumed that the antenna was located at a fixed origin.




Figure 52: An early version of EL-E used a fixed antenna. By rotating the mobile base, we could obtain RFID sensor values for each bearing angle. Note: For this early work, we used the power-scaling as an approximation to RSSI. A lower minimum reader (transmit) power equates to a stronger RSSI reading; thus, we use an argmin optimization

4.1.4.3 Other Advanced Techniques

For each full panning scan (bearing estimation), the robot obtained a single estimate of the direction toward the tag. This opens up the possibility of using bearing-only SLAM algorithms (a Bayesian inference technique) where each measurement is actually the result of many relative RSSI measurements. This approach would be slower than using each measurement independently (eg. the 2-DoF probabilistic models discussed in Chapter 3); however, the bearing estimates may yield more robust real-world performance. This is an interesting concept worthy of future study.

One compelling alternative to mechanically actuating the RFID reader antennas is to use a multiple-antenna setup called a *phased array* to electronically steer the antenna’s radiation pattern

[98]. This could eliminate actuation all together and provide a bearing estimate for each positive tag detection. Furthermore, phased arrays can obtain the phase information about incident RF signals (as opposed to the RSSI, which is only magnitude information). The phase information can be used to estimate the backscatter-modulated signal's direction of arrival [8]. Having the phase information also opens up a host of other modern radar techniques: phase difference of arrival [94], frequency-modulated continuous wave radar, multi-frequency continuous wave radar, or even Doppler velocity tracking [136]. These techniques are commonplace in modern radars; however, they are currently beyond the capability of commercial UHF RFID readers.

4.2 RFID Servoing

In the previous section, we performed bearing estimation by performing a (time-intensive) rotation of the RFID reader’s antenna. The robot located a maximum in the expected RSSI from this scan to yield an estimate of bearing angle toward the tag. In this section, we look at how two antennas at the same location (co-located) but at different angles (off-pointing, called *left* and *right*) can provide an estimate of the gradient in RSSI so that the robot can incrementally rotate its mobile base to face in the direction of a tagged object. We call this process *bearing-only RFID servoing*. This type of controller is often referred to as servoing, as it uses a sensed value of error along with feedback control to improve the performance of a mechanism. This technique draws significant inspiration from classic radar techniques such as *sequential lobing* or *lobe switching* [136].

Adjusting rotational velocity according to bearing-only *and* moving forward at a constant velocity, a robot can successfully approach a tagged object of interest. We call this more general technique *RFID servoing* and we report on experiments in both simulation and in home / office environments. Researchers have used similar techniques to guide a mobile robot toward a battery-powered (active) target tag in a cluttered environment [82]; however, we believe our implementation and evaluation are the first to use passive (UHF) RFID tags. The trade-off between forward velocity and rotational velocity resembles a classic Braitenberg vehicle formulation [23], but the complex interplay between model parameters complicates closed-form theoretical analysis.

4.2.1 Bearing-Only RFID Servoing

In bearing estimation, the function of RSSI versus rotation angle for a patch antenna, $g(\theta)$ as depicted in Figure 43, is unimodal (convex) and smoothly varying (differentiable), so we should be able to perform a more efficient series of measurements to locate the bearing toward the tag. Assume we have a single “virtual antenna” rigidly affixed to the robot and aligned with the robot’s heading (coordinate frame, F_{robot} , which is rotated by Θ_n relative to some arbitrary map frame), as shown in Figure 53. If we could compute the derivative of g at our current heading, $\left. \frac{\partial g}{\partial \theta} \right|_{\Theta_n}$, then we could perform gradient ascent to iteratively refine our heading rather than performing a time-intensive articulation of the virtual antenna (a brute-force search over Θ). To formalize this, we could rotate to a new heading (Θ_{n+1}) from our previous heading (Θ_n) according to the gradient and a fixed step

size (γ),

$$\Theta_{n+1} = \Theta_n + \gamma \cdot \left. \frac{\partial g}{\partial \theta} \right|_{\Theta_n}. \quad (54)$$

If we performed this process every t seconds, then we have effectively defined a feedback controller that alters the robot's angular velocity ($\dot{\Theta} = \frac{\partial \Theta}{\partial t} \approx \frac{\Theta_{n+1} - \Theta_n}{t}$) in proportion to the gradient,

$$\dot{\Theta} \propto \left. \frac{\partial g}{\partial \theta} \right|_{\Theta_n}. \quad (55)$$

In fact, assume that we can obtain instantaneous measurements of $\frac{\partial g}{\partial \theta}$. Let

$$h(\Theta) = \dot{\Theta} = \kappa \cdot \left. \frac{\partial g}{\partial \theta} \right|_{\Theta_n}, \quad (56)$$

and consider the Lyapunov function, $V(\Theta) = |\Theta - \hat{\Theta}|$, then

$$\begin{aligned} \frac{dV}{dt} &= \frac{dV}{d\Theta} \cdot \frac{d\Theta}{dt} \\ &= \text{sign}(\Theta - \hat{\Theta}) \cdot h(\Theta) \\ &= - |h(\Theta)| < 0 \quad \forall \Theta \neq \hat{\Theta}, \end{aligned} \quad (57)$$

which implies that the true heading toward the tag ($\Theta = \hat{\Theta}$) is asymptotically Lyapunov stable [78, 157], so this feedback controller should converge to the true tag heading.

Unfortunately, we cannot obtain multiple RSSI measurements from our single virtual antenna². However, we could use two physical antennas offset at fixed angles ($\pm\alpha^\circ$) relative to our virtual antenna to “sample” the virtual antenna's RSSI function at those two angles, as shown in Figure 53³. In this case, we can approximate the derivative by the first-order difference equation:

$$\left. \frac{\partial g}{\partial \theta} \right|_{\Theta_n} \approx \frac{g(\Theta_n + \alpha) - g(\Theta_n - \alpha)}{2\alpha}, \quad (58)$$

where $g(\Theta_n + \alpha)$ is the RSSI of the left antenna ($RSSI_{left}$) and $g(\Theta_n - \alpha)$ is the RSSI of the right antenna ($RSSI_{right}$). Substituting this approximation into our feedback controller, we define bearing-only RFID servoing as:

$$\Delta RSSI = RSSI_{left} - RSSI_{right} \quad (59)$$

$$\dot{\Theta} = \kappa \cdot \Delta RSSI. \quad (60)$$

²This may be possible using more advanced techniques such as phased-array beam steering, but these are not supported by our UHF RFID hardware.

³These are sketches to illustrate the overall concept, not real data.

The left antenna will have a maximum RSSI when $\Theta = \hat{\Theta} - \alpha$, the right when $\Theta = \hat{\Theta} + \alpha$. The difference between left and right RSSI, $\Delta RSSI$, gives an indication of which direction the robot should turn so as to face the tag. If the left antenna receives a stronger signal (as when $\Theta < \hat{\Theta}$), the robot should turn left (an angular velocity $\dot{\Theta} > 0$); if the right antenna receives a stronger signal (as when $\Theta > \hat{\Theta}$), then the robot should turn right (an angular velocity $\dot{\Theta} < 0$).

In practice RSSI measurements are noisy, may not be obtained for certain tag configurations, and are subject to non-ideal propagation effects (eg. multipath, shadowing, etc). We reformulate RFID servoing, treating the left and right RSSI measurements as random variables ($RSSI_{left}$ and $RSSI_{right}$) so that we alter the robot's angular velocity in proportion to the expectation of the difference random variable, $\Delta RSSI = RSSI_{left} - RSSI_{right}$ according to:

$$E(\Delta RSSI | \theta = \theta) = E(RSSI_{left} | \theta = \theta) - E(RSSI_{right} | \theta = \theta) \quad (61)$$

$$\dot{\Theta} = \kappa \cdot E(\Delta RSSI | \theta = \theta) \quad (62)$$

to perform bearing-only RFID servoing, where we use a feedback controller to adjust the robot's heading so that the robot is eventually oriented toward the RFID tag.

4.2.2 RFID Servoing

Unlike bearing estimation, which required the tag to remain at a fixed position while articulating the transmit antenna (a time-intensive process), bearing-only RFID servoing will continuously update the robot's heading even as a tagged object moves. When the tagged object stops moving, the controller stabilizes to the equilibrium point, at which time the robot will be oriented facing the new tag position. Stabilizing to a new equilibrium point would also occur if the *robot* were to move to a new location and then resume bearing-only RFID servoing. If the robot were to just move forward slowly enough (such that the servoing stabilized between each discretized "forward step"), then the robot would approach the tag; we employed a similar technique during some of our previous work [37]. Unfortunately, our analysis provides no guarantees about how *long* it takes for the bearing to stabilize after each forward step, only that it does so *asymptotically*. Instead of waiting for stabilization at each forward step, we employ an alternative approach dubbed *RFID*




Figure 53: Consider a robot with a single (virtual) antenna mounted to its chest (A). If the robot rotates in place, this antenna would obtain the RSSI signal (C) and with derivative (D). We can approximate the derivative by taking the difference in RSSI between two antennas at fixed, offset angles ($\pm\alpha^\circ$) (B) – their signals are shown in (E).

servoing, where the robot simply moves forward at a constant velocity while updating its angular velocity in proportion to the expected difference in left and right RSSI readings; the robot halts when its forward path is impeded by an obstacle. This technique is akin to Braitenberg vehicles [23]. We know there will be subtle interplays between the forward and angular velocities; however, we benefit from an algorithm that can approach tagged objects in real-time – and even track or follow tagged objects that move.

4.2.3 RFID Servoing in Simulation

Using the simulated RFID sensor model from Section 2.4 in the PR2 Gazebo simulation environment, we can simulate RFID servoing (under relatively ideal conditions) in a variety of different configurations: antenna mountings, starting conditions, type of reader transmit/receive antenna, type of tag, angular / forward velocity, and proportionality constants.

For example, Figure 54 shows the trajectories of a simulated robot performing RFID servoing in a simulated version of the same room that was used to capture data in Chapter 2 (Figures 45 and 25). The simulated PR2 robot was using two S9025P reader patch antennas mounted to the shoulders, with off-pointing angle(s) $\alpha = 40^\circ$, servoing to an Alien Squiggle tag in the “vertical” position from Figure 23. The current servoing behavior requires open space between the robot and the tagged object, as it stops as soon as the on-board range sensors (eg. laser rangefinders) detect a potential collision in front of the robot.

We have highlighted several simulated paths in Figure 54. In the context of Braitenberg vehicles, which operate on differential equations, these cases illustrate examples where the initial conditions and sensor characteristics result in a dynamic system that converges-to, diverges-from, and (nearly) circles the simulated tag. Experimenting with different system parameters, we can explore the trade-offs (eg. time to complete versus divergent cases) without extensive on-robot testing.

Due to the unmodeled variations that we have explored in the Friis-based model (eg. how it degrades when a tag is placed on an object, or multipath), it is still important to vet RFID servoing’s ability to approach tagged objects on the actual robot.

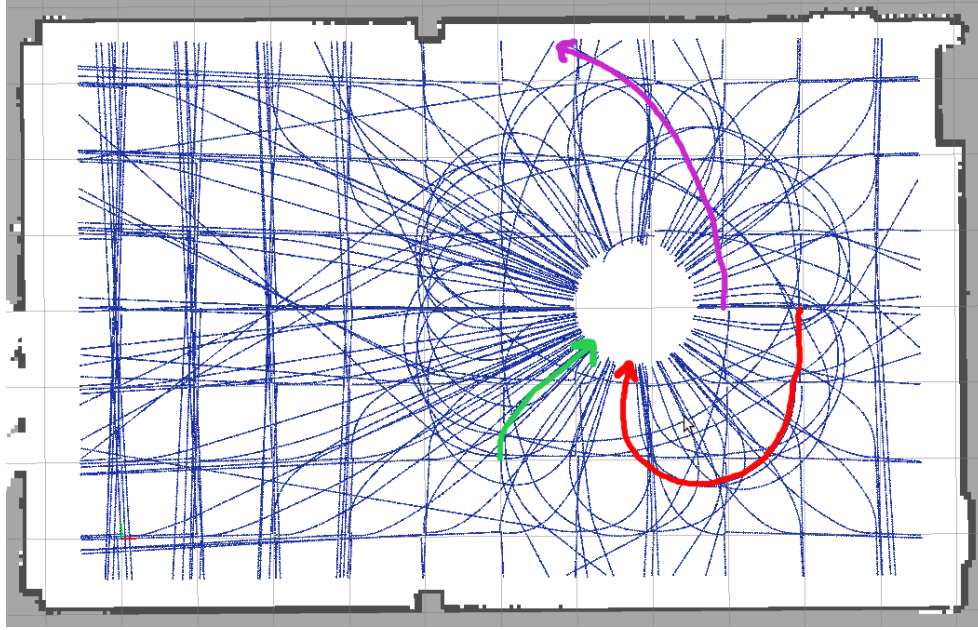


Figure 54: We evaluate RFID servoing in simulation by initially placing the PR2 at each of the grid locations and executing the servoing algorithm. Here, we can witness the subtle interplay between angular and forward velocities, where the PR2’s trajectory converges to the tag (green), diverges from the tag (purple), and nearly circles the tag (red). [The tag is located at the center.]

4.2.4 RFID Servoing on the PR2

Our bearing-only RFID servoing system is analogous to a conical scan radars; from the radar literature, suggested squint angles (α) are 28-45% of the antennas’ 3-dB beamwidth [24]. For our system, we selected $\alpha = 40^\circ$, which is 40% of the antennas’ (100°) 3-dB beamwidth. As shown in Figure 55, our antennas are not colocated at the same origin owing to physical constraints and robot mounting considerations; instead, they are mounted on the robot’s shoulders. We employ a running average filter to compute the expected RSSI for the left and right antennas (and thus the expected difference). The ThingMagic reader returns RSSI values between 72 to a maximum of 105; we assign a RSSI value of 69 to negative tag detections so that the system can operate in the presence of noise (eg. spurious negative readings). Thus, $|\Delta RSSI| \in [0, 36]$. We selected a constant forward velocity of 10cm/sec and manually tuned the proportionality constant $\kappa = 1.15^\circ/sec$ to yield angular velocities $|\dot{\Theta}| \in [0, 40]^\circ/sec$; these settings provide a good trade-off between forward velocity and angular velocity such that the robot approaches tags in a timely fashion while still promptly responding to bearing updates. These are the same settings used for the simulated servoing results

shown in Figure 54.

We evaluated RFID servoing in a large, uncluttered environment. We performed two sets of experiments. In the first experiment, we positioned a tag (vertically mounted) on a large cardboard box in near the center of the large room (in the same fashion as the relatively ideal data captures from Chapter 2, Figure 25). Using ROS's *navstack* package, we positioned the PR2 at various initial positions and orientations, and then the robot performed RFID servoing until its forward progress was impeded – either by the room's walls (a failure) or the tag's cardboard box (a success). An example of the a successful servoing attempt is shown in Figure 55. For initial headings within $\pm 90^\circ$ of the ground-truth heading, the PR2 successfully approached the tag in 152 / 170 (89.4%) of the trials using *only* RFID servoing in this large open room; however, the trajectories taken by the robot (Figure 56) can sometimes be more circuitous than predicted by RFID servoing in simulation.




Figure 55: The PR2 uses its two shoulder-mounted RFID antennas to servo to an Alien Squiggle tag mounted vertically on the front face of a white cardboard box.

In the second experiment, we tagged an orange medication bottle (with the tag spiraled around the bottle) and placed the tagged object in clutter on a small table in the middle of the same room. Again, we positioned the PR2 at various initial positions and orientations, and then the robot performed RFID servoing until its forward progress was impeded – either by the room's walls (a

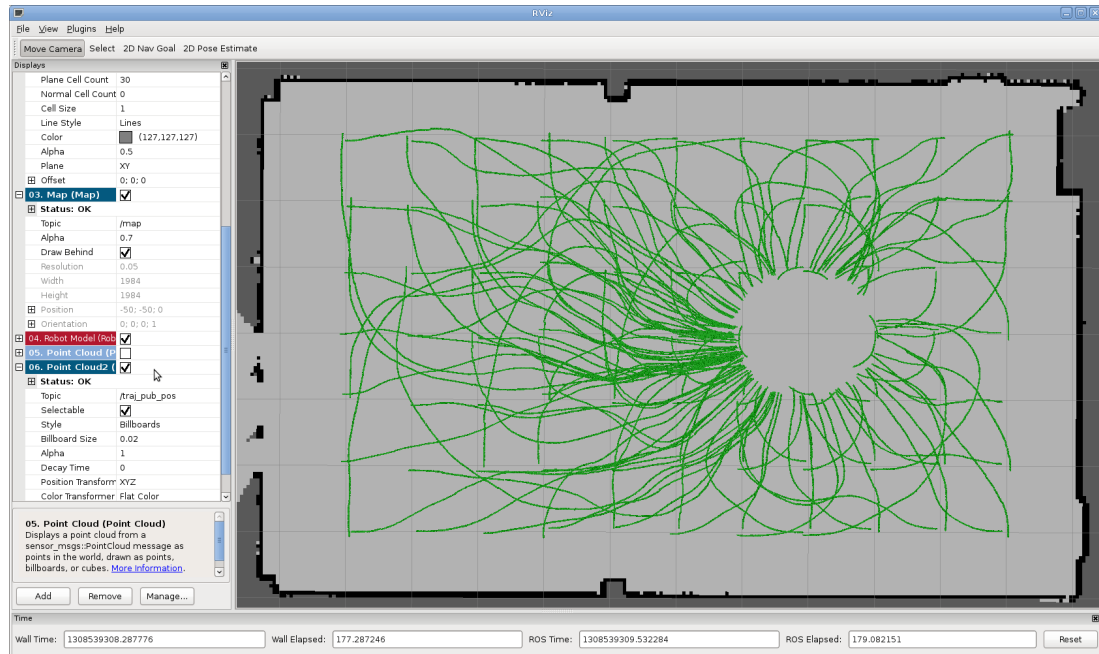


Figure 56: The PR2 robot uses its two shoulder-mounted RFID antennas to servo to the vertically-oriented RFID tag. These figures illustrate the robot’s trajectories for initial conditions that resulted in successful (top) and unsuccessful (bottom) approaches to the tag.

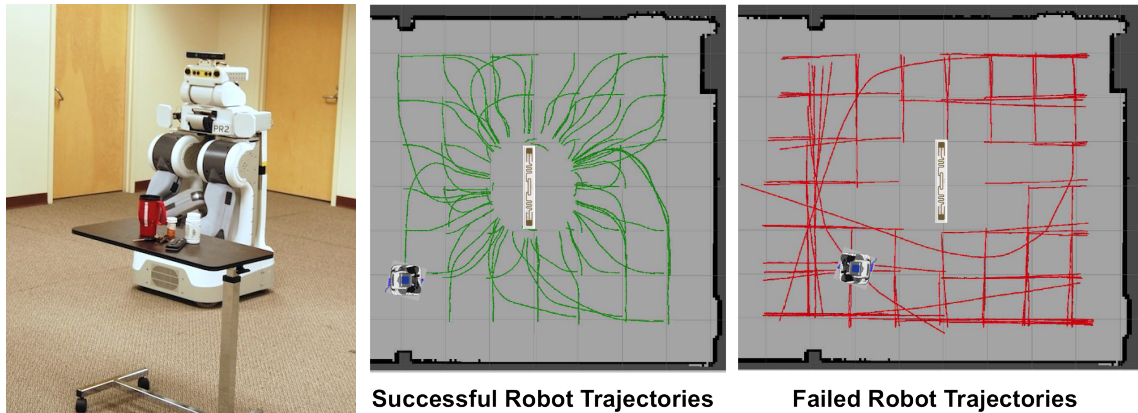


Figure 57: **Left:** The PR2 with two shoulder-mounted antennas performs RFID servoing to approach the tagged orange medication bottle. **Middle:** The PR2 was placed at each grid location, and each of four cardinal directions. This plot shows the trajectories taken by the robot for successful attempts. **Right:** The trajectories of unsuccessful servoing attempts.

failure) or the table upon which the tagged object sat (a success). For initial headings within $\pm 90^\circ$ of the ground-truth heading, the PR2 successfully approached the tag in 77 / 84 (91.6%) trials using *only* RFID servoing in this large, open room. We show the trajectories taken by the robot during RFID servoing for successful and failed attempts in Figure 57.

In both of these scenarios (a tag in relative isolation and a tagged medication bottle in clutter), the initial conditions that resulted in successful RFID servoing attempts (ie. $\Theta_{init} \in \pm 90^\circ$) are well-matched to the final conditions afforded by bearing estimation in the last section. We will use this observation later in this chapter to develop a hybrid global-local RFID search algorithm; we also show how they can be combined (in Chapter 6) to perform medication delivery.

In another set of experiments, we evaluated EL-E’s ability to approach a tagged object on a bookshelf from a grid of 36 distinct starting locations in a $3.7 \times 7.3m$ room – see Figure 58. At each location, we initially oriented EL-E so that the robot faced the bookshelf (as though having performed bearing estimation). We considered a trial successful if EL-E stopped less than 1 meter from the tagged object, and if that object was fully visible from EL-E’s on-board camera. This definition of success is well-matched to other close-range methods we have developed on EL-E (discussed later in Chapters 5 and 6). EL-E was successful in 32 of the 36 initial conditions, as shown in Figure 58 [40].



Figure 58: *Top Row:* In a large room (right), a bookshelf contains 37 tagged objects. EL-E has been tasked with approaching the tagged medication box, boxed in red (left). *Bottom Row:* The results of EL-E approaching the tagged object (red circle). White circles indicate success; black circles indicate failure.

For both the PR2 and EL-E, this approaching behavior is a valuable foundation but has limitations. The current servoing behavior requires open space between the robot and the tagged object, as the robot stops as soon as its forward progress is impeded (ie. its on-board laser range sensors detect an impending collision). Since RFID servoing does not maintain an explicit representation of the tag’s location, the robot must rely on other methods to determine if it has successfully approached the desired tagged object. Often, this will be accomplished through the application of other, complementary sensing technologies (such as vision). We explore this idea later in Chapter

6.

4.2.5 Discussion

The trade-off between forward velocity and angular velocity is not the only major design decision for RFID servoing. The off-pointing angle, α , is also of crucial importance. In the radar literature related to conical scan radars (radars that take multiple readings centered about a radar target), α is sometimes called the “squint angle.” To minimize angular error, the optimal squint angle is 45% of the antennas’ 3-dB beamwidth, the condition where the two antennas’ radiation patterns are just barely overlapping. In practice, squint angles can range anywhere from 45% to 28% of the 3-dB beamwidth, depending on how conservative the designer wants to be in ensuring that both antennas can perceive the target [24]. For our choice of Laird Technologies S9025P patch antennas (with 100° 3-dB beamwidth) and $\alpha = 40^\circ$, we are using a squint angle of 40% of the 3-dB beamwidth – well within the range suggested by the literature.

The choice of patch antenna (specifically, its 3-dB beamwidth) is an important consideration for RFID servoing. Unlike bearing estimation, high-directivity antennas (eg. pencil-beams) are not necessarily a desirable trait. Having a wider beamwidth allows the two antennas to initially perceive the tag over a broader set of initial conditions, but at the expense of tag detection range (antennas with lower directivity have shorter detection ranges per the Friis equation) and possibly poorer angle estimation due to increased noise from multipath.

4.3 Bearing Estimation – Azimuth (Θ) and Elevation (Φ)

We can extend bearing estimation to two dimensions, estimating both azimuth (Θ) and elevation (Φ) angles toward a stationary tag by taking RSSI measurements while panning ($\theta \in [-180^\circ, 180^\circ]$) and tilting ($\phi \in [-90^\circ \dots 90^\circ]$) the reader antenna. This corresponds quite closely to an acquisition radar used to locate (and track) aircraft in the early days of radar [136].

We develop a method of displaying the RSSI values as images, where the value at each pixel corresponds to the expected RSSI measurement at the associated pan and tilt angles. We dub these visualizations “RSSI images.” They provide a natural and intuitive way to visualize the otherwise invisible RFID sensor values. Later, in Chapter 6, we show this representation can be used to fuse data from complementary sensors and facilitate mobile manipulation.

4.3.1 Similarity to Azimuth-Only Bearing Estimation

Because we employ patch antennas with a unimodal radiation pattern and global maximum in both pan/tilt, the function $RSSI = g(\theta, \phi)$ will have a unique maximum when the antenna is pointed directly at the tag under ideal, free-space conditions. Thus, we employ a similar optimization,

$$\hat{\Theta}, \hat{\Phi} = \underset{\theta, \phi}{\operatorname{argmax}} g(\theta, \phi). \quad (63)$$

Again, we assume that RSSI measurements are composed of a noise-free component, $g(\theta, \phi)$ and a zero-mean noise signal represented by its second-order (Gaussian) statistics, $\mathcal{N}(0, \Sigma)$, so that the distribution of the random variable $RSSI$ is expressed as

$$P(RSSI | \theta = \theta, \phi = \phi) = g(\theta, \phi) + \mathcal{N}(0, \Sigma). \quad (64)$$

The optimization then seeks the angles that yield the maximum expected RSSI.

$$\hat{\Theta}, \hat{\Phi} = \underset{\theta, \phi}{\operatorname{argmax}} E(RSSI | \theta = \theta, \phi = \phi). \quad (65)$$

4.3.2 RSSI Images

Dealing with a 2-dimensional quantity related to pan-tilt angles suggests a natural image-based representation where each $\langle x, y \rangle$ pixel corresponds to a discretized pan-tilt steradian $\langle \theta \pm \epsilon, \phi \pm \epsilon \rangle$

and has an intensity I given by the expected RSSI,

$$I(x = \theta, y = \phi) = E(\text{RSSI} \mid \theta = \theta, \phi = \phi). \quad (66)$$

Essentially, this treats the pan-tilt antenna as a spherical camera, which yields an “RSSI image.” When co-located (and calibrated) with a visible light camera, RSSI images yield insights into the otherwise invisible RF environment. For example, consider the scenario depicted in Figure 59. Here, a tagged object (outlined in red) is moved from left-to-right across a scene. The RSSI image roughly tracks this motion and suggests a region of interest where other sensors should look for the desired tagged object. In many respects, the RSSI image “focuses” attention in the complementary (visual) sensing modality. This suggests a rather straight forward method of sensor fusion that we revisit in Chapter 6.




Figure 59: The RSSI images (bottom) *roughly* track the position of the tagged bottle (boxed in red) in the raw camera images (top) and yield insights into the non-visible RF sensor measurements. Using this early (off-robot) test rig, the RSSI images are roughly calibrated with the corresponding camera images.

Sometimes objects that are occluded in the visible spectrum are still readily sensed by our 900 MHz RF measurements. In the case of Figure 60, a tagged object is located in one of the closed cabinet drawers. While we cannot visually locate the tagged object, the RSSI image suggests that the object is indeed hidden in the cabinet.




Figure 60: Under some environmental conditions, RSSI images can detect the presence of a tagged object even when it is visually occluded. Using this early (off-robot) test rig, the RSSI images (bottom) are roughly calibrated with the corresponding camera images (top).

4.3.3 RSSI Images in Practice

As shown in Figure 61, we employ an actuated Laird Technologies *S9028PC12NF* patch antenna (see Section 2.1.3.2). This antenna has a greater directivity (65° 3dB-beamwidth) than the *S9025P* (100° 3dB-beamwidth) used for earlier bearing estimation. In the RSSI images we show, the scale is roughly 7° per 45 pixels (or 0.155° per pixel) corresponding to the 640×480 pixel visual image. Using this antenna, the stronger RSSI image regions (colored red) are generally 200-250 pixels wide / tall – corresponding to $30\text{-}40^\circ$ in either dimension. This corresponds nicely with half the 3dB-beamwidth of our antenna, and seems to reinforce the discussions in Section 4.1 regarding the accuracy of the pan-tilt estimation using our argmax optimization. The key design take-away: the more directional the antenna, the better the angular estimates and more precise (peaked) the RSSI image.

To create an RSSI image, we pan and tilt the directional patch antenna, used for both transmit and receive (monostatic mode), through various angles – as shown in Figure 61. The RSSI images in




Figure 61: An early version of the mobile, manipulating robot EL-E (left) pans and tilts a high-directivity patch antenna (right – ordered left-to-right, top-to-bottom) to capture a RSSI image.

Figure 59 and 60) were *not* created by panning and tilting the transmit / receive antenna to each of the 300,000+ (640×460) pan-tilt locations as governed by the visual image – this would have required an unreasonable amount of time (and the angular resolution of the Robotis Dynamixel servos used to actuate the antennas could not achieve the necessary image resolution). Plus, we know from our choice of antenna (and its 3dB-beamwidth) that the effective angular resolution of our RFID sensor is much lower than the resolution of the visual camera. We use this insight to develop a more expedient and sparse sampling technique to build the RSSI images you see. We pan the antenna through $\theta = \pm 60^\circ$ over 10 seconds for each of 12 tilt angles, which are uniformly spaced between $\phi = [-60^\circ, +30^\circ]$. At a 12Hz sampling rate, this results in 120 RFID measurements per panning scan or 1440 total RFID reads per RSSI image. For each reading, we store the RSSI value in the corresponding 640×460 pixel location. This results in sparse, sub-sampled RSSI image. We then apply a computationally-efficient Gaussian convolution filter with standard deviation, $\sigma = 45$ pixels ($\sigma = 7^\circ$) to create the final RSSI images – essentially interpolating between and smoothing adjacent RFID measurements. Unfortunately, the Gaussian filter transforms the pixel's intensity $I(x, y)$ units from RSSI into an unspecified unit whose range varies depending on the sampling granularity (scan time and number of pan / tilt angles) and the convolution filter's width. To make the resulting units

consistent across different configurations, we normalize the RSSI image’s intensities to occupy the range between $[0.0, 1.0]$. Relative to the maximum pixel value I_{max} and the minimum pixel value I_{min} , we used a linear scaling given by,

$$I(x, y) \leftarrow \frac{I(x, y) - I_{min}}{I_{max} - I_{min}}. \quad (67)$$

RSSI images are quite illustrative, as they let us visualize the otherwise-invisible RFID sensor measurements superimposed next to a camera image of the environment. A number of RSSI images are shown under various conditions in Figures 63-65, spanning the spectrum from nearly-ideal results, to others that are decent but have significant multipath, and finally a select few that produced lackluster results.

We use data to understand the relationship between the RSSI image intensity $I(x, y)$ and the likelihood, $P(I(x, y) | \theta, \phi, tag_{present})$, that we obtain such a measurement when the tag is *present* (denoted $tag_{present}$) at the corresponding pan/tilt angles $I(x = \theta, y = \phi)$. From *sixty* hand-labeled ground-truth RSSI images, we can build a histogram sensor model of this distribution. The ground-truth observations were recorded by selecting the center of the tagged object from a (registered) visual camera image, and selecting all pixels in the RSSI image within a 10-pixel radius; this process, along with photos of the surrounding environment is described in Section 6.2. From the same data, we also build a distribution of the likelihood, $P(I(x, y) | x, y, tag_{absent})$, that the measured RSSI image intensity would be obtained if the tag were *absent* (denoted tag_{absent}) at the given location. Both of these distributions are shown in Figure 62.

We can make two observations about the distributions depicted in Figure 62. First, a pixel with strong normalized RSSI image intensity ($I(x, y) \approx 1.0$) is more than *fifteen* times as likely to have the tag present ($tag_{present}$) than a pixel with weaker intensity ($I(x, y) < 0.5$), or

$$P(I(x, y) \approx 1.0 | x, y, tag_{present}) \geq 15 \cdot P(I(x, y) < 0.5 | x, y, tag_{present}). \quad (68)$$

In the context of the RSSI images depicted in Figures 63-65, where red and blue correspond to intensities of 1.0 and 0.0 respectively, the deep red regions are *fifteen* times more likely to contain the tag than the yellow, green, or blue regions. Second, a pixel with *zero* intensity ($I(x, y) = 0.0$) is *five* times as likely to **not** contain the tag than a pixel that possesses some non-zero value ($I(x, y) > 0.0$),




Figure 62: From hand-labeled ground-truth data, we can build a Bayesian sensor model akin to those in Chapter 3 that will be used in Chapter 6 to fuse RSSI image data with data from other sensors.

or

$$P(I(x, y) = 0.0 \mid x, y, tag_{absent}) \geq 5 \cdot P(I(x, y) > 0.0 \mid x, y, tag_{absent}). \quad (69)$$

In the context of Bayesian estimation from Chapter 3, these distributions form a data-driven sensor model for RSSI images. In Chapter 6 we will show how these RSSI image intensity distributions can be used to fuse RSSI image data with data from other sensors.




Figure 63: We created numerous RSSI images under various conditions using the robot-mounted system in Figure 61. The RSSI images (right) correspond to the environmental conditions depicted in the raw camera images (left); the tagged objects are boxed in red: medication box on a shelf (1st row), medication box on a shelf that is barely detectable (2nd row), an empty water bottle on the floor (3rd row), and a TV remote on a shelf (4th row). In these RSSI images, the artificial coloring represents a scale from 0.0 (blue) to 1.0 (red).




Figure 64: Additional RSSI images captured using the robot-mounted system in Figure 61. These show conditions that result in various amounts of multipath, particularly ground-bounce. The tagged objects are boxed in red: medication box on a table (1st row), medication box on a shelf (2nd row), a spirally-tagged medication bottle on the floor (3rd row), and an empty water bottle placed up high (4th row). In these RSSI images, the artificial coloring represents a scale from 0.0 (blue) to 1.0 (red).




Figure 65: Additional RSSI images captured using the robot-mounted system in Figure 61 showing borderline failure cases. The tagged objects are boxed in red: empty water bottle on a table results in multipath (1st row), TV remote on the floor is at the edge of a fortuitous detection nearby (2nd row), the ground-bounce multipath dominates when this empty water bottle was on this particular table (3rd row), and multipath again dominated when measuring the empty water bottle on the floor (4th row). In these RSSI images, the artificial coloring represents a scale from 0.0 (blue) to 1.0 (red).

4.4 Global RFID Search

In the previous sections, we examined several local optimization-based RFID behaviors that inform the robot how to move relative to its current pose to obtain increased RSSI measurements and thus get closer to the tag. But a mobile robot can do more than just pan and tilt its antennas; it can use its mobile base to travel to new locations, where it can make new RFID measurements of the tag and its surrounding environment. In this section, we develop a global optimization-based technique for RFID search where the robot explores its broader environment to capture RFID measurements from diverse positions and orientations. The goal of RFID search is to determine a robot position and orientation (pose, \hat{P}) that yields a “good” view of the tag, where “good” is (1) near the tag and (2) oriented in the direction of the tag.

In this section, we show how global RFID search can be accomplished by maximizing the expected RSSI per the following optimization,

$$\hat{P} = \underset{P}{\operatorname{argmax}} \operatorname{E} (\mathbf{RSSI} \mid \mathbf{P} = P). \quad (70)$$

We discuss how this optimization relates to RSSI measurements predicted by the Friis model, and we discuss some of the practical constraints of real-world robot operation. Later, in Section 4.5, we show how to perform a *sparse* global RFID search combined with local optimization-based behaviors to discover, locate, and approach tagged objects in a real home environment.

4.4.1 Related Work

We define search as the task of locating a desired (often transient) object in an environment. This is a foundational robot capability upon which many other robot behaviors are built, such as object fetching and retrieving – which actually requires search in at least two instances: locating the object to grasp, and then locating the person for delivery. Borrowing insights from Lilienthal *et al* , search algorithms can be broadly classified into two types: tracing-based methods and analytical-model based methods [95]. Tracing-based methods seek to reposition the robot as close as possible to the desired object *without* explicitly yielding the object’s location as a final estimate. Meanwhile, analytical-model based methods seek to explicitly estimate the 2-DoF, 3-DoF, or 6-DoF pose parameters of the object (either on a map or relative to the robot). Both methods provide valuable

capabilities and can sometimes be used in concert; often, the choice of algorithms is largely dependent upon the robot's complement of sensors.

4.4.1.1 Tracing-Based Methods

Again, the goal of tracing-based methods is to reposition the robot as close as possible to a desired object. These methods are especially common when the sensor (or system) is low-bandwidth and/or low-acuity. For example, numerous tracing-based methods stem from the odor source localization literature, which relies on a low-bandwidth, low-acuity (scalar) odor sensors to locate an odor source [93]. Sometimes these tracing-based methods are called “chemical plume tracing” and there are many different approaches: gradient-following, upwind, SPIRAL algorithm, etc [93]. In fact, RFID servoing (Section 4.2) and RFID search in this section are both examples of a tracing-based method.

The gradient-based tracing methods suffer from local minima, and may have difficulties dealing with large (global) environments. This was something we observed with RFID servoing, where obstacle detection was the stopping condition; this could easily prevent the robot from exploring an entire building. Other tracing-based methods use exploration algorithms to create evidence grids that contain a history of the sensor values obtained when the robot was at each location in a map. A number of exploration strategies exist to explore the global environment. Brute force search is certainly one option (visiting every single location and orientation); however, other methods may provide more efficient (or sparse) coverage. For example, frontier-based exploration queues up locations at the boundary of known empty space and unexplored space to ensure complete coverage is obtained by the robot (within some sensing radius) [162]; this method is used to build maps using the PR2 [118], but could also be applied to search a room for RFID tags. Meanwhile, other exploration algorithms provide provably complete exploration strategies to ensure that an entire environment is covered [81].

In this section, we develop a refined RFID search algorithm that combines an evidence grid approach (using a minimally-defined “snaking” exploration pattern) to select a good starting location, which then gets refined using RFID servoing. To the best of our knowledge, this is the first tracing-based robot search algorithm using long-range UHF RFID.

4.4.1.2 Analytical-Model Based Methods

High-bandwidth and/or high-acuity sensors, such as cameras and point cloud sensors naturally lend themselves to analytical-models owing to their high-resolution, largely-geometric data. Search algorithms using these sensors typically focus on object recognition (ie. detecting the desired object in the stream of high-bandwidth data), from which the location or pose of the object can be immediately calculated using the intrinsic properties of the sensor. With the availability of low-cost depth cameras (sensors that provide both a RGB image and a depth reading for each pixel), researchers are building dense 3D maps of the environment and large databases of 3D object models [58, 89], which help build “semantic maps” – maps that contain information about the identity and location of objects within the room. Semantic mapping is essentially another form of search. For example, work by Quigley *et al* uses dense point clouds to discover, detect, and locate cups in an office environment using both visual and 3D features [121]. These algorithms rely on object recognition; solving the object recognition is a oft-examined, but largely unsolved problem, especially in large environments with diverse (potentially unique) objects. Reducing the number of false positives is one of the key challenges in this space.

We believe that the complementary nature of UHF RFID sensing and (depth) cameras offers a compelling value proposition for robotic search, where the refined RFID search quickly positions the robot nearby the desired tag (with high certainty that the signals belong to the desired object) and then the high-acuity sensors can be used to recognize and localize the desired object from this advantageous viewpoint. We believe this complementary nature would greatly benefit existing work that uses (false positive prone) color histograms for long-range sensing and visual SIFT features for close-up object detection and localization [97]. We explore this concept in more detail in Chapter 6.2.

Of course, analytical-model based methods can also be used for RFID search. In fact, the Bayesian localization techniques discussed in Chapter 3 are examples of analytical-model methods. While not the focus of this chapter, it should be possible to marry these two techniques – either using them to produce independent vantage points, or perhaps to combine their respective estimates. This remains an area for further research.

4.4.2 Exhaustive Brute-Force Search

A robot could perform an exhaustive, brute-force search by moving its RFID reader antenna to all $\langle X, Y, Z \rangle$ Cartesian locations (relative to the unknown tag location), and at each location the antenna could be panned and tilted through all possible orientations. At each location, we already know from bearing estimation (Section 4.3) that the Friis equation predicts a maximum expected RSSI when the reader's patch antenna is directly pointed at the tag; since these poses are already the result of an *argmax* optimization, we can consider just this subset of antenna poses. Now, the question becomes: which of these locations (with the reader antenna pointed directly toward the tag) yields maximum expected RSSI measurements?

Consider just a single ray emanating from the tag, as shown in Figure 66. For an infinitesimally fine search granularity, the robot would have obtained exhaustive pan-tilt measurements for each point along this ray (ie. different radii, r), and the local *argmax* operation would have selected the reader antenna orientation(s) directly pointing toward the tag. For this subset of poses, the only 6-DoF pose parameter that differs from point-to-point is the radius from the tag, so that the Friis equation predicts,

$$RSSI \propto \left(\frac{1}{r}\right)^4, \text{ or} \tag{71}$$

$$RSSI \propto -r \text{ (dB)}. \tag{72}$$

Specifically, the RSSI increases monotonically for decreasing distances (radii) along this ray. Thus, the maximum expected RSSI (along this ray) occurs when the RFID reader antenna is infinitesimally close to the tag and oriented directly toward it. This is true for all rays emanating from the tag; any one of these poses would satisfy our goal for global RFID search.

4.4.3 Practical Constraints

In practice, exhaustive brute-force search is intractable; the robot may not be able to achieve certain poses owing to: its navigation capabilities, antenna mounting considerations, nearby obstacles or clutter, and even the physical extent of the tagged object. For each ray (i) emanating from the tag, there will be some minimum viable distance, d_i , attainable by the robot's RFID reader antenna where the antenna is pointing directly toward the tag. Again, since RSSI increases monotonically

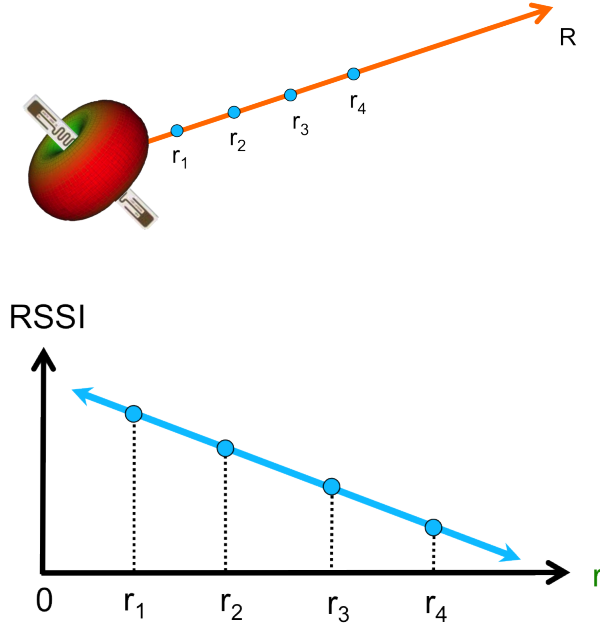


Figure 66: For each location along a ray emanating from a tag (eg. r_1), an RFID reader with a patch antenna will obtain maximum RSSI readings when directly pointed toward the tag. For these configurations, the only 6-DoF pose parameter that varies from point-to-point along the ray is the radius (distance) from the tag; the Friis equation predicts that RSSI is a linear function of radius (in logarithmic units), so the maximum expected RSSI (along this ray) when the reader antenna is infinitesimally close to the tag and oriented directly toward it.

for decreasing distances along each ray (eg. a local *argmax* along each ray), we can consider just the subset of antenna poses corresponding to d_i .

In Figure 67, we show the positions corresponding to d_i for five rays emanating from the tag. For global RFID search, the robot will select the pose that yields the maximum expected RSSI. In general, this pose may not be closest to the tag; the tag's radiation pattern plays a crucial role. Illustrated in Figure 67, the robot may have received the maximum expected RSSI measurement at the position indicated along R_3 , so that the robot's distance to the tag is given by $d = d_3$. This position is the closest possible to the tag along R_3 , but it is not the closest achievable position (which occurs along R_4 where $d_4 = d_{best} = \min d_i$). In general, the distance error (defined as $d_{err} = |d - d_{best}|$) is bounded by

$$d_{err} \leq \max(d_i) - \min(d_i) \quad (73)$$

It is worth noting that, if the tag's radiation pattern is isotropic, then it does not play a role; the distance from the tag will be the only pertinent factor so that $d = d_{best} = \min d_i$. This suggests




Figure 67: Along each ray emanating from the tag, the robot can only get the RFID reader antenna so close. In general, the ray with the maximum RSSI reading is unknown, so all that we can say is that the distance error is bounded (less than or equal to) the maximum distance along any of the rays.

that (on average), as the tag antenna’s radiation pattern becomes increasingly non-isotropic, global RFID search’s distance error will degrade.

4.4.4 On-Robot Implementation

For our on-robot implementation, which we discuss in detail in Section 4.5, we used two (patch) reader antennas mounted on the PR2’s shoulders. During our global RFID search trials, we kept the antennas at fixed heights off the ground, and we only articulated the antennas through pan angles rather than both pan and tilt. Effectively, this means that we performed *planar* global RFID search over location and orientation $\langle X, Y, \Theta \rangle$ (rather than full 6-DoF pose) given by the following optimization,

$$X, Y, \Theta = \underset{x, y, \theta}{\operatorname{argmax}} \mathbb{E} (\text{RSSI} \mid X = x, Y = y, \Theta = \theta). \quad (74)$$

Despite being an approximation to the more general 6-DoF global RFID search, we will shown

(in the next section) that our implementation performs well in practice.

4.5 Hybrid Global-Local RFID Search

Moving to all 6-DoF positions and orientations during global RFID search is a time-intensive, and oft-intractable proposition. Instead, we developed a hybrid global-local RFID search algorithm for discovering, locating, and approaching UHF RFID tagged objects in a real home environment by combining sparse global RFID search with the local RFID behaviors we described previously. The general premise behind this technique is that sparse global RFID search can obtain a good (approximate) robot configuration (position and orientation) within the neighborhood of the tagged object, followed by several local optimization-based behaviors to *climb the gradient* within the local basin of attraction to improve the robot's configuration: first, bearing estimation to refine the angular estimate (point the robot toward the tag); followed by RFID servoing to approach the tagged object until impeded by an obstacle. The combination of these behaviors results in a search algorithm that is easy to implement, easy to generalize, does not require training data or sensor models, and results in final robot states that are both (1) near the desired tagged object, and (2) oriented towards it. Later, we will show that this method yields superior results compared to state-of-the-art Bayesian tag localization techniques, with reasonable sensor models.

4.5.1 Evaluation

We evaluated UHF RFID search in a real home environment (Georgia Tech's Aware Home). We defined *nine* locations in the room to place each of *nine* tagged objects for RFID search. We have indicated the nine locations in Figure 68); the nine tagged objects, along with an overhead (map view) of the Aware Home are shown in Figure 69.

We executed global RFID search *nine* times. For each execution, we placed one tagged object in each location (so that across all nine global search executions, each tagged object was placed in each location). The robot continuously pans its antennas and queries for all tags in the environment (makes RFID measurements of tag detection and RSSI) while following a search pattern with a particular search granularity. For the purposes of this work, the robot programatically generated a sparse search pattern consisting of zig-zag waypoints (shown in Figure 69) that span a human-specified rectangle encompassing the three rooms to search (kitchen, dining room, and living room); the robot uses SLAM to determine it's own pose in the home and a Dynamic Window Approach




Figure 68: *Top:* We evaluated UHF RFID search at Georgia Tech’s Aware Home, a realistic home environment. *Middle and Bottom:* We define nine different locations (numbered zero through eight) in the Aware Home where tagged objects will be placed for RFID search.





Figure 69: *Top:* The *nine* tagged objects used during the RFID search evaluation. *Bottom:* The *nine* different locations demarcated on an overhead map of the Aware Home.

(DWA) to move between waypoints and avoid obstacles [99]. We selected a search granularity of 1.5 meters; at this granularity, the global search (for all tags simultaneously) requires approximately 3 minutes to execute. For every positive tag detection, the robot records the RFID measurement, the pose of the detecting antenna, and the pose of the robot.

After performing sparse global search (for all tags), the robot used the recorded measurements to perform hybrid global-local search for each of the *nine* separate tagged objects (for a total of 81 search trials). For each trial, the robot used Equation 74 to determine the position and orientation with maximum expected RSSI (global RFID search). From this initial pose in the neighborhood of the tagged object, the robot pans its directional (patch) antennas and uses bearing estimation (Equation 49) to refine its heading estimate. Finally, the robot adjusts its two shoulder-mounted antennas to $\alpha = \pm 40^\circ$ and performs RFID servoing, moving forward at a constant rate of 10 cm/sec while adjusting its angular velocity in proportion to the expected difference in RSSI between left and right antennas. The robot halts when impeded by an obstacle (or the tagged object).

During global RFID search, the robot obtained positive tag detections for the desired tagged object in 69 / 81 trials. In 12 instances, the desired tagged object was undetected. The failures occurred for the TV remote (7 failures), the keys (4 failures), and the hair brush on the floor (1 failure). The failures for the TV remote and keys are likely attributable to the heavy presence of metallic components interfering with RF signals. Despite using on-metal tags, these objects were not detected; designing tag antennas for use on metal, near liquids, and affixed to people is an ongoing area of research [10, 9, 3]. Additionally, we chose a 1.5 meter search granularity for our experiments; re-searching with a finer granularity would likely result in positive tag detections. In the future, search at multiple granularities (or over time while performing other tasks) would probably prove fruitful. However, these 12 failure cases demonstrate an important limitation of RFID-based methods: if the tag is undetected, the methods are not applicable.

In the remaining 69 trials with at least one positive tag detection, the robot returned to the position and orientation with maximum expected RSSI, and executed the local optimization-based behaviors: bearing estimation and RFID servoing. Photographs of the final robot positions for several trials are shown later in Figures 72-80. In Figure 70, we show the final robot poses for all 69 trials.






Figure 70: The final robot positions (and orientations) after performing hybrid global-local RFID search for tagged objects in Georgia Tech’s Aware Home.

Because of the arrangement of furniture and other obstacles, the best distance achievable for each test location is highly variable and also depends on the robot’s shape and navigation capabilities. The set of achievable positions within the Aware Home using our navigation system is shown in light gray in Figure 70 (the light blue corresponds to keepout zones due to nearby obstacles). For example, in location #1 (the fireplace mantle) the robot is *at best* capable of getting within 0.54 meters (planar distance) of the tagged object using this navigation system; in location #5 (where the object is blocked by the kitchen table), the robot is *at best* capable of getting within 1.53 meters of the tagged object. We define a distance error measure that takes into account the locational variation; specifically, we define the distance error as the difference between the best possible distance achievable and the distance achieved after performing hybrid RFID search. A sample calculation is shown in Figure 71.

The overall performance of hybrid global-local RFID search using these error measures is shown in Table 22. Several illustrative examples of hybrid global-local RFID search are shown in Figures



Distance from Tag: $d_{\text{best}} = 0.83\text{m}$
 Angular Estimate: $\theta_{\text{best}} = 0.0^\circ$



Distance from Tag: $d = 1.14\text{m}$
 Angular Estimate: $\theta = +2.4^\circ$

Distance Error: $|d_{\text{best}} - d| = 0.31\text{m}$
 Angular Error: $|\theta| = 2.4^\circ$

Figure 71: Top Left: Owing to obstacles and the robot's size, the PR2 is incapable of navigating to all positions and orientations in the Aware Home. Red positions are valid positions (all orientations feasible at these positions); blue positions are unachievable. **Top Right:** The final position of the PR2 robot after performing hybrid RFID search is shown in blue. **Bottom:** The distance and angular errors are defined as the difference between the best achievable for this tagged object location and that achieved after performing RFID search.

72 through 80. Detailed statistics for all 81 trials (on a tag-by-tag and location-by-location basis) are reported in Tables 4 through 21.



Figure 72: The results for RFID search with tagged objects at location 0. *Top:* The locations of positive and negative reads (grey and red, respectively) during RFID search (blue robot trajectory), the “best” location given by global search (orange), and the final location after servoing (teal). *Bottom:* A photograph of the PR2 post-servoing and a 3D point cloud captured by the PR2’s head-mounted Kinect showing the robot’s perspective.





Figure 73: The results for RFID search with tagged objects at location 1. *Top:* The locations of positive and negative reads (grey and red, respectively) during RFID search (blue robot trajectory), the “best” location given by global search (orange), and the final location after servoing (teal). *Bottom:* A photograph of the PR2 post-servoing and a 3D point cloud captured by the PR2’s head-mounted Kinect showing the robot’s perspective.





Figure 74: The results for RFID search with tagged objects at location 2. *Top:* The locations of positive and negative reads (grey and red, respectively) during RFID search (blue robot trajectory), the “best” location given by global search (orange), and the final location after serving (teal). *Bottom:* A photograph of the PR2 post-servicing and a 3D point cloud captured by the PR2’s head-mounted Kinect showing the robot’s perspective.




Figure 75: The results for RFID search with tagged objects at location 3. *Top:* The locations of positive and negative reads (grey and red, respectively) during RFID search (blue robot trajectory), the “best” location given by global search (orange), and the final location after serving (teal). *Bottom:* A photograph of the PR2 post-servoing and a 3D point cloud captured by the PR2’s head-mounted Kinect showing the robot’s perspective.



Figure 76: The results for RFID search with tagged objects at location 4. *Top:* The locations of positive and negative reads (grey and red, respectively) during RFID search (blue robot trajectory), the “best” location given by global search (orange), and the final location after servoing (teal). *Bottom:* A photograph of the PR2 post-servoing and a 3D point cloud captured by the PR2’s head-mounted Kinect showing the robot’s perspective.





Figure 77: The results for RFID search with tagged objects at location 5. *Top:* The locations of positive and negative reads (grey and red, respectively) during RFID search (blue robot trajectory), the “best” location given by global search (orange), and the final location after servoing (teal). *Bottom:* A photograph of the PR2 post-servoing and a 3D point cloud captured by the PR2’s head-mounted Kinect showing the robot’s perspective.






Figure 78: The results for RFID search with tagged objects at location 6. *Top:* The locations of positive and negative reads (grey and red, respectively) during RFID search (blue robot trajectory), the “best” location given by global search (orange), and the final location after serving (teal). *Bottom:* A photograph of the PR2 post-servoing and a 3D point cloud captured by the PR2’s head-mounted Kinect showing the robot’s perspective.





Figure 79: The results for RFID search with tagged objects at location 7. *Top:* The locations of positive and negative reads (grey and red, respectively) during RFID search (blue robot trajectory), the “best” location given by global search (orange), and the final location after servoing (teal). *Bottom:* A photograph of the PR2 post-servoing and a 3D point cloud captured by the PR2’s head-mounted Kinect showing the robot’s perspective.




Figure 80: The results for RFID search with tagged objects at location 8. *Top:* The locations of positive and negative reads (grey and red, respectively) during RFID search (blue robot trajectory), the “best” location given by global search (orange), and the final location after servoing (teal). *Bottom:* A photograph of the PR2 post-servoing and a 3D point cloud captured by the PR2’s head-mounted Kinect showing the robot’s perspective.

Table 4: RFID search results (summary) for tagged objects at location #0.

	Global RFID Search <i>only</i> mean (std dev)	Hybrid Global-Local RFID Search mean (std dev)
Tag Detections per Object	23.4 (36.8)	33.4 (53.0)
Robot-Tag Distance	2.0 m (0.5)	1.9 m (0.3)
Distance Error	0.6 m (0.5)	0.4 m (0.3)
Angular Error	16.9° (2.2°)	19.5° (12.6°)

Table 5: Hybrid RFID search results (detailed) for tagged objects at location #0.

	Orange Meds	TV Remote	Red Bottle	Keys
Tag Detections	22 / 3144 (0.7%)	0 / 2454 (0.0%)	19 / 2621 (0.7%)	0 / 3155 (0.0%)
Robot-Tag Distance	1.795 m	--	1.733 m	--
Distance Error	0.337 m	--	0.275 m	--
Angular Error	-11.1°	--	-15.9°	--

	Med Bottle	Med Box	Teddy Bear	Cordless Phone	Hair Brush
...	3 / 2965 (0.1%)	76 / 2343 (3.2%)	169 / 3125 (5.4%)	12 / 3263 (0.4%)	0 / 3247 (0.0%)
	2.554 m	1.831 m	1.712 m	1.806 m	--
	1.096 m	0.373 m	0.254 m	0.348 m	--
	46.7°	-14.6°	9.4°	19.5°	--

Table 6: RFID search results (summary) for tagged objects at location #1.

	Global RFID Search <i>only</i> mean (std dev)	Hybrid Global-Local RFID Search mean (std dev)
Tag Detections per Object	97.9 (101.3)	171.0 (169.6)
Robot-Tag Distance	1.4 m (0.5)	0.9 m (0.1)
Distance Error	0.9 m (0.5)	0.4 m (0.1)
Angular Error	25.9° (12.1°)	15.8° (9.6°)

Table 7: Hybrid RFID search results (detailed) for tagged objects at location #1.

	Orange Meds	TV Remote	Red Bottle	Keys
Tag Detections	103 / 3626 (2.8%)	0 / 2978 (0.0%)	240 / 2668 (9.0%)	19 / 2755 (0.7%)
Robot-Tag Distance	1.085 m	--	0.818 m	0.846 m
Distance Error	0.541 m	--	0.274 m	0.302 m
Angular Error	11.8°	--	8.6°	6.7°

	Med Bottle	Med Box	Teddy Bear	Cordless Phone	Hair Brush
...	500 / 3535 (14.1%)	181 / 2844 (6.4%)	408 / 2399 (17.0%)	29 / 3272 (0.9%)	59 / 3290 (1.8%)
	1.088 m	0.916 m	1.017 m	0.855 m	0.918 m
	0.544 m	0.372 m	0.474 m	0.312 m	0.374 m
	-7.0°	-18.7°	20.1°	37.7°	15.6°

Table 8: RFID search results (summary) for tagged objects at location #2.

	Global RFID Search mean (std dev)	Hybrid Global-Local RFID Search mean (std dev)
Tag Detections per Object	118.3 (146.4)	233.4 (248.5)
Robot-Tag Distance	2.4 m (0.8)	1.6 m (0.6)
Distance Error	1.4 m (0.8)	0.6 m (0.6)
Angular Error	19.8° (10.3°)	32.4° (18.8°)

Table 9: Hybrid RFID search results (detailed) for tagged objects at location #2.

	Orange Meds	TV Remote	Red Bottle	Keys
Tag Detections	598 / 3506 (17.1%)	0 / 3247 (0.0%)	368 / 3348 (11.0%)	0 / 2454 (0.0%)
Robot-Tag Distance	1.180 m	--	1.408 m	--
Distance Error	0.171 m	--	0.399 m	--
Angular Error	15.6°	--	-2.4°	--

	Med Bottle	Med Box	Teddy Bear	Cordless Phone	Hair Brush
...	234 / 2680 (8.7%)	688 / 3524 (19.5%)	143 / 2979 (4.8%)	2 / 2399 (0.1%)	68 / 3615 (1.9%)
	1.568 m	1.875 m	2.842 m	1.123 m	1.176 m
	0.559 m	0.866 m	1.833 m	0.114 m	0.167 m
	-20.3°	55.7°	-44.7°	-34.4°	53.5°

Table 10: RFID search results (summary) for tagged objects at location #3.

	Global RFID Search <i>only</i> mean (std dev)	Hybrid Global-Local RFID Search mean (std dev)
Tag Detections per Object	161.2 (189.6)	247.1 (223.8)
Robot-Tag Distance	1.7 m (0.5)	1.2 m (0.1)
Distance Error	0.8 m (0.5)	0.3 m (0.1)
Angular Error	28.3° (16.5°)	20.0° (15.4°)

Table 11: Hybrid RFID search results (detailed) for tagged objects at location #3.

	Orange Meds	TV Remote	Red Bottle	Keys
Tag Detections	137 / 3170 (4.3%)	0 / 3119 (0.0%)	195 / 3450 (5.7%)	0 / 2978 (0.0%)
Robot-Tag Distance	1.209 m	--	1.137 m	--
Distance Error	0.298 m	--	0.226 m	--
Angular Error	27.1°	--	-2.8°	--

	Med Bottle	Med Box	Teddy Bear	Cordless Phone	Hair Brush
...	479 / 2925 (16.4%)	519 / 2692 (19.3%)	637 / 3318 (19.2%)	73 / 2997 (2.4%)	184 / 2498 (7.4%)
	1.169 m	1.129 m	1.098 m	1.345 m	1.468 m
	0.258 m	0.218 m	0.187 m	0.434 m	0.557 m
	1.0°	-41.8°	21.6°	-38.1°	-7.6°

Table 12: RFID search results (summary) for tagged objects at location #4.

	Global RFID Search <i>only</i> mean (std dev)	Hybrid Global-Local RFID Search mean (std dev)
Tag Detections per Object	142.7 (145.5)	294.0 (253.2)
Robot-Tag Distance	1.5 m (0.2)	0.8 m (0.1)
Distance Error	0.9 m (0.2)	0.2 m (0.1)
Angular Error	22.1° (14.9°)	34.5° (21.4°)

Table 13: Hybrid RFID search results (detailed) for tagged objects at location #4.

	Orange Meds	TV Remote	Red Bottle	Keys
Tag Detections	204 / 2367 (8.6%)	0 / 2984 (0.0%)	705 / 3443 (20.5%)	248 / 3667 (6.8%)
Robot-Tag Distance	0.886 m	--	0.789 m	0.776 m
Distance Error	0.294 m	--	0.197 m	0.184 m
Angular Error	-19.6°	--	-3.4°	58.1°

	Med Bottle	Med Box	Teddy Bear	Cordless Phone	Hair Brush
...	543 / 3336 (16.3%)	608 / 2882 (21.1%)	299 / 2940 (10.2%)	24 / 3421 (0.7%)	15 / 2963 (0.5%)
	0.667 m	0.900 m	0.875 m	0.826 m	0.989 m
	0.075 m	0.309 m	0.284 m	0.234 m	0.397 m
	-8.5°	33.3°	-38.0°	49.1°	66.2°

Table 14: RFID search results (summary) for tagged objects at location #5.

	Global RFID Search mean (std dev)	Hybrid Global-Local RFID Search mean (std dev)
Tag Detections per Object	55.8 (110.1)	110.3 (191.8)
Robot-Tag Distance	2.3 m (0.5)	1.9 m (0.2)
Distance Error	0.8 m (0.5)	0.4 m (0.2)
Angular Error	16.7° (23.9°)	25.1° (26.1°)

Table 15: Hybrid RFID search results (detailed) for tagged objects at location #5.

	Orange Meds	TV Remote	Red Bottle	Keys
Tag Detections	52 / 2822 (1.8%)	0 / 2193 (0.0%)	132 / 3233 (4.1%)	0 / 3119 (0.0%)
Robot-Tag Distance	1.752 m	--	1.801 m	--
Distance Error	0.221 m	--	0.270 m	--
Angular Error	-19.5°	--	-11.2°	--

	Med Bottle	Med Box	Teddy Bear	Cordless Phone	Hair Brush
...	95 / 3446 (2.8%)	638 / 3408 (18.7%)	5 / 2706 (0.2%)	6 / 2643 (0.2%)	65 / 3400 (1.9%)
	2.057 m	2.286 m	2.272 m	1.655 m	1.749 m
	0.526 m	0.755 m	0.741 m	0.124 m	0.218 m
	-20.4°	-9.8°	-87.9°	-8.5°	-18.1°

Table 16: RFID search results (summary) for tagged objects at location #6.

	Global RFID Search mean (std dev)	Hybrid Global-Local RFID Search mean (std dev)
Tag Detections per Object	149.0 (110.3)	234.0 (159.0)
Robot-Tag Distance	0.9 m (0.1)	0.8 m (0.0)
Distance Error	0.2 m (0.1)	0.1 m (0.0)
Angular Error	27.8° (15.8°)	23.6° (15.1°)

Table 17: Hybrid RFID search results (detailed) for tagged objects at location #6.

	Orange Meds	TV Remote	Red Bottle	Keys
Tag Detections	137 / 3305 (4.1%)	40 / 2793 (1.4%)	325 / 2370 (13.7%)	107 / 3188 (3.4%)
Robot-Tag Distance	0.811 m	0.832 m	0.784 m	0.826 m
Distance Error	0.096 m	0.117 m	0.069 m	0.111 m
Angular Error	-34.0°	8.7°	26.7°	-8.1°

	Med Bottle	Med Box	Teddy Bear	Cordless Phone	Hair Brush
...	430 / 3351 (12.8%)	446 / 3556 (12.5%)	416 / 3274 (12.7%)	53 / 2595 (2.0%)	152 / 2688 (5.7%)
	0.764 m	0.839 m	0.806 m	0.815 m	0.797 m
	0.048 m	0.123 m	0.091 m	0.099 m	0.081 m
	3.1°	-32.6°	-15.2°	-31.1°	-52.9°

Table 18: RFID search results (summary) for tagged objects at location #7.

	Global RFID Search <i>only</i> mean (std dev)	Hybrid Global-Local RFID Search mean (std dev)
Tag Detections per Object	100.4 (122.7)	277.2 (267.8)
Robot-Tag Distance	2.4 m (0.3)	1.5 m (0.5)
Distance Error	1.6 m (0.3)	0.7 m (0.5)
Angular Error	19.7° (16.5°)	16.7° (20.7°)

Table 19: Hybrid RFID search results (detailed) for tagged objects at location #7.

	Orange Meds	TV Remote	Red Bottle	Keys
Tag Detections	127 / 2858 (4.4%)	0 / 3155 (0.0%)	405 / 3145 (12.9%)	64 / 2513 (2.5%)
Robot-Tag Distance	1.737 m	--	1.444 m	1.389 m
Distance Error	0.905 m	--	0.613 m	0.558 m
Angular Error	4.7°	--	10.1°	3.3°

	Med Bottle	Med Box	Teddy Bear	Cordless Phone	Hair Brush
...	452 / 3366 (13.4%)	717 / 3580 (20.0%)	657 / 3707 (17.7%)	32 / 3286 (1.0%)	41 / 2601 (1.6%)
	1.145 m	1.137 m	1.107 m	1.605 m	2.570 m
	0.313 m	0.305 m	0.276 m	0.774 m	1.739 m
	2.7°	2.4°	6.2°	57.6°	46.5°

Table 20: RFID search results (summary) for tagged objects at location #8.

	Global RFID Search <i>only</i> mean (std dev)	Hybrid Global-Local RFID Search mean (std dev)
Tag Detections per Object	140.0 (124.8)	188.1 (146.7)
Robot-Tag Distance	1.0 m (0.0)	0.9 m (0.0)
Distance Error	0.2 m (0.0)	0.2 m (0.0)
Angular Error	23.4° (21.8°)	21.2° (17.1°)

Table 21: Hybrid RFID search results (detailed) for tagged objects at location #8.

	Orange Meds	TV Remote	Red Bottle	Keys
Tag Detections	125 / 2603 (4.8%)	78 / 2649 (2.9%)	133 / 3301 (4.0%)	75 / 2822 (2.7%)
Robot-Tag Distance	0.963 m	0.966 m	0.910 m	0.937 m
Distance Error	0.167 m	0.170 m	0.114 m	0.140 m
Angular Error	23.7°	-1.9°	5.0°	47.3°

	Med Bottle	Med Box	Teddy Bear	Cordless Phone	Hair Brush
...	260 / 2354 (11.0%)	454 / 3154 (14.4%)	426 / 3327 (12.8%)	102 / 3409 (3.0%)	40 / 3125 (1.3%)
	0.955 m	0.952 m	0.982 m	0.932 m	0.948 m
	0.159 m	0.156 m	0.186 m	0.136 m	0.152 m
	-17.4°	-13.4°	-22.4°	-6.6°	53.3°

4.5.2 Comparison to Bayesian Estimation

Using the *same* data captured during our RFID search experiments, we used the 2-DoF, probabilistic tag detection and RSSI model from Figure 38 to compute the posterior probability distribution over 2-DoF tag pose (in the map frame) using data acquired during our RFID search experiments. We opted to represent this distribution as a probability map rather than a resampled particle filter to ensure that model over-confidence did not cause particle deprivation [143]. Using brute force, we determined the best *location* and *orientation* in the room for the robot to view the desired tagged object given this posterior distribution (the one that minimized the expected error).

Unfortunately, we know from the Friis radar equation that the relationship between tag and reader antenna is 6-DoF (not just 2-DoF); furthermore, from the RFID literature, we know that tagging objects can dramatically affect tag detection and RSSI. All of these factors introduce variation, and may be exacerbated as objects are relocated to new environments with (potentially changing) RF properties. In Figure 41, we showed several different tag configurations that have dramatically different sensor models compared to the controlled conditions and compared to one another.

In Figure 81, we show the results of Bayesian tag localization using a probability map (a distribution of tag likelihood with the state discretized at 5cm resolution) using the sensor model from Figure 38. The cells in the distribution represent the top 95% of the probability mass; we can see that the true tagged object locations (green) are actually located in areas with relatively low probability. Additionally, we observe that much of the probability mass is further away from the robot than the actual tag (since the relatively-ideal model predicts a much greater read range compared to the tagged object models).

4.5.2.1 Determining a “Good” Robot Pose

As we mentioned earlier, our goal for RFID search is to determine a final robot state that is both (1) near the tagged object and (2) oriented toward it. Thus, we pose the following two questions: (1) Given this distribution over tag pose, where should we position the robot to minimize the expected distance error? (2) At that position, in which direction should the robot face to minimize the expected angular error given the distribution over tag pose. To answer these questions, we performed a brute-force search over all possible (discretized) robot positions (Figure 81), and selected the one




Figure 81: Posterior probability distribution over 2-DoF tag pose using the data from global RFID search measurements only (left) and global RFID search plus local RFID search measurements (right) for a tagged object on the kitchen countertop (top), on the floor (middle), and on the couch (bottom). The robot position and orientation that yields the minimum distance and angular error (given the distribution) is shown as an arrow (orange or blue) in each image.

that minimized the expected distance error. Once we have this position, we performed another brute-force search over all possible robot orientations and chose the orientation that minimized the expected angular error. Collectively, these two quantities yielded position and orientation estimates that are analogous to those yielded by hybrid RFID search. In Table 22, we show the overall statistics for RFID search using Bayesian tag localization. In terms of position, Bayesian localization search

has comparable (not statistically significant) performance compared to our (optimization-based) hybrid global-local RFID search. However, our approach achieves superior (statistically-significant) performance in terms of orienting the robot toward the tag.

Table 22: RFID Search Using Hybrid Optimization-Based Approach Versus Bayesian Localization

	Distance Error mean (std dev)	Angular Error mean (std dev)
Hybrid Global-Local RFID Search	0.36 m (0.33 m)	23.2° (19.0°)
Bayesian Localization	0.31 m (0.42 m)	39.8° (37.6°)
p-value	0.4900	0.0014
Significant for $\alpha < 0.05$	No	Yes

Since our hybrid global-local optimization-based approach to RFID search is easy to implement, easy to generalize, does not require training data, and yields superior results, we believe that it is the superior approach to discover, locate, and approach tagged *objects* in human environments.

4.6 Conclusions

In this chapter we presented a new, general approach for perceiving UHF RFID tags affixed to locations, people, and objects of interest that is easy to implement, easy to generalize, and does not require training data or sensor models. We formulated our approach as an optimization problem, where the robot actuates its directional (patch) reader antennas and/or uses mobility to provide opportunistic views of the RF landscape and estimate a desired quantity of interest (eg. bearing toward the tag) and seeks to maximizing the expected RSSI measurements associated with a desired tag ID. We unify our approaches using the following optimization,

$$\hat{x} = \operatorname{argmax}_x E (\text{RSSI} \mid \mathbf{X} = x).$$

We developed a series of *local* optimization-based behaviors that allow the robot to estimate the direction toward a tag (bearing estimation) or to approach a tag (RFID servoing) given some initial position in the (local) neighborhood of the tag:

- **Bearing Estimation – Azimuth (Θ) Only:**

$$\hat{\Theta} = \underset{\theta}{\operatorname{argmax}} E (RSSI | \theta = \theta)$$

- **Bearing Estimation – Azimuth (Θ) and Elevation (Φ):**

$$\hat{\Theta}, \hat{\Phi} = \underset{\theta, \phi}{\operatorname{argmax}} E (RSSI | \theta = \theta, \phi = \phi)$$

- **RFID Servoing – Angular Velocity ($\dot{\Theta}$):**

$$\dot{\Theta} = \kappa \cdot E (\Delta RSSI | \theta = \theta)$$

From an optimization perspective, these techniques perform local search (eg. gradient ascent [19]) to refine the robot’s state (position and/or orientation) within a local basin of attraction, thereby yielding (progressively) greater expected RSSI measurements.

We also developed a global RFID search behavior, where the robot determines the best position and orientation (X, Y, Θ) to observe the desired tag by moving throughout the entire environment and selecting the pose that yields the maximum expected RSSI. This technique is akin to sampling-based global search techniques in the optimization and planning literature [91, 149]; in the preceding framework, global RFID search is expressed as:

- **Global RFID Search:**

$$\hat{X}, \hat{Y}, \hat{\Theta} = \underset{x, y, \theta}{\operatorname{argmax}} E (RSSI | X = x, Y = y, \theta = \theta)$$

We developed a hybrid global-local RFID search algorithm that first performs a sparse global RFID search to determine an acceptable initial robot state (position and orientation), followed by a series of local RFID search behaviors (bearing estimation and RFID servoing) to find a “good” final robot state that is both (1) near the desired tagged object, and (2) oriented towards it. We evaluated hybrid RFID search using nine tagged objects in nine locations in a real home environment for a total of 81 trials.

We compared the final robot pose obtained using our optimization-based techniques to those obtained when using state-of-the-art Bayesian tag localization techniques (Chapter 3). We showed

that our hybrid global-local RFID search technique achieved comparable performance in positioning the robot, but superior performance (statistically-significant) in orienting the robot toward the tag. Since our optimization-based approach is easy to implement, easy to generalize, does not require training data, and yields superior results, we believe that it is the superior approach to discover, locate, and approach tagged *objects* in human environments.

CHAPTER V

SHORT-RANGE UHF RFID

The robotics community uses RFID predominantly for navigation and localization; comparatively little work uses RFID for manipulation. In this chapter, we present two short-range UHF RFID systems mounted-near or embedded-in a mobile robot’s manipulator to aid in manipulation: to locate tagged objects from short distances (less than 1 meter), and to provide high-fidelity object identification of tagged objects that the robot has already grasped. The antennas in these two systems operate in the near-field RF regime (magnetostatic coupling), providing short-range capabilities analogous to low-frequency (LF) or high-frequency (HF) RFID. Using UHF RFID in the near-field is an atypical mode of operation, which allows us to interact with the *same* UHF RFID tags at both short range (via these new antennas) and at long-range via far-field electromagnetic coupling – as with the patch antennas for long-range RFID localization (Chapter 3) and long-range RFID behaviors (Chapter 4). The design and evaluation of these antennas and robotic behaviors are the primary contribution of this chapter. Later in Chapter 6, we examine mobile manipulation applications that leverage UHF RFID sensing in these two regimes (long-range and short-range). We believe the future combination of these two operating regimes will be a powerful tool for robots operating in unstructured human environments, particularly when the consequences of mis-identification are dire (eg. when dealing with medication).

In this chapter, we describe the design of *two* short-range UHF RFID sensor systems mounted on the autonomous mobile manipulator, EL-E. Both are shown in Figure 82. The first system is mounted to EL-E’s wrist; we show that this system can be used during manipulation to selectively grasp the correct tagged object even in the presence of other (visually-identical) objects. The other system is embedded in EL-E’s fingers; we refer to this as the *in-hand* RFID reader. We evaluate the in-hand reader’s ability to detect a variety of tags applied to a diverse set of objects being held in the robot’s manipulator. The results demonstrate that the system is capable of detecting a wide variety of tags, but that care must be taken to select a tag appropriately matched to an object’s material




Figure 82: *Left:* A loaded microstrip antenna (green circuit board) is mounted to EL-E’s wrist. We call this the *wrist-mounted* antenna. *Right:* Two ceramic microstrip antennas (small white rectangles) are embedded inside each of EL-E’s fingers, comprising the *in-hand* RFID reader system. Both of these short-range UHF RFID systems was capable of detecting any of the 32 UHF RFID tag variants shown in Figure 6. This was confirmed by performing read attempts with each tag variant within 1cm of the antennas (under free-space conditions).

composition.

Using these two systems, we develop several short-range robot behaviors. Again, these behaviors leverage a robot’s ability to reposition itself for additional, advantageous views of the tag and its environment. For short-range UHF RFID, we observe that closer tag-antenna distances produce greater RSSI readings. We use this property to develop the several short-range UHF RFID capabilities:

- **Locating a Tagged Object:** Moving the wrist-mounted antenna in front of tagged objects, we can determine which of the objects likely possesses the desired tag.
- **Short-Range RFID Servoing:** Using the differential signal between in-hand antennas on opposing fingers, we can servo the robot’s manipulator to try and maintain the object’s position between the robot’s fingers.
- **Raster-Scan RSSI Images:** Scanning the in-hand reader in the xy-plane above a tagged object, which we dub “raster scanning,” we can build RSSI Images’ that show the RSSI signal strength as a function of xy-position. These raster-scan RSSI images provide insights into the otherwise invisible RFID signals.

All three short-range behaviors above leverage the *same* optimization-based framework developed in Chapter 4 for long-range operation. Using the in-hand reader, we develop one additional robot behavior:

- **Verifying a Tagged Object’s Identity:** We can monitor RSSI readings as a grasped object is lifted to determine (1) if the grasped object is tagged, and (2) the tagged object’s identity. We show that this behavior works even in the presence of nearby (distracting) tagged objects.

5.1 *Related Work*

Short-range low frequency (LF) and high frequency (HF) RFID are used extensively for access control [48], relying on a fundamental premise: tags are a unique token that is *only* detected when near the reader. This has been employed extensively in human-computer interaction (HCI) for situated story telling [115], children’s games [84], and tangible interfaces [5]. In all of these examples, a tag’s presence provides strong object recognition. This same principle has been also employed in robotics, where cabinets and drawers are aware of their tagged contents [18]. However, these systems require many readers distributed throughout the environment. In the system(s) we develop, the robot carries the short-range sensor on-board, which reduces infrastructure costs and makes our systems more portable. Furthermore, since the short-range RFID antennas operate on the *same* UHF RFID tags as the long-range antennas, we further reduce the infrastructure requirements (one tag serves multiple purposes) while benefiting from both the long-range and short-range capabilities.

There are relatively few examples of robotic systems using RFID for manipulation purposes. Short-range high-frequency (HF) RFID systems have been used to locate books on shelves for subsequent manipulation [163], identify Lego blocks for autonomous construction [156], and to identifying surgical tools [50]; however, there are very few additional examples. Meanwhile, there are numerous examples of RFID readers being employed for human manipulation. For example, researchers have used wrist-mounted RFID readers to determine which object is being held in a person’s hand, which provides contextual awareness for activity recognition systems [116, 49, 137]. Later, we demonstrate analagous capabilities using our in-hand, short-range UHF RFID system. To the best of our knowledge, the systems described in this chapter are the first instance of *UHF RFID* used for manipulation, human or robotic.

In the RFID literature, there are many examples of near-field UHF antennas [111, 119]; however, most are too bulky to be contained on or near the robot’s manipulator. We design and evaluate several new near-field UHF RFID antenna systems mounted-near or embedded-in a mobile robot’s manipulator, and we use these antennas to develop several robot behaviors – these are the principle contributions of this chapter. Later, in Chapter 6, we will show that this capability is useful for applications where knowing the identity of a grasped object with high confidence is important (eg. medication delivery).

5.2 Design of a Wrist-Mounted Reader




Figure 83: A loaded-microstrip antenna (*left*) is mounted to EL-E’s wrist (*middle*). EL-E moves the antenna in front of a tagged medication bottle (*right*) while querying for the specific tag ID and monitoring RF signals.

The wrist-mounted RFID system (shown in Figure 83) consisted of two main components: a ThingMagic Mercury 4e (M4e) RFID reader *and* a near-field, magnetostatically coupled antenna. By means of a flexible coaxial cable running along the robot’s arm, the RFID reader unit was connected to the near-field antenna located at the manipulator’s wrist joint. This near-field antenna was designed by our collaborator (Dr. Matt Reynolds) and was distributed with the M4e reader system; it is a printed circuit loaded-microstrip antenna designed to emphasize the production of closeby RF magnetic fields. Specifically, the wrist-mounted antenna is a specially-designed printed circuit board made of FR4 laminate that contains a terminated microstrip section $1/4$ -wavelength (8cm) in

length. When the passive tag is in proximity to the microstrip section, the fringing magnetic fields from the microstrip section couple to the small loop sections a UHF RFID tag's antennas. The resulting near-field coupling exhibits a far-field gain of less than -13dBi , so it radiates approximately 20dB less (1%) of the energy radiated by the far-field antennas. This results in a correspondingly shorter detection range for the near-field antenna. In practice, the detection range was less than 12cm when using a Alien ALN-9640 Gen2 Squiggle Tag in free space. A far-field radiation pattern measurement is not useful for characterizing the near-field performance of the microstrip antenna, so we performed a near-field measurement by systematically moving the tag across the microstrip antenna, scaling the transmit power, and recording the minimum transmit power required to read the tag – recall that this is a proxy for RSSI, as was explained in Chapter 4. A lower minimum read power corresponds to a larger RSSI on modern RFID readers (eg. the Mercury M5e).

Figure 84 shows measured near-field coupling, as determined by the minimum power required to read the tag, at each 10mm increment of translation from left to right across the near-field antenna, at varying tag-to-antenna separations. The near-field antenna exhibits a spatially selective coupling. Though tag-antenna orientation certainly plays a role, we note that the minimum read power (maximum RSSI) corresponds to the location nearest the tag. This suggests that argmax techniques (like those from Chapter 4) may be equally-valuable for short-range operation; we will subsequently explore this insight.

By moving the manipulator left-to-right across the sample region, we can determine the position of maximal coupling between the tag and the near-field antenna. Using arguments analogous to RFID search (Section 4.4), the configuration with minimum read power (maximum RSSI) is most likely to be the configuration where the wrist-mounted antenna is closest to the tagged object. As seen in Figure 84, this position of alignment corresponds to a spatial error of $\pm 1.0\text{cm}$ at a tag-to-coupler separation of 2cm. We expect that improved positioning accuracy can be achieved by further optimizing the near-field antenna geometry; for example, custom antenna geometries that account for nearby RF interference (eg. the robot arm) might yield more directive near-field coupling.

In subsequent sections, we will explore robot behaviors that leverage this property of spatial RSSI measurements to infer the location of tagged objects. The loaded microstrip coupler suffers from several major drawbacks:

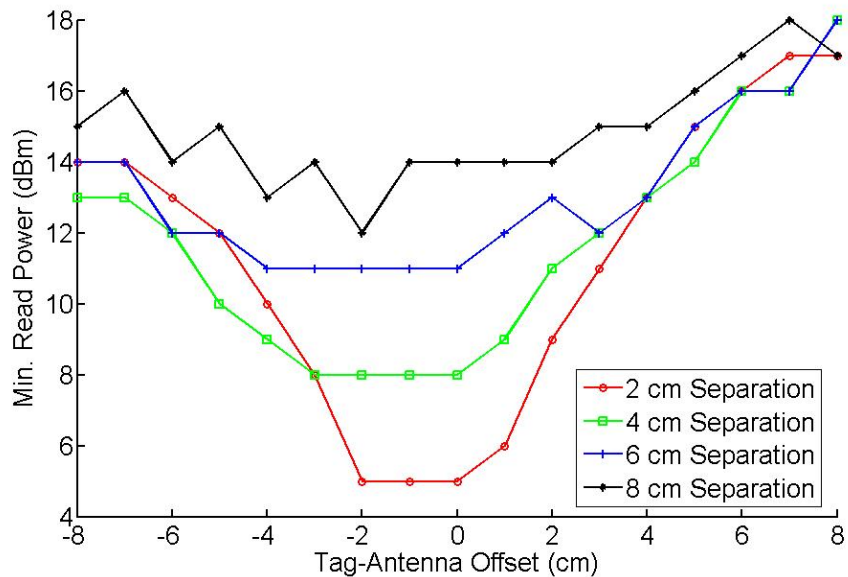


Figure 84: As EL-E moves the short-range wrist-mounted antenna in front of a tagged medication bottle (as in Figure 83) the minimum required reader transmit power to detect the tag (a proxy for RSSI) is minimized when the short-range antenna is closest to the tagged object. This is analogous to a maximum in RSSI.

- The loaded microstrip antenna’s radiation pattern is strongest “in front” of the antenna. It is incapable of detecting tagged objects in the robot’s manipulator.
- The physical size of a quarter-wavelength microstrip is directly related to the substrate’s dielectric constant [14]. For the choice of FR4 printed circuit board, the physical size of the antenna cannot be reduced without significantly hampering the antenna’s detection range. While this could be addressed through alternate substrates, it was beyond the scope of our research.
- The mounting location of the antenna on EL-E’s wrist causes kinematic constraints that (1) reduce the manipulator’s effective workspace; and (2) make it difficult to scan antenna objects laying flat on the floor or on tables (ie. it is difficult to scan the antenna parallel to the ground).
- Since the antenna operates on magnetic coupling, nearby metallic structures will adversely affect coupling efficiency. This means that the short-range antenna radiation pattern is highly affected by nearby metallic structures. These changes are difficult to model even when the environment is static and known; however, the environment will likely be unknown (*a priori*),

and it will change as the robot moves around and interacts with objects.

Because of these drawbacks, we sought an alternative design that could be (1) small enough to be directly integrated into EL-E's fingers, and (2) be used to detect grasped objects. Ultimately, this precipitated the design of the more compact *in-hand* RFID reader.

5.3 *Design of an In-Hand Reader*

Initial efforts to adapt the loaded microstrip design (from the previous section) for in-hand operation proved unfruitful. The resulting antennas did not yield sufficient read range and/or lacked a form factor that was amenable for embedding in the robot's hand. Instead, we employed *ceramic* microstrip antennas from Johanson Technology (part number 0920AT50A080E) to form the basis of an *in-hand* reader. These antennas provide superior performance and advantageous form-factor (measuring just 5×11 mm). The in-hand RFID system consists of two main electronic components, mounted at the base of the robot's arm: a Thing Magic Mercury5e (M5e) RFID reader, and a custom circuit board with communications (USB) and a four-port RF switch. Each of the four RF switch outputs is connected to a ceramic microstrip antenna from Johanson Technology via a flexible coaxial cable extending up the robot arm. The RFID reader can cycle through all four of the ceramic microstrips, providing RFID sensor measurements (both tag detection and RSSI, owing to the newer RFID reader module) at a rate of 3Hz. A photograph of the entire system is shown in Figure 85; a close-up of the in-hand system components is shown in Figure 86.

We placed two microstrip antennas on each finger. Like all microstrip antennas, the ceramic microstrips possess strong polarization bias (strong dependence on tag-antenna orientation). To help address this issue, a pair of antennas are placed at 90° offsets to provide comprehensive RF coverage from each finger. Like the loaded microstrip, the small ceramic microstrips are highly sensitive to nearby metal structures. We replaced the existing robot fingers (predominantly metal) with custom fingers that were 3D-printed with ABS plastic. We tuned the antennas' return loss over the UHF spectrum for each antenna using a π -type inductor-capacitor (LC) matching network. The matching network for each antenna is situated just in front of the antenna's RF feed and shares the printed circuit board on which the ceramic microstrip antennas are mounted. The matching network achieves better than ≈ -10 dB return loss over the UHF spectrum, as shown in Figure 86.




Figure 85: A later version of EL-E possesses two long-range body-mounted UHF RFID antennas for long-range (6+ meter) operation and four short-range ($\approx 40\text{cm}$) in-hand UHF RFID antennas in the manipulator.

These antennas provide superior performance compared to the loaded-microstrip from the previous section. We observed consistent tag detection distances in excess of 40cm under ideal conditions (a tag mounted on cardboard) and consistent detection distances of 15cm under practical conditions (a tag spiral-wrapped on a medication bottle).

5.4 Characterizing In-Hand Reader Detection Performance

A key contribution of this work is characterizing the ability of the in-hand reader system to positively detect a tag affixed to an object being held in the robot's manipulator. To test the system, we focused on ten objects: TV remote, cellular phone, fork, medication bottle, metal soda can, book, cordless phone, wallet, medication box, and water bottle. These objects were chosen for a number of reasons. First, many of these items appear on a list of objects for robotic retrieval prioritized by people with motor impairments, so they are immediately applicable to healthcare robots [29].

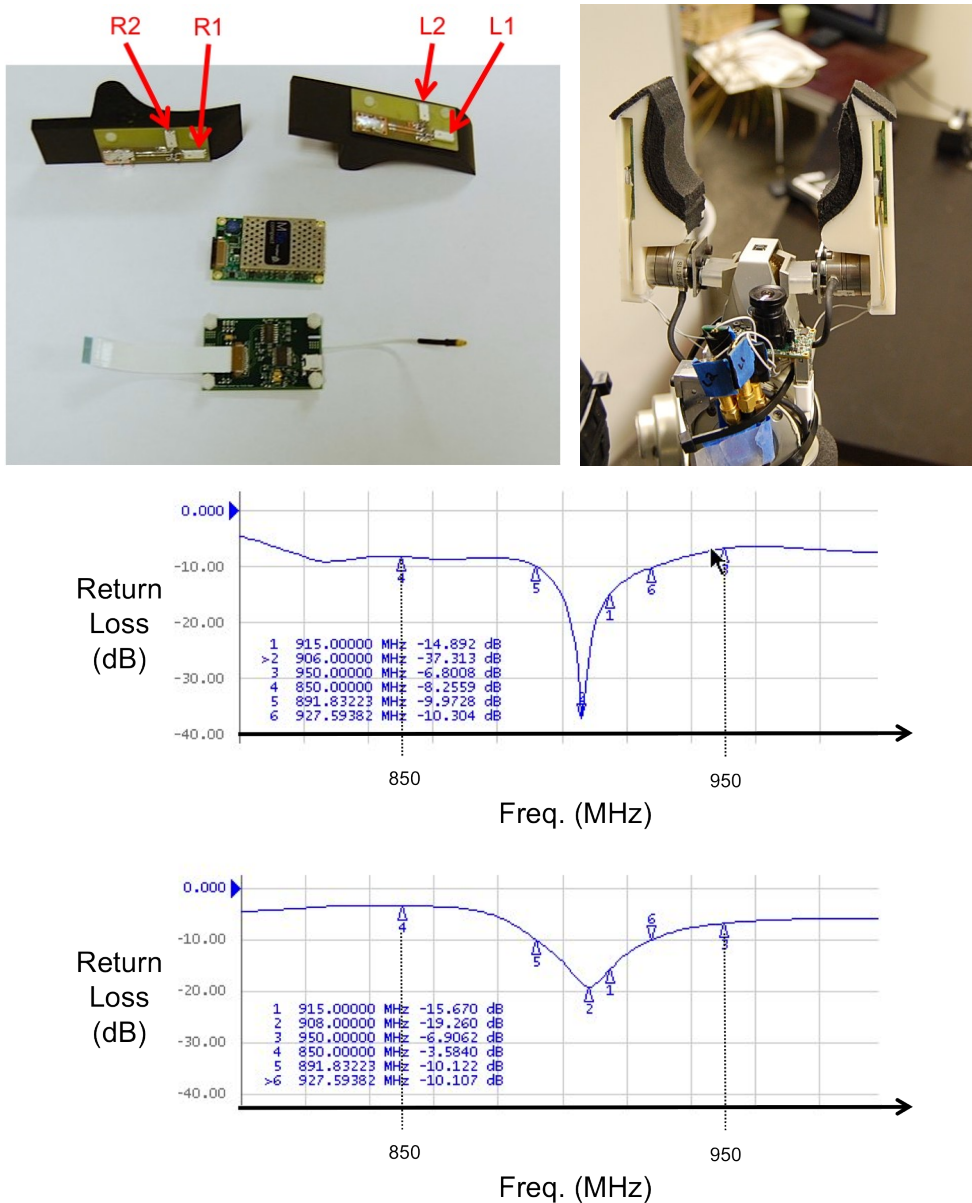


Figure 86: *Top Left:* Photos of the system components: two antenna PCB's each with a pair of short-range ceramic microstrip antennas (labeled L1, L2, R1, and R2), a ThingMagic M5e UHF RFID reader module, and a custom carrier board with integrated RF switch. *Top Right:* A later version of the fingers with the antennas and printed circuit board embedded inside EL-E's fingers. *Bottom:* Network analyzer plots show antenna return loss of less than -10dB for a pair of antennas (L1 on top, L2 on bottom). The measurements were taken with the antennas embedded inside EL-E's fingers; the tag and finger were located in free space.

Second, the medication-centric items cannot be handled by service animals such as service dogs or helper monkeys. This makes the objects appealing to service robot applications. Finally, some of these objects have material properties that have a high likelihood of interfering with RF signals (ie. metal). These objects will require careful consideration of tag selection, and illustrate challenges that may be encountered during real-world deployments.

We performed a total of 140 trials, wherein each object was affixed with a UHF RFID tag, placed in the robot’s manipulator, and then queried *once* by each of the four short-range antennas (a process taking less than 300ms). This was repeated for two different UHF tag variants (Alien Technologies Gen2 Squiggle tag and MetalTag Metal-Mount Flex tag) and seven canonical orientations within the manipulator (*ie.* as though a successful grasp had already been performed). A trial was considered successful if any one of the four antennas correctly received a positive tag detection. In all trials, the robot’s manipulator was positioned in relative isolation, at least 50cm from nearby objects (*ie.* in free space).

The results of these experiments are summarized in Table 23. The detailed results appear in Tables 24 and 25. In the detailed results, there are four marks for each trial, which correspond to each of the four RFID antennas. In order, the marks correspond to antennas L_1 , L_2 , R_1 , and R_2 (referring to Figure 86). For each antenna, a “√” represents a successful tag detection, whereas a “×” represents a failed detection attempt.

The in-hand reader positively detected tagged objects in 124 out of 140 trials, for an overall success rate of 88.5%. Trials involving the On-Metal tag were successful in 70 out of 70 trials (100%). However, trials involving the Squiggle tag were only successful for 54 out of 70 trials (77%). The Squiggle tag performed well on objects with predominantly RF-transparent materials such as plastic and cardboard (49 out of 49 trials, 100%), but was virtually undetected on metal (1 out of 14 trials, 7%). This is a known issue for generic UHF RFID tags (such as the “Squiggle” tag used) [10]. The water bottle is something of a special case, being detected 4 out of 7 trials (57%); in this case, a majority of the RF energy is absorbed by the water [9], but the tag was still occasionally detected. Tags designed to work on or near metal (eg. the on-metal tag) are specifically meant to address this limitation of normal dipole-like tags (eg. the Squiggle tag); however, such functionality comes at a price. The metal-mount tag is more than an order of magnitude more expensive per tag

Table 23: Tagged Object Detection for Grasped Objects – Overall Results

	Alien Technologies Squiggle Tag	MetalTag Metal-Mount Flex Tag
TV Remote	7/7 = 100%	7/7 = 100%
Cellular Phone	7/7 = 100%	7/7 = 100%
Fork	1/7 = 14%	7/7 = 100%
Medication Bottle	7/7 = 100%	7/7 = 100%
Metal Soda Can	0/7 = 0%	7/7 = 100%
Book	7/7 = 100%	7/7 = 100%
Cordless Phone	7/7 = 100%	7/7 = 100%
Wallet	7/7 = 100%	7/7 = 100%
Medication Box	7/7 = 100%	7/7 = 100%
Full Water Bottle	4/7 = 57%	7/7 = 100%

(\$2.50 each) compared to the Squiggle tag (sub-\$0.20).

It is also worth noting that all four antennas did not always obtain consensus. Occasionally multipath, material properties, polarization, or a number of other considerations may prevent some antennas from receiving a positive detection where the others prevail. In these cases, antenna diversity serves an important role; a single positive detection on any one of the antennas is sufficient to infer the object’s presence nearby the short-range antenna. We will use this insight to develop various short-range robot behaviors.

Table 24: Tagged Object Detection for Grasped Objects (Under Various Orientations) Using In-Hand RFID Reader (1 of 2)








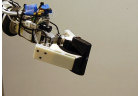



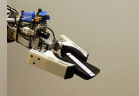


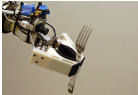

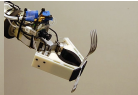
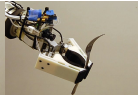
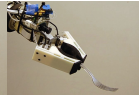
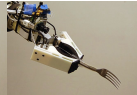
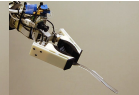














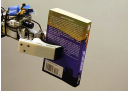











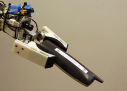
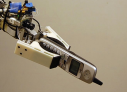
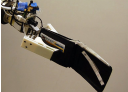
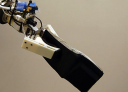

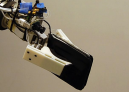

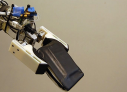




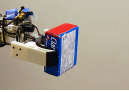










TV Remote:							
Squiggle:	✓✓✓✓	✓✓✓✓	✓✓✓✓	×✓✓✓	✓✓✓✓	✓✓✓×	✓✓✓✓
On-Metal:	✓✓✓✓	✓✓✓✓	✓✓×✓	✓✓✓×	✓✓✓✓	✓✓✓×	✓✓✓✓
Cell Phone:							
Squiggle:	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓×	✓✓✓✓	✓✓✓✓
On-Metal:	✓×××	✓✓×✓	✓×××	✓×✓✓	✓×✓×	✓✓×✓	✓✓✓✓
Fork:							
Squiggle:	××××	××××	××××	×××✓	××××	××××	××××
On-Metal:	✓×✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓
Med. Bottle:							
Squiggle:	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓
On-Metal:	×✓✓✓	×✓✓✓	✓✓✓×	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓
Metal Can:							
Squiggle:	××××	××××	××××	××××	××××	××××	××××
On-Metal:	×××✓	××✓×	×✓××	✓×××	✓×××	✓×××	✓✓✓×

Table 25: Tagged Object Detection for Grasped Objects (Under Various Orientations) Using In-Hand RFID Reader (2 of 2)

Book:							
Squiggle:	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓
On-Metal:	✓✓✓✓	✓✓✓✓	✓✓✓×	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓
Cordless Phone:							
Squiggle:	✓✓✓✓	✓✓✓✓	✓✓✓✓	×✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓
On-Metal:	✓✓✓✓	✓✓✓✓	×✓✓✓	×××✓	✓✓✓✓	✓×✓✓	✓✓✓✓
Wallet:							
Squiggle:	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓×	✓✓✓✓	✓✓✓✓	✓✓✓×
On-Metal:	✓✓✓✓	✓✓✓×	✓×✓×	✓✓✓✓	✓✓×✓	✓×××	✓✓✓×
Med. Box:							
Squiggle:	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓
On-Metal:	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓	✓✓✓✓
Water Bottle:							
Squiggle:	×××✓	××××	××✓×	××××	××××	××××	✓✓×✓
On-Metal:	✓✓✓✓	✓×✓✓	×××✓	✓✓×✓	✓✓×✓	✓✓✓✓	×××✓

5.5 *Short-Range UHF RFID Robot Behaviors*

In this section, we will look at a number of robot behaviors that use the wrist-mounted and in-hand UHF RFID systems to provide functionality in the intermediate regime between long-range navigation and actually making contact to manipulate (grasp) objects. Sometimes this regime is referred to as “pre-touch” or “pre-grasp” and can include manipulator or finger alignment [159], pre-contact shaping of the hand [102], and / or object recognition and pose estimation [128].

In particular, we are going to leverage our observation that short-range UHF RFID sensing provides maximum RSSI readings when the antenna-tag distances are minimized. Moving the robot’s end effector (and attached antennas) to various positions above or in front of the RFID tagged object while capturing RFID sensor measurements, we can estimate the location of the tagged object. By designing particular motions (eg. 1-dimensional motions along a line or 2-dimensional scans about a plane), we will show that the same argmax methods from Chapter 4 are equally-applicable for short-range UHF RFID. Further, early experiments suggest that the difference in signal between antennas on opposing figures are amenable to RFID servoing techniques.

5.5.1 **Locating a Tagged Object**

One foundational capability for robot manipulation is to determine the identity and location of a *desired* object located in front of the robot, perhaps with several other objects nearby (ie. clutter). Traditionally, computer vision and 3D point cloud perception are commonly used to perform this task [128]. While computer vision and 3D point cloud perception can provide precise pose estimation, object recognition algorithms do not always provide reliable identification, especially in the presence of virtually identical objects. In Chapter 1 we discussed how visual fiducials could serve this role, but short-range UHF RFID sensing can be used too. Specifically, moving the robot’s manipulator while monitoring the RFID sensor signals, we can determine a tagged object’s identification and approximate location without relying on visual information. This is similar to work by Josh Smith that used electric field sensing [159], except that our work uses RFID tags for precise identification information and the RSSI values obtained are related to the *magnetic* coupling between the antenna(s) and tag(s).

For the moment, we assume that the candidate objects lie along a straight line – let’s call this

the x -dimension. Using other sensors (eg. cameras, laser rangefinders, or RGB-D cameras), we compute a collision-free arm trajectory that is parallel to this line with the antenna(s) facing the candidate objects. Now, we can record RFID sensor measurements (detection and RSSI) specific to the desired RFID tag (using singulation), producing a 1-dimensional function that yields RSSI versus the arm's x -position,

$$f : x \rightarrow RSSI. \quad (75)$$

From Figure 84, we know that the minimum required transmit power (maximum RSSI) will occur when the distance between the tag and the wrist-mounted read antenna are minimized. Thus, the desired object's location along the line (x_o) will correspond to the maximum RSSI,

$$x_o = \operatorname{argmax}_x f(x). \quad (76)$$

If we assume that the RFID sensor measurements have zero-mean Gaussian noise, we can take multiple samples at each location and compute the expectation (as in Chapter 4), so that

$$P(RSSI | x) = f(x) + \mathcal{N}(0, \sigma), \text{ and} \quad (77)$$

$$x_o = \operatorname{argmax}_x E(RSSI | X = x). \quad (78)$$

$$\dot{x}_o = \kappa \cdot E(\Delta RSSI | X = x). \quad (79)$$

To test this capability, we positioned EL-E (with the wrist-mounted reader system) in front of three visually-identical bottles on a table. The bottles were positioned in a line within EL-E's manipulation workspace with ≈ 2 cm separation between adjacent bottles. We placed RFID tags on two of the bottles and left the third untagged. This setup is shown in Figure 87.

For each trial, we varied the ordering of the bottles and randomly selected one of the two tag ID's (corresponding to one of the two tagged bottles) for EL-E to grasp using an existing (slightly-modified) grasp controller developed at Georgia Tech's Healthcare Robotics Lab [76]. We used EL-E's planar laser rangefinder to determine the location of all three objects on the table and compute the arm scan trajectory. We moved EL-E's arm while capturing RFID sensor readings (minimum required transmit power in this case). We associated the location of the minimum reader transmit power (maximum RSSI) with the nearest laser-segmented object location, which is then passed to the existing grasp controller. Sample data for two scans (one for each tag ID) is shown in

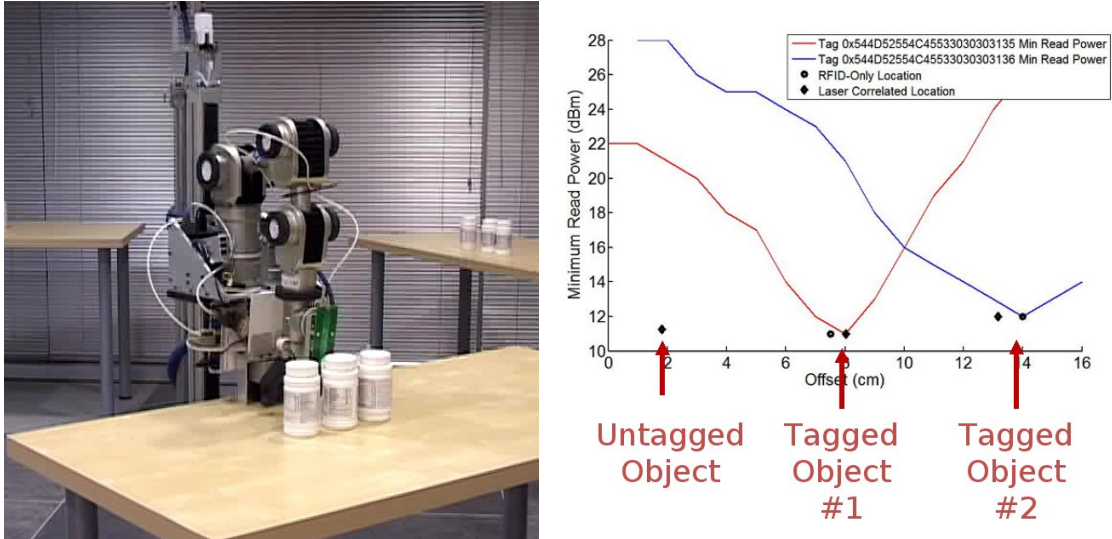


Figure 87: When scanned in front of three medication bottles (left), the signals from the wrist-mounted short-range antenna can determine the identity of the otherwise visually-identical bottles (right). The offset with minimum required read power closely coincides with the center of the objects, as determined by a planar laser rangefinder.

Figure 87. In 5 out of 5 trials (100%), EL-E successfully associated the RFID sensor measurements with the desired tagged object. In all trials, EL-E used the laser-segmented object locations to successfully grasp the desired tagged object [76]. Agreement between the RFID-determined position and the laser segmentation position was within 1cm in all of the experiments.

In general, objects are unlikely to be positioned in a perfect line, so a single 1-dimensional scan is unlikely to pass by all tagged objects with equal tag-antenna distances; indeed, the location of the tag on the object can also affect tag-antenna distances. There are several possible ways to account for this. For example, it may be possible to scan along a line tangent to each object individually to determine its identity, or even just move the short-range antenna to the same distance (say, 2cm) from each object and just associate an object's identity with the tag that yields the maximum RSSI for that location. Both of these techniques are in the spirit of relative sensing methods, like those examined in Chapter 4. As we will see later, this latter approach will be insufficient to determine the ID of an object under all circumstances; environmental conditions can create cases where RSSI is higher for a nearby tagged object (eg. with superior RF material properties), even though the antenna is closer to the desired tagged object. However, as long as the identification-and-grasp behavior is sufficient to whittle down the number of candidate objects to a manageable number,

the robot can simply grasp each possible candidate object and verify its identity post-grasp before terminating; we will discuss post-grasp identification in Section 5.6.

5.5.2 Short-Range RFID Servoing

The in-hand RFID reader system places two ceramic microstrip antennas in each of EL-E's opposing fingers. When the robot's gripper is entirely open (10cm gap), the antennas on the left and right fingers are separated by 12cm, with $\pm 6\text{cm}$ between the finger tips and the gripper's center. When we move the robot's end-effector along the same 1-dimensional trajectory as the wrist-mounted reader in the last section, The RSSI measurements obtained by the left and right antennas will be $f(x - 6\text{cm})$ and $f(x + 6\text{cm})$ (respectively) when the gripper's center is located at x , as shown in Figure 88. A data capture for all four antennas illustrating this relationship is shown in Figure 88.

We can make a few observations about the data:

- Scans that pass closer to the tag have greater peak RSSI measurements and exhibit greater slope. For each scan, the location that achieves the maximum RSSI (for a particular antenna) will be the location where that particular antenna is closest to the RFID tag. We observed both of these effects for the wrist-mounted RFID system that used power-scaling as a proxy to RSSI in Figure 84.
- Antennas L_1 and R_1 in Figure 86 exhibit greater RSSI measurements compared to L_2 and R_2 . This is due to polarization mismatch; microstrip and dipole antennas both have strong linear polarization, which makes antenna-tag coupling highly dependent on orientation. In this data capture, antennas L_1 and R_1 are aligned in orientation with the dipole tag, so they have strong polarization matching. Meanwhile antennas L_2 and R_2 are rotated 90° with respect to the tag, resulting in polarization mismatch and weaker RFID sensor measurements.
- The curves in Figure 88 have a similar flavor to those in Figure 53 from Chapter 4, where we performed bearing-only RFID servoing. Intuitively, if we receive a stronger signal from the left finger's antenna, we could move left; if we receive a stronger signal from the right finger's antenna, we could move right. The difference between the signals obtained from opposing antennas (L_1 and R_1) has a sign inflection when the tag is positioned in the middle of the





Figure 88: *Top:* EL-E moves her gripper with embedded RFID antennas along a linear trajectory in front of a tagged object. EL-E's gripper is opened, with its fingers 12cm apart. When the tagged object is centered about the object, this corresponds to a position $x_{gripper} = 0cm$, so that the left finger is located at $x_{left} = -6cm$ and $x_{right} = +6cm$. *Middle:* Filtered RSSI values over a two-second sliding window taken while moving the arm along a trajectory in front of a tagged object. Each curve corresponds to one of EL-E's four antennas. In this plot, the distance between the gripper and the tagged object at $x_{gripper} = 0cm$ is 4cm. *Bottom:* In this plot, the distance between the gripper and the tagged object at $x_{gripper} = 0cm$ is 9cm.

gripper. This position represents a stable equilibrium for short-range RFID servoing.

The latter observation prompted us to briefly explore 1-dimensional RFID servoing using the in-hand RFID reader system (using antennas L_1 and R_1). Akin to bearing-only RFID servoing from Section 4.2, when the gripper's center ($x = x_{gripper}$) deviates from the location on the scan trajectory closest to tag (by δ), the RSSI measurements obtained from the left and right antennas will be given by

$$RSSI_{left} = f(x = \delta - 6cm), \text{ and} \quad (80)$$

$$RSSI_{right} = f(x = \delta + 6cm). \quad (81)$$

Treating these measurements with zero-mean Gaussian error, the left and right antenna measurement probabilities and their difference can be defined as

$$P(RSSI_{left} | \delta) = f(x = \delta - 6cm) + \mathcal{N}(0, \sigma) \text{ and} \quad (82)$$

$$P(RSSI_{right} | \delta) = f(x = \delta + 6cm) + \mathcal{N}(0, \sigma) \text{ so} \quad (83)$$

$$E [P(\Delta RSSI | \delta)] = E [P(RSSI_{left} | \delta)] - E [P(RSSI_{right} | \delta)]. \quad (84)$$

To servo along the linear trajectory, we move the arm according to

$$\dot{x} = \kappa_3 \cdot E [P(\Delta RSSI | \delta)]. \quad (85)$$

Figure 89 shows the response of RFID servoing. A tagged medication bottle is moved from the right of the frame to the left. EL-E's hand tracks this movement by servoing along a line parallel to the robot. After servoing (tracking) the tagged medication bottle for 10 seconds, EL-E robot reaches out a fixed distance and closes her gripper. This sort of behavior might be useful for handing off objects to the robot – the RFID servoing allows the robot to correct for small mis-alignments in the gripper's initial position. Similar techniques have been used for (non-tagged) electric field sensing [159].

5.5.3 Raster-Scan RSSI Images

We are not limited to 1-dimensional motions for the wrist-mounted and in-hand RFID systems. We could just as easily scan the short-range antennas through 2-dimensions – moving the end effector




Figure 89: EL-E servos relative to a tagged medication bottle by using the difference in RSSI between opposing fingers’ antennas. After servoing for a fixed amount of time, EL-E reaches out and closes her gripper to perform a grasp.

back and forth in the plane above or in front of tagged objects. These back and forth motions, commonly called a raster scan, allow us to capture RSSI values for each xy -location on the plane. An image-based representation, where each $\langle x, y \rangle$ pixel corresponds to a discretized $\langle x, y \rangle$ location on the plane, is a natural way to visualize this spatial data. This is similar to the RSSI images from Section 4.3; we dub these “raster-scan RSSI images” to signify that we are scanning along a plane rather than in pan-tilt. Figure 90 shows raster-scan RSSI images for each of the four RFID antennas embedded in the in-hand reader. The tagged object (black arrow) is indicated at the center of each of the raster-scan RSSI images. In this position, the gripper is centered directly over the tagged object. The maximum RSSI values for the finger antennas occur when the gripper is offset from this position by $\pm 6\text{cm}$ – where the finger antennas are closest to the tagged object.

These raster-scan images provide insights into the otherwise invisible RFID sensor signals. For example, visualizing the difference between the left and right finger antennas signals (L_1 and R_1) in Figure 90, we are able to visualize the region where servoing would converge. Much like RSSI images, we believe this technique could also prove valuable for object localization or sensor fusion.

5.6 Additional Robot Behavior: Verifying a Grasped Object’s Identity

In Section 5.4, we showed that the in-hand RFID system was capable of detecting objects in the robot’s hand. However, this is subtly different from identifying *which* tagged object is currently located *in* the robot’s hand. Recall that the in-hand reader can detect tags out to 40cm, so multiple






Figure 90: Rastering the manipulator on a plane above the tagged-object yield contour plots of received signal strength (RSSI) for all four antennas. Using the difference between RSSI on opposing antennas yields useful signals for servoing.

tagged objects nearby the end effector could be positively detected when querying for the identity of nearby tags. Further, it is not necessarily the case that the closest object yields the strongest measurement. The RSSI measurement is a complex function of: relative tag pose, tag type, object materials, and the RF environment around the reader antennas. For example, an object with “good” RF materials may result in a stronger measurement even though it is farther away than a tagged object with “bad” RF material properties being held directly in the robot’s gripper. This observation led us to develop a robot behavior that can *verify* a grasped object’s identity, not *just* detect the tag.

One method to verify a grasped object’s identity is to monitor RSSI values (or even presence / absence of a tag) over time while repositioning the robot’s end effector (and the grasped object). In Figure 91, we show a scenario where EL-E is grasping a cellphone near a number of other tagged objects (all objects using Squiggle tags) on a tabletop – a scenario that might be indicative of a grasp attempt, perhaps using the other short-range RFID behaviors. Once EL-E has picked up the cellphone, the robot lifts the gripper up and away from all of the other objects on the table. The final location of the in-hand reader antennas is more than 40cm (the maximum read distance) from the initial grasp location. This should mean that any tagged objects that were nearby during the grasp attempt should no longer be near the in-hand reader and should no longer be detected. The cellphone is the only object still being detected, so we can infer that the cellphone is indeed the object in EL-E’s gripper.

From this single illustrative example, we can see a number of interesting signal characteristics:

- The cellphone has one of the lowest initial RSSI values of any of the tagged objects, likely owing to extensive metallic components. As EL-E moves far away from the other tagged objects, the cellphone’s RSSI signal is the only one that remains present. This indicates that the cellphone is the object contained in EL-E’s manipulator.
- The RSSI signals (for all the objects) do not remain constant. A variety of factors come into play as the in-hand reader antennas are moved through the environment. In particular, the environmental conditions can change drastically: nearby objects, multipath, and even interference from the robot itself. All of these factors can change the RFID read strength.
- Similarly, two of the in-hand antennas were unable to detect the Squiggle tag on the cellphone








Figure 91: Monitoring RSSI values over time (*bottom plots*) while moving the arm up and away from distracting objects (*top photos*) allows EL-E to correctly identify the grasped object: a cellular telephone.

after EL-E's arm pulled away from the table. As in Section 5.4, the diversity provided by the four in-hand antennas is quite useful for detecting the tag under marginal conditions.

- In the previous section, we suggested that this method could be used to grasp and then identify tagged objects in scenarios where the precise tag ID could not be ascertained. This capability speaks to the benefits of a mobile manipulator, which can not only change its position to get various RFID readings, but also manipulate the world for its own benefit.

In our example, we only relocated the in-hand reader by approximately 60cm by lifting EL-E's arm up and away from the table. Alternatively, the robot could completely relocate to a new location its mobile base. Taken to the extreme, EL-E could move to a different room or to a wide-open, empty space to obtain strong confirmation about the grasped object's identity. This approach may facilitate analogous grasp verification behaviors using the long-range antennas (eg. holding the tagged object in front of the long-range antennas to detect its presence), which would require large translations to escape from the initial detection position. The in-hand reader's short-range operation facilitates quick, in-place object verification. Depending on the result of this grasped object verification behavior, EL-E could either continue on to subsequent tasks (for success) or promptly replace the object and attempt a re-grasp (for failure).

We believe that grasped object verification is an important capability for systems that require high-confidence identification and tracking. In some cases, such as medication fetching and delivery, the consequences of a mis-identification could be dire. The ability to at-will verify the identity of a object in a robot's gripper may be a foundational capability for these types of applications.

5.7 Conclusions

In this chapter, we described the design of *two* short-range UHF RFID sensor systems utilizing near-field RF coupling that were used by the autonomous mobile manipulating robot, EL-E. The first system utilized a loaded microstrip affixed to EL-E's wrist. The second system employed four ceramic microstrip antennas embedded in EL-E's fingers to for an in-hand antenna capable of detecting grasped objects.

Using these two unique systems, we examined several short-range robot behaviors that drew inspiration from analogous relative behaviors in Chapter 4 to detect the location of tagged objects

at short reader-tag distances:

- **Locating a Tagged Object:** Moving the wrist-mounted antenna in front of (visually-identical) tagged objects, we determined which of the objects likely possesses the desired tag.
- **Short-Range RFID Servoing:** Using the differential signal between in-hand antennas on opposing fingers, we servo the robot’s manipulator to try and maintain the object’s position between the robot’s fingers.
- **Raster-Scan RSSI Images:** Scanning the in-hand reader in the xy-plane above a tagged object, we built “raster scan RSSI Images” that showed the RSSI signal strength as a function of xy-position. These raster-scan RSSI images provided insights into the otherwise invisible RFID signals.
- **Verifying a Tagged Object’s Identity:** We monitored RSSI readings as a grasped object was lifted to determine (1) if the grasped object is tagged, and (2) the tagged object’s identity. We showed that this behavior works even in the presence of nearby (distracting) tagged objects.

CHAPTER VI

UHF RFID SENSING FOR MOBILE MANIPULATION

In this chapter, we will explore several complete mobile manipulation systems that utilize UHF RFID sensing:

- **Physical, Perceptual, and Semantic (PPS) Tags:** We develop a new type of tag (dubbed a PPS-tag) that combines a UHF RFID tag with additional forms of augmentation (eg. compliant materials with visually-distinct properties) to provide physical, perceptual, and semantic assistance to robots. We present five exemplar PPS-tags along with a set of robotic behaviors that utilize UHF RFID perception and optimization-based RFID behaviors (eg. bearing estimation and RFID servoing) to accomplish various tasks, such as: turning on and off light switches, opening and closing drawers, operating lamps, and disposing of trash in a waste bin.
- **Multi-Sensor Fusion and Mobile Manipulation Using RSSI Images:** We develop a framework for multi-sensor fusion that eliminates the need for the physical and perceptual augmentation beyond just the UHF RFID tag. We employ RFID perception and optimization-based RFID behaviors to locate and approach the tagged objects, which terminate with the tagged object close-to and in front of the robot. This allows us to bring to bear additional sensing modalities (eg. cameras and laser rangefinders that produce 3D point clouds) to perceive the tagged object. We use the UHF RFID tag's unique identifier as an index into a semantic database, where we store and retrieve information about the tagged object's appearance (as opposed to the PPS-tag's appearance). Finally, we demonstrate a complete mobile manipulation system that is capable of locating, approaching, perceiving, and grasping tagged objects in relative isolation on the floor.
- **UHF RFID for Medication Delivery and Adherence:** We apply RFID perception and optimization-based RFID behaviors to medication delivery. We develop a system wherein

a robot confirms the identity of a grasped tagged medication bottle, uses optimization-based RFID behaviors to locate and approach the intended (tagged) recipient, and then hands off the medication.

6.1 *Physical, Perceptual, and Semantic (PPS) Tags*

Autonomous mobile manipulation within human environments represents both an exciting opportunity for new robotic applications and a grand challenge for robotics [77]. Although researchers continue to make progress in this area, autonomous mobile manipulators do not yet exhibit sufficiently robust performance to support many promising applications. For example, if assistive mobile manipulators could robustly operate within real homes for extended periods of time, they could provide valuable in-home assistance. We see the critical deficiencies of current robots as falling into the following three inter-related categories:

Physical:

The robot's mechanical structure may be poorly matched to the task. For example, a robot with a primitive gripper may be unable to pull on a recessed handle, or a small mechanism may be too difficult to grasp reliably.

Perceptual:

The robot may be unable to reliably perceive the task-relevant features required for consistent success at the task. For example, a thin pull chain may be too small for a robot's laser range finder to detect, or the robot may be unable to reliably detect drawer handles due to the wide variety of handles found in human environments.

Semantic:

The robot may be unable to infer the task-relevant semantics, such as what actions it can perform with a particular mechanism or the implications of those actions. For example, the robot may not realize that it can pull on a chain to operate a lamp and that this should either increase or decrease the light from the lamp.

Many approaches seek to address one or more of these shortcomings. In this section, we propose augmenting environments to directly help robots with these three challenges. Specifically, we present PPS-tags, which stands for physical, perceptual, and semantic tags¹. We have designed

¹This was joint work with Hai Nguyen from the Healthcare Robotics Lab. Hai was the lead student on this project and was responsible for most of the robot behaviors and success / failure detection. My primary contribution were the RFID-based robot behaviors and sensing.



Figure 92: *Top Row:* PPS-tags affixed to a flip-type light switch, a rocker-type light switch, and a cabinet drawer. *Bottom Row:* EL-E interacting with each of the PPS-tags.

these tags to be affixed to sparse task-relevant locations in the environment in order to help the robot physically interact with the location, perceive the location, and understand the location’s semantics. Figure 92 shows EL-E interacting with three different types of PPS-tags.

While we ultimately hope to develop robots that will not require modifications of the environment, we believe PPS-tags offer several advantages at this time. For example, PPS-tags have the potential to accelerate the deployment of autonomous mobile manipulators in real-world applications. This could have societal and economic benefits. It could also benefit robotics research by providing data from real-world usage scenarios. Also, PPS-tags could represent a beneficial path

for system development and research. One can imagine first developing a robotic system that uses PPS-tags and then gradually removing them or altering them in conjunction with the development of improved mechanical, perceptual, or semantic capabilities. Similarly, researchers can use PPS-tags to immediately explore system-level questions, rather than waiting for the solutions to long standing problems such as object recognition. Additionally, we believe PPS-tags might enable simple inexpensive robots to perform complex tasks.

6.1.1 Related Work

Robots in Augmented Environments:

People often alter environments for robots. For example, in factories people create robotic work cells matched to the tasks performed by the robot. There are also many examples of environmental modification for robots outside of industrial settings. For example, most high performing systems in RoboCup competitions depend on environments that are easy to perceive with color vision [26, 155]. Also, many robots have depended on perceptual augmentation of the environment, such as with ARTags and QR tags [73, 74]. Roomba owners “roombaize” their homes, a process that often involve changing furniture layouts, cleaning up wires, and tucking in rug tassel [141]. People sometimes attach fabric to the handles of doors and drawers so that service dogs can operate them. Previous work from Healthcare Robotics Lab demonstrated the use of towels as a physical and perceptual aid for a robot [109]. In contrast to prior work, PPS-tags combine physical, perceptual, and semantic assistance to enable a robot to perform a variety of tasks using similar behaviors. Other installations, such as smart homes, distribute cameras or RFID readers throughout the environment [158, 18]. In contrast, PPS-tags can be simple, sparsely distributed, inexpensive, and independent from one another. In our current implementation, the robot does not require detailed models of the environment nor the tagged objects, and instead uses sparse task-relevant information.

RFID-Assisted Robots

As explored in this thesis, Radio-frequency identification (RFID) represents another example of environmental augmentation for robots. Due to the low cost of tags and the opportunity for non line-of-sight perception, RFID tags have enjoyed a great amount of attention in robotics –

particularly for navigation, localization, and mapping [86, 55, 83]. Researchers have also explored opportunities for associating semantic information with the unique identifier provided by an RFID-tag. Using XML profiles, the authors of [79] defined object properties such as weight, and grip force for a table mounted robot. Ha *et al* [53] proposed a knowledge architecture based on the semantic web language, OWL-S, to describe objects, possible actions, and the expected effects of actions. Baeg [13] described a smart home environment with interoperating devices such as RFID enabled tables, shelves, and mobile robots. Jang, Sohn, and Cho [68] presented an architecture for associating semantic labels and properties such as indicating what areas are restricted within a physical space. Hidaya et. al. [134] proposed that objects should be tagged with their affordances.

Although many researchers have previously suggested that RFID-indexed databases could be used by robots, there is a lack of published results describing real robots making use of the proposed information. The authors of [103] may be the first to have implemented their architecture on a mobile manipulator, and is the only work we have found that describes a real mobile manipulator making use of RFID-indexed semantic information. Their robot moved a cup and a chair using object properties loaded from a RFID-indexed database. We have implemented our system with a mobile manipulator and tested it with three different devices. We have found that a relatively simple semantic structure with only a handful of entries is sufficient to support these tasks. We pursued a bottom-up design strategy in which the goal of the robot performing specific, well-defined tasks dictated the contents of the semantic database. PPS-tags also combine this semantic assistance with physical, and perceptual help.

6.1.2 Three Examples of PPS-Tags

The PPS-tag concept is general and could take many forms. In this section, we present three illustrative examples of PPS-tags, shown in Figure 92, that we have evaluated on a mobile manipulator.

6.1.2.1 Physical: Manipulating High-Friction, Compliant Materials

Each of the three tags shown in Figure 92 provides a different form of physical assistance. All of them are compliant with high friction. Two of the PPS-tags (attached to the flip-type light switch



Figure 93: *Left:* Compliant, slip-resistant foam tubing is used to make eating utensils easier to grasp. *Middle:* High-friction Dycem polymer is used as an assistive device to prevent slip. *Right:* Towels are used by service dogs to interact with doors and drawers [109].

and the drawer) also increase the target volume over which the robot can successfully manipulate the tag.

For the first type of PPS-tag, we use a non-slip, compliant red foam tube attached to the flip-type light switch. This tubing is normally used to make eating utensils and cylindrical objects, such as toothbrushes and pencils, easier for people with motor-impairments to manipulate (Figure 93). We purchased this foam tube from an online store (RehabMart.com) that supplies materials and devices to assist people with physical disabilities.

For the second PPS-tag, we use patches of red Dycem polymer attached to a rocker-type light switch. This light switch complies with the Americans with Disabilities Act, so we will commonly refer to it as the ADA Light Switch. The Dycem polymer is a high friction, compliant material commonly affixed to wheel chairs and walkers to prevent grip slippage (Figure 93). We purchased this material from the same online store.

For the third PPS-tag, we use a red towel, which is affixed to drawers and doors to facilitate manipulation. This interaction is inspired by service dogs that assist motor impaired individuals. Service dogs do not like to grasp (bite) metal door and drawer handles. Animal owners commonly tie soft, compliant towels onto door and drawer handles to make it easier and more pleasant for the service dog to grasp [109]. Physically, the towel provides a large compliant target for grasping – this makes it much easier to grasp using simple grippers, which may not be well-suited to manipulate the vast array of handles found in human environments.

6.1.2.2 *Perceptual: RFID Sensing, Color Images, and 3D Point Clouds*

All three PPS-tags have features that simplify perception for various robot sensing modalities. For the version of EL-E used in this work (shown in Figure 94), the relevant sensors and perceptual features are:

UHF RFID:

A UHF RFID provides both perceptual (through UHF RFID sensing) and semantic (through a ID-indexed database) information about the PPS-tag. While other technologies (ie. QR codes) might be able to serve a similar role, our research focuses on UHF RFID. For our examples, the RFID tag is placed on or near the object or mechanism; further refinements could incorporate the tag into the dycem polymer, foam tube, or towel. The UHF RFID tag made it possible for EL-E to interact with the PPS-tag using it's onboard long-range and short-range UHF RFID systems using any of the techniques developed in previous chapters.

Color Cameras:

All three illustrative PPS-tags are comprised of a bright red material. This makes the PPS-tags stand out in images captured by any of EL-E's on-board cameras: a Point Grey Firefly camera (640×480 resolution) mounted in EL-E's hand, and a higher resolution Point Grey Dragonfly 2 camera (1024×768 resolution) near the base of EL-E's arm. Also, the larger physical size of the PPS tags aids visual recognition from a distance.

3D Point Clouds:

EL-E has a tilting, infrared laser rangefinder capable of producing dense 3D point clouds. Unlike highly-reflective metal handles, PPS-tags materials do not produce specular reflections, which allows the laser rangefinder to acquire accurate sensor readings. Further, the large size of the PPS-tags help EL-E reliably detect small objects, such as flip-type light switch or thin drawer handles, from a distance.

6.1.2.3 *Semantic: An RFID-Indexed Database with Grounded Semantics*

The UHF RFID tag's unique identifier, sensed by long-range or short-range RFID systems, can serve as an index into a database containing grounded, semantic information pertinent to interacting




Figure 94: *Left:* The version of EL-E used for the PPS-tag experiments. *Right:* The downward-facing laser rangefinder is used as a light curtain to determine when to halt during RFID servoing.

with the tag. This may include such information as: which actions can be performed with this PPS-tagged object, how to perform them (behaviors and their parameters), conditions for success or failure, and the robot’s past history of interaction with the device. We implemented the semantic database as a series of nested hashables (in the Python programming language) with the first level of the hashables indexed by the RFID tag’s unique identifier. A sample top-level entry in the database for the rocker-type light switch PPS-tag is shown in Figure 95.

We focus on grounded semantics for robot manipulation. By this, we mean that we restrict the semantic database to hold information that directly informs the robot’s manipulation behaviors. Each object-specific entry in the database contains three main components: properties, actions, behaviors. “properties” is a place for information about the object that is not specific to an action (Figure 95). The “actions” hashtable maps user-friendly names for a PPS-tag’s available actions to associated robot behaviors. Finally, for each robot behavior in “behaviors,” we define separate hashables (e.g. “push_bottom”, “push_top”) where we store parameters used by each of the robot’s behaviors.

To illustrate the semantic database and its use by the robot, we now describe the example entry

```

'rocker-switch': {
  'properties': {
    'type': 'ada light switch',
    'name': 'A D A light switch 1',
    'pps_tag': 'dycem',
    'change': 'overall brightness',
    'switch_travel': 0.02,
    'height': 1.22,
    'on_plane': True,
    'direction': 'up',
    'ele': {'color_segmentation':
            [[34, 255],
             [157, 255],
             [0, 11]]},
  },
  'actions': {
    'off': 'push_bottom',
    'on': 'push_top'
  },
  'behaviors': {
    'push_bottom': {
      'force_threshold': 3.0,
      'height_offset': -0.02
      'ele': { 'gripper': 5}
    },
    'push_top': {
      'force_threshold': 3.0,
      'height_offset': 0.02,
      'ele': {'gripper': 5}
    }
  }
}

```

Figure 95: A semantic database entry for operating a rocker-type light switch. The current database is implemented as a hashtable using Python (programming language) dictionaries.

(Figure 95) in more detail:

'properties'

In 'type', we store the class of object, such as 'ada light switch', 'light switch', or 'drawer'. In 'name', we store a unique name that is specific to this particular object instance, such as 'A D A light switch 1', 'light switch 1', or 'drawer 1'. Both of these levels of naming, class and instance, could potentially be useful to the robot, such as when collecting data from experience which may relate to the specific instance or the class of object being used. 'pps_tag' defines the type of tag being used, such as 'dycem', 'towel', or 'foam tube'. 'change' describes the change in state that should be observed upon using the object successfully, such as 'overall_brightness' for lighting and 'location' for the drawer. 'direction' tells the robot where to look to observe this state change, such as 'up' for the light switches. 'ele' contains a hash table with information specific to the robot EL-E. In this case it holds the HSV color segmentation boundaries that segment red color comprising the tag when using EL-E's camera(s).

'actions'

For the 'ada light switch', the two associated actions are turning the light 'on' and turning the light 'off'. These map to the 'push_top' and 'push_bottom' behaviors respectively.

'behaviors'

Each of these behaviors also has an entry which stores information important to performing the action. For example, 'push_bottom' holds information critical to pushing the bottom of the rocker switch in order to turn the light off. It has the entries 'height_offset' with a value of -0.02 meters, and 'force_threshold' with a value of 3 Newtons. These describe how far below the center of the PPS-tag to push and the force to apply when pushing. The 'ele' entry for the 'push_bottom' behavior how close the opening angle that should be used by EL-E's gripper when performing this action. The value (5 degrees) places EL-E's gripper in a pinching configuration that is useful for pushing the button.

6.1.3 Interacting with PPS-Tags

Now we are going to describe how EL-E interacts with our three PPS-tags. We assume that EL-E is located nearby the tag, with an unobstructed path to the tag. This assumption is well-matched to the end conditions provided by RFID search from Section 4.4, which did not exist at the time this work was performed.

6.1.3.1 *Selecting a PPS-tag and Action*

First, we must select a nearby tag and action for EL-E to perform. Using the long-range UHF RFID system, EL-E can prompt a user to supply this input. Unlike previous work (described later), the user in this situation is not required to be near the robot; the user could be located at a remote location.

EL-E pans her long-range, RFID transmit antennas through $\pm 135^\circ$ while querying for nearby RFID tags. The resulting list of tag IDs is presented to the user via a graphical user interface, shown in Figure 96. We call this a “contextually-aware user interface,” since only tag IDs that EL-E knows how to interact with are displayed to the user. In Figure 96, several nearby tags are detected – the user has selected the drawer containing a PPS-tag. The contextually-aware user interface presents the user with the actions available for that PPS-tag: ‘open’ and ‘close’. The user has selected the ‘open’ action, so EL-E will execute the robot behavior associated with this action.

It is worth noting that tag IDs unknown to EL-E are not displayed in the user interface. Conceivably, EL-E could query remote databases to obtain information about how to interact with unknown tags. This could support on-demand retrieval of robot behaviors, ie. the tags could act as uniform resource locators (URLs) for behaviors or applications (apps) stored on the internet. We think this is a compelling way of adding new functionality to robots, as described in Chapter 1:

During a trip to a department store, an individual purchases a robot and a box of standard labels: “dish”, “dish washer”, “clothing”, “washing machine”, “toy”, “litter box”, “scoop”, and “storage bin”. The individual returns home, applies the labels as directed, unboxes the robot, and turns it on. The robot connects to the web, downloads the appropriate behaviors, and is instantly able to operate in the labeled world: loading the dishwasher with labeled dishes, putting away labeled toys, cleaning the cats’ litter

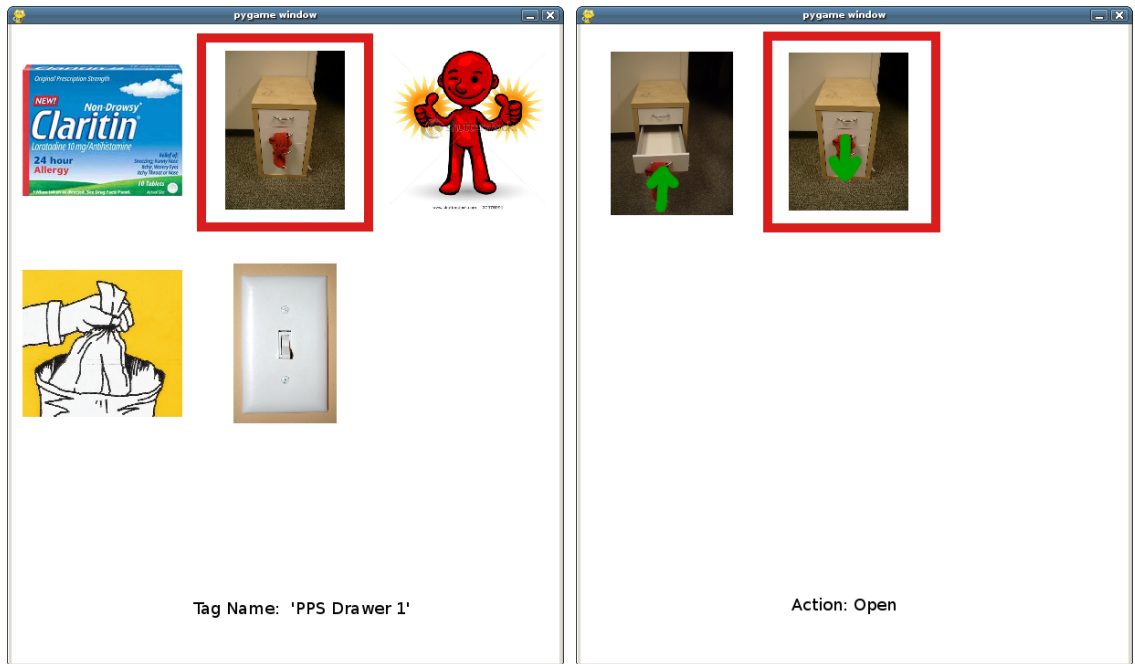


Figure 96: The contextually-aware graphical user interface displays the tags (left) that were detected nearby and that EL-E knows how to interact with. After selecting a tag, the interface displays the actions associated with the tag (right). In this example, the user has selected the drawer with a PPS-tag, and has selected the 'open' action. EL-E will now execute the appropriate robot behavior.

box, and washing labeled clothing. Improved or new functionality, such as delivering medicine, is just a few labels and an internet connection away.

For the purposes of this work, the semantic database and robot behaviors are stored locally on the robot.

6.1.3.2 *Initially Approaching the PPS-Tag*

Now that EL-E has been instructed to perform a specific task on a specific PPS-tag, she needs to move in closer to the PPS-tag. At this point, EL-E does not know where the PPS-tag is located in 3D space. EL-E uses the relative RFID behaviors from Chapter 4 to approach the PPS-tag, estimate it's position using other sensors, and then move to a position nearby in preparation for performing the desired action.

First, EL-E orients herself toward the PPS-tag. EL-E pans her long-range antennas, acquires RFID measurements targeting just the selected tag, performs bearing estimation according to Section 4.1, and then orients herself accordingly. This process puts EL-E in a starting position roughly

oriented toward the tag. From Section 4.1, we know that bearing estimation provides low-fidelity results from longer distances; therefore, this initial orientation may not align the PPS-tag with EL-E's other sensors. Thus, EL-E does not perform tag pose estimation at this time. Instead, EL-E moves in closer to the tag so that a shorter-distance bearing estimation can yield better sensor alignment.

EL-E closes the distance between herself and the PPS-tag using RFID servoing from Section 4.2. The previous orientation step put EL-E into favorable initial conditions (roughly oriented toward the tag) for RFID servoing, so EL-E positions her antennas (at $\alpha = 40^\circ$) and begins servoing in bearing ($\kappa_3 = 0.02$) while moving forward at 10cm per second. By assumption, there are no obstacles between EL-E and the PPS-tag, so the first obstacle encountered will be the PPS-tagged object. Using her downward-facing laser rangefinder as a light curtain (Figure 94), EL-E moves forward until an obstacle (exceeding 5cm in height from the floor) "breaks" the light curtain within 30cm of the robot, indicating that unobstructed forward movement is no longer safe [67]. This obstacle may be the PPS-tagged object, cabinet, wall, or an undesired nearby object.

From this closer vantage, EL-E re-orientes herself using bearing estimation. This process results in conditions where EL-E is facing in the direction of the PPS-tag. The semantic database entry for the PPS-tag has a 'properties' attribute that indicates the height of the PPS-tag. EL-E uses her linear actuator to reposition her arm, tilting laser rangefinder, and high-resolution camera to this height. Now, the conditions are such that the PPS-tag is located in the field of view of the high-resolution camera and tilting laser rangefinder, which can be used to compute the pose of the PPS-tag through sensor fusion.

6.1.3.3 Estimating the PPS-tag's 3D Position

With the PPS-tag in front of EL-E's sensors, our algorithm acquires a color image from the high resolution camera and a dense 3D point cloud using the tilting laser rangefinder. These two sensors are calibrated and registered with one another. We use a series of perceptual algorithms to determine the PPS-tag's 3D position and a waypoint (pose) for the robot that is appropriate for manipulating the PPS-tag.

First, we need to determine the PPS-tag's 3D location. All three PPS-tags share a very similar red color that occupies a large area in the camera image. Our algorithm performs a 2D segmentation of red patches in the image using minimum and maximum thresholds defined in HSV space using the parameters defined in EL-E's color segmentation properties ('properties' \rightarrow 'el' \rightarrow 'color_segmentation') from the semantic database [70]. The next step post processes this raw segmented image with a series of morphological operations: hole filling, closing, then opening. Using the known 3D rigid body transformation between the point cloud and the camera, the algorithm projects the 3D point cloud into the color segmented image. Points projected onto red-segmented regions are kept and used to construct a 3D occupancy grid (resolution 1 cm^3). We use 3D connected components to compute the centroid of the largest connected component. The 3D location of this centroid is our estimate of the PPS-tag's 3D position.

Next, we need to determine a waypoint (pose) where we can position the robot to interact with the PPS tag. We define a cylindrical volume of interest (VOI) around the PPS-tag's 3D position, such that the axis of the cylinder is parallel to gravity and passes through the selected location. Using points from the previously-captured 3D point cloud that fall within the VOI, we use MLSAC (a variant of RANSAC) to find all planes in the VOI. For MLSAC we use the implementation provided through the ROS personal-robots repository [126, 122]. We throw out all of the planes with fewer than 100 points and select the remaining plane whose member points come closest to the the PPS-tag's 3D position. Now, given the estimated location and orientation of the plane and the estimated location of the PPS-tag, we calculate a waypoint 50 cm from the selected tag in the direction perpendicular to the plane. We then check if the robot can be centered at the waypoint facing the PPS-tag with its arm extended without colliding with points in the point cloud. If this test passes, EL-E drives to the waypoint (position and orientation) using the Robot Operating System's (ROS) navigation stack.

An example of this perceptual algorithm for the drawer PPS-tag is shown in Figure 97.²

²The data in this figure was obtained when using the laser-pointer interface PPS-tag implementation discussed in Section 6.1.5. Still, it is indicative of the conditions experienced in this implementation using long-range RFID behaviors for navigation.




Figure 97: The camera image (left) and point cloud (right) used to estimate the PPS-tag’s 3D position and the robot waypoint (position and orientation) for subsequent manipulation. (Key for point cloud colors: Cyan points lie within the volume of interest. Red points were found to correspond to the PPS-tag. Dark green points denote the plane found closest to the PPS-tag’s centroid. The yellow-green bounding boxes check for potential collisions with EL-E’s body and arm.)

6.1.3.4 Close-Range Alignment and Tag ID Verification

While driving towards the waypoint, EL-E updates her estimate of the PPS-tags location using odometry. This results in substantial accumulated error both in orientation and distance to the tag, so the EL-E now performs two additional navigation steps so that it can read the PPS-tags RFID and manipulate the object.

First, EL-E unfolds her arm into a forward-facing position (see Figure 92) and visually servos her arm so that the red PPS-tag is centered in the eye-in-hand camera mounted to EL-E’s arm. While servoing, EL-E monitors the forces on her fingers, and stops if a collision is detected. If necessary, EL-E updates her orientation to keep the tag in her manipulation workspace. Upon completion, EL-E’s end effector is approximately 10cm away from the PPS-tag.

Second, to refine the distance estimate to the PPS-tag, which also may suffer from odometric errors during navigation, EL-E reaches forward with her arm until she detects contact with the force-torque sensors at the base of her fingers or until the IR range sensor in the gripper’s palm detects

an obstacle. EL-E then uses her in-hand RFID system to query for nearby tags to verify that the tag ID with maximum RSSI measurement matches the user-selected tag ID. Since the PPS-tags are sparsely distributed throughout the environment, there is little chance of obtaining stronger reads from distracting PPS-tags (see Chapter 5). If the tag IDs match, then EL-E continues and executes the user-selected action. If the tag IDs do *not* match, the action is aborted and failure recorded.

EL-E then pulls back her arm a fixed distance (approximately 10cm) before executing the user-selected action.

6.1.3.5 *Execute the User-Selected Action*

The remainder of the interaction with the user-selected PPS-tag depends on which action (robot behavior) was selected by the user. We examine each action individually; the parameters for these operations are stored in the semantic database.

Flip-Type Light Switch: 'On' and 'Off'

Both actions for the flip-type light switch execute the *light switch behavior*. In this behavior, EL-E moves her carriage up (or down depending on the command), closes her gripper (but not all the way), moves the gripper forward until contact has been made with the wall (using a 2 N threshold), moves the gripper away from the wall by 2 cm (to clear the plate on which the light switch is mounted), then moves the carriage down (or up) using torque control with the gripper extended, and stops when the maximum force on the finger-mounted force-torque sensor is greater than 12 N or when the carriage has traveled in excess of 15 cm.

To monitor the effects of using the light switch, EL-E takes a picture of the expected location of the light source with its stereo camera prior to, and after moving the carriage up or down. To determine the effect of its attempt to use the light switch, EL-E takes the average intensity of the image before the action and subtracts the average intensity of the image after the action. If the magnitude of the change is greater than a threshold and the result is positive, EL-E concludes that the light has been turned off, and if it is negative EL-E concludes that the light has been turned on. An example of the signals obtained during light switch operation are shown in Figure 98.³

³The data in this figure was obtained when using the laser-pointer interface PPS-tag implementation discussed in

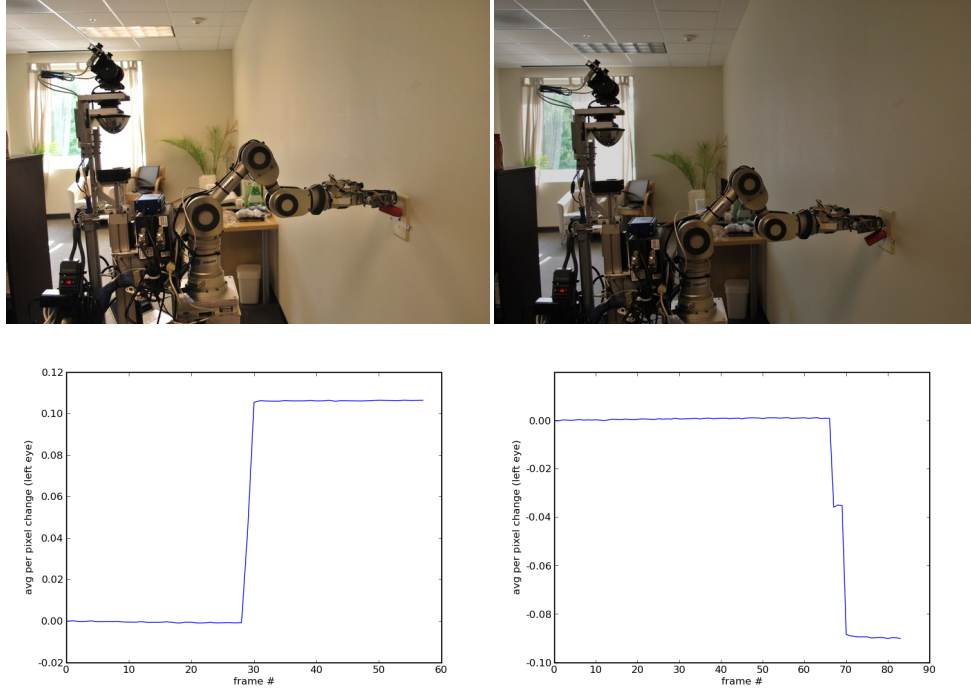


Figure 98: *Left Pair:* EL-E turning off a light switch. *Right Pair:* Brightness changes resulting from the light being switched from on to off.

Rocker-Type Light Switch: 'On' and 'Off'

Both actions for the rocker-type light switch execute the *push behavior*. The goal of the push behavior is for the robot to apply a force normal to the detected plane upon which the PPS-tag rests. This same behavior is also used to close a drawer. EL-E first sets its gripper to the settings specified in the semantic database – 5° for the rocker-type light switch, as shown in Figure 95. This setting puts the grippers' fingers close together for more appropriate “poking” motions. Next, EL-E moves its end-effector forward stopping if a force greater than 3 N is detected or the IR sensor in the palm reports an obstruction. The same brightness change detector from the flip-type light switch is used to determine success or failure.

Drawer: 'Open' and 'Close'

To open the drawer, EL-E executes the *pull back behavior*. To pull back on the towel portion of the PPS-tag, EL-E first moves its arm forward to grasp the towel. As there is only one

Section 6.1.5. Still, it is indicative of the conditions experienced in this implementation using long-range RFID behaviors for navigation.

degree of freedom for both of EL-E's fingers, the towel grasping behavior in our previous work would sometimes cause large forces to accumulate if EL-E's grippers were not positioned close to the center of the towel [109]. To some extent, this issue was mitigated by the compliance provided by the towel. In this work, we have implemented a new grasping behavior that uses the force-torque sensors to move the end effector laterally while grasping to correct for small misalignments. This grasping behavior has the effect of centering the towel in the middle of EL-E's gripper making it more likely that forces on EL-E's fingers will be distributed evenly across the two fingers as the towel is pulled backwards. If the grasping behavior detects that it has been successful (forces on both of EL-E's fingers exceed 2 N), then EL-E proceeds to pull on the towel by moving backwards with its mobile base. EL-E moves back in steps of 20 cm, stopping when either a force threshold is exceeded (drawer is fully opened), the force on the fingers drops below a threshold (fingers lose their grip), or the robot has moved back farther than 25cm (drawer is fully open). At the end of each complete pulling step, EL-E runs the towel grasping behavior again to maintain its grip on the towel. During this pull back behavior, EL-E records the displacement of its end effector between the time when it first successfully grasps the towel and when the robot either loses its grip or finishes pulling. If this distance is greater than 10cm, EL-E declares its action successful. Otherwise it declares a failure. The result (and the displacement) is recorded in the semantic database for future use.

To close the drawer, EL-E executes the same *push behavior* as the rocker-type light switch, but with slightly different configuration options (as stored in the semantic database). EL-E sets its gripper to 50° – a fully-open configuration that maximizes the chances that the end-effector will make contact with the drawer's surface. EL-E moves its end-effector forward stopping if a force greater than 3 N is detected or the IR sensor in the palm reports an obstruction. Now, EL-E pushes forward with her mobile base for a distance of 35cm or until the fingers report a force exceeding 20N (the push-forward distance on the rocker-type light switch was 0cm). The mobile base is used to push the door closed, as the total drawer travel (25cm) exceeds the reach of EL-E's arm alone. If the recorded push distance exceeded 10cm from the point of first-contact, EL-E concludes that the door has been successfully closed.

Table 26: Detailed results of EL-E operating the three PPS-tags from Figure 92

PPS-Tagged Object	Scenario	Success Rates
Flip-type light switch	Switching on	5/5 = 100%
	Switching off	4/5 = 80%
Rocker-type light switch	Switching on	4/5 = 80%
	Switching off	5/5 = 100%
Cabinet Drawer	Pull open	4/5 = 80%
	Push closed	5/5 = 100%

6.1.4 Evaluating EL-E’s Interaction with PPS-tags

We now present results from our tests of EL-E’s effectiveness in operating PPS-tagged devices. Our first goal was to test the behaviors multiple times to estimate their reliability. Our second goal was to evaluate the system’s dependence on the relative orientation of the robot to the object being operated. Thirdly, we wanted to test EL-E’s ability to recognize when it failed to operate a device. All the trials that we report here were performed in the Healthcare Robotics Lab using standard office fluorescent lighting.

We evaluated EL-E’s interaction with the three PPS-tags pictured in Figure 92. For each PPS-tagged object, we conducted ten trials – one trial for each action at for five initial locations evenly spaced by 35cm along a line 1.7 meters from the tagged object, running parallel to the wall. In all trials, EL-E’s initial orientation was facing the wall.

For each trial, EL-E first generated the user interface, then approached the user-selected tag, and finally manipulated the tag. In all 30 trials, the PPS-Tag of interest was correctly verified prior to manipulation using the finger-mounted antennas, and EL-E correctly determined success or failure at the task using information from the semantic database, such as observing whether or not the lighting changed. EL-E succeeded in 27/30 trials for an overall success rate of 90%. These results are summarized in Table 26. Further, because EL-E was able to successfully detect failures, she could have tried again after the three recognized failures.

The two failures on the light switches were a direct result of the relative RFID behaviors. When orienting the robot after RFID servoing, bearing estimation produced a poor estimate that resulted in the PPS-tag being beyond the field of view of the camera and laser rangefinder. This caused the

algorithm in Section 6.1.3.3 to fail, which EL-E correctly detected and recorded. The single cabinet drawer failure was *not* a result of the relative RFID behaviors. After approaching and orienting toward the PPS-tag, EL-E successfully segmented the tag from the camera image and point cloud. However, the largest plane detected in the volume of interest (Section 6.1.3.3) was the cabinet's side rather than the front face that contained the PPS-tag. EL-E recognized the lack of red points in front of this plane and aborted. This was a known failure mode for earlier implementations of PPS-tags [107]. EL-E also detected and recorded this failure.

6.1.5 Extending Clickable World Using PPS-Tags

In 2008, researchers in the Healthcare Robotics Lab developed a new behavior selection system for human-robot interaction that mapped virtual buttons overlaid on the physical environment to the robot's behaviors, thereby creating a clickable world [108]. In the clickable world framework, a user clicks on a virtual button by briefly illuminating a 3D location with an off-the-shelf green laser pointer. Using an omnidirectional camera and a pan/tilt stereo camera, EL-E could detect the "click" and estimate the 3D location of the laser point [76]. Inferring context based on perception surrounding the click's location (eg. proximal object, table, floor, or wall), and the state of a robot (eg. holding an object or not), EL-E would execute an appropriate behavior: pick-up a designated object from a floor or table, deliver an object to a designated person, place an object on a designated table, go to a designated location, and touch a designated location with its end effector. This interaction is shown in Figure 99.




Figure 99: A clickable world interface enables a user to trigger appropriate robotic behaviors by clicking on virtual buttons using a green laser-pointer [108].

In this section, we examine an extension to the clickable world system that uses PPS-tags (Section 6.1) to provide additional context to a laser-designated location. This context, in the form of the semantic database, facilitates additional robot capabilities: turning on and off a pull-chain lamp and disposing trash into a waste bin. The physical and perceptual features of the PPS-tags again help to facilitate these capabilities.

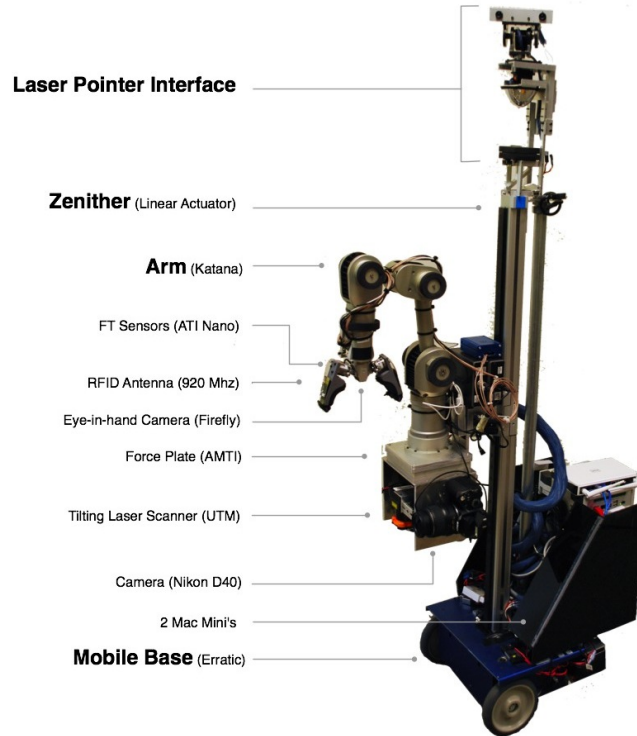


Figure 100: *Top:* Our robot EL-E (pronounced "Ellie"). *Bottom:* Fingers with short range RFID antennas.

6.1.5.1 Deviations from Previously-Explored PPS-Tags

This implementation of PPS-tags used an earlier version of EL-E that did not possess long-range UHF RFID antennas (Figure 100). Instead, EL-E *only* had the short-range, finger-mounted UHF RFID antennas. Without the long-range UHF RFID, it is not possible to construct the context-aware user interface, perform the RFID bearing estimation, or execute RFID servoing. Thus, the initial procedures for all five behaviors will be consistent, but differ slightly from those explored in Section 6.1.3.

Instead, in the context of the clickable world framework, a user located in close proximity to the robot designates a nearby PPS-tag with a green laser pointer. The robot perceives the 3D location of the selection using the laser pointer interface [76]. If the 3D location of the tag is more than 1.0 meters away, EL-E will drive towards the given location. If EL-E travels for more than a threshold distance (1.5 meters), it will ask the user to re-designate the location to reduce accumulated odometric errors. If EL-E is within 1.0 meters of the PPS-tag, then EL-E stops and

estimates the PPS-tag's 3D position using the technique described in Section 6.1.3.3. Specifically, EL-E uses the distinctive red color in a high-resolution camera image and a dense 3D point cloud obtained from the tilting laser rangefinder to estimate the PPS-tag's 3D position based on sensor data instead of user input. EL-E also estimates the waypoint for the robot to approach the tag perpendicular to the largest plane in the corresponding volume of interest. Again, EL-E drives to the waypoint (position and orientation) using the Robot Operating System's (ROS) navigation stack.

At this point, EL-E follows the procedure for close-range alignment and tag ID detection from Section 6.1.3.4. However, when EL-E reaches out to touch the PPS-tag (a procedure similar to the "poke" behavior in clickable [108]), she does not verify the ID. At this point, EL-E still does not know which tag she is interacting with, or which action to perform. Instead, after reaching out, EL-E queries for the identity of nearby PPS-tags using her in-hand UHF RFID reader system. We conclude that the tag ID with maximum RSSI measurement corresponds to the tag ID indicated by the user's earlier laser pointer selection. This is a reasonable conclusion; assuming that the PPS-tags are sparsely distributed throughout the environment, there is little chance of obtaining stronger reads from distracting PPS-tags (see Chapter 5).

Now that EL-E knows the identity of the PPS-tag, she uses the semantic database to determine which actions are available for this tag. If there is only one possible action, EL-E performs it. If there are two, EL-E audibly asks the user to select between the two choices by pointing the laser pointer either up or down – such that the laser point is either more than 50cm off of the ground for the first option, or less than 50cm off the ground for the second. For all of our PPS-tags, the number of available actions is either one or two, so this binary indication is sufficient. At this point, EL-E executes the selected action using the same tag-specific behaviors already examined in Section 6.1.3.5.

For the laser-pointer interface version of PPS-tags, we implemented two additional PPS-tag examples: a pull-chain lamp and a trashcan. These two new PPS-tags and the three older ones used in the clickable world framework are shown in Figure 101.

The tag-specific behaviors for these two new PPS-tags are as follows:

Pull-Chain Lamp: 'On' and 'Off'

For the *pull lamp behavior*, EL-E tugs down on a PPS-tag affixed to the pull chain of a



Figure 101: *Top Row:* PPS-tags affixed to a flip-type light switch, a rocker-type light switch, a pull-chain lamp, a cabinet drawer, and a trashcan. *Bottom Row:* EL-E interacting with each of these PPS-tags.

commonly available (IKEA) free standing living room lamp. For this behavior, EL-E reaches forward until it detects the PPS-tag with the infrared proximity sensor in the gripper’s palm. It then backs up by 8cm, placing the PPS tag in the center of the gripper, and closes its gripper until a threshold force is exceeded, upon which it stops closing. After gripping the pull chain, EL-E uses the same carriage control as in the flip-type light switch operation to apply a downward force directly on the lamp’s chain, stopping when EL-E’s fingers either detect a force greater than 10 N or have moved down more than 7 cm. As with the other lighting related behaviors, the pull lamp behavior monitors the change in lighting to determine success or failure. Unlike the light switches, which operate ceiling lights, the lamp pull chain points the stereo cameras forward to detect the change in lighting and indicate success or failure.

Trashcan: ‘Dispose of Trash’

In the *dispose of trash* behavior, the goal is for EL-E to take an object from its hand and drop that object into a PPS-tagged waste bin. The object in EL-E’s hand often obstructs the eye-in-hand camera, so visual servoing is not used to refine the robot’s pose prior to executing




Figure 102: *Left:* The trashcan PPS-tag. *Right:* Visualization of the drop behavior. The segmented planar front face of the trash can is in red. The large red dot is the centroid of the detected PPS-tag, and the large blue dot is the location at which EL-E will attempt to release the object.

this behavior (Section 6.1.3.4). Instead, EL-E uses the PPS-tag’s 3D location and the robot waypoint to calculate a drop location: the 3D position that EL-E should hold out its gripper (with the object) before releasing it to fall into the waste bin. An example of this calculation is shown in Figure 102. For this behavior, EL-E detects success based on whether or not it senses the object in its grasp using its finger-mounted force-torque sensors and the palm-mounted IR range sensor.

6.1.5.2 *Evaluating EL-E’s Interaction with PPS-Tags Designated by the Laser-Pointer*

We now present results from our tests of EL-E’s effectiveness in operating PPS-tagged devices in the clickable world framework. Our first goal was to test the behaviors multiple times to estimate their reliability. Our second goal was to evaluate the system’s dependence on the relative orientation of the robot to the object being operated. Thirdly, we wanted to test EL-E’s ability to recognize when it failed to operate a device. All the trials that we report here were performed in the Healthcare Robotics Lab using standard office fluorescent lighting.

In the first set of trials, we varied the tagged device used by the robot, the robot’s position with

Table 27: PPS-Tag Experimental Results

PPS-Tagged Object	Scenario	Success Rates
Flip-type light switch	Switching on	4/4 = 100%
	Switching off	4/4 = 100%
ADA-Compliant rocker-type light switch	Switching on	4/4 = 100%
	Switching off	4/4 = 100%
Lamp	Switching on	4/4 = 100%
	Detect switching failure	4/4 = 100%
Trash can	Drop object into	4/4 = 100%
	Detect dropping failure	4/4 = 100%
Cabinet Drawer	Pull open	3/4 = 75%
	Push closed	3/4 = 75%
	Detect pulling failure	2/4 = 50%
	Detect pushing failure	3/4 = 75%

respect to the device, and the action selected for a total of 32 trials. At the beginning of each trial, we positioned EL-E 1.5 meters away from the device’s PPS-tag in one of four directions. The robot was always facing towards the tag at the beginning of the trial. We then provided EL-E with a 3D location via the laser pointer interface. If multiple actions were available, we would also select the action to perform using the laser pointer interface.

In detail, for the regular and ADA light switch (rocker) we placed EL-E on evenly spaced locations along the 1.5 meter radius half-circle centered at the light switch. We placed the lamp, drawer, and trashcan next to a wall and performed the same procedure. However, we also placed the lamp such that its pull chain faced outwards in the direction perpendicular to the wall, as required by our current implementation.

In the second set of trials, we tested EL-E on an unplugged lamp, a stuck drawer, and a sticky object to test the robot’s ability to detect failure. The sticky object used in this case was a sphere of double-sided tape in order to simulate potential failures in releasing normal objects. In this case, we defined success as EL-E attempting to perform the task, performing what would usually be a successful action, and reporting that it was not able to perform the task as indicated.

We present the results from these two sets of trials in Table 27. In the first set, EL-E was able to carry out all tasks with the exception of a pulling and a pushing trial on the drawer.

In the drawers opening and closing experiments, our implementation failed in several experiments (75% for the RFID-only PPS-tags, 80-90% for the laser-pointer designated PPS-tags) [109]. These errors corresponded with initial conditions where the robot was at oblique angles relative to the front face of the drawer, which caused the side of the drawer (rather than the front) to be selected during planar segmentation.

This implementation of PPS-tags used an earlier version of EL-E that does not possess the long-range UHF RFID antennas (Figure 100). Instead, EL-E *only* has the short-range, finger-mounted UHF RFID antennas. This suggests that the methods described in this section could be equally-applicable to other short-range RFID technologies, such as the more ubiquitous LF and HF RFID.

6.1.6 Discussion and Conclusions

We have only presented five PPS-tag types. It is not hard to imagine tags that have different and potentially better properties. For example, tags with more compact and aesthetic designs would be beneficial. Also, it could be useful for the tags to provide a 6D frame of reference, like ARTags, rather than only a position with an implied orientation coming from a nearby plane; this would be one advantage of incorporating a visual fiducial instead of (or in addition to) UHF RFID tags. In addition, with tasks that involve multiple task-relevant locations, such as carrying a two-handled tray or performing tasks in the kitchen, we expect that richer semantic information that enables tags to reference one another could be valuable.

With the semantic database, we now have the ability to gather information about the robot's interaction with each device over time, which could potentially serve as a resource for self-guided learning. More generally, multi-robot coordination through the tags could be valuable, such as forms of stigmergy [17]. Having grounded, hierarchical semantic information that can be tailored to individual robots or abstracted to groups of robots might also prove valuable. One could imagine a smarter, more sensor rich robot traveling through the environment, tagging locations, and recording relevant information for use by less sophisticated robots. More generally, we expect that exploring the potential for simple robots to operate PPS-tagged devices would be worthwhile.

The work in this section makes two main contributions. First, we have presented the concept of PPS-tags which provide physical, perceptual, and semantic help to robots. Second, we presented

five examples of PPS-tags along with a set of robotic behaviors utilizing UHF RFID sensing that enabled us to evaluate their performance on a set of tasks.

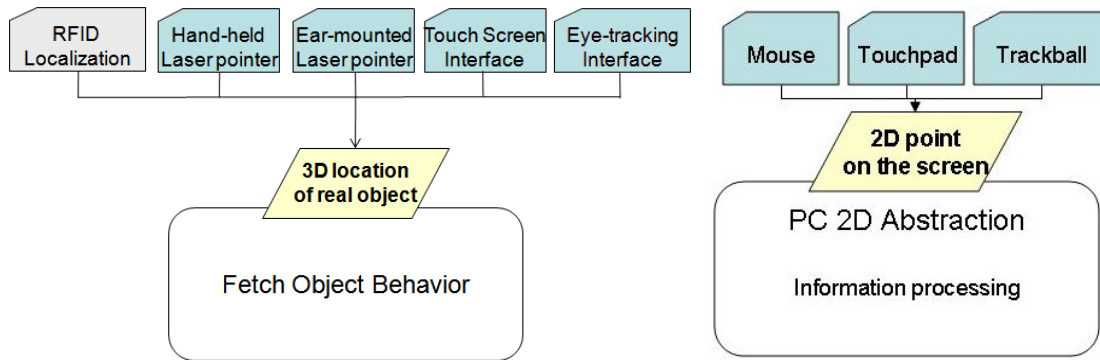


Figure 103: *Left:* A 3D point can serve as a natural interface (abstraction layer) between user interfaces and mobile manipulation algorithms for object fetching and retrieval [28]. *Right:* This is analogous to 2D interfaces used on modern personal computers.

6.2 Multi-Sensor Fusion and Mobile Manipulation Using RSSI Images

In the previous section, we examined PPS-tags that provided physical, perceptual, and semantic assistance through environmental augmentation. However, ideally robots would *not* require any environmental augmentation. In this section, we examine a framework for multi-sensor fusion that eliminates the need for the physical and perceptual augmentation, beyond just the UHF RFID tag. As with PPS-tags, we use RFID-perception (eg. bearing estimation) to bring to bear additional sensing modalities (eg. cameras and laser rangefinders that produce 3D point clouds). Again, we use the tag as a unique index into a semantic database, where we store and retrieve information about the tagged object’s appearance (as opposed to the PPS-tag’s appearance).

In this section, we develop a probabilistic sensor fusion framework that uses information from a semantic database to fuse information from three different sensing modalities: RSSI images, camera images, and laser rangefinders. Ultimately the sensor fusion yields a 3D point corresponding to the tagged object’s position. In previous work, we showed that such a 3D point was a natural abstraction layer between user interfaces and existing mobile manipulation algorithms for object fetching and retrieval, similar to personal computers’ 2D point abstraction layer (eg. mouse-based interfaces – see Figure 103) [28]. We show that the 3D point generated from our sensor fusion framework, combined with a RFID-generated context-aware user interface and RFID-perception (bearing estimation), is sufficient to seed the existing mobile manipulation algorithm for object fetching.




Figure 104: The mobile manipulator, EL-E used in this section has two articulated, long range RFID antennas (top) and short-range near-field RFID antennas on the end effector (bottom).

We evaluated these methods using a robot with actuated, long-range RFID antennas and finger-mounted short-range antennas, shown in Figure 104. The robot scans its environment to discover which tagged objects are within range, creates a user interface, orients toward the user-selected object using bearing estimation, estimates the 3D location of the object using sensor fusion (including an RSSI image), approaches and grasps the object, and then uses its finger-mounted antennas to confirm that the desired tagged object has been grasped. In our tests, the sensor fusion system (with an RSSI image) correctly located the requested object in 17 out of 18 trials (94.4%), an 11.1% improvement over the system's performance when not using an RSSI image. The robot correctly oriented to the requested object in 8 out of 9 trials (88.9%), and in 3 out of 3 trials the entire system successfully grasped the object selected by the user.

6.2.1 Related Work

Mobile manipulators typically have many sensors: cameras, ultrasonic ranging, laser rangefinders, depth cameras (RGB-D cameras), tactile sensors, bump sensors, etc. An active area of research is how to combine (fuse) data from these disparate sensing modalities to reason about the state of the

world and achieve a greater degree of intelligence or increase the mobile manipulator’s capabilities. There are many different types of multi-sensor integration and fusion techniques; for example: Weighted average, Kalman filter, Bayesian estimation, consensus sensors, decision theory, evidential reasoning, and fuzzy logic [96]. These techniques have been applied to mobile robot navigation [71, 160], localization [105], people tracking [32], object detection [114], and planning / perception [12].

For our mobile manipulation system, we transform the data from other sensors into an image-based representation so that we can use a pixel-based sensor fusion framework. The goal of our pixel-based solution is to determine salient or interesting regions in an image by scoring each pixel [135]; specifically, to determine pixels that are likely to correspond to the desired tagged object. Similar techniques have been used to fuse data from cameras and laser rangefinders to detect (safe) roadways [123], detect obstacles [41, 42, 90], or to segment images into regions of interest [87].

However, our technique is slightly different – we already know the identity of the sought after object. This allows us to load properties or classifiers specific to our desired tagged object; for example, we use color histograms features that are specific to each tagged object [87]. While we do not currently use other object-specific features, there are plenty from the related literature to choose from, such as: SIFT features [131], 3D models [127], or even features about RFID-specific sensors (eg. RSSI Image features). Relative to the related work, our system yields the following capabilities:

- The presence of the tag’s unique ID provides strong evidence that the tagged object is nearby, which allows us to build context-aware user interfaces.
- We show that the relative RFID behaviors from Chapter 4 allow the robot to bring to bear other (complementary) sensing modalities, including RSSI images, to perceive the tag without having to sift through many other scenes (which could produce false-positives).
- We can load object-specific sensor models from a semantic database indexed by the tag ID. We can fuse sensor data on a pixel-level to locate pixels likely belonging to the tag ID. We also use our new RSSI image sensing modality. Our algorithm yields a 3D point belonging to the tagged object.

- Finally, we use the output of our sensor fusion (a 3D point) as a parameter to existing mobile manipulation algorithms, which enables the robot to approach and grasp the object, and then uses its finger-mounted antennas to confirm that the desired object has been grasped.

6.2.2 Sensor Fusion

The goal of our approach to sensor fusion is to combine object-specific features extracted from multiple sensing modalities to produce a single, maximum-likelihood 3D location that can be used by a mobile manipulation system to retrieve an object with a UHF RFID tag. We provide a probabilistic framework for fusing these sensing modalities, which is accomplished through *five* steps:

1. Capture Raw Data
2. Extract Object-Specific Features
3. Load Object-Specific Probabilistic Feature Models
4. Infer a Tagged Object's 2D Image Location
5. Infer a Tagged Object's 3D Location

We will each examine in turn.

6.2.2.1 Capturing Raw Data

In this work, we consider the output of three approximately coincident sensors with overlapping fields of view: an RSSI image (from Section 4.3), a low resolution (640x480) camera image from a rectified camera, and a 3D point cloud from a tilting laser rangefinder. A representative data capture from each sensor is shown in Figure 109.

6.2.2.2 Extract Object-Specific Features

In order to fuse the output of these three sensors, we first geometrically register them with one another and create a common (*raw*) image-based representation. From each of these raw sensor images, we extract one or more spatially-varying features F_i to form feature images I_i . We build a fused image I that consists of a set of n feature images $I_0...I_n$. The number (and type) of feature images generated from the raw data will depend on the desired tagged object – more specifically,

the information available in the semantic database relating to that tagged object. For our system, we use a common set of three feature images for each tagged object: a camera feature image, a range feature image, and a RSSI feature image.

Camera Feature Image

The raw data from a camera image is a three-tuple value, or red-green-blue (RGB) *color*, for each pixel. Using this raw value directly as the (single) camera feature image,

$$I_{cam}(x, y) : \mathbb{R}^2 \rightarrow color \text{ (RGB tuple)}. \quad (86)$$

There are many other possible feature images that could be generated or extracted from a camera image, for example: edge features, Haar-like features, SIFT features, or even the output of an object detection algorithm centered about the pixel (ie. the output of a face detector might be very useful if we were trying to locate a tagged person). While all of these features are possible, we only use the RGB color feature from the camera image.

Range Feature Image

EL-E's tilting laser rangefinder generates a 3D point cloud. We estimated the 6DOF transformation from the laser rangefinder to the camera by hand measurements, and then refined this estimate using visualization software that displays the transformed 3D point cloud on the corresponding camera image. Transforming the 3D point cloud results in a range image, $I_{range}(x, y)$, that is registered with the camera image, $I_{cam}(x, y)$. An example range image superimposed on top of the camera image is shown in Figure 105. Since the laser rangefinder is approximately coincident with the visual camera and it's horizontal and vertical resolution is less than the camera, each range measurement will correspond (uniquely) to one of the camera pixels. If there exists a 3D point ($p3d$) from the point cloud that maps into a particular pixel, the pixel is assigned a value of 1.0; otherwise, the pixel is assigned a value of 0.0,

$$I_{range}(x, y) : \mathbb{R}^2 \rightarrow \begin{cases} 1.0 & \text{if } \exists p3d \in \text{cloud that transforms to } x, y \\ 0.0 & \text{otherwise.} \end{cases} \quad (87)$$

Again, this is just one type of 3D point cloud feature. Other possible features include: spin images [?], 3D segmentations [66], and point feature histograms [127]. However, use just this simple range feature image.




Figure 105: We transform a 3D point cloud (left) into the camera’s reference frame to generate a range feature image. The range feature image is displayed superimposed on the camera image (right).

RSSI Feature Image

Finally, we create an RSSI image as described in Section 4.3 by panning and tilting a highly-directive patch transmit / receive antenna. The RSSI image yields the (interpolated) RSSI value at a particular pan-tilt angle (corresponding to a particular x-y pixel). This raw data forms the RSSI feature image,

$$I_{rssi}(x, y) : \mathbb{R}^2 \rightarrow \mathbb{R} \text{ (RSSI)}. \quad (88)$$

All three sensing modalities (camera image, point cloud, and RSSI image) now share a common image-based representation. We have defined three different feature images from the raw data. Before performing sensor fusion, we load the object-specific probabilistic models from the tag-indexed semantic database.

6.2.2.3 Load Object-Specific Probabilistic Feature Models

The goal of our sensor fusion method is to locate a desired object in the sensor data. To accomplish this, we load object-specific models from the semantic database (indexed by the object’s tag ID). The models are probabilistic and describe the likelihood of obtaining the acquired sensor value if the tag is present (or absent) at a particular pixel location. We show how to fuse the distributions and locate the tagged object shortly; right now, we will look at the *form* of the probability distributions stored in the database.

For a feature image I_i , we model each feature F_i as being generated with some probability given by $p_{f_i|tag}(F_i, \text{True})$, if a tag is at the bearing associated with the location. If a tag is not at the bearing associated with the location, we model the probability of a given feature value as $p_{f_i|tag}(F_i, \text{False})$. Our current measurement of the feature value F_i is stored in the feature image I_i and indexed by x, y pixel locations, $I_i(x, y)$. Thus, we equate the distributions,

$$p_{f_i|tag}(F_i, \text{True}) = p_{f_i|tag}(I_i(x, y), \text{True}), \text{ and} \quad (89)$$

$$p_{f_i|tag}(F_i, \text{False}) = p_{f_i|tag}(I_i(x, y), \text{False}) \quad (90)$$

with the conditional distribution property that

$$p_{f_i}(I_i(x, y)) = p_{f_i|tag}(I_i(x, y), \text{True}) + \quad (91)$$

$$p_{f_i|tag}(I_i(x, y), \text{False}).$$

We assume that adjacent pixels are independent. We store the probability distributions from Equations 89-90 on an object-by-object basis in the semantic database. We derive the distributions from data on a per-sensor (per feature image) basis.

Camera Feature Image Distributions

From the camera image, we employed color histograms as the probability distributions. We selected color histograms for their simplicity; again, other visual features could be integrated into this framework and may be more discriminative. For the color histogram, the object probability, $p_{color|tag}(I_{cam}(x, y), \text{True})$, is obtained from one (or more) image of the tagged object stored in the tag-indexed database, as is shown in Figure 106. Meanwhile, the non-object background probability, $p_{color|tag}(I_{cam}(x, y), \text{False})$, is generated from a color histogram accumulated over the set of images of the environment collected during navigation.

Range Feature Image Distribution

No distributions are stored in the semantic database for the range feature. We treat the range feature image as a special case that we will discuss when we infer that tagged object's 2D / 3D location.




Figure 106: The probability distributions for the camera images are represented as color histograms. *Left:* One of the tagged objects used in our sensor fusion experiments. *Right:* The color histogram, corresponding to $p_{color|tag}(I_{cam}(x, y), \text{True})$, for this object.

RSSI Feature Image

The feature from the RSSI image consists of the RSSI value from $I_{rssi}(x, y)$. The associated probabilities, $p_{rssi|tag}(RSSI, \text{True})$ and $p_{rssi|tag}(RSSI, \text{False})$, were obtained as a histogram from 60 hand-labeled ground-truth observations. The ground-truth observations were recorded by selecting the center of the tagged object from the (registered) visual camera image, and selecting all pixels in the RSSI image within a 10-pixel radius. The ground-truth observations were recorded in the environments shown in Figures 110. The resulting RSSI feature image distributions are shown in Figure 107.

Before fusing the sensor data, we load the appropriate object-specific probability distributions from the tag-indexed semantic database.

6.2.2.4 Sensor Fusion: Infer a Tagged Object's 2D Image Location

To recap, the fused image I consists of a set of n feature images $I_0 \dots I_n$, where each feature image I_i represents the spatially varying value of feature F_i . We model each of these features as being




Figure 107: RSSI feature probability distributions determined from 60 hand-labeled training examples. *Left:* $p_{rssi|tag}(RSSI, True)$. *Right:* $p_{rssi|tag}(RSSI, False)$.

generated with some probability $p_{f_i|tag}(F_i, True)$, if a tag is at the bearing associated with the location. If a tag is not at the bearing associated with the location, we model the probability of a given feature value as $p_{f_i|tag}(F_i, False)$. We further model these feature values as being conditionally independent given the presence or absence of the tag at the bearing associated with the location, and as independent from one another. Given these assumptions, we can find the probability that a tag is at a given location using Bayes' rule:

$$p_{tag|f_0...f_n}(V, F_0...F_n) = \frac{p_{f_0...f_n|tag}(F_0...F_n, V)p_{tag}(V)}{p_{f_0...f_n}(F_0...F_n)} \quad (92)$$

$$= \frac{(\prod_{i=1}^n p_{f_i|tag}(F_i, V)) p_{tag}(V)}{\prod_{i=1}^n p_{f_i}(F_i)} \quad (93)$$

$$= p_{tag}(V) \prod_{i=1}^n \frac{p_{f_i|tag}(F_i, V)}{p_{f_i}(F_i)} \quad (94)$$

We assume a uniform prior on the position of each tag, $p_{tag}(V)$. Assuming independence of the feature vectors for each x, y location of the fused image I ,

$$p_{image}(I) = \prod_{i=0}^n p_{f_0...f_n|tag}(I_i(x, y), V(x, y)). \text{ And} \quad (95)$$

$$p_{f_i}(I_i(x, y)) = p_{f_i|tag}(I_i(x, y), True) + p_{f_i|tag}(I_i(x, y), False). \quad (96)$$

The maximum likelihood (ML) estimate of the location of the tag (at pixel location x_{ml} and y_{ml}) is given by,

$$x_{ml}, y_{ml} = \underset{x,y}{\operatorname{argmax}} \left\{ \prod_{i=1}^n \frac{p_{f_i|tag}(I_i(x, y), True)}{p_{f_i}(I_i(x, y))} \right\}. \quad (97)$$

To compute the ML estimate, we load the object-specific probabilistic models from the tag-indexed semantic database and compute (or look up, since our models are histogram distributions)

$$\frac{p_{f_i|tag}(I_i(x, y), \text{True})}{p_{f_i}(I_i(x, y))} = \frac{p_{f_i|tag}(I_i(x, y), \text{True})}{p_{f_i|tag}(I_i(x, y), \text{True}) + p_{f_i|tag}(I_i(x, y), \text{False})} \quad (98)$$

for each feature image I_i . For the range feature image, we employ a special distribution (not in the semantic database),

$$\frac{p_{range|tag}(I_{range}(x, y), V)}{p_{range}(I_{range}(x, y))} = \begin{cases} 1.0 & \text{if } I_{range}(x, y) = 1.0 \\ 0.0 & \text{otherwise.} \end{cases} \quad (99)$$

This distribution treats the point features as a “masking” distribution and ensures that the 2D image pixel selected as the maximum likelihood estimate (x_{ml} and y_{ml}) has a unique 3D point from the point cloud associated with it.

The result of the argmax computation thus yields the most likely location of the tagged object in the 2D image. This process is graphically depicted in Figure 108, with an example using real sensor data shown in Figure 109.

6.2.2.5 Sensor Fusion: Infer a Tagged Object’s 3D Location

Our choice of 3D point features and the “masking” distribution for the range image ensured that the ML estimate (x_{ml} and y_{ml}) has a unique 3D point from the point cloud associated with it. To transform the 2D image location into a 3D location, we simply recall which of the 3D points from the point cloud was transformed to this pixel location. This yields a 3D location which likely belongs to the tagged object.

The result of this five-step multi-sensor fusion process yields a candidate 3D location of the tagged object. This process is summarized in Figure 109. The 3D location can be used by other mobile manipulation algorithms, such as fetching and/or retrieving objects.

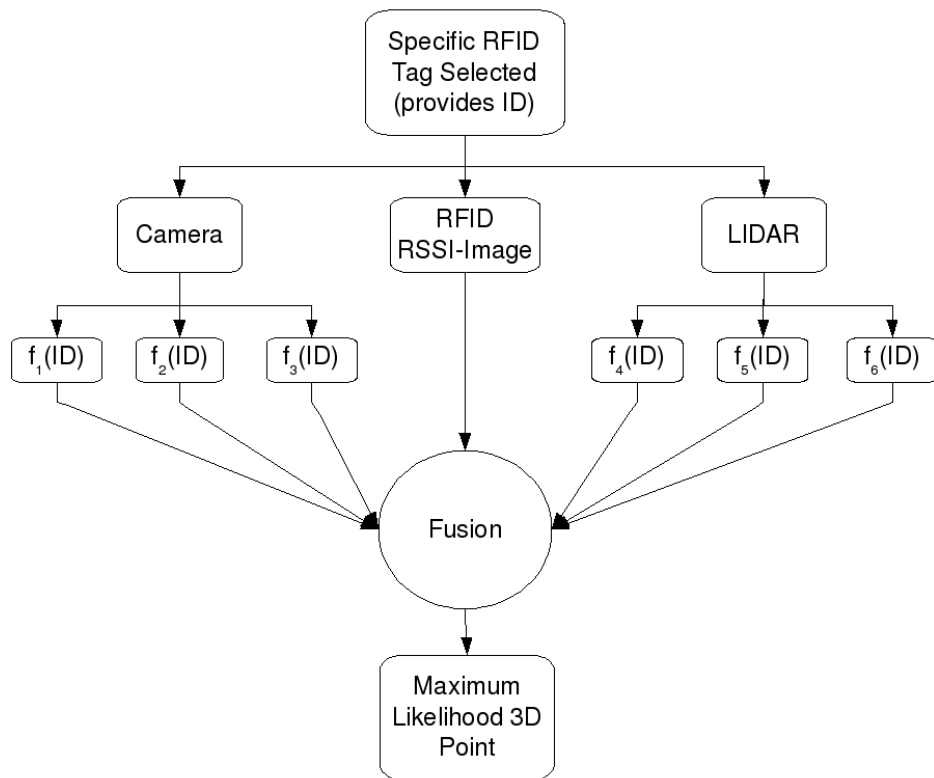


Figure 108: Method for producing a maximum-likelihood 3D point estimate for the location of an RFID-tagged object

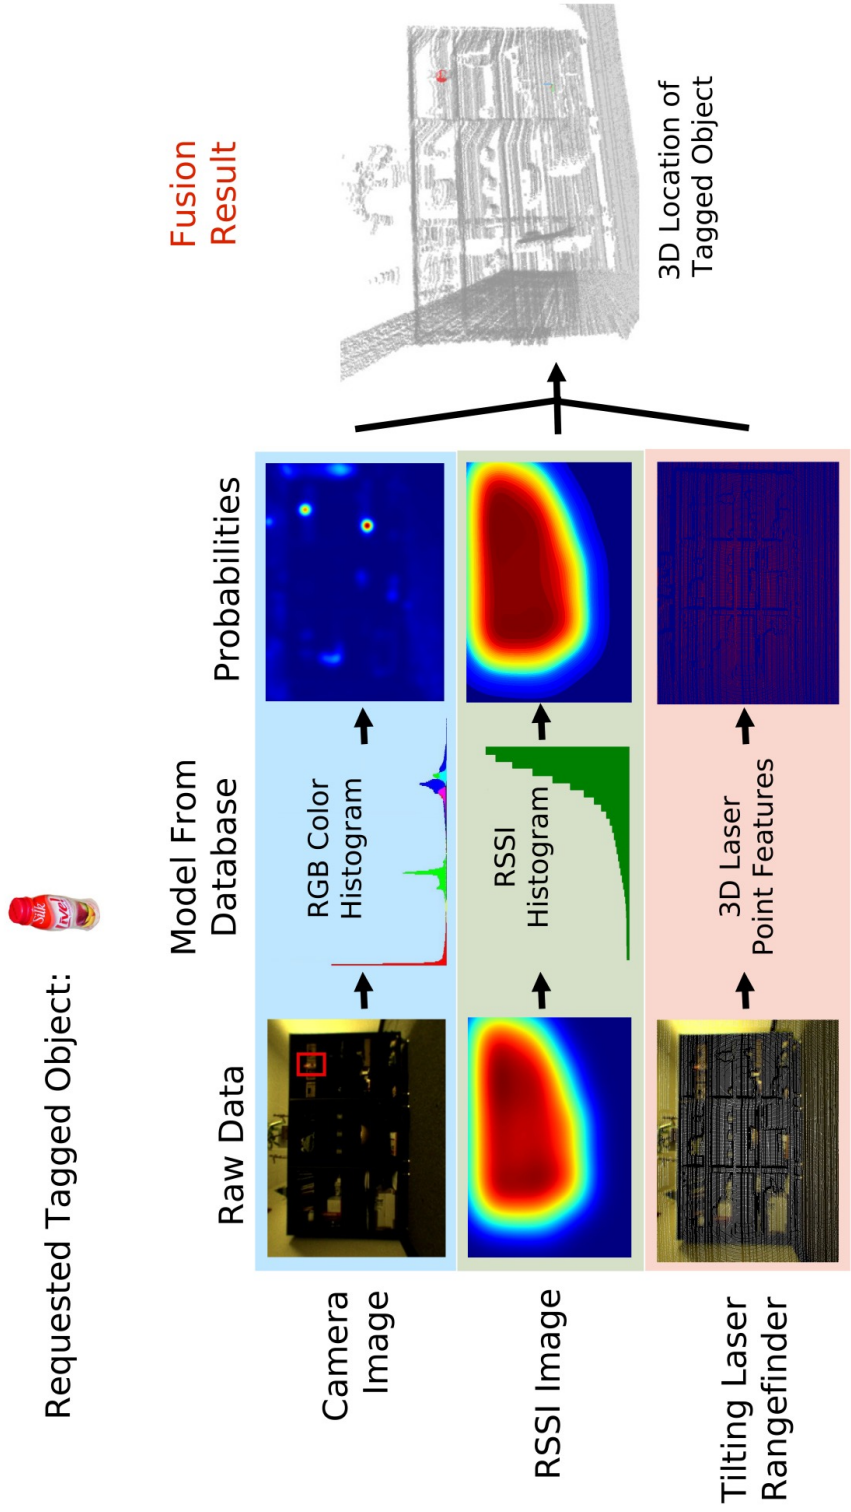


Figure 109: Multi-Sensor Fusion using RSSI Images.

6.2.3 Sensor Fusion Evaluation

We performed a number of tests of the sensor fusion system’s accuracy when estimating tagged object locations in 3D. For our test scenario, we chose three objects with distinct color histograms: a red water bottle, a blue medication box, and an orange disposable beverage bottle. We chose two cluttered but unobstructed scenes and three locations within each scene where each of the three objects was tested, resulting in a total of 18 3D location estimation trials – shown in Figure 110.

The algorithm from Figure 109 was executed for each trial and was deemed successful if the 3D point derived from the fused image belonged to the desired object. The 3D location estimation was successful in 17 of the 18 trials (94.4%), with the only failure occurring for the orange disposable drink bottle due to a nearby orange object that had a similar color histogram. The success rate *without* the RSSI image on the same dataset was 15 of 18 (83.3%); thus, incorporation of the RSSI image resulted in an 11.1% improvement in the system’s performance.

Table 28: Sensor fusion evaluation results: Three tagged objects were placed at each of three labeled locations in two different scenes. The sensor fusion algorithm was executed, yielding a 3D point corresponding to the tagged object.





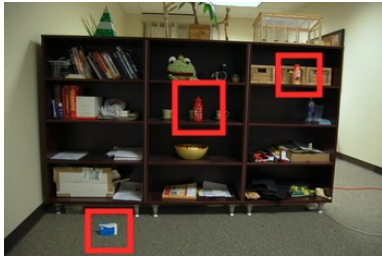
Scenes	Locations	Object to Locate		
				
	1	✓	✓	✓
	2	✓	✓	✓
	3	✓	✓	✓
	1	✓	✓	×
	2	✓	✓	✓
	3	✓	✓	✓



Figure 110: The two scenes for sensor fusion experimentation. The sole failure occurred when the orange disposable bottle was placed in the upper-right location in the bottom scene.

6.2.4 Mobile Manipulation System

Our goal in designing a multi-sensor fusion framework is to facilitate mobile manipulation. Specifically, we designed a mobile manipulation system using our mobile manipulator EL-E (in Figure 104) to fetch a nearby tagged object. First, EL-E uses the RFID antennas to scan the environment for tagged objects in the environment and builds a graphical, context-aware interface for a remote user. The interface is constructed using object names and photos obtained from the semantic database. The remote user selects an object to be grasped. EL-E estimates the bearing to the tag of interest using RFID bearing estimation from Section 4.1.4.1. EL-E rotates to that bearing, which places the tagged object within the field of view of EL-E's other sensors (ie. camera and tilting laser rangefinder). Next, EL-E performs sensor fusion as previously described, which results in a 3D estimate of the object's location. EL-E then uses the 3D estimate of the object's location to approach and grasp the object using an overhead grasp algorithm developed at Georgia Tech's Healthcare Robotics Lab [28, 106, 66]. Finally, after the grasp attempt is completed, short-range RFID antennas in EL-E's end effector verify that the desired tag ID has been grasped, indicating success or failure, which was described in Section 5.

6.2.5 Mobile Manipulation Evaluation

For object approaching and grasping to be successful using our sensor fusion technique, the bearing estimation must terminate with the tagged object in the other sensors' field of view; otherwise, there the sensor fusion (and resulting 3D location) will not correspond to the desired tagged object, and grasping is guaranteed to fail.

We evaluated bearing estimation (as described in Section 4.1.4.1) with the same three objects from the sensor fusion experiments, each positioned in three different locations in a laboratory environment (as shown in Figure 111). Bearing estimation was successful in 8 of 9 trials (88.9%), where success was defined by halting with the desired object within fused image's field of view. In the one failure, the red bottle was located on the chair; bearing estimation yielded a poor direction estimate, so EL-E turned such that the bottle was not in the camera's field of view.

We performed three tests of the entire mobile manipulation system: from user interface generation, bearing estimation, sensor fusion, approaching and grasping, and tagged object verification.




Figure 111: Three different bearing estimation scenarios, with object locations highlighted. Bearing estimation was attempted for all three objects, each in the three different locations (9 total attempts). In 8 of 9 instances, the robot correctly achieved a bearing that placed the tagged object in the fused sensor image.

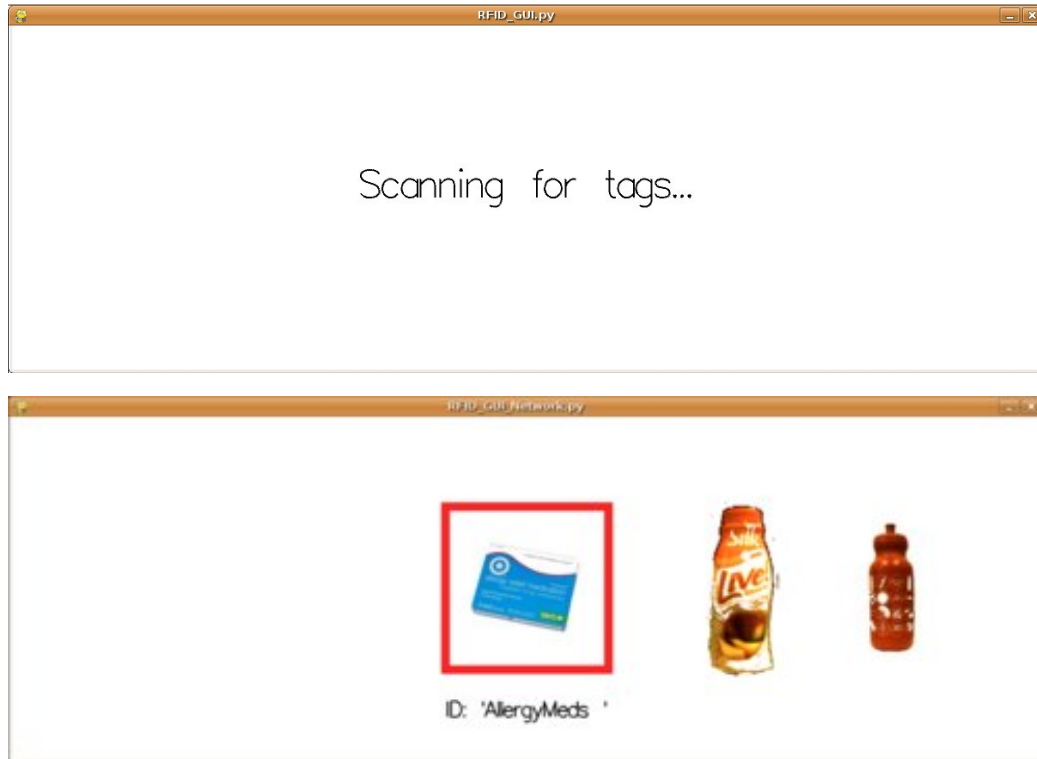


Figure 112: Dynamically generated user interface presenting a menu of tagged objects available to be grasped by the robot.

In all three trials, the robot successfully grasped the correct object and verified the ID of the object post-grasp using the short-range RFID antennas.

6.2.6 Discussion and Conclusions

We have presented an integrated set of methods that enable a mobile manipulator to grasp an object to which a self-adhesive UHF RFID tag has been affixed. Among the contributions made in this section:

- We showed that the presence of a tag’s unique ID provides strong evidence that the tagged object is nearby, which allowed us to build context-aware user interfaces.
- We showed that the relative RFID behaviors from Chapter 4 allow the robot to bring to bear other (complementary) sensing modalities (including RSSI images) to perceive the tagged object. We employed object-specific sensor models from a semantic database indexed by the unique tag ID.

- We developed a framework for fusing data from multiple (complementary) sensing modalities, including RSSI images. This algorithm yields a 3D point belonging to the tagged object.
- We showed that the 3D points output from our sensor fusion framework were sufficient to seed existing mobile manipulation algorithms. We showed that the robot was able to use this 3D location to approach and grasp the object, and then uses its finger-mounted antennas to confirm that the desired object has been grasped.

This work demonstrates that RFID-based perception can potentially be useful in all aspects of a mobile manipulation task: from discovering of what objects are available, to the production of customized user interfaces, to navigating up to tagged objects, manipulation objects, and ultimately verifying that the correct object has been grasped.

6.3 UHF RFID For Medication Delivery and Adherence

Poor medication adherence is a serious problem; many people do not take their medications as prescribed, which undermines the benefits of drug treatments [57]. Research shows that several short-term approaches to medication adherence can be effective, for example: counseling, written information, and personal phone call reminders; however, for long-term medication adherence, no simple intervention, and only some complex ones actually lead to improvements in health outcomes. Yet many approaches have been considered: convenient care, education programs, counseling, reminders, self-monitoring, and family therapy [57]. This is problematic. In fact, it has been suggested that the full benefits of medications cannot be realized at currently achievable levels of adherence; therefore, more studies of innovative approaches to assist patients to follow prescriptions for medications are needed [101].

Like other robotics researchers, we believe that autonomous mobile robots providing timely, in-person reminders could provide both the motivation and the means to improve adherence without significant caregiver or care-receiver burden [144]. In essence, an autonomous mobile robot could be tasked with delivering the *right* medication to the *right* person at the *right* time. Along with the medication, the robot could also provide the means to take the medication (eg. a glass of water) and significant motivation through established psychological bonds (eg. reciprocity) that are commonly developed between social robots and their owners [16].

We believe that UHF RFID sensing is well-matched to the challenges of robotic medication delivery. By tagging medication bottles (as with the FDA E-Pedigree program) and having care-receivers carry UHF RFID tags, the robot can use the optimization-based approaches to RFID search to acquire the medication and then discover, approach, and administer it in a timely fashion. Several properties of UHF RFID make it particularly beneficial in this scenario. First, UHF RFID sensing is not subject to the lighting limitations of cameras; the algorithms could locate recipients equally well in both darkness or lighted rooms. Second, UHF RFID has a low false positive rate – a crucial feature when the consequences of misidentification are dire (as with medication delivery). Coupled with short-range RFID verification and perhaps other forms of identity confirmation (eg. facial recognition), robots can ensure that the *right* medication is given to the *right* person.




Figure 113: EL-E uses UHF RFID behaviors to locate, approach, and deliver medication to a seated recipient.

Deploying and evaluating a robotic system for medication adherence is beyond the scope of this work; however, we did demonstrate *two* systems for UHF RFID medication delivery using the local search methods from Chapter 4: one using EL-E and one using the PR2. In this section, we describe the details of our implementation – the first of its kind to use UHF RFID for medication delivery. Georgia Tech’s Human Factors and Aging Lab is currently testing our UHF RFID medication delivery implementation with older adults to gauge their the robot’s perceived usefulness and perceived ease of use. The results of this study are forthcoming.

6.3.1 UHF RFID Medication Delivery on the EL-E Platform

We developed a system whereby EL-E could deliver a tagged object to a person wearing a UHF RFID tag on his or her wrist – a setup akin to a hospital’s identification wrist band. Using the version of EL-E shown in Figure 94, we are able to deliver medication to a seated recipient as shown in Figure 113).

For this implementation of medication delivery, we assume EL-E is already nearby the intended recipient (ie. as though global search has already been carried out). EL-E executes a series of local,

optimization-based RFID search behaviors to more precisely locate the person, approach, and then hand off the object in its gripper. First, EL-E performs bearing estimation (Section 4.1) followed by RFID servoing (Section 4.2), stopping when the downward-facing laser rangefinder indicates an impeding obstacle. At this time, EL-E reaches out its hand to a fixed position and monitors the forces and torques at the base of its fingers, waiting for the recipient to reach out and grab the object. When a fixed force-torque threshold is exceeded, EL-E releases the object. Before departing, EL-E asks the user to (optionally) place a new object in its gripper. If EL-E receives an object, the robot uses its short-range in-hand UHF RFID reader to ascertain the identity of the grasped object (Chapter 5). At this point, no further action is taken; rather, this is a stub into future behaviors where EL-E could perform additional tasks.

To evaluate EL-E's ability to deliver tagged objects to a tagged person, we performed 10 delivery trials (with ourselves as subjects) in a laboratory environment. Each trial began with EL-E positioned approximately 2 meters away from the recipient, with one of two possible tagged objects in the robot's hand (a TV remote or medication bottle, split equally among trials). For a trial to be deemed successful, EL-E needed to correctly identify the object initially in its hand, approach the recipient, and deliver the tagged object in such a way that the recipient could remain seated. Furthermore, in all 10 trials the recipient placed an object back in EL-E's hand. In half of the trials, recipients simply replaced the delivered object; in the other half, recipients handed EL-E the alternate object. Detecting the identity of this object was also required for a trial to be deemed successful. EL-E succeeded in all 10 trials.

6.3.2 UHF RFID Medication Delivery on the PR2

We ported the medication delivery system from EL-E to the PR2. Georgia Tech's Human Factors and Aging (HFA) Lab is currently using our UHF RFID medication delivery implementation with older adults in Georgia Tech's Aware Home (Figure 115) to study the robot's perceived usefulness and perceived ease of use, as well as peoples' impressions of robot medication delivery.

For the HFA Lab studies, the robot is tasked with delivering a tagged medication bottle to a recipient wearing a tagged neck-worn lanyard. By design, we assume that the recipient is seated in the Aware Home's living room. The medication delivery algorithm is preceded by a navigation

step, where the PR2 moves from any starting location in the Aware Home to the center of the living room⁴. The navigation is performed using a variant of FastSLAM localization coupled with an A* global planner and dynamic window local planner afforded by the open source Robot Operating System (ROS) navigation stack [99]. From this vantage in the center of the room, we assume that the intended recipient is nearby, as though global RFID search had already been performed. The PR2 then executes the same medication delivery steps as EL-E. First, the robot uses local, optimization-based RFID search behaviors to more precisely locate the recipient through bearing estimation and approach them using RFID servoing. Then the PR2 reaches out its hand (with the medication bottle) to a fixed position and monitors the values of its finger-mounted tactile sensors. When the recipient grasps the medication bottle and the tactile sensor values exceed a threshold, the robot opens its gripper and releases the object, completing the delivery process. This process is depicted in Figure 114. It is worth noting that the PR2 UHF RFID system does not possess a short-range in-hand reader (as does EL-E). Thus, some other method may be required for short-range identification of recipients and / or grasped objects. Currently, no short range verification exists for the PR2; this remains an issue for future refinement.

Before progressing to the (now in progress) user studies, we tested this system using ourselves as subjects. We performed 10 trials, with random initial robot positions. In all 10 trials, the robot was able to successfully navigate to the living room's center and use the local optimization-based UHF RFID behaviors to deliver the medication to a seated recipient. More informally, this demonstration has been run dozens of times with dozens of recipients and in the presence of many spectators. Anecdotally, the demonstration (as formulated) seems to be quite robust.

Thus far the HFA Lab has processed 14 study participants. With one participant, the robot was not functioning due to a high-level system failure, thus medication delivery was not attempted; this failure was beyond the purview of the medication delivery demonstration. Of the remaining 13 participants, the medication delivery demonstration was successful in 12 / 13 trials. In Figure 115, we can see an older adult from one of the successful attempts interacting with the PR2. In the one medication delivery failure, the robot stopped prematurely (3 feet from the intended recipient),

⁴The PR2 cannot safely and autonomously navigate through doors in the Aware Home, as the clearance is less than 2cm on either side of the robot. Thus, the starting locations were always selected so as to avoid intervening doors.




Figure 114: To deliver medication to a user seated in the Aware Home’s living room, the PR2 first navigates (top left) into the center of the room (top right). Then, the robot uses local RFID search behaviors to more precisely locate the recipient (middle left), approach them using its mobile base (middle right), stop when forward motion is impeded (bottom left), and then hand out and deliver the medication to the recipient (bottom right).

and reached out to hand off the medication. Of course, this was much too far away to support comfortable delivery to a seated user. We believe this latter case occurred due to a hallucinated obstacle in the robot's sensor data, which the robot detected and thus prematurely halted forward motion to prevent a collision. Since the robot never registered any changes in its tactile sensor measurements during the hand-off interval (where a user would normally grasp the object being held in the robot's manipulator), the robot never actually released the object. Though it is not currently supported in the code base, it is conceivable that the robot could have autonomously detected this failure condition and re-executed the medication delivery process. However, the study administrators simply halted and re-executed the demonstration manually; it completed successfully on the second try.

The detailed results and analysis from the HFA Lab's study (including the results for UHF RFID medication delivery using the PR2) are forthcoming.



Figure 115: The Human Factors and Aging Lab at Georgia Tech is currently running user studies to understand how older adults perceive personal robots. UHF RFID medication delivery is one of the demonstrations and technologies being examined. In this demonstration (not actually part of the user study), the PR2 has successfully demonstrated UHF RFID medication delivery to an older adult seated in Georgia Tech's Aware Home. [Photo used with user permission.]

6.3.3 Discussion and Conclusions

We believe that autonomous mobile robots may be a compelling technology to improve medication adherence, and that this topic is worthy of further study. By delivering the *right* medication to the *right* person at the *right* time, autonomous mobile robots can provide the means and motivation to properly administer medication, all while reducing caregiver and/or care-receiver burden. Further, we believe that UHF RFID sensing is well-matched to the challenges of robotic medication delivery. By tagging medication bottles (as with the FDA E-Pedigree program) and having care-receivers carry UHF RFID tags, the robot can use the optimization-based approaches to RFID search to acquire medication and then discover, approach, and administer it in a timely fashion. Several properties of UHF RFID make it particularly beneficial in this scenario. First, UHF RFID sensing is not subject to the lighting limitations of cameras; the algorithms could locate recipients from a distance equally-well in both darkness or lighted rooms. Second, UHF RFID has a low false positive rate – a crucial feature when the consequences of misidentification are dire (as with medication delivery). Coupled with short-range RFID verification and perhaps other forms of identity confirmation (eg. facial recognition), robots can ensure that the *right* medication is given to the *right* person at the *right* time.

We tested UHF RFID medication delivery using two different robot platforms (EL-E and the PR2). On EL-E, our tests were successful in 10 / 10 trials in a laboratory environment. More recently, the PR2 medication delivery was tested in a realistic home environment with older adults, where it was successful in 12/13 trials.

6.4 Conclusions

In this chapter, we explored several complete mobile manipulation systems that utilize UHF RFID sensing, and made several key contributions:

Physical, Perceptual, and Semantic (PPS) Tags:

We developed a new type of tag (dubbed a PPS-tag) that combines a UHF RFID tag with additional forms of augmentation (eg. compliant materials with visually-distinct properties) to provide physical, perceptual, and semantic assistance to robots. We presented five exemplar PPS-tags along with a set of robotic behaviors that utilize UHF RFID perception and

optimization-based RFID behaviors (eg. bearing estimation and RFID servoing) to accomplish various tasks, such as: turning on and off light switches, opening and closing drawers, operating lamps, and disposing of trash in a waste bin.

Multi-Sensor Fusion and Mobile Manipulation Using RSSI Images:

We developed a framework for multi-sensor fusion that eliminated the need for the physical and perceptual augmentation beyond just the UHF RFID tag. We employed RFID perception and optimization-based RFID behaviors to locate and approach the tagged objects, which terminated with the tagged object close-to and in front of the robot. This allowed us to bring to bear additional sensing modalities (eg. cameras and laser rangefinders that produce 3D point clouds) to perceive the tagged object. We used the UHF RFID tag's unique identifier as an index into a semantic database, where we stored and retrieved information about the tagged object's appearance (as opposed to the PPS-tag's appearance). Finally, we demonstrated a complete mobile manipulation system that is capable of locating, approaching, perceiving, and grasping tagged objects in relative isolation on the floor.

UHF RFID for Medication Delivery and Adherence:

We applied RFID perception and optimization-based RFID behaviors to medication delivery. We developed a system wherein an autonomous mobile manipulator confirmed the identity of a grasped tagged medication bottle, used optimization-based RFID behaviors to locate and approach the intended (tagged) recipient, and then handed off the medication to a seated user. We tested UHF RFID medication delivery using two different robot platforms (EL-E and the PR2). On EL-E, our tests were successful in 10 / 10 trials in a laboratory environment. More recently, the PR2 medication delivery was tested in a realistic home environment with older adults, where it was successful in 12/13 trials.

CHAPTER VII

CONCLUSIONS

Personal robots with mobility, autonomy, and manipulation capabilities have the potential to dramatically improve quality of life for various user populations, such as older adults, people suffering from disabilities, and ultimately *you* or *me*. Unfortunately, unstructured human environments pose myriad challenges that hinder home robot deployment. In this thesis we addressed some of these challenges through a new robotic sensing modality that leverages a small amount of environmental augmentation in the form of low-cost, passive Ultra High Frequency (UHF) Radio-Frequency Identification (RFID) tags. Specifically, we explored robot behaviors and radio frequency perception techniques using robot-mounted UHF RFID readers, allowing an autonomous mobile manipulator (AMM, a robot with manipulation capabilities) to discover, locate, and interact with UHF RFID tags applied to objects, people, or locations of interest. Unlike UHF RFID sensing techniques commonly employed in factory settings, where readers are statically fixed in the environment, the behaviors and algorithms we explored in this thesis explicitly relied on the robot’s ability to change its configuration or that of the world through mobility or manipulation, thereby providing multiple opportunistic views of the RF landscape.

We employed commercial UHF RFID readers that can sense tags beyond *six* meters under idealistic conditions; however, they do *not* provide precise tag location information. Rather, they provide a binary indication of tag presence or absence (a detection); in the event of a positive tag detection they also provide a receive signal strength indicator (RSSI) measurement, a scalar value that indicates the strength of the tag’s response. We develop a RFID sensor model from first principles based on the Friis transmission equation, which is commonly used in the RF community to elucidate system design considerations under ideal, free-space operation. It is well-known that the Friis model is a crude approximation to real-world radar system performance [136]; however, we use insights gleaned from the Friis model to describe the pertinent relationships between reader-tag properties (eg. relative pose, antenna properties, and environmental conditions) and expected RF sensor values

(tag detection and RSSI). Extensive empirical testing confirms that the sensor system is stochastic in nature, where tags may or may not be detected due to a variety of factors: stochasticity purposefully introduced in the reader-tag communication protocols for anti-collision, changing environmental influences, and on-board energy storage that affects a tag's ability to power-up. Due to the stochastic nature of the RFID measurements, researchers (ourselves included) have developed extensive data-driven, probabilistic sensor models to encapsulate some of the variability experienced in real-world operation. Using state-of-the-art *probabilistic* techniques such as Bayesian filters (eg. particle filters), researchers can use the data-driven models to localize tags. These methods have proven fruitful when we can control the placement of tags (nearby material properties and tag orientation), which is useful for infrastructure tags that remain fixed in the environment. These techniques use a robot motion model and a RF sensor model that considers readings one at a time (ie. in isolation with an *absolute* measurement of RSSI) to localize RFID tags.

However, empirical testing also suggests a complication. The simple act of tagging an object can dramatically affect the received signals by altering the tag antenna properties or local RF propagation. For example, the object's material properties (eg. metal) and the tag mounting (eg. wrapped around an object) can both dramatically affect the received signals. Furthermore, the object's position and orientation can also dramatically affect the received signals; when objects are relocated, the tag's orientation changes as does the surrounding environment (and its RF properties). Combined, these two factors may cause substantive deviations from the ideal Friis model and/or data-driven probabilistic models. To compensate, the models may require *significant* alterations to retain their utility; the alterations can be non-trivial, unknown apriori, change over time, or be intractable to sense *in situ*. We show that the data-driven probabilistic models undergo substantially changes due to these factors, and that this has an adverse affect on tagged object localization in home environments.

Instead, we develop a series of optimization-based UHF RFID behaviors that glean insights from the ideal Friis sensor model that use successive *relative* RFID measurements during specially-crafted antenna motions to estimate derived quantities of interest rather than trying to estimate an exact tag location. For example, by panning a highly-directive RFID reader antenna through various angles, we can estimate the heading (azimuth angle) toward a tag – it will be the direction that

yields the maximum RSSI reading. Other techniques include panning and tilting to estimate both azimuth and elevation angles, RFID servoing, and RFID-based search. All of these techniques draw inspiration from the early radar literature, as well as the literature on local and global optimization strategies. When using patch antennas with unimodal radiation patterns, we unify these algorithms under a common optimization framework and evaluate the capabilities and limitations of these new algorithms. The optimization-based UHF RFID behaviors for locating and approaching tags offer several key benefits: they are easy to implement; easy to generalize; do not require training data; and when compared to state-of-the-art probabilistic techniques for object localization, they achieve comparable performance in positioning the robot near the tagged object and superior performance in orienting towards the tagged object.

We develop *two* novel UHF RFID systems mounted-on or embedded-in a mobile robot's manipulator that operate over short range using the *same* tags as used for long-range operation. These systems operate on magnetostatic rather than electromagnetic coupling, which is an atypical mode of operation for UHF RFID tags. The first system is mounted to the robot's wrist; we show that this system can be used during manipulation to selectively grasp the correct tagged object even in the presence of other (visually-identical) objects using optimization techniques analogous to those developed for long-range operation. The other system is embedded in the robot's fingers; we refer to this as the *in-hand* RFID reader. We evaluate the in-hand reader's ability to detect a variety of tags applied to a diverse set of objects being held in the robot's manipulator. The results demonstrate that the system is capable of detecting a wide variety of tags, but that care must be taken to select a tag appropriately matched to an object's material composition.

This thesis also examines three complete mobile manipulation systems that utilize both long-range and short-range optimization-based UHF RFID behaviors.

1. We developed a new type of tag (dubbed a PPS-tag) that combines a UHF RFID tag with additional forms of augmentation (eg. compliant materials with visually-distinct properties) to provide physical, perceptual, and semantic assistance to robots. We presented five exemplar PPS-tags along with a set of robotic behaviors that utilize UHF RFID perception and optimization-based RFID behaviors (eg. bearing estimation and RFID servoing) to accomplish various tasks, such as: turning on and off light switches, opening and closing drawers,

operating lamps, and disposing of trash in a waste bin.

2. We developed a framework for multi-sensor fusion that eliminated the need for the physical and perceptual augmentation beyond just the UHF RFID tag. We employed RFID perception and optimization-based RFID behaviors to locate and approach the tagged objects, which terminated with the tagged object near and in front of the robot. This allowed us to bring to bear additional sensing modalities (eg. cameras and laser rangefinders that produce 3D point clouds) to perceive the tagged object. We used the UHF RFID tag's unique identifier as an index into a semantic database, where we stored and retrieved information about the tagged object's appearance (as opposed to the PPS-tag's appearance). Finally, we demonstrated a complete mobile manipulation system that is capable of locating, approaching, perceiving, and grasping tagged objects in relative isolation on the floor.
3. We applied RFID perception and optimization-based RFID behaviors to medication delivery. We developed a system wherein an autonomous mobile manipulator confirmed the identity of a grasped tagged medication bottle, used optimization-based RFID behaviors to locate and approach the intended (tagged) recipient, and then handed off the medication to a seated user. We tested UHF RFID medication delivery using two different robot platforms (EL-E and the PR2). On EL-E, our tests were successful in 10 / 10 trials in a laboratory environment. More recently, the PR2 medication delivery was tested in a realistic home environment with older adults, where it was successful in 12/13 trials.

In this thesis we show that autonomous mobile manipulators can leverage this modest form of environmental augmentation using short-range and long-range UHF RFID perception and special robot behaviors to facilitate navigation and manipulation – an early approximation to this scenario:

During a trip to a department store, an individual purchases a robot and a box of standard labels: “dish”, “dish washer”, “clothing”, “washing machine”, “toy”, “litter box”, “scoop”, and “storage bin”. The individual returns home, applies the labels as directed, unboxes the robot, and turns it on. The robot connects to the web, downloads the appropriate behaviors, and is instantly able to operate in the labeled world: loading

the dishwasher with labeled dishes, putting away labeled toys, cleaning the cats' litter box, and washing labeled clothing. Improved or new functionality, such as delivering medicine, is just a few labels and an internet connection away.

Of course, this is just an interim vision; ultimately we would like robots that are capable of robustly operating in *unaugmented* home environments. We believe that RFID-assisted robots may provide sufficiently-compelling applications (economics) to bootstrap a personal robot revolution, paving the way for unaugmented operation.

REFERENCES

- [1] “Helping hands: Monkey helpers for the disabled inc.,” in <http://www.helpinghandsmonkeys.org/>, December 2006.
- [2] “National spinal cord injury statistical center (nscisc), supported by the national institute on disability and rehabilitation research,” in <http://www.spinalcord.uab.edu/>, 2007.
- [3] “Performance analysis of alien squiggle tag in human presence,” in *Antennas and Propagation Society International Symposium (APSURSI), 2010 IEEE*, pp. 1–4, July 2010.
- [4] ABRAMSON, S. and LEVI, E., “System and method for determining the location of a machine,” Sept. 26 2007. US Patent App. 11/861,978.
- [5] AFRICANO, D., BERG, S., LINDBERGH, K., LUNDHOLM, P., NILBRINK, F., and PERS-SON, A., “Designing tangible interfaces for children’s collaboration,” in *CHI’04 extended abstracts on Human factors in computing systems*, pp. 853–868, ACM, 2004.
- [6] AGRANOFF, A., GODBOUT, C., and JOHNS, J., “Stroke motor impairment,” in <http://www.emedicine.com/pmr/topic189.htm>, 2007.
- [7] ALLEN, K. and BLASCOVICH, J., “The value of service dogs for people with severe ambulatory disabilities. a randomized controlled trial.,” *Journal of the Americal Medical Association*, vol. 275, no. 13, 1996.
- [8] ANGERER, C., LANGWIESER, R., and RUPP, M., “Direction of arrival estimation by phased arrays in rfid,” in *Workshop on RFID Technology*, 2010.
- [9] ARUMUGAM, D. and ENGELS, D., “Characteristics of passive uhf rfid tags on liquids,” in *Antennas and Propagation Society International Symposium, 2009. APSURSI’09. IEEE*, pp. 1–4, IEEE.
- [10] ARUMUGAM, D. and ENGELS, D., “Characteristics of passive uhf rfid tags on metal slabs,” in *Antennas and Propagation Society International Symposium, 2009. APSURSI’09. IEEE*, pp. 1–4, IEEE.
- [11] ASFOUR, T., AZAD, P., VAHRENKAMP, N., REGENSTEIN, K., BIERBAUM, A., WELKE, K., SCHRÖDER, J., and DILLMANN, R., “Toward humanoid manipulation in human-centred environments,” *Robotics and Autonomous Systems*, 2007.
- [12] AUFRERE, R., MERTZ, C., and THORPE, C., “Multiple sensor fusion for detecting location of curbs, walls, and barriers,” in *Intelligent Vehicles Symposium, 2003. Proceedings. IEEE*, pp. 126–131, IEEE, 2003.
- [13] BAEG, S.-H., PARK, J.-H., KOH, J., PARK, K.-W., and BAEG, M.-H., “Robomaidhome: A sensor network-based smart home environment for service robots,” in *16th IEEE International Conference on Robot and Human Interactive Communication*, 2007.

- [14] BALANIS, C. A., *Antenna Theory: Analysis and Design*. Wiley and Sons, third ed., 2005. pp. 816–843.
- [15] BAMBHA, N., BHATTACHARYYA, S., TEICH, J., and ZITZLER, E., “Hybrid global/local search strategies for dynamic voltage scaling in embedded multiprocessors,” in *Hardware/Software Codesign, 2001. CODES 2001. Proceedings of the Ninth International Symposium on*, pp. 243–248, IEEE, 2001.
- [16] BECKER, B., “Social robots—emotional agents: Some remarks on naturalizing man-machine interaction,” *Ethics in Robotics*, vol. 6, p. 37, 2006.
- [17] BECKERS, R., HOLLAND, O., and DENEUBOURG, J., “From local actions to global tasks: Stigmergy and collective robotics,” in *Artificial life IV*, vol. 181, p. 189, Citeseer, 1994.
- [18] BEETZ, M., STULP, F., RADIG, B., BANDOUC, J., BLODOW, N., DOLHA, M. E., FEDRIZZI, A., JAIN, D., KLANK, U., KRESSE, I., MALDONADO, A., MARTON, Z.-C., MSENLECHNER, L., RUIZ, F., RUSU, R. B., and TENORTH, M., “The Assistive Kitchen – A Demonstration Scenario for Cognitive Technical Systems,” in *International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2008.
- [19] BERTSEKAS, D., *Nonlinear programming*. Athena Scientific Belmont, MA, 1999.
- [20] BLOSS, R., “Mobile hospital robots cure numerous logistic needs,” *Industrial Robot: An International Journal*, vol. 38, no. 6, pp. 567–571, 2011.
- [21] BLUETHMANN, W., AMBROSE, R., FAGG, A., ROSENSTEIN, M., PLATT, R., GRUPEN, R., BREZEAL, C., BROOKS, A., LOCKERD, A., PETERS, R., JENKINS, O., MATARIC, M., and BUGAJSKA, M., “Building an autonomous humanoid tool user,” in *International Conference on Humanoid Robots*, (Santa Monica, Los Angeles, CA, USA.), IEEE Press, 2004.
- [22] BOHN, J. and MATTERN, F., “Super-distributed RFID Tag Infrastructures,” *LECTURE NOTES IN COMPUTER SCIENCE*, pp. 1–12, 2004.
- [23] BRAITENBERG, V., *Vehicles: Experiments in synthetic psychology*. The MIT press, 1986.
- [24] BROOKER, G., “Conical-scan antennas for w-band radar systems,” in *Radar Conference, 2003. Proceedings of the International*, pp. 406–411, IEEE, 2003.
- [25] BROOKS, R., ARYANANDA, L., EDSINGER, A., FITZPATRICK, P., KEMP, C., O’REILLY, U.-M., TORRES-JARA, E., VARSHAVSKAYA, P., and WEBER, J., “Sensing and manipulating built-for-human environments,” *International Journal of Humanoid Robotics*, 2004.
- [26] BRUCE, J., BALCH, T., and VELOSO, M., “Fast and Inexpensive Color Image Segmentation for Interactive Robots,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2000.
- [27] BUSTILLO, M., “Wal-mart radio tags to track clothing,” *Wall Street Journal*, p. A1, 2010.
- [28] CHOI, Y. S., ANDERSON, C. D., GLASS, J. D., and KEMP, C. C., “Laser pointers and a touch screen: Intuitive interfaces to an autonomous mobile robot for the motor impaired,” in *ACM SIGACCESS*, 2008.

- [29] CHOI, Y. S., DEYLE, T., CHEN, T., GLASS, J. D., and KEMP, C. C., “Benchmarking assistive mobile manipulation: A list of household objects for robotic retrieval prioritized by people with als,” in *International Conference on Rehabilitation Robotics (ICORR)*, 2009.
- [30] CONNELL, J., “A behavior-based arm controller,” *IEEE Transactions on Robotics and Automation*, vol. 5, pp. 784–791, December 1989.
- [31] CUDRNAK, J. A., *RFID Modeling, Characterization, and Optimization*. PhD thesis, California Polytechnic State University, 2009.
- [32] CUI, J., ZHA, H., ZHAO, H., and SHIBASAKI, R., “Multi-modal tracking of people using laser scanners and video camera,” *Image and Vision Computing*, vol. 26, no. 2, pp. 240–252, 2008.
- [33] DALLAWAY, J. and JACKSON, R., “RAID: A vocational robotic workstation,” in *Proceedings of the 1992 International Conference on Rehabilitation Robotics. Keele University (England)*, 1992.
- [34] DALLAWAY, J., JACKSON, R., and TIMMERS, P., “Rehabilitation robotics in europe,” *Rehabilitation Engineering, IEEE Transactions on*, vol. 3, pp. 35–45, Mar 1995.
- [35] DARIO, P., “MOVAID: a personal robot in everyday life of disabled and elderly people,” *Technology and Disability*, vol. 10, no. 2, pp. 77–93, 1999.
- [36] DEAVOURS, D., RAMAKRISHNAN, K., and SYED, A., “Rfid performance tag analysis,” tech. rep., Technical Report, Information and Telecommunication Technology Center, The University of Kansas, 2005.
- [37] DEYLE, T., ANDERSON, C., KEMP, C. C., and REYNOLDS, M. S., “A foveated passive UHF RFID system for mobile manipulation,” *Intelligent Robots and Systems, 2008 (IROS 2008) IEEE/RSJ International Conference on*, pp. 3711–3716.
- [38] DEYLE, T., KEMP, C. C., and REYNOLDS, M. S., “Probabilistic UHF RFID tag pose estimation with multiple antennas and a multipath RF propagation model,” *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*, pp. 1379–1384, 2008.
- [39] DEYLE, T., NGUYEN, H., REYNOLDS, M., and KEMP, C. C., “RF Vision: RFID Receive Signal Strength Indicator (RSSI) Images for Sensor Fusion and Mobile Manipulation,” *Intelligent Robots and Systems, 2009 (IROS 2009) IEEE/RSJ International Conference on*.
- [40] DEYLE, T., NGUYEN, H., REYNOLDS, M., and KEMP, C., “RFID-Guided Robots for Pervasive Automation,” *IEEE Pervasive Computing*, vol. 9, pp. 37–45, 2010.
- [41] DIMA, C., VANDAPPEL, N., and HEBERT, M., “Sensor and classifier fusion for outdoor obstacle detection: an application of data fusion to autonomous off-road navigation,” 2003.
- [42] DIMA, C., VANDAPPEL, N., and HEBERT, M., “Classifier fusion for outdoor obstacle detection,” in *Robotics and Automation, 2004. Proceedings. ICRA’04. 2004 IEEE International Conference on*, vol. 1, pp. 665–671, IEEE, 2004.
- [43] DOLLAR, A. M. and HOWE, R. D., “Towards grasping in unstructured environments: Grasper compliance and configuraton optimization,” *Advanced Robotics*, vol. 19, no. 5, pp. 523–543, 2005.

- [44] EPC GLOBAL US INC., “Class 1 Generation 2 UHF RFID protocol for operation at 860MHz-960MHz, version 1.0.9,” January 2005. Available online at: <http://www.epcglobalus.org/>.
- [45] ERNST, H., *MH-1, a computer-operated mechanical hand*. PhD thesis, Massachusetts Institute of Technology, 1962.
- [46] FIALA, M., “Vision guided control of multiple robots,” 2004.
- [47] FINKENZELLER, K., *RFID handbook*. Wiley Hoboken, NJ, 2003.
- [48] FINKENZELLER, K. and MÜLLER, D., *RFID Handbook: Fundamentals and Applications in Contactless Smart Cards, Radio Frequency Identification and Near-Field Communication*. Wiley, 2010.
- [49] FISHKIN, K., PHILIPPOSE, M., REA, A., RES, I., and SEATTLE, W., “Hands-on RFID: wireless wearables for detecting use of objects,” in *Wearable Computers, 2005. Proceedings. Ninth IEEE International Symposium on*, pp. 38–41, 2005.
- [50] FRIEDMAN, D., DOSHER, J., KOWALEWSKI, T., ROSEN, J., and HANNAFORD, B., “Automated tool handling for the trauma pod surgical robot,” in *Robotics and Automation, 2007 IEEE International Conference on*, pp. 1936–1941, IEEE.
- [51] GOLDBERG, D. and VOESSNER, S., “Optimizing global-local search hybrids,” *Urbana*, vol. 51, p. 61801, 1999.
- [52] GRAF, B., HANS, M., and SCHRAFT, R. D., “Care-o-bot ii - development of a next generation robotic home assistant,” *Autonomous Robots*, vol. 16, no. 2, pp. 193–205, 2004.
- [53] HA, Y.-G., SOHN, J.-C., and CHO, Y.-J., “Service-oriented integration of networked robots with ubiquitous sensors and devices using the semantic web services technology,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2005.
- [54] HAHNEL, D., BURGARD, W., FOX, D., FISHKIN, K., and PHILIPPOSE, M., “Mapping and localization with rfid technology,” in *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, vol. 1, pp. 1015–1020, IEEE, 2004.
- [55] HAHNEL, D., BURGARD, W., FOX, D., FISHKIN, K., and PHILIPPOSE, M., “Mapping and localization with RFID technology,” in *Proceedings of IEEE International Conference on Robotics and Automation*, vol. 1, pp. 1015–1020, 26 April-1 May 2004.
- [56] HAIGH, K. Z. and YANCO, H. A., “Automation as caregiver: A survey of issues and technologies,” *AAAI-02 Workshop on Automation as Caregiver: The Role of Intelligent Technology in Elder Care*, pp. 39–53, 2002.
- [57] HAYNES, R., YAO, X., DEGANI, A., KRIPALANI, S., GARG, A., and McDONALD, H., “Interventions for enhancing medication adherence (review),” *Cochrane Database of Systematic Reviews*, 2005.
- [58] HENRY, P., KRAININ, M., HERBST, E., REN, X., and FOX, D., “Rgb-d mapping: Using depth cameras for dense 3d modeling of indoor environments,” in *the 12th International Symposium on Experimental Robotics (ISER)*, 2010.

- [59] HERNANDEZ, M., ECHEVERRY, S., GUERRA, R., KIM, D., HAMILTON, C., , and BEHAL, A., “User-guided vision-based control of a wheelchair mounted assistive robotic arm in an unstructured environment,” in *RESNA*, (Arlington, VA), May 2008.
- [60] HILLENBRAND, U., OTT, C., BRUNNER, B., BORST, C., and HIRZINGER, G., “Towards service robots for the human environment: the Robutler,” in *Proceedings Mechatronics & Robotics (MechRob)*, pp. 1497–1502, 2004.
- [61] HODGES, S., THORNE, A., MALLINSON, H., and FLOERKEMEIER, C., “Assessing and optimizing the range of uhf rfid to enable real-world pervasive computing applications,” *Pervasive Computing*, pp. 280–297, 2007.
- [62] HSU, J., 2011. PR2 Simulator (pr2_simulator) ROS Package. Willow Garage, Robot Operating System.
- [63] HUTCHINSON, C., *The ARRL handbook for radio amateurs*. Amer Radio Relay League, 2000.
- [64] HUTCHINSON, C., “The kalman filter applied to aerospace and electronic systems,” *Aerospace and Electronic Systems, IEEE Transactions on*, no. 4, pp. 500–504, 1984.
- [65] INAMURA, T., OKADA, K., TOKUTSU, S., HATAO, N., INABA, M., and INOUE, H., “HRP-2W: A humanoid platform for research on support behavior in daily life environments,” *Robotics and Autonomous Systems*, 2008.
- [66] JAIN, A. and KEMP, C. C., “Behavior-based door opening with equilibrium point control,” in *RSS Workshop: Mobile Manipulation in Human Environments*, 2009.
- [67] JAIN, A. and KEMP, C. C., “EL-E: An Assistive Mobile Manipulator that Autonomously Fetches Objects from Flat Surfaces,” *Autonomous Robots*, 2010.
- [68] JANG, M., SOHN, J.-C., and CHO, Y., “Building semantic robot space based on the semantic web,” in *IEEE International Conference on Robot and Human Interactive Communication*, 2007.
- [69] JOHO, D., PLAGEMANN, C., and BURGARD, W., “Modeling RFID signal strength and tag detection for localization and mapping,” in *Proceedings of the IEEE international conference on robotics and automation (ICRA), Kobe, Japan*, 2009.
- [70] KAKUMANUA, P., MAKROGIANNISA, S., and BOURBAKIS, N., “A survey of skin-color modeling and detection methods,” in *Pattern Recognition*, 2007.
- [71] KAM, M., ZHU, X., and KALATA, P., “Sensor fusion for mobile robot navigation,” *Proceedings of the IEEE*, vol. 85, no. 1, pp. 108–119, 1997.
- [72] KAMON, I., FLASH, T., and EDELMAN, S., “Learning to grasp using visual information,” in *Robotics and Automation, 1996. Proceedings., 1996 IEEE International Conference on*, vol. 3, 1996.
- [73] KATO, H., BILLINGHURST, M., I. POUPYREV, K. I., and TACHIBANA, K., “Virtual object manipulation on a table-top ar environment,” in *International Symposium on Augmented Reality*, 2000.

- [74] KATSUKI, R., OTA, J., TAMURA, Y., MIZUTA, T., KITO, T., ARAI, T., UHEYAMA, T., and NISHIYAMA, T., "Handling of object with marks by a robot," in *International Conference on Intelligent Robots and Systems*, 2003.
- [75] KEMP, C., "Ramona nichols, founder and operator of georgia canines for independence," in *personal communication*, November 2007.
- [76] KEMP, C. C., ANDERSON, C. D., NGUYEN, H., TREVOR, A. J., and XU, Z., "A point-and-click interface for the real world: Laser designation of objects for mobile manipulation," in *International Conference on Human-Robot Interaction*, 2008.
- [77] KEMP, C. C., EDSINGER, A., and TORRES-JARA, E., "Challenges for robot manipulation in human environments," *IEEE Robotics & Automation Magazine*, vol. 14, pp. 20–29, March 2007.
- [78] KHALIL, H. and GRIZZLE, J., *Nonlinear systems*, vol. 3. Prentice hall New Jersey, 2002.
- [79] KIM, B. K., MIYAZAKI, M., OHBA, K., HIRAI, S., and TANIE, K., "Web services based robot control platform for ubiquitous functions," in *International Conference on Robotics and Automation (ICRA)*, 2005.
- [80] KIM, D., *Propagation Measurements and System Design for Long-Range RF Tags*. PhD thesis, Georgia Institute of Technology, 2002.
- [81] KIM, J., ZHANG, F., and EGERSTEDT, M., "A provably complete exploration strategy by constructing voronoi diagrams," *Autonomous Robots*, vol. 29, no. 3-4, pp. 367–380, 2010.
- [82] KIM, M., KIM, H. W., and CHONG, N. Y., "Automated robot docking using direction sensing RFID," in *Proceedings of IEEE International Conference on Robotics and Automation*, pp. 4588–4593, 10-14 April 2007.
- [83] KLEINER, A., PREDIGER, J., and NEBEL, B., "Rfid technology-based exploration and slam for search and rescue," in *IROS*, 2006.
- [84] KONKEL, M., LEUNG, V., ULLMER, B., and HU, C., "Tagaboo: a collaborative childrens game based upon wearable rfid technology," *Personal and Ubiquitous Computing*, vol. 8, no. 5, pp. 382–384, 2004.
- [85] KRAGIĆ, D., PETERSSON, L., and CHRISTENSEN, H., "Visually guided manipulation tasks," *Robotics and Autonomous Systems*, vol. 40, no. 2-3, pp. 193–203, 2002.
- [86] KULYUKIN, V., GHARPURE, C., and NICHOLSON, J., "Robocart: Toward robot-assisted navigation of grocery stores by the visually impaired," in *International Conference on Intelligent Robots and Systems (IROS)*, 2005.
- [87] KURUGOLLU, F., SANKUR, B., and HARMANCI, A., "Color image segmentation using histogram multithresholding and fusion," *Image and Vision computing*, vol. 19, no. 13, pp. 915–928, 2001.
- [88] KWEE, H., DUIMEL, J., SMITS, J., DE MOED, A., and VAN WOERDEN, J., "The MANUS Wheelchair-Borne Manipulator: System Review and First Results," in *Proc. IARP Workshop on Domestic and Medical & Healthcare Robotics, Newcastle*, 1989.

- [89] LAI, K., BO, L., REN, X., and FOX, D., “A large-scale hierarchical multi-view rgb-d object dataset,” in *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2011.
- [90] LALLIER, E. and FAROOQ, M., “A real time pixel-level based image fusion via adaptive weight averaging,” in *Information Fusion, 2000. FUSION 2000. Proceedings of the Third International Conference on*, vol. 2, pp. WEC3–3, IEEE, 2000.
- [91] LAVALLE, S., *Planning algorithms*. Cambridge Univ Pr, 2006.
- [92] LAZARO, A., GIRBAU, D., and SALINAS, D., “Radio link budgets for uhf rfid on multipath environments,” *Antennas and Propagation, IEEE Transactions on*, vol. 57, no. 4, pp. 1241–1251, 2009.
- [93] LI, J., MENG, Q., WANG, Y., and ZENG, M., “Odor source localization using a mobile robot in outdoor airflow environments with a particle filter algorithm,” *Autonomous Robots*, pp. 1–12, 2011.
- [94] LI, X., ZHANG, Y., and AMIN, M., “Multifrequency-based range estimation of rfid tags,” in *RFID, 2009 IEEE International Conference on*, pp. 147–154, IEEE, 2009.
- [95] LILIENTHAL, A., LOUTFI, A., and DUCKETT, T., “Airborne chemical sensing with mobile robots,” *Sensors*, vol. 6, no. 11, pp. 1616–1678, 2006.
- [96] LUO, R. and KAY, M., “Multisensor integration and fusion in intelligent systems,” *Systems, Man and Cybernetics, IEEE Transactions on*, vol. 19, no. 5, pp. 901–931, 1989.
- [97] MA, J., CHUNG, T., and BURDICK, J., “A probabilistic framework for object search with 6-dof pose estimation,” *The International Journal of Robotics Research*, 2011.
- [98] MAILLOUX, R. and BOOKS24X7, I., *Phased array antenna handbook*. Artech House, 2005.
- [99] MARDER-EPPSTEIN, E., BERGER, E., FOOTE, T., GERKEY, B., and KONOLIGE, K., “The office marathon: Robust navigation in an indoor office environment,” in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pp. 300–307, IEEE, 2010.
- [100] MAYET, R., ROBERZ, J., SCHMICKL, T., and CRAILSHEIM, K., “Antbots: A feasible visual emulation of pheromone trails for swarm robots,” *Swarm Intelligence*, pp. 84–94, 2011.
- [101] MCDONALD, H., GARG, A., and HAYNES, R., “Interventions to enhance patient adherence to medication prescriptions,” *JAMA: the journal of the American Medical Association*, vol. 288, no. 22, p. 2868, 2002.
- [102] MILLER, A., KNOOP, S., CHRISTENSEN, H., and ALLEN, P., “Automatic grasp planning using shape primitives,” in *Robotics and Automation, 2003. Proceedings. ICRA'03. IEEE International Conference on*, vol. 2, pp. 1824–1829, IEEE, 2003.
- [103] MING, A., XIE, Z., YOSHIDA, T., YAMASHIRO, M., TANG, C., and SHIMOJO, M., “Home service by a mobile manipulator system - system configuration and basic experiments,” in *International Conference on Information and Automation*, 2008.
- [104] NATALE, L. and TORRES-JARA, E., “A sensitive approach to grasping,” in *Proceedings of the Sixth International Workshop on Epigenetic Robotics*, pp. 87–94, 2006.

- [105] NEIRA, J., TARDÓS, J., HORN, J., and SCHMIDT, G., “Fusing range and intensity images for mobile robot localization,” *Robotics and Automation, IEEE Transactions on*, vol. 15, no. 1, pp. 76–84, 1999.
- [106] NGUYEN, H., ANDERSON, C. D., TREVOR, A. J., JAIN, A., XU, Z., and KEMP, C. C., “El-e: An assistive robot that fetches objects from flat surfaces,” in *Robotic Helpers, Int. Conf. on Human-Robot Interaction*, 2008.
- [107] NGUYEN, H., DEYLE, T., REYNOLDS, M., and KEMP, C. C., “PPS-Tags: Physical Perceptual and Semantic Tags for Autonomous Mobile Manipulation,” *Intelligent Robots and Systems, 2009 (IROS 2009) IEEE/RSJ International Conference on, Semantic Perception for Mobile Manipulation Workshop*.
- [108] NGUYEN, H., JAIN, A., ANDERSON, C. D., and KEMP, C. C., “A clickable world: Behavior selection through pointing and context for mobile manipulation,” in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2008.
- [109] NGUYEN, H. and KEMP, C. C., “Bio-inspired Assistive Robotics: Service Dogs as a Model for Human-Robot Interaction and Mobile Manipulation,” in *IEEE RAS/EMBS International Conference on Biomedical Robotics and Biomechatronics (BIOROB)*, 2008.
- [110] NIKITIN, P. and RAO, K., “Effect of gen2 protocol parameters on rfid tag performance,” in *RFID, 2009 IEEE International Conference on*, pp. 117–122, IEEE.
- [111] NIKITIN, P., RAO, K., and LAZAR, S., “An Overview of Near Field UHF RFID,” in *Proc. IEEE Int’l. Conf. RFID*, pp. 167–74, 2007.
- [112] OKADA, K., OGURA, T., HANEDA, A., FUJIMOTO, J., GRAVOT, F., and INABA, M., “Humanoid Motion Generation System on HRP2-JSK for Daily Life Environment,” *planning*, vol. 6, p. 7, 2005.
- [113] OSHITA, M., “Motion-capture-based avatar control framework in third-person view virtual environments,” in *Proceedings of the 2006 ACM SIGCHI international conference on Advances in computer entertainment technology*, pp. 2–es, ACM, 2006.
- [114] PALETTA, L. and ROME, E., “Learning fusion strategies for visual object detection,” in *Intelligent Robots and Systems, 2000.(IROS 2000). Proceedings. 2000 IEEE/RSJ International Conference on*, vol. 2, pp. 1446–1452, IEEE, 2000.
- [115] PARK, M. and JUNG, K., “A Novel Infrastructure of Digital Storytelling Theme Museums Based on RFID Systems,” *LECTURE NOTES IN COMPUTER SCIENCE*, vol. 4551, p. 893, 2007.
- [116] PHILIPOSE, M., FISHKIN, K., PERKOWITZ, M., PATTERSON, D., FOX, D., KAUTZ, H., and HÄHNEL, D., “Inferring Activities from Interactions with Objects,” *IEEE PERSVASIVE COMPUTING*, pp. 50–57, 2004.
- [117] PHILIPOSE, M., SMITH, J., JIANG, B., MAMISHEV, A., ROY, S., and SUNDARA-RAJAN, K., “Battery-free wireless identification and sensing,” *IEEE Pervasive Computing*, pp. 37–45, 2005.

- [118] PITZER, B., KAMMEL, S., DUHADWAY, C., and BECKER, J., “Automatic reconstruction of textured 3d models,” in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pp. 3486–3493, IEEE, 2010.
- [119] QING, X., GOH, C., and CHEN, Z., “Segmented loop antenna for uhf near-field rfid applications,” *Electronics letters*, vol. 45, no. 17, pp. 872–873, 2009.
- [120] QT-P1, “Age groups and sex: 2000 (data set: Census 2000 summary file 1 (sf 1) 100 percent data) for the united states,” in <http://factfinder.census.gov/>, 2007.
- [121] QUIGLEY, M., BATRA, S., GOULD, S., KLINGBEIL, E., LE, Q., WELLMAN, A., and NG, A., “High-accuracy 3d sensing for mobile manipulation: Improving object detection and door opening,” in *Robotics and Automation, 2009. ICRA’09. IEEE International Conference on*, pp. 2816–2822, IEEE, 2008.
- [122] QUIGLEY, M., GERKEY, B., CONLEY, K., FAUST, J., FOOTE, T., LEIBS, J., ERIC BERGER, R. W., and NG, A., “ROS: an open-source Robot Operating System,” in *Open-Source Software workshop of (ICRA)*, 2009.
- [123] RASMUSSEN, C., “Combining laser range, color, and texture cues for autonomous road following,” in *Robotics and Automation, 2002. Proceedings. ICRA’02. IEEE International Conference on*, vol. 4, pp. 4320–4325, IEEE, 2002.
- [124] REMAZEILLES, A., LEROUX, C., and CHALUBERT, G., “SAM: A robotic butler for handicapped people,” in *Robot and Human Interactive Communication, 2008. RO-MAN 2008. The 17th IEEE International Symposium on*, pp. 315–321, 2008.
- [125] REYNOLDS, M., 2011. Personal communication.
- [126] RUSU, R. B., MEEUSSEN, W., CHITTA, S., and BEETZ, M., “Laser-based perception for door and handle identification,” in *Proceedings of International Conference on Advanced Robotics*, 2009.
- [127] RUSU, R., BLODOW, N., and BEETZ, M., “Fast point feature histograms (fpfh) for 3d registration,” in *Robotics and Automation, 2009. ICRA’09. IEEE International Conference on*, pp. 3212–3217, IEEE, 2009.
- [128] RUSU, R., BRADSKI, G., THIBAUX, R., and HSU, J., “Fast 3d recognition and pose using the viewpoint feature histogram,” in *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pp. 2155–2162, IEEE.
- [129] SAXENA, A., DRIEMEYER, J., and NG, A., “Robotic Grasping of Novel Objects using Vision,” *The International Journal of Robotics Research*, vol. 27, no. 2, p. 157, 2008.
- [130] SAXENA, A., WONG, L., and NG, A., “Learning Grasp Strategies with Partial Shape Information,” *AAAI*, 2008.
- [131] SE, S., LOWE, D., and LITTLE, J., “Vision-based mobile robot localization and mapping using scale-invariant features,” in *Robotics and Automation, 2001. Proceedings 2001 ICRA. IEEE International Conference on*, vol. 2, pp. 2051–2058, IEEE, 2001.
- [132] SEO, D., WON, D., YANG, G., CHOI, M., KWON, S., and PARK, J., “A probabilistic approach for mobile robot localization under rfid tag infrastructures,” in *2005 International Conference on Control, Automation, and Systems (ICCAS 2005)*, pp. 1797–1801, 2005.

- [133] SICILIANO, B. and KHATIB, O., eds., *Springer Handbook of Robotics*, ch. 53 Rehabilitation and Health Care Robotics. Springer, July 2008.
- [134] SIDIQ S. HIDAYAT, B. K. and OHBA, K., “Affordance-based ontology design for ubiquitous robots,” in *IEEE International Symposium on Robot and Human Interactive Communication*, 2008.
- [135] SINGHAL, A., LUO, J., and BROWN, C., “A multilevel bayesian network approach to image sensor fusion,” in *Information Fusion, 2000. FUSION 2000. Proceedings of the Third International Conference on*, vol. 2, pp. WEB3–9, IEEE, 2000.
- [136] SKOLNIK, M., *Introduction to Radar Systems*. McGraw-Hill New York, New York, USA, 1962.
- [137] SMITH, J., FISHKIN, K., JIANG, B., MAMISHEV, A., PHILIPPOSE, M., REA, A., ROY, S., and SUNDARA-RAJAN, K., “RFID-based techniques for human-activity detection,” *Communications of the ACM*, vol. 48, no. 9, pp. 39–44, 2005.
- [138] SOLDATOVA, L., CLARE, A., SPARKES, A., and KING, R., “An ontology for a robot scientist,” *Bioinformatics*, vol. 22, no. 14, p. e464, 2006.
- [139] SRINIVASA, S., FERGUSON, D., WEGHE, M., DIANKOV, R., BERENSON, D., HELFRICH, C., and STRASDAT, H., “The Robotic Busboy: Steps Towards Developing a Mobile Robotic Home Assistant,” in *International Conference on Intelligent Autonomous Systems*, 2008.
- [140] STANGER, C. A., ANGLIN, C., HARWIN, W. S., and ROMILLY, D. P., “Devices for assisting manipulation: a summary of user task priorities,” *IEEE Transactions on Rehabilitation Engineering*, vol. 2, p. 10, December 1994.
- [141] SUNG, J.-Y., GRINTER, R., CHRISTENSEN, H., and GUO, L., “Housewives or technophiles?: Understanding domestic robot owners,” in *Proceedings of the ACM Conference on Human Robot Interaction*, 2008.
- [142] THINGMAGIC, 2009. Personal communication.
- [143] THRUN, S., BURGARD, W., and FOX, D., *Probabilistic Robotics*. MIT Press, 2005.
- [144] TIWARI, P., WARREN, J., DAY, K., MACDONALD, B., JAYAWARDENA, C., KUO, I., IGIC, A., and DATTA, C., “Feasibility study of a robotic medication assistant for the elderly,” 2011.
- [145] TOPPING, M. and SMITH, J., “The development of Handy 1, a rehabilitation robotic system to assist the severely disabled,” *Industrial Robot*, vol. 25, no. 5, pp. 316–20, 1998.
- [146] TSUI, K., YANCO, H., KONTAK, D., and BELIVEAU, L., “Development and evaluation of a flexible interface for a wheelchair mounted robotic arm,” in *Proceedings of the 3rd international conference on Human robot interaction*, pp. 105–112, ACM New York, NY, USA, 2008.
- [147] VAN DER LOOS, H., “VA/Stanford rehabilitation robotics research and development program: lessons learned in the application of robotics technology to the field of rehabilitation,” *Rehabilitation Engineering, IEEE Transactions on [see also IEEE Trans. on Neural Systems and Rehabilitation]*, vol. 3, no. 1, pp. 46–55, 1995.

- [148] VAN DER LOOS, H., WAGNER, J., SMABY, N., CHANG, K., MADRIGAL, O., LEIFER, L., and KHATIB, O., "ProVAR assistive robot system architecture," in *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*, vol. 1, 1999.
- [149] VANDERBILT, D. and LOUIE, S., "A monte carlo simulated annealing approach to optimization over continuous variables," *Journal of Computational Physics*, vol. 56, no. 2, pp. 259–271, 1984.
- [150] VORST, P., SCHNEEGANS, S., YANG, B., and ZELL, A., "Self-localization with RFID snapshots in densely tagged environments," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, pp. 1353–1358, Citeseer, 2008.
- [151] VORST, P., SOMMER, J., HOENE, C., SCHNEIDER, P., WEISS, C., SCHAIRER, T., ROSENSTIEL, W., ZELL, A., and CARLE, G., "Indoor positioning via three different rf technologies," in *Proc. 4th European Workshop on RFID Systems and Technologies (RFID SysTech 2008)*, Freiburg, Germany, Citeseer, 2008.
- [152] VORST, P. and ZELL, A., "Semi-autonomous learning of an rfid sensor model for mobile robot self-localization," in *European Robotics Symposium 2008*, pp. 273–282, Springer, 2008.
- [153] VORST, P., KOCH, A., and ZELL, A., "Efficient self-adjusting, similarity-based location fingerprinting with passive UHF RFID," in *IEEE International Conference on RFID-Technology and Applications (RFID-TA2011)*, (Sitges, Barcelona, Spain), September 15-16 2011.
- [154] WAARSING, B., NUTTIN, M., and BRUSSEL, H. V., "Introducing robots into a human-centred environment - the behaviour-based approach," in *International Conference on CLAWAR*, May 2001.
- [155] WASIK, Z. and SAFFIOTTI, A., "Robust color segmentation for the RoboCup domain," in *International Conference on Pattern Recognition*, 2002.
- [156] WERFEL, J., BAR-YAM, Y., RUS, D., and NAGPAL, R., "Distributed construction by mobile robots with enhanced building blocks," in *Proceedings IEEE International Conference on Robotics and Automation*, pp. 2787–2794, 2006.
- [157] WIKIPEDIA, 2011. Lyapunov Function: Example.
- [158] WILLIAMS, A., XIE, D., OU, S., GRUPEN, R., HANSON, A., and RISEMAN, E., "Distributed smart cameras for aging in place," 2006.
- [159] WISTORT, R. and SMITH, J., "Electric field servoing for robotic manipulation," *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*.
- [160] XINGYONG, S., LEE, H., and CHO, H., "A sensor fusion method for mobile robot navigation," in *SICE-ICASE, 2006. International Joint Conference*, pp. 5310–5316, IEEE, 2006.
- [161] YAMAMOTO, Y., PIRJANIAN, P., MUNICH, M., DIBERNARDO, E., GONCALVES, L., OSTROWSKI, J., and KARLSSON, N., "Optical sensing for robot perception and localization," in *Advanced Robotics and its Social Impacts, 2005. IEEE Workshop on*, pp. 14–17, IEEE.
- [162] YAMAUCHI, B., "A frontier-based approach for autonomous exploration," in *Computational Intelligence in Robotics and Automation, 1997. CIRA'97., Proceedings., 1997 IEEE International Symposium on*, pp. 146–151, IEEE, 1997.

- [163] YUAN, K., HONG, A., ANG, M., and PENG, G., “Unmanned library: an intelligent robotic books retrieval & return system utilizing RFID tags,” in *Systems, Man and Cybernetics, 2002 IEEE International Conference on*, vol. 4, 2002.
- [164] ZETTER, K., “Hackers clone e-passports,” *Wired News (3 August 2006) Referenced*, pp. 71521–0, 2006.
- [165] ZHOU, Y., LIU, W., and HUANG, P., “Laser-activated rfid-based indoor localization system for mobile robots,” in *Robotics and Automation, 2007 IEEE International Conference on*, pp. 4600–4605, IEEE, 2007.
- [166] ZLLNER, R., ASFOUR, T., and DILLMANN, R., “Programming by demonstration: Dual-arm manipulation tasks for humanoid robots,” in *International Conference on Intelligent Robots and Systems*, (Sendai, Japan), 2004.