

NORTHWESTERN UNIVERSITY

Motion and Sensing in Electrosensory Systems

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Mechanical Engineering

By

James R Solberg

EVANSTON, ILLINOIS

December 2007

© Copyright by James R Solberg 2007

All Rights Reserved

ABSTRACT

Motion and Sensing in Electrosensory Systems

James R Solberg

Sensing is a fundamental operation for almost any motion-based system. We have chosen electrosensory systems as a platform to explore sensing and control in both artificial and biological systems. Electrosense is a convenient sensing modality because artificial electrosensory systems are relatively simple to implement, and weakly electric fish have some of the most well-documented sensorimotor pathways. In particular, active electrolocation is investigated, where the task is to estimate the location of a target using measurements from a self-generated electric field. The fundamentals of electrolocation are described first with a finite-element numerical approximation of the governing equations, and then simple models are used to predict electrosensory observations. Several belief maintenance schemes are employed to fuse sensor data and explicitly account for uncertainties in the position of the target. In the biological realm, a protocol for simulating the sensory acquisition and belief maintenance during prey-capture behavior in the weakly electric fish was developed. Using these simulations optimal sensing was investigated, and results provide insight into the interdependencies and co-evolution of sensing and motion

systems of the weakly electric fish. In the artificial realm, an electrosensory robot capable of actively locating underwater targets by measuring perturbations in a self-generated electric field was built. Using seven different control algorithms, the robot can successfully locate nearby targets, as well as localize itself when placed in a pre-mapped environment, in both fresh and saltwater.

Acknowledgements

This work would not be possible without the tremendous amount of guidance and support that I have been fortunate to receive from my friends, family, colleagues, and advisors.

My first round of thanks has to go to my advisors, Professors Kevin Lynch and Malcolm MacIver. Their direction and expertise kept me excited and on track. Their generous contributions—both time and ideas—to this work went well beyond the expectations of the standard advisor roll, so much so that I felt compelled to use “we” as the first-person pronoun instead of “I” when writing this document.

I would also like to thank the rest of my PhD committee, Professors Mitra Hartmann, Randy Freeman, and Steve LaValle, for their participation and valuable insights. In particular, I would like to thank Mitra for her involvement and contributions early on.

The Laboratory for Intelligent Mechanical Systems (LIMS) has been my intellectual and social home for the past four and half years. I could not possibly imagine a better group of professors and students. To the co-directors of LIMS—Michael, Ed, and Kevin—I ask you to maintain this environment so future LIMSians could be as fortunate as I. In particular I would like to thank Professor Michael Peshkin for his help with the electronics on my robot. Also, I am forever in debt to Mike Taylor; without his “sleepovers” at the lab working on the robot, it may have taken 5 and half years instead of 4 and a half years to finish.

Most importantly, I would like to thank my family for their love, support, and babysitting:

- My in-laws—Wylie, Erica, Ioanna, and Corinna—and my parents have provided thousands of hours of daycare, allowing me to complete this work in finite time.
- My parents have supported me through all my crazy endeavors even though they could not understand why someone would spend fourteen years in college (sorry, I don't have an answer).
- My sons—Andrew and Keith—have made every day exciting.
- Above all I thank my wife, Emily. Without her emotional (and financial) support, I would not have been able to pursue a PhD.

In memory of the original Keith Solberg.

Table of Contents

| | |
|---|-----------|
| ABSTRACT | 3 |
| Acknowledgements | 5 |
| List of Tables | 13 |
| List of Figures | 14 |
| Chapter 1. Introduction | 19 |
| 1.1. Thesis Overview | 21 |
| 1.2. Weakly Electric Fish Use Active Sensing for Active Sensing | 22 |
| 1.3. Background on Electric Field Sensing | 23 |
| 1.4. Electrosense as a Possible Modality for AUVs | 26 |
| 1.5. Belief Maintenance and Planning Under Uncertainty | 28 |
| 1.6. Preview | 31 |
| Part 1. Modeling Electrolocation | 33 |
| Chapter 2. Modeling Electrosensory Observations | 35 |
| 2.1. Numerical Construction of the EEV in 2-D | 35 |
| 2.2. An Analytical Solution to the EEV in 3-D | 46 |
| 2.3. Dipole Model | 50 |

| | |
|---|------------|
| | 9 |
| 2.4. Ellipse Model | 54 |
| Chapter 3. Belief Maintenance During Electrolocation | 59 |
| 3.1. Possibilistic Belief Maintenance | 60 |
| 3.2. Belief Maintenance via Histogram Filter | 71 |
| 3.3. Belief Maintenance via Particle Filter | 83 |
| 3.4. Uncertainty maps | 96 |
| Part 2. Electrosense Optimization for Prey Capture in Weakly Electric Fish | 101 |
| Chapter 4. Electrosense Optimization for Prey Capture: | |
| Introduction and Formulation | 103 |
| 4.1. Introduction | 103 |
| 4.2. Optimal Sensing Formulation | 108 |
| 4.3. Methods | 117 |
| Chapter 5. Electrosense Optimization for Prey Capture: | |
| Results and Discussion | 128 |
| 5.1. The Tradeoff Between Number of Sensors and Sensor Noise | 128 |
| 5.2. Why is the Electric Field Stronger Near the Tail? | 134 |
| 5.3. Determining the Best Sensor Layout | 138 |
| 5.4. Interdependencies Between Sensing and Motion | 144 |
| Part 3. The Electrosenster: A Robotic Electrosensory System | 150 |
| Chapter 6. Experimental Setup | 152 |

| | |
|---|-----|
| | 10 |
| 6.1. Related Systems | 152 |
| 6.2. Aquatic Environment | 154 |
| 6.3. Hardware | 156 |
| 6.4. Electrodes | 157 |
| 6.5. Electrical Circuitry | 157 |
| 6.6. Computer Hardware Structure | 162 |
| 6.7. Electric Field | 163 |
| Chapter 7. Sensor Characterization | 164 |
| 7.1. Empirical Construction of Sensor Model | 164 |
| 7.2. Empirical EEVs | 166 |
| 7.3. Detection Range Analysis for a Single Sensor | 170 |
| 7.4. Computed vs Measured Detection Distances | 177 |
| Chapter 8. Probabilistic Target Electrolocation | 179 |
| 8.1. Probabilistic-Based Control Algorithms | 179 |
| 8.2. Experimental Set 1: Cylinders | 188 |
| 8.3. Experimental Set 2: Eight Different Conditions for Spheres | 197 |
| 8.4. Experimental Set 3: Comparison of Six Controllers | 205 |
| 8.5. Conclusions Across Experiment Sets | 211 |
| Chapter 9. More Capabilities of the Electrosenster | 212 |
| 9.1. A Global Localization Scenario | 212 |
| 9.2. A Non-Probabilistic-Based Controller | 216 |
| 9.3. Performance of Non-Probabilistic-Based Controller | 219 |

| | |
|---|-----|
| | 11 |
| 9.4. Extensions of the Current Setup | 223 |
| Chapter 10. Conclusions and Future Directions | 227 |
| 10.1. The Driving Objectives of this Thesis | 227 |
| 10.2. Future Directions: Modeling the Sensing and Motion Systems of the Weakly Electric Fish | 230 |
| References | 234 |
| APPENDICES | 256 |
| Appendix A. A Simple Example Using a Particle Filter | 257 |
| Appendix B. Electroplating the Electrodes | 259 |
| Appendix C. The Electric Field of the Electrodes | 260 |
| Appendix D. Some Examples of Measured EEVs of Cylinders | 265 |
| Appendix E. Closed-Form Solution for Voltage Perturbation | 271 |
| Appendix F. More Complex Models of the Electrosensory System of Weakly Electric Fish | 274 |
| F.1. Modeling the Weakly Electric Fish | 274 |
| F.2. Models of Neural Coding in Primary Electrosensory Afferents | 276 |
| F.3. Electrosensory Processing Models | 280 |
| F.4. Prey Localization | 281 |
| Appendix G. Simple Models of Sensorimotor Transformations in Fish and Robots | 283 |
| G.1. The Kalman filter | 283 |

| | |
|--|-----|
| | 12 |
| G.2. SLAM | 286 |
| G.3. The Utility of Uncertainty in Simple Models | 287 |
| G.4. Sensor-Based Control with Simple Models | 287 |
| Appendix H. More Generalized Models for Sensorimotor Transformations in Fish and Robots | 289 |
| H.1. (Active) Markov Localization | 289 |
| H.2. Markov Decision Processes | 291 |
| H.3. Reinforcement Learning | 295 |
| Appendix I. Feedback-Feedforward Hybrid Controllers | 299 |
| Appendix J. Predictive Controllers in Biological Systems | 303 |
| J.1. Relevant Models of the Primate Motor Control System | 304 |
| J.2. Relevance to Weakly Electric Fish | 309 |
| Appendix K. Information in the Brain | 311 |
| Appendix. Vita | 313 |

List of Tables

| | | |
|-----|--|-----|
| 2.1 | Definition of terms for Maxwell’s equations. | 37 |
| 3.1 | Confidence intervals of standard deviations. | 76 |
| 4.1 | The sixteen sensor layouts. | 124 |
| 5.1 | The sensing performance for each combination of sensory layout and trajectory type. | 142 |
| 5.2 | Performance of fish-like and non-fish-like trials. | 147 |
| 7.1 | Empirically measured and theoretical detection distance. | 177 |
| 8.1 | Statistics for electrolocation data set 1. | 195 |
| 8.2 | Statistics for electrolocation data set 2. | 204 |
| 8.3 | Statistics for electrolocation data set 3. | 209 |
| 9.1 | Statistics for global localization. | 214 |
| 9.2 | Electrolocation statistics using zero-cross (experimental data set 1). | 221 |
| 9.3 | Electrolocation statistics using zero-cross with additional pass (experimental data set 2). | 223 |

List of Figures

| | | |
|------|---|----|
| 1.1 | The weakly electric fish | 20 |
| 2.1 | Unperturbed electric field for the simulation world. | 40 |
| 2.2 | A conductor perturbs the electric field. | 42 |
| 2.3 | An insulator perturbs the electric field. | 43 |
| 2.4 | Numerical EEV for conducting disk. | 45 |
| 2.5 | Electric field for dipole model. | 47 |
| 2.6 | An illustration of the dipole emitter/detector layout. | 52 |
| 2.7 | EEV for dipole model. | 53 |
| 2.8 | Comparison of multipole model with measured data from an actual fish. | 55 |
| 2.9 | Electric field of ellipse model. | 56 |
| 2.10 | An example of four EEVs for the ellipse model. | 58 |
| 3.1 | Possibilistic noise model with the dipole model | 61 |
| 3.2 | Possibilistic noise model with the ellipse model. | 63 |
| 3.3 | Four possibilistic beliefs from dipole model | 65 |
| 3.4 | Four possibilistic beliefs superimposed for dipole model | 66 |

| | | |
|------|--|----|
| | | 15 |
| 3.5 | Fused belief after each observation in the dipole model. | 67 |
| 3.6 | Single-observation beliefs for each sensor in the ellipse model. | 68 |
| 3.7 | Possibilistic sensor fusion for the ellipse model. | 69 |
| 3.8 | A simple example of a histogram belief scheme. | 71 |
| 3.9 | The entropy of two beliefs. | 73 |
| 3.10 | Single-observation histogram beliefs for ellipse model. | 75 |
| 3.11 | Determining the likelihood of an observation for a cell in a histogram filter. | 77 |
| 3.12 | Probabilistic sensor fusion for the ellipse model. | 79 |
| 3.13 | Probabilistic sensor fusion for the ellipse model with different sensor order. | 80 |
| 3.14 | Comparing beliefs from possibilistic and probabilistic belief schemes for the ellipse model. | 81 |
| 3.15 | Single-observation histogram beliefs for dipole model. | 82 |
| 3.16 | Probabilistic sensor fusion for the dipole model. | 84 |
| 3.17 | Comparing beliefs from possibilistic and probabilistic belief schemes for the dipole model. | 85 |
| 3.18 | Converting a particle belief representation to a histogram belief representation. | 87 |
| 3.19 | Single-observation particle beliefs for dipole model. | 89 |
| 3.20 | Probabilistic sensor fusion for the dipole model with particle belief. | 91 |

| | | |
|------|--|-----|
| 3.21 | Comparing beliefs from possibilistic and probabilistic belief schemes (both histogram and particle) for the dipole model | 92 |
| 3.22 | Single-observation particle beliefs for ellipse model. | 93 |
| 3.23 | Probabilistic sensor fusion via particle filtering for the ellipse model. | 94 |
| 3.24 | Comparing beliefs from possibilistic and probabilistic belief schemes (both histogram and particle) for the ellipse model. | 95 |
| 3.25 | Uncertainty maps for dipole model. | 97 |
| 3.26 | Uncertainty maps for ellipse model. | 99 |
| 4.1 | A 1-D example of prey localization using a Bayes filter. | 111 |
| 4.2 | The black ghost knifefish fit to an ellipsoid. | 117 |
| 4.3 | The motivation for using the <code>parts-std</code> uncertainty norm rather than entropy. | 121 |
| 4.4 | An example of a sensor layout using our parametrization. | 123 |
| 5.1 | Setup for toy problem (sensing equivalence) | 129 |
| 5.2 | Sensing equivalence for toy problem. | 131 |
| 5.3 | Sensing equivalence for the 3-D ellipsoid. | 133 |
| 5.4 | Uncertainty map for normal and flipped e-field. | 136 |
| 5.5 | Three examples of sensor layouts. | 139 |
| 5.6 | The normalized <code>sense-utility</code> of the mech-opt trajectories versus the normalized <code>sensor-dist</code> metric. | 148 |

| | | |
|-----|--|-----|
| | | 17 |
| 6.1 | The Electrosenster. | 153 |
| 6.2 | Top view photo of the hardware of the Electrosenster. | 155 |
| 6.3 | Block Diagram of the signal processing circuitry of the Electrosenster. | 159 |
| 6.4 | Synchronous detection. | 161 |
| 7.1 | Electrosenster workspace for two different spheres. | 167 |
| 7.2 | EEVs of spherical targets for eight different experimental conditions. | 169 |
| 7.3 | Sensor observation as a function of target position for theoretical detection analysis. | 172 |
| 7.4 | Sensor observation as a function of target position along diagonal. | 173 |
| 7.5 | Dimensionless detection distance as a function of dimensionless noise. | 175 |
| 8.1 | Motivation for <code>max-grad</code> controller. | 184 |
| 8.2 | Way points for <code>spiral</code> controller. | 188 |
| 8.3 | Frames from a typical electrolocation sequence taken from experimental set 1. | 192 |
| 8.4 | Final frame of the electrolocation sequence. | 193 |
| 8.5 | Snapshots of a typical electrolocation sequence from experimental set 2. | 200 |
| 8.6 | Final frame of electrolocation sequence from experimental set 2. | 201 |
| 8.7 | Performance of the <code>rnd-wlk</code> verses the <code>min-var</code> controllers for data set 2 | 203 |
| 8.8 | EEV for experimental set 3. | 206 |

| | | |
|------|---|-----|
| 8.9 | Performance of the six controllers for data set 3. | 208 |
| 8.10 | The mean-min controller gets stuck. | 210 |
| 9.1 | Global localization environment. | 213 |
| 9.2 | Snapshots of a typical global localization sequence. | 215 |
| 9.3 | Final snapshot of the global localization sequence. | 216 |
| 9.4 | Electrolocation using zero-cross controller. | 218 |
| 9.5 | Two target trajectories relative to the robot using the zero-cross controller. | 222 |
| A.1 | A simple example using a particle filter. | 258 |
| B.1 | Electrode electroplating setup. | 259 |
| C.1 | Electric field of the Electrosenster. | 264 |
| D.1 | The EEVs for the three conditions in experimental set 1. | 267 |
| D.2 | EEVs for plastic cylinders under 4 different experimental conditions. | 268 |
| D.3 | EEVs for brass cylinders of four different diameters. | 270 |
| F.1 | Modeling voltage to spike trains in the fish. | 279 |
| F.2 | Schematic of spike trains to prey location. | 281 |
| J.1 | A very crude block diagram of biological controller | 305 |

CHAPTER 1

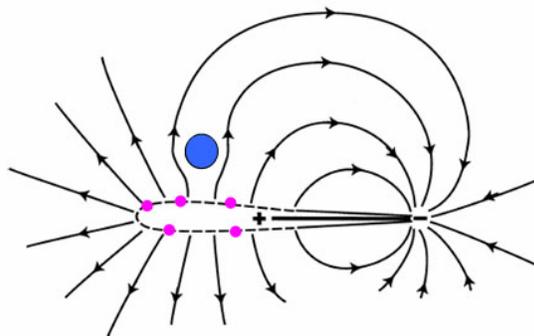
Introduction

The black ghost knifefish (*Apteronotus albifrons*) is a peculiar species of predatory fish that live in the rivers of Amazon basin (Figure 1.1(A)). They hunt at night in turbid waters where vision is of little use for capturing the small insect larvae and water fleas that they often feast on. These knifefish utilize a self-generated electric field for sensing nearby objects. As a potential prey enters the fish's electric field, the prey distorts the nominal electric field (assuming the electrical impedance of the prey differs from the surrounding water). These distortions are detected by the knifefish's voltage-sensitive sensory organs that are scattered over the surface of its body. The fish uses the voltage measurements from the perturbed electric field to estimate the location and other properties of the prey [149]. The black ghost knifefish, along with several other species of South American and African fish that emit an electric field for sensing purposes, are collectively referred to as *weakly* electric fish to differentiate them from other aquatic species that generate much stronger electric fields to stun prey or for self-defense.

Figure 1.1(B) depicts a 2-D simplified electrostatic model of a weakly electric fish in the presence of a circular object (blue disk in Figure 1.1(B)) with higher resistance than the surrounding water. Dipole charges create an electric field (black arrows are field lines), and the pink circles represent volt sensors. The insulating blue disk diverts the electric field causing the sensors to measure a voltage that is different than if there was no object



(A) The black ghost knifefish: *Apteronotus albifrons*



(B) A simple model of the weakly electric fish

Figure 1.1. The weakly electric fish. (A): A photo of the weakly electric fish. (B): A simple model of the electric field of the weakly electric fish. The black curves are flow lines of the electric field. The blue disk is an electric insulator. The pink dots are voltage-sensitive receptors.

present. In general the electric field is time-varying, and the objects in this field exhibit a complex impedance (from both resistance and capacitance).

The electrosensory system of weakly electric fish is quite impressive. Their self-generated electric field is only about 1 mV/cm near the body of the fish [187]. Yet, their approximately 14,000 voltage-sensitive organs (electroreceptors) can detect a 0.1 %, or about $1 \text{ }\mu\text{V/cm}$, change in their nominal electric field [157]. This level of sensitivity allows the black ghost knifefish to detect a water flea (*Daphnia magna*) about 2.8 cm away in low conductivity water [154].

The weakly electric fish plays two very important roles in this document—inspiration and simulation. A major theme of this document is the development of an electrosensory system that rivals the performance of the weakly electric fish. The weakly electric fish is used only for inspiration in this case, and details of our system are not constrained

to exactly match the specifics of the weakly electric fish. In particular electrolocation—which is the task of estimating the position of target from voltage measurements—is investigated. Secondly, the weakly electric fish is used to investigate the interdependencies between sensing and motion. Sensory acquisition and prey localization is simulated for several different types of fish-like trajectories. By quantifying the sensing performance of these trajectories, conclusions are drawn about the sensing effectiveness of the trajectory choice and structure of the electrosensory system.

1.1. Thesis Overview

In almost any system with both sensing and motion control, it is necessary to recognize their interdependencies. We have chosen electrosensory systems as a platform to explore sensing and control in both artificial and biological systems. Electrosense is a convenient sensing modality because artificial electrosensory systems are relatively simple to implement, and weakly electric fish have some of the most well documented sensorimotor pathways. In the artificial realm, we have built a robotic electrosensory robot capable of actively locating underwater targets using perturbations in a self-generated electric field. In the biological realm, we have simulated the sensory acquisition and belief maintenance of the weakly electric fish during the prey-capture task. In both cases—as in most systems with both sensing and motion control—characteristics of the sensing system influence the optimal motion plan.

This document addresses three major objectives:

- (1) Developing models of electrosense and complementary methods of target belief maintenance.

- (2) Investigating interdependencies between motion and sensing in weakly electric fish.
- (3) Implementing an active, robotic electrosensory system for locating underwater targets.

These objectives are realized through focusing on two electrosense-based systems: the weakly electric fish and the Electrosenster. The weakly electric fish has been described above. The Electrosenster is an XY robot equipped with a self-generated electric field and sensors to measure perturbations in the field (just like the weakly electric fish). The Electrosenster is a much simpler system than the fish. Nonetheless, it is capable of performing electrolocation of targets using the same sensing modality.

1.2. Weakly Electric Fish Use Active Sensing for Active Sensing

Active sensing is a reoccurring theme throughout this document. Coincidentally, weakly electric fish exhibit two different meanings of the word “active sensing,” which are differentiated below [174].

Active Sensing definition 1. During the task of electrolocation weakly electric fish strategically control the trajectory of their body so their electroreceptors can receive the necessary information to determine the location of the target. In this context *active sensing* refers to the control strategy of moving to maximize the knowledge of the world [51].

Active sensing is a fundamental task for both robots and animals wishing to increase their knowledge of their surroundings [26, 106, 165, 174]. Active sensing usually takes the form of a sensor-based control scheme and is often synonymous with exploration.

Active sensing strategies have been used for both linear and nonlinear [117] systems as well as for stochastic [151] systems. A few applications include target tracking [135, 216], sensor planning [60, 137], haptic exploration [27, 190, 198, 219], localization [241, 257], and locating vapor sources [183]. The chapters to come will use electrosense-based systems to demonstrate various active sensing techniques.

Active Sensing definition 2. Weakly electric fish sense changes in their *self*-generated electric field to infer the state of their surroundings. In this context *active sensing* refers to actively emitting the energy that will be sensed. Since weakly electric fish generate the electric field that is subsequently sensed, their electrosense can be called an active sensing modality.

The first definition of active sensing is used more frequently in this document than the second, thus, we use the term “active sensing” to denote definition 1 above unless otherwise noted.

1.3. Background on Electric Field Sensing

Since sensing electric fields is a fundamental concept throughout this document, it is useful to understand how it is used in other applications.

As a potential sensing technology for both marine and freshwater applications, electrosense combines some of the advantages of somatosensation, in that it is easily distributed across a body surface to achieve omnidirectionality, and of vision, in that it occurs at a distance [155]. It is also useful in dark, cluttered or turbid environments where vision is useless, as evidenced by the weakly electric fish’s ability to hunt in these conditions [155]. As a leading model system in neurobiology for how animals process

sensory information, a great deal is known about biological electrolocation [20, 246] facilitating the development of artificial electrolocation systems [39].

There are a few examples of engineered electric field sensing systems. One of the earliest was the Theremin, a musical instrument that made its first public appearance in 1921 and measures the player's capacitance to determine the pitch of the output tone [224]. The Theremin is the first device that could measure the position of an object (the body parts of the player in this case) by measuring changes of an emitted electric field [206]. Using principles similar to a Theremin, a system has been developed by the MIT Media Lab that is able to extract the 3-D position of a user's hand [207]. In both systems the capacitance of a nearby hand changes the nominal electric field.

In a previous study an artificial electrosensory array was built to study the feasibility of using such a system in underwater robots [156, 155]. This system was able to determine the distance of a submerged 10 mm diameter plastic sphere up to a distance of about 12 mm away from the sensor. The distance estimation algorithm was based on the spatial distribution of the sensor measurements. A related study built an artificial electrosensory system to investigate the possibility of using such a system for obstacle avoidance in underwater robots [74]. This system could detect either a conducting or insulating sphere 25 mm in diameter at a range of 5 mm. In preliminary experiments the robot could perform obstacle avoidance using this electrosensory system.

The biological and man-made electric field sensing approaches described above are a specific instance of *impedance imaging* [48]. In impedance imaging the goal is to determine the spatial distribution of electrical resistance and permittivity (a measure of how readily charges separate under an imposed electric field) of a specified region given voltage and/or

current measurements at the boundary. Two popular variants of impedance imaging are resistivity imaging and electrical impedance tomography.

Resistivity imaging began its widespread use in the 1920s by geophysicists who used arrays of metal electrodes inserted into the ground to estimate a map of the subsurface strata. Data are collected by injecting current across pairs of electrodes and then measuring the voltages at the other electrodes. Variants of the basic technique are still used today by geophysicists for imaging the subterranean and extracted core samples [86].

In the late 1970s tomographic impedance measurement techniques were developed for non-invasive imaging of impedance variations inside the human body [15]. This technique is now referred to as *electrical impedance tomography* (EIT, see [72], [111], and [148] for reviews). Internal impedance images can be used to diagnose such medical conditions as pulmonary emboli or blood clots in the lungs [195]. The problem of EIT is to estimate the spatial distribution of impedances from a series of electrical measurements at the boundary. In medical imaging applications, electrodes are placed on the surface of the skin, and an internal impedance image is deduced. The basic mathematical model used to construct an impedance image is Laplace's equation with complex impedance,

$$(1.1) \quad \nabla \cdot \gamma(\mathbf{x}, \omega) \nabla V = 0,$$

where $\mathbf{x} \in \mathbb{R}^3$ is the spatial coordinate, ω is the frequency of the applied current, V is the electric potential, and the inverse of electrical impedance (called admittivity) is represented by $\gamma(\mathbf{x}, \omega) = \sigma(\mathbf{x}, \omega) + i\omega\epsilon(\mathbf{x}, \omega)$, where σ is the electrical conductivity and ϵ is the electric permittivity. The forward problem associated with Equation (1.1) is deducing the relationship between the electric potentials and currents at the boundary for a given

impedance. The inverse problem, sometimes known as the Calderón Problem [57], is to determine the impedance of the interior of the body given the applied electric potentials and current measurements on the boundary [179]. This nonlinear inverse problem is typically severely ill-posed [87, 118]. In practice, the number of degrees of freedom of a parameterized impedance map that can be determined is limited by the number and precision of the independent observations made. In special cases, simplifying assumptions can reduce the complexity of the solution [244].

1.4. Electrosense as a Possible Modality for AUVs

We have just seen some applications of electric field sensing, but none give strong motivation for developing an active robotic electrosensory system. In this section the case is made for using electrosense in underwater vehicles. By developing a simple robotic electrosensory system, we are taking the first steps to implementing electrosense on more complex (man-made) systems including autonomous underwater vehicles.

Autonomous underwater vehicles (AUVs) are routinely used by oil and gas companies, the military, and scientists for tasks such as developing seafloor maps, mine detection, and studying the ocean and ocean floor (e.g. aquatic flora and fauna monitoring), respectively. Many new designs of AUVs that are currently being developed promise to perform much more challenging tasks such as reconnaissance, search and rescue operations, and weapon deployment. Because of tight constraints on power and space, AUV designs are often streamlined to perform a specific task. An effective AUV design incorporates complementary sensing and propulsion systems for its intended task. For example, it would be inefficient for an AUV to have a long-range sonar system if it was only traversing tight

underwater caves since the sound waves could not penetrate the walls of the caves. Thus, when choosing a sensing system for an AUV it is important to find the best technology for both the AUV's motion capabilities and for collection of the pertinent information.

Most contemporary AUVs emphasize long range sensing. The majority of these sensing systems can be classified as either optical or acoustical. The optical class includes video camera and laser range finders. These systems are able to provide high spatial resolution for shorter range applications and when the optical properties of the water are favorable. Sonar systems, on the other hand, have a longer range, but provide poor spatial resolution. Another advantage of sonar is that it is unaffected by poor visibility.

Less mainstream underwater sensing technologies could be utilized to better complement the AUV's motion capabilities and designated task. A few examples of potential AUV sensing technologies are ultrasound, whisker systems inspired by those of seals and rodents, and the use of emitted low power electric fields and field perturbation sensing as is utilized by nocturnal electrosensory fish. Ultrasonic technologies offer the potential of higher resolution than their more commonly used audible and sub-audible counterparts, but their higher frequency acoustic waves are more readily dissipated. Underwater whiskers, inspired by whiskers on seals, can be used to sense fluid current patterns and possibly even fine details of solid objects based on mechanical interaction with the whiskers [215]. Electrosense is also a viable option, depending on the task of the AUV. A self-generated electric field (i.e. active electrosense) allows for control of the intensity, direction, timing, and spectral characteristics of the stimulus energy [174]. An electrosensory system on an AUV could construct an impedance image of its surroundings, which could be used for navigation or surveying. Active (sensing self-generated energy) electrosense

could potentially offer high spatial resolution, but because of the quartic power-law dependence of geometric spreading [174], electrosense should only be used for short-range applications. Electrosense is omnidirectional, and is unaffected by water visibility. A good application for electrosense would be in an AUV working in cramped caverns with an omnidirectional locomotion system. In this situation the AUV needs a good image of its nearby surroundings to make reactive control choices.

1.5. Belief Maintenance and Planning Under Uncertainty

Any real sensing system has noise. Thus, we have chosen to address the uncertainties that arise from noisy sensors (specifically, from electrosense). These uncertainties should be accounted for in both the estimation scheme and controller [24].

Conceptually, a *belief* is any reasonable representation of the state of a system [212]. This is a very vague and abstract definition, so, for the purposes of this document we limit a belief to a mathematical (and usually statistical) description of the parameters that define the system. Many times the belief will take the form of probability distributions over a state space or parameter space. For example, if we wish to keep track of the position of a point on a line, then we could choose the coordinate position along that line, say x , as the state. If we knew with absolute certainty where the point was, then we could just use that single number as the belief (e.g., $x = 1$ is the belief). But, because the sensors used to estimate x are noisy, the position is uncertain, and the uncertainties should be reflected in the belief representation (for example the belief could be a probability distribution over all possible values of x). *Belief maintenance* is the process of updating the belief with new information (e.g., sensor observation).

The classical approach to motion planning under uncertainty is to first estimate the state of the system based on all previous motion and sensor information and then execute a control policy based on that state estimate. If the *history information space*, \mathcal{I}_{hist} , is defined as the collection of all the actions and all the sensor observations (past and present) along with the initial conditions of the system, then this classical control scheme can be written as the following two step process:

$$(1.2) \quad \kappa_{s.e.} : \mathcal{I}_{hist} \rightarrow \mathcal{X}$$

$$(1.3) \quad \pi : \mathcal{X} \rightarrow \mathcal{U}$$

where $\kappa_{s.e.}$ is a mapping from information space to state space (s.e. = state estimate). The mapping π is a control policy that generates a control vector, $u \in \mathcal{U}$, based on the state estimate, $\hat{x} \in \mathcal{X}$. This decoupling of the estimation and control components facilitates an efficient and intuitive flow of information. But, crucial information could potentially be lost in the transformation in Equation (1.2), depriving the control policy of sufficient information to act optimally.

Alternatively, a control policy can be an explicit function of the history information space:

$$(1.4) \quad \pi : \mathcal{I}_{hist} \rightarrow \mathcal{U}$$

This guarantees that the control policy operates with the maximum amount of information available. In general \mathcal{I}_{hist} is a very large dimensional space that grows by the number

of observations plus control actions at each time step. In almost any scenario a control policy that is an explicit function of \mathcal{I}_{hist} becomes impractical. Additionally, \mathcal{I}_{hist} does not lend itself to an intuitive representation of the knowledge of the system's state. An intuitive representation of the state can facilitate the design of a control policy.

This document employs an intermediate approach where the control policy is a function of the belief, b , of the system. We define *belief space*, \mathcal{B} , as a type of *derived information space* [142]. An information mapping, κ_{bel} , transforms an element in \mathcal{I}_{hist} to an element in \mathcal{B} . Now, the control policy, π , operates on the belief space (derived information state), as expressed here:

$$(1.5) \quad \kappa_{bel} : \mathcal{I}_{hist} \rightarrow \mathcal{B}$$

$$(1.6) \quad \pi : \mathcal{B} \rightarrow \mathcal{U}$$

Equations (1.2) and (1.3) are a specific implementation of Equations (1.5) and (1.6), where \mathcal{B} becomes \mathcal{X} for the case of directly estimating the system's state for the control policy to use.

Recall the example of tracking the x coordinate of a point on a line. A PDF over all possible x values is a valid belief, and thus an element in belief space. All of the raw sensor data would be contained in \mathcal{I}_{hist} . An information mapping would map all the sensor data into a PDF over x . The controller uses this PDF to generate the control command.

1.6. Preview

An overview of this document was given in Section 1.1, which outlined three objectives. Each of these objectives has been assigned to one of the three parts of this document.

In **Part 1** models of electrolocation and target belief maintenance are presented. In **Chapter 2** two methods of modeling electrosensory observations are examined. The first method is a numerical simulation that is general but computationally expensive. In the second method an analytic solution is provided for a single spherical target in the electric field. In **Chapter 3** the noise in the sensors is explicitly accounted for with several different belief maintenance schemes.

Part 2 investigates interdependencies between motion and sensing in weakly electric fish. We use the models developed from Part 1 to simulate the sensory acquisition and belief maintenance of the weakly electric fish during the prey-capture task. Using these simulations we can quantify the sensing performance of different types of trajectories. Results reveal how well the fish's sensing system has evolved for the prey-capture task.

Part 3 is dedicated to the implementation and evaluation of a robotic electrosensory system for locating underwater targets (a.k.a. the Electrosenster). In **Chapter 6** the experimental setup is described. **Chapter 7** develops a sensor model as well as characterizes the detection distance. Next, **Chapter 8** evaluates the performance of the several active controllers whose task is electrolocation. Finally, in **Chapter 9** a few more capabilities of the Electrosenster are demonstrated including global localization and a non-probabilistic-based controller.

Chapter 10 1) states the driving objectives of this work, and 2) discusses extensions of this work that include developing accurate models of sensorimotor transformations in

the weakly electric fish. Throughout this document we have chosen very simple models to investigate various aspects of weakly electric fish. While these models have provided some valuable results, they cannot accurately replicate sensorimotor transformations of the weakly electric fish. **Appendix F** discusses models of the afferent and early stages of electrosensory information that are grounded in the morphology, anatomy, and neurobiology of the weakly electric fish. **Appendices I and J** address the complexities of modeling controllers in biology. Once the intricacies of the fish's sensorimotor transformations are exposed, we may discover that more parsimonious models of sensor-based control (discussed in **Appendix G**) may accurately model the salient components. If not, then more complex approaches may be needed (discussed in **Appendix H**).

Part 1

Modeling Electrolocation

Part 1 investigates the fundamentals of electrolocation. We examine how voltage measurements of an electric field can be used to estimate the location of a target perturbing the electric field. In Chapter 2 two types of models are considered. The first employs a finite-element numerical approximation of the governing equations of the electric field. While accurate and general, this method is computationally expensive. The other method is only valid for spherical targets but is fast to compute since closed-form solutions can be found. This second method will be used in the remainder of this document for simulating electrosensory observations.

Two different electrosensory systems are presented. The first model—called the *dipole model*—represents the emitter/detector layout of the electrosense-based robot that will be investigated in Part 3. The second—called the *ellipse model*—is a simplified model of the weakly electric fish.

Once the fundamentals of electrolocation are established, Chapter 3 investigates sensor fusion techniques with noisy sensors. First, we consider the case when statistics about the sensor noise are not known, but the noise is bounded (possibilistic models). Second, Gaussian noise is assumed, and we explore belief maintenance techniques for target electrolocation using both a histogram filter and a particle filter. All the sensor fusion techniques are demonstrated on both the dipole and ellipse models.

CHAPTER 2

Modeling Electrosensory Observations

Electrolocation is the task of estimating the location of a target based on electrical measurements (i.e. voltage and/or current) near the target. In this chapter the electrolocation of circular and spherical targets based on voltage measurements from a self-generated electric field is examined.

Two methods of modeling electrosensory observations are considered in this chapter. The first method uses differential equations derived from Maxwell's Equations to describe the electric field. Then, a finite element solver numerically approximates the electric field. This method is very general but is computationally expensive. The second method uses an analytical model to predict electrosensory observations. This model is only valid for simple objects (spheres or ellipsoids), but is very fast to compute.

The ensemble of electrosensory viewpoints (EEV) maps the configuration of the target to the expected sensor observation. The EEV, along with a sensor model and sensor fusion scheme, can be used to determine the configuration of the target. This chapter demonstrates how to compute EEVs using both numerical methods and an analytical model for simple targets.

2.1. Numerical Construction of the EEV in 2-D

The purpose of this section is to illustrate how voltage measurements of an electric field can be used to estimate the location of a nearby target. We first consider a general

formulation derived from governing differential equations. A simple example of a conductor in an electric field is used. The example used is made as simple as possible while still providing the general intuition behind electric field sensing. The example consists of a single target near an electrostatic dipole producing DC current. In the interest of simplicity the following assumptions are made:

- (1) linear and isotropic media
- (2) time-invariant fields
- (3) ignore the presence of magnetic fields
- (4) two-dimensional (planar)

2.1.1. Formulation

We begin with the fundamental equations for electric fields and then use the assumptions above to simplify the model.

2.1.1.1. Fundamental Equations. The physics of all electric and magnetic waves are mathematically described by a set of four differential equations commonly referred to as Maxwell's Equations [120] (terms are defined in Table 2.1):

$$(2.1) \quad \nabla \cdot \mathbf{D} = \rho$$

$$(2.2) \quad \nabla \cdot \mathbf{B} = 0$$

$$(2.3) \quad \nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}$$

$$(2.4) \quad \nabla \times \mathbf{H} = \mathbf{J} + \frac{\partial \mathbf{D}}{\partial t}$$

| | |
|----------|---|
| E | electric field |
| H | magnetic field |
| D | electric displacement field |
| B | magnetic flux density |
| ρ | free electric charge density, not including dipole charges bound in a material |
| J | free current density, not including polarization or magnetization currents bound in a material |

Table 2.1. Definition of terms for Maxwell's equations (Equations (2.1) through (2.4)).

Also of interest is the equation of continuity of charge. To obtain the continuity equation, begin with Equation (2.4), then take the divergence of both sides. Because the divergence of a curl is zero, the result is:

$$(2.5) \quad \nabla \cdot \nabla \times \mathbf{H} = \nabla \cdot \mathbf{J} + \frac{\partial \nabla \cdot \mathbf{D}}{\partial t} = 0$$

If Maxwell's Equation (2.1) is substituted into Equation (2.5), then the continuity equation for current density is obtained. This relation states that the divergence of the current density is equal to the negative rate of change of the charge density, and shown here:

$$(2.6) \quad \nabla \cdot \mathbf{J} = -\frac{\partial \rho}{\partial t}$$

Next, it is shown that by applying the four simplifying assumptions to Equations (2.1) through (2.6) Poisson's equation will emerge. Poisson's equation will serve as the basis for the numerical simulations.

2.1.1.2. Application of the Four Assumptions. Each of the four assumptions listed above will be applied to the fundamental equations.

(1) **Linear and isotropic media.** By invoking this assumption, the following constituent relations can be assumed (σ is the electrical conductivity; ε is the electrical permittivity; and μ is the magnetic permeability):

$$D = \varepsilon E$$

$$B = \mu H$$

$$(2.7) \quad J = \sigma E$$

where σ , ε , and μ are time-independent scalars.

(2) **Time-invariant fields.** One problematic complication of Maxwell's equations is that the time rate of change of \mathbf{B} affects the electric field (see Equation (2.3)), and the time rate of change of \mathbf{D} affects the magnetic field (See Equation (2.4)). Thus, Maxwell's equations must be solved simultaneously in the case of time-varying fields. When these fields become time-invariant, then this type of coupling disappears, and the two fields can be determined independent of each other. This is why static-field problems are relatively simple. Therefore, to greatly reduce the complexity of the problem, time-invariant fields are assumed here.

(3) **Ignore the presence of magnetic fields.** Even with time-invariant fields, there is still coupling between electric and magnetic fields through Equation (2.3), which arises from conduction currents. Since $J = \sigma E$, Equation (2.3) can be written as $\nabla \times \mathbf{H} = \sigma E$,

which means the magnetic field is dependent on the (time-invariant) electric field. But, because of combination of Equations (2.3), (2.6), and (2.7), the sources of the electric field are independent of the magnetic field. Thus, the coupling between electric and magnetic fields (under these strict assumptions) resulting from conduction currents is unidirectional. Furthermore, it will be assumed that these magnetic fields will not interfere in any other way (i.e. The presence of magnetic fields will be ignored).

(4) **Two-dimensional (planar).** The dimension of the example will be kept as low as possible for easy computation and visualization.

Given the above four assumptions, Equations (2.1) through (2.6) are reduced to Poisson's equation for pure conduction current:

$$(2.8) \quad \boxed{-\nabla \cdot (\sigma \nabla V) = \rho}$$

This is the equation that is used in the next section to numerically determine the electric field with a target nearby.

2.1.2. Numerical Simulation of Electric Field

Equation (2.8) is now used to develop a numerical simulation to demonstrate how targets of impedances different than their surrounding medium perturb an electric field.

Simulations are implemented with MATLAB's Partial Differential Equation Toolbox (The Mathworks, Natick MA, USA), which is a finite-element solver. The simulation uses Poisson's equation (Equation (2.8)) with both Dirichlet and Neumann type boundary conditions. For electrostatic problems, the Dirichlet type boundary condition takes the

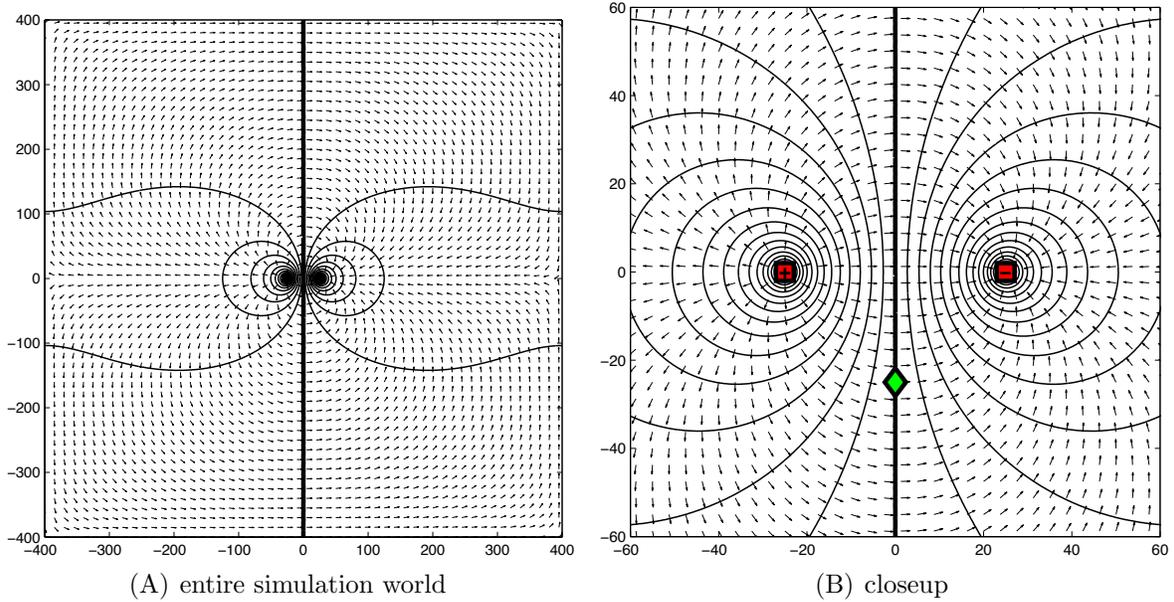


Figure 2.1. Unperturbed electric field for the simulation world. Arrows show the direction of the electric field. Contours are isopotentials (constant voltage) with the heavy line at zero volts. Contours are at 50 mV increments. In subfigure (B) the emitters are represented with red squares. The center of the green diamond represents the point of voltage observation.

form:

$$(2.9) \quad hV = r$$

and the Neumann type boundary condition takes the form:

$$(2.10) \quad \mathbf{n} \cdot (\sigma \nabla V) + qV = g$$

where \mathbf{n} is a vector normal to the boundary surface; h , r , q , and g are parameters that define the nature of the constraint (these will be assigned later).

The ideal simulation would compute the electric field generated by a simple dipole in an infinite medium. But, since the finite element simulation requires a finite domain, a “very large” medium is used. The modeled electric field is generated by two 1 mm diameter disks that are centered at $y = 0$ and $x = \{-25, 25\}$. For the disk at $(x = -25, y = 0)$ Dirichlet boundary conditions with $h = 1$ and $r = 1$ are used (see Equation (2.9)). This constrains the electrostatic potential at the disk’s surface to 1 volt (essentially making it a voltage source). Similarly, the Dirichlet boundary conditions of $h = 1$ and $r = -1$ are imposed at the disk centered at $(x = +25, y = 0)$ (making this a voltage source of -1 volt).

The “infinite” medium is approximated as an 800 mm by 800 mm square centered at $(x = 0, y = 0)$. Neumann boundary conditions with $g = 0$ and $q = 0$ (see Equation (2.10)) are imposed at the four edges of the square. In electrostatic problems with $q = 0$, Neumann boundary conditions define the normal component of the electric field. With $g = 0$ this particular constraint says there is no normal component of the electric field at the square’s edges. This is equivalent to the square’s edges being ideal electrical insulators. Figure 2.1 shows the electric field for this setup.

2.1.3. Electrosensory Viewpoint

In Figure 2.1 no target is present to distort the electric field generated by the electric dipole (two red squares). In this situation, a sensor located at position $(x = 0 \text{ mm}, y = -25 \text{ mm})$ —which is symbolized by the green diamond—lies on the iso-potential contour of 0 V. In Figure 2.2 a circular perfect conductor of radius 3 mm is centered at coordinates $(x = 15 \text{ mm}, y = -15 \text{ mm})$. The presence of the conductor distorts the electric field, and

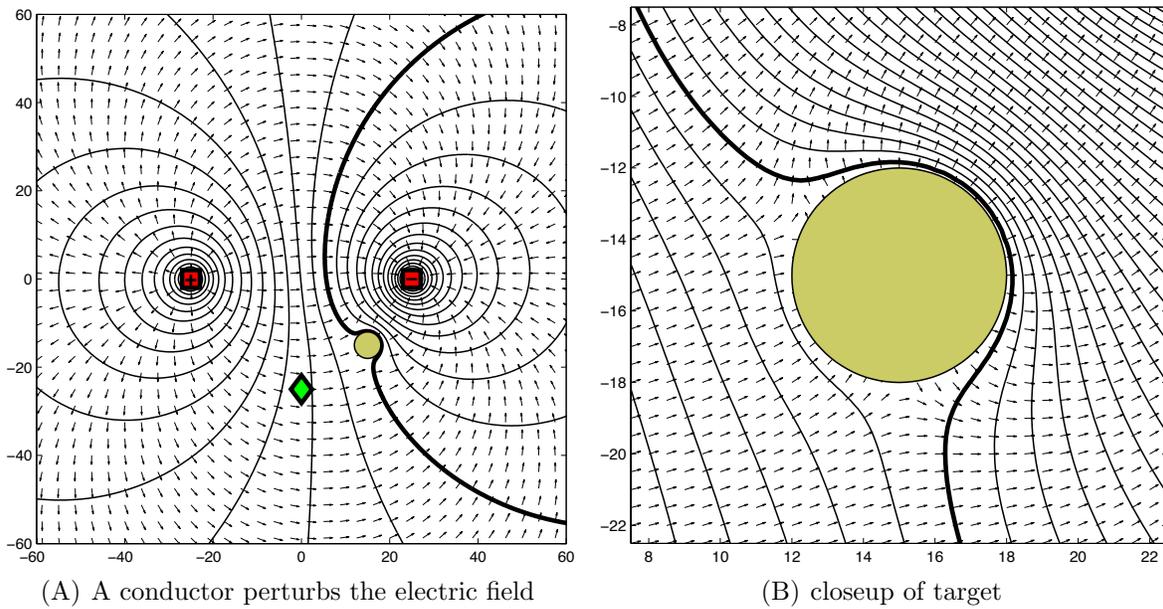


Figure 2.2. The electric field is perturbed by an electrically conducting disk of 3 mm radius. The emitters of the electric field are located at the two red squares. The green diamond sensor reads 112 mV. Since the target is a perfect conductor, no electric field can exist inside it. Thus, the surface of the target is an equipotential, and the electric field is perpendicular to the surface. The heavy contour line is zero volts. In subfigure (A) the contours are at 50 mV increments and are at 10 mV increments in subfigure (B).

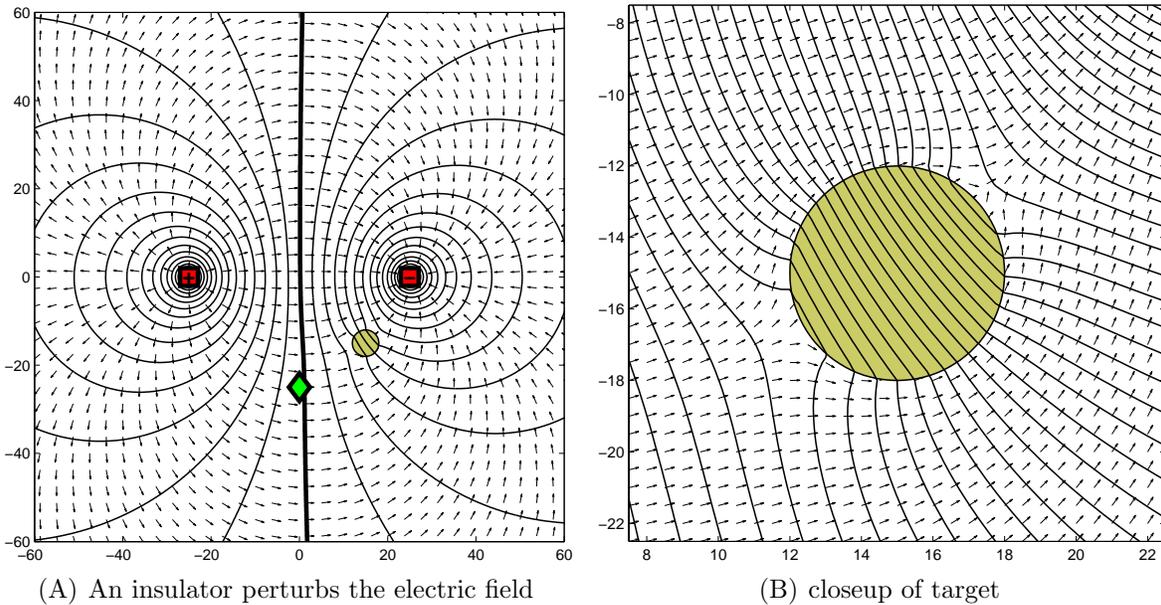


Figure 2.3. The electric field is perturbed by an electrically insulating disk of 3 mm radius. The emitters of the electric field are located at the two red squares. The green diamond sensor reads 9.5 mV. Since the target is a perfect insulator, the electric field is tangent to the surface. The heavy contour line is zero volts. In subfigure (A) the contours are at 50 mV increments and are at 10 mV increments in subfigure (B).

thus perturbs the original electrostatic potential. The sensor is now at $V_{ct} = 112$ mV, and the observation for this target at this location is $w = 112$ mV. We define the combination of target position and the observation the *electrosensory viewpoint* (EV) for these sets of conditions. If \mathbf{x} is the location of the target relative to the sensor frame, and w is the measured voltage, then the EV is simply $\{\mathbf{x}, w\}$. In this example, the EV is $\{(15 \text{ mm}, -15 \text{ mm}), 112 \text{ mV}\}$.

In Figure 2.3 the conducting target from the previous example is replaced with an insulating target. As before, its presence perturbs the nominal electric field. The green diamond sensor will observe 9.5 mV in this case. Thus, the EV is $\{(15 \text{ mm}, -15 \text{ mm}), 9.5 \text{ mV}\}$.

2.1.4. The Ensemble of Electrosensory Viewpoints

The collection of EVs for every possible target location composes the *ensemble of electrosensory viewpoints* (EEV). The functional form of the EEV maps the target position to the observed sensor measurement, and can be written as:

$$(2.11) \quad \text{EEV}_{\mathbf{e}}(\mathbf{x}) = w$$

The subscript \mathbf{e} denotes unchanging aspects of the environment that the EEV depends on, which here includes the size, shape, and material properties of the target, the applied electric field, and the conductivity of the medium.

An EEV can be created for the 6 mm diameter conducting disk used previously. The disk is placed at each of the 144 points of a 12 by 12 grid and the simulated voltage observed by the green diamond is recorded. For example, above it was shown that if the disk is placed at location $(x = 15, y = -15)$, then the sensor observes 112 mV. In Figure 2.4 this location is represented as the large yellow dot, and the value of +112 mV is assigned to the cell centered at $(15, -15)$. Once this process is repeated for the remaining 143 yellow dots, the EEV shown in Figure 2.4 emerges.

Figure 2.4 can be used to localize the 6 mm diameter disk. Assume it is known that only the 6 mm conducting disk is perturbing the electric field, and the sensor is noise-free. Also, the disk is constrained to be only at one of the 144 grid locations. If the sensor observes 112 mV, and if the cell centered at $(15, -15)$ in Figure 2.4 is the only disk location that rendered 112 mV, then it could be concluded with certainty that the disk is located at position $(15, -15)$. Of course this very simple case does not accurately portray

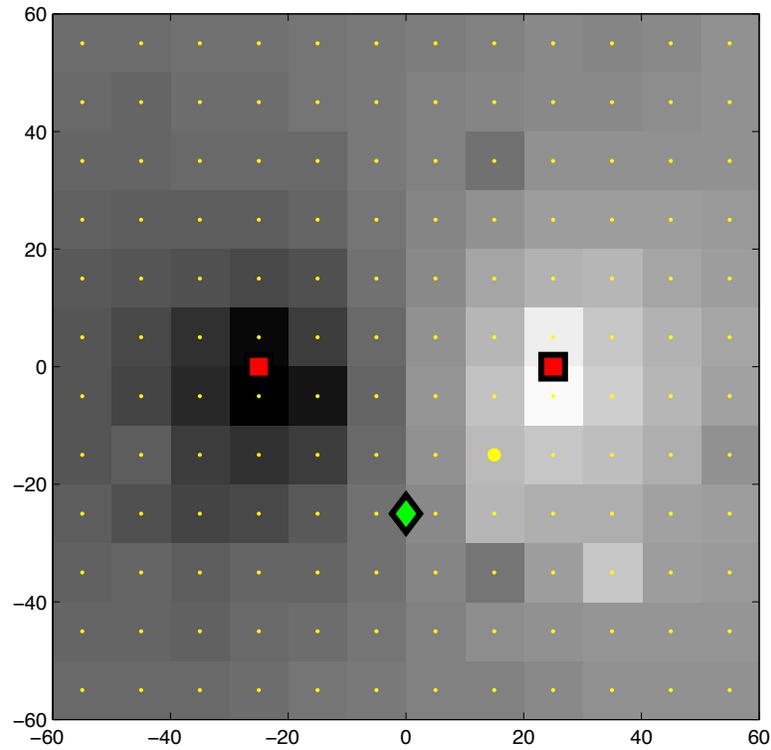


Figure 2.4. The EEV for the 6 mm diameter conducting disk. The two red squares generate the electric field. The green diamond is the voltage sensor. Yellow dots represent all the locations the target was placed during the construction of the EEV. The color of each cell represents the observed voltage, where white is the largest positive observation and black is the largest negative value.

real-world target localization situations. In forthcoming chapters it will be shown how to electrolocate targets under more realistic conditions.

2.2. An Analytical Solution to the EEV in 3-D

A finite-element numerical simulation of electric fields is a computationally expensive method for estimating the voltage observation, making it impractical to use in 3-D simulations. Luckily, there is an analytical solution for predicting electrosensory observations induced by simple targets perturbing electric fields. Also, we now only consider electric fields that are generated by point charges. This assumption allows for quick calculation of the electric field.

2.2.1. Electric Field Generated by Point Charges

In the previous section the 2-D electric field needed to be computed for all nodes, even if only the electric field for a single point in space was needed. Fortunately, the electric fields examined in this document can be closely approximated by a computationally-efficient, analytical solution, which uses point charges to generate the electric field. All models of electric fields throughout the remainder of this document will be assumed to be generated by ideal point charges.

If $\mathbf{x}_q \in \mathbb{R}^3$ represents the charge-centered relative coordinates, then an isolated point charge defines an electric field for every point in its vicinity by the Equation [235]:

$$(2.12) \quad \mathbf{E}_f(\mathbf{x}_q) = \frac{kQ}{|\mathbf{x}_q|^2} \hat{\mathbf{x}}_q$$

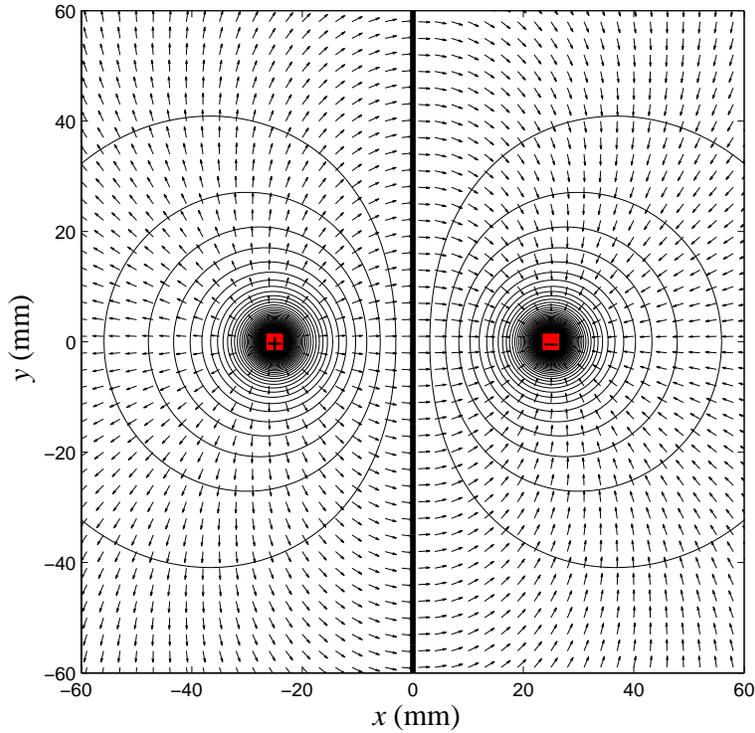


Figure 2.5. Electric field for dipole model. Contours are at 10 mV increments. $Q_+ = 1.11 \cdot 10^{-13}$ C and $Q_- = -1.11 \cdot 10^{-13}$ C.

where Q is the magnitude of the charge (in Coulombs), $\hat{\mathbf{x}}_q$ is a unit vector from the charge to \mathbf{x}_q , $|\mathbf{x}_q|^2$ is the square of the magnitude of \mathbf{x}_q ; and k is the Coulomb constant ($8.988 \times 10^9 \text{ N} \cdot \text{m}^2/\text{C}^2$).

The electric field at each point in space is a vector and obeys the superposition principle. That is, the net electric field due to a system of charges can be found by computing the electric field due to each charge in the system separately and then adding these vectors to obtain the net electric field. Simply, $\mathbf{E} = \sum_i \mathbf{E}_i$. Figure 2.5 shows an example of the resultant electric field generated from two point charges (an electric dipole).

2.2.2. An Analytical Model for Electrosensory Observations

We now consider an idealized 3-D model for the observed voltage perturbation created by a sphere of known diameter and conductivity in water of known conductivity [187]. This model gives an analytical solution of the $EEV_e(\cdot)$ for spheres.

If we let $\mathbf{x}_t \in \mathbb{R}^3$ represent the target-centered relative coordinates, then the model is represented as:

$$(2.13) \quad \boxed{\delta V(\mathbf{x}_t) = \frac{a^3 \mathbf{E}_f \cdot \mathbf{x}_t}{|\mathbf{x}_t|^3} \left(\frac{\sigma_{\text{target}} - \sigma_{\text{water}}}{\sigma_{\text{target}} + 2\sigma_{\text{water}}} \right)}$$

where $\delta V(\mathbf{x}_t)$ represents the change in potential (mV) at position \mathbf{x}_t (cm); a is the target's radius (cm); \mathbf{E}_f is the electric field vector at the location of the target (mV/cm) when no target is present; σ_{target} is the conductivity of the target ($\mu\text{S}/\text{cm}$), and σ_{water} is the conductivity of the water. The right term of Equation (2.13) within the parentheses is called the *electrical contrast factor* and denoted χ . The electrical contrast factor varies from $\chi = -1/2$ (for perfect insulators) to $\chi = +1$ (for perfect conductors). When the conductivity of the target and water are equal, the electrical contrast is zero and the target becomes electrically invisible. It should be noted, however, that this model ignores phase shifts due to capacitive components of the target impedance. In weakly electric fish, such phase shifts are detected by a different electrosensory system [175]. Such phase shifts are likely used to differentiate between inanimate objects and live objects, which have capacitance due to biological membranes [245]. This simple model also assumes the electric field is uniform across the target. As the target approaches close to the source of the electric field, this assumption will be violated.

Equation (2.13) represents the model of electrosensory observations that will be used for the remainder of this document and will be referred to extensively. Thus, we will restrict the simulations to locating spheres.

We now have the tools to construct $\text{EEV}_{\mathbf{e}}(\cdot)$ for spheres. First, Equations (2.12) and (2.13) will be converted from the charge-centered and target-centered coordinate systems to a more convenient sensor-centered coordinate system. If \mathbf{x}_s^q is the position of the charge in the sensor-centered coordinate frame, then in Equation (2.12) let $\mathbf{x}_q = \mathbf{x}_s - \mathbf{x}_s^q$. To convert coordinates in Equation (2.13) let $\mathbf{x}_t = -\mathbf{x}_s$. Now, let $\text{EEV}_{\mathbf{e}}(\mathbf{x}_s) = w = \delta V(\mathbf{x}_s)$. Thus, for a single sensor and single point charge, the EEV in the sensor frame can be defined as:

$$(2.14) \quad \text{EEV}_{\mathbf{e}}(\mathbf{x}_s) = \left(\frac{a^3 k Q}{|\mathbf{x}_s|^3 |\mathbf{x}_s - \mathbf{x}_s^q|^3} \right) \left(\frac{\sigma_{\text{target}} - \sigma_{\text{water}}}{\sigma_{\text{target}} + 2\sigma_{\text{water}}} \right) \left(-\mathbf{x}_s \cdot (\mathbf{x}_s - \mathbf{x}_s^q) \right)$$

Equation (2.14) is a function of \mathbf{x}_s , \mathbf{x}_s^q , Q , k , a , σ_{target} , and σ_{water} . The parameter k is a universal constant; and \mathbf{x}_s^q , Q , a , σ_{target} , and σ_{water} are defined by the \mathbf{e} vector, which denotes unchanging aspects of the environment that the EEV depends on. It is assumed that the charge is rigidly connected to the sensor (this is the case for both systems examined in this document), meaning \mathbf{x}_s^q is constant and included in \mathbf{e} . Thus, $\text{EEV}_{\mathbf{e}}(\mathbf{x}_s)$ is a function of only \mathbf{e} and \mathbf{x}_s .

Of course if the point charge in Equation (2.14) is truly electrically isolated, then no current will flow to or from the charge. Without flowing current it is impossible to record a voltage. But, this model would work very well for a point charge in a large, electrically-grounded box.

For the remainder of this document two electrosensory systems will be investigated. In the next two sections these two different emitter/detector layouts are examined using the analytical model for sensor observation (Equation (2.13)). The first is a model of the robotic electrolocation system, and the second is a simplified model of the weakly electric fish.

2.3. Dipole Model

The first emitter/detector setup considered is that of the robotic electrolocation system (*The Electrosenster*) that will be thoroughly discussed in forthcoming chapters. The electric field is generated by two (approximately) point emitters, thus this setup is referred to as the “dipole model”. With the dipole model there are two sensors, and the model lives in full 3-D space, although only a 2-D slice will be examined.

The Electrosenster is a 2-DOF XY robot (similar to a plotter) that moves its 2 emitters and 2 detectors through a tank of water. The emitter/detector layout can be seen in Figure 2.6. It is able to electrolocate targets by sensing voltage changes in its self-generated electric field. We use the simple 3-D models of electric fields (Equation (2.12)) and observations (Equation (2.13)) induced by spheres. A single observation, w , is the difference of the two sensor voltages.

Both the emitters and detectors are constrained to move in the plane $z = 0$. Each emitter and detector is approximately a 1 mm radius metal sphere. In the Electrosenster, the potential across the two emitters is 2 volts. But, recall that the model of electric field uses point charges. If we assume the electric field is generated by two point charges of $+q$

and $-q$ located at positions \mathbf{x}_p and \mathbf{x}_n , respectively, then the electric field at position \mathbf{x} can be computed as

$$(2.15) \quad \mathbf{E}(\mathbf{x}) = \frac{q k}{|\mathbf{x} - \mathbf{x}_p|^3}(\mathbf{x} - \mathbf{x}_p) - \frac{q k}{|\mathbf{x} - \mathbf{x}_n|^3}(\mathbf{x} - \mathbf{x}_n)$$

where k is Coulomb's constant. If the emitters are assumed to be spherical conductors, then the equivalent charge for an emitter of radius r_{emitter} and voltage $\pm V_{\text{emitter}}$ is $q_{\text{eq}} = V_{\text{emitter}} r_{\text{emitter}} / k$. Thus, if $r_{\text{emitter}} = 1$ mm, then $q_{\text{eq}} = 1.11 \cdot 10^{-13}$ C.

To prevent the emitters and detectors from making contact with the target, the center of the target is placed at a depth of $z = -(r_{\text{target}} + 3 \text{ mm})$, where r_{target} is the radius of the target. With $r_{\text{emitter}} = 1$ mm, there is 2 mm clearance between the bottom of the emitters and the top of the sphere. Complete details of the Electrosenser and its operating conditions can be found in Chapter 6.

We first consider a hypothetical example to illustrate the dipole model. In Figure 2.6(A) no target is present to distort the electric field generated by the two red squares. In this situation both the top sensor (sensor 1) and bottom sensor (sensor 2) lie on the iso-potential contour of 0 V. In Figure 2.6(B) a spherical perfect conductor of radius 5 mm is centered at coordinates $(x = 20 \text{ mm}, y = -20 \text{ mm}, z = -8 \text{ mm})$. The presence of the conductor distorts the electric field, and thus perturbs the original electrostatic potential. Sensor 1 is now at $V_1 = 0.62$ mV, and sensor 2 is now at $V_2 = 0.97$ mV. We define $\text{EEV}_{\mathbf{e}}(\mathbf{x}) = V_1 - V_2 = w$, where now $\mathbf{x} \in \mathbb{R}^3$. Since $V_1 - V_2 = -350 \mu\text{V}$, the electrosensory viewpoint is $\{(20 \text{ mm}, -20 \text{ mm}), -350 \mu\text{V}\}$.

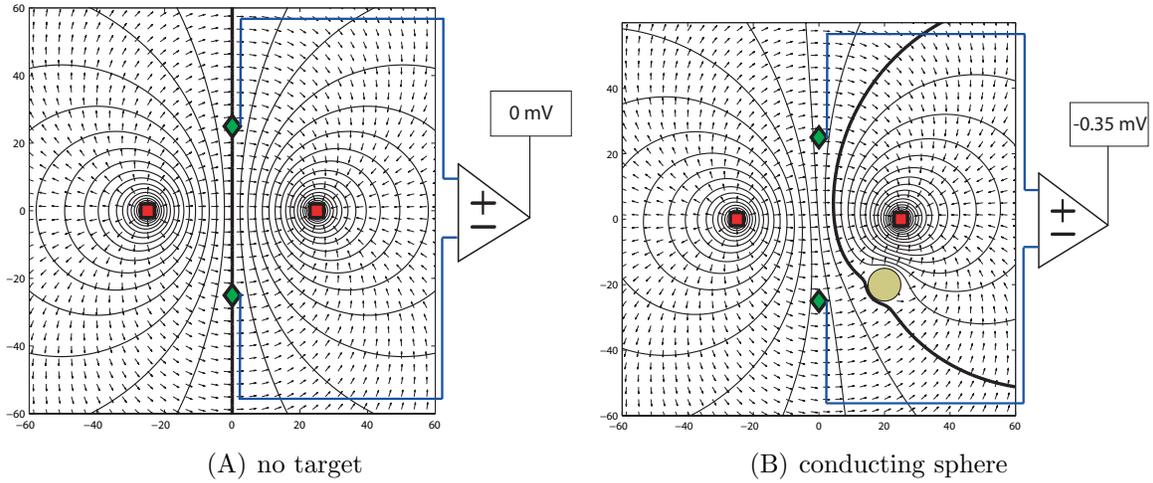


Figure 2.6. An illustration of the dipole emitter/detector layout. The left red square is $+1$ V and the right is at -1 V. Isopotential contours are shown has black lines. The bold line represents the 0 V contour. (A) Unperturbed electric field. Both green sensors observe 0 V. (B) A circular target of diameter 10 mm is centered at $x = 20$, $y = -20$. For example if the top sensor observes 0.62 mV, and the bottom sensor observes 0.97 mV, then, $w = -0.35$ mV.

In Figure 2.7 the EEV of a 10 mm diameter spherical ideal conductor is depicted as a filled contour plot with a gray colormap for the observation w as a function of the x and y coordinates of the center of the target. Figure 2.7 is a 2-D slice of the EEV at a height of $z = +8$ mm relative to the center of the target ($+z$ points out of the x - y plane). As before, an electric field is applied at the red squares, and two sensors (green diamonds) measure voltage. If a single measurement gives $w = -350$ μ V, the target must lie somewhere on one of the two yellow contours in Figure 2.7 (target not shown in this figure). These contours are the set $\{\mathbf{x} \mid \text{EEV}_e(\mathbf{x}) = -350$ μ V $\}$.

The isopotential perturbation contours of Figure 2.7 therefore represent the theoretical limit on how well a target of known properties can be localized by a single noiseless observation of the perturbation in 2-D. Each contour in Figure 2.7 represents the 1-D

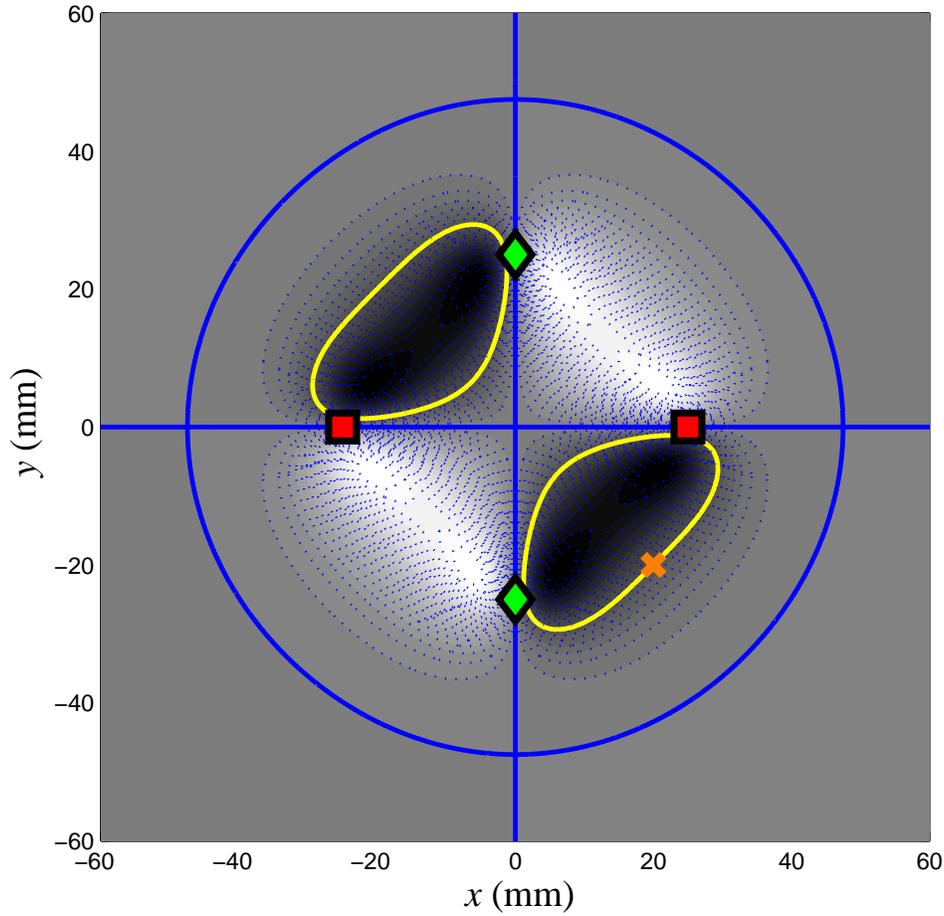


Figure 2.7. An example of an ensemble of electrosensory viewpoints (EEV). The target is a 10 mm diameter ideal conductor. The $EEV_e(\mathbf{x})$ is computed from Equation (2.13). To construct this EEV, we successively place the simulated sphere at each position on a grid covering the illustrated domain. For each place on this grid, we compute the perturbation *at each of the green diamonds (sensors)* according to Equation (2.13), given the field applied at the red squares, and take the difference between these perturbations. Off-grid values are interpolated. White represents the largest positive observation, and black is the largest negative observation. Intermediate grays are linear interpolations of these extreme values. The solid blue contour is $w = 0$ V. The yellow contour is the localization subspace for the single observation of $w = -350 \mu\text{V}$.

“localization subspace” of a target for the given sensor observation. It is impossible to disambiguate the location of any single target from the location of any other single target (observed at a different time) located on the same contour. The length of the 1-D subspace is a measure of the “localization quality”; shorter subspaces are better for localization than longer ones.

2.4. Ellipse Model

The other emitter/detector setup examined in this document is inspired by the weakly electric fish. It was shown by [12] (and later verified by [8] and [71]) that the electric field of weakly electric fish can be modeled as a line of uniformly distributed electric poles along the rostro-caudal axis of the model fish¹. The model has one negative pole located at the “tail” of the fish, and the remaining m are positive poles. The net charge of the poles is zero, thus, the magnitude of the single negative pole is m times greater than the m positive poles. The model of the electric field with n total poles, one negative pole, and $m = n - 1$ positive poles is [71]:

$$(2.16) \quad \mathbf{E}(\mathbf{x}) = -\frac{q}{|\mathbf{x} - \mathbf{x}_p^n|^3}(\mathbf{x} - \mathbf{x}_p^n) + \sum_{i=1}^m \frac{q/m}{|\mathbf{x} - \mathbf{x}_p^i|^3}(\mathbf{x} - \mathbf{x}_p^i)$$

\mathbf{x}_p^i is the position of the i th pole. The quantity q is analogous to electric charge in an electrostatic model and is distributed such that the first m poles have a “charge” of q/m and the remaining poles have a charge of $-q/(n - m)$, resulting in a total net charge of zero. The poles are uniformly distributed along the midline from head to tail. This model has been experimentally verified, and is depicted in Figure 2.8. A charge of $q = 10$

¹Other models are marginally more similar to the actual electric field of the fish [11]. But, these models requires numerical methods, which are much more computationally expensive.

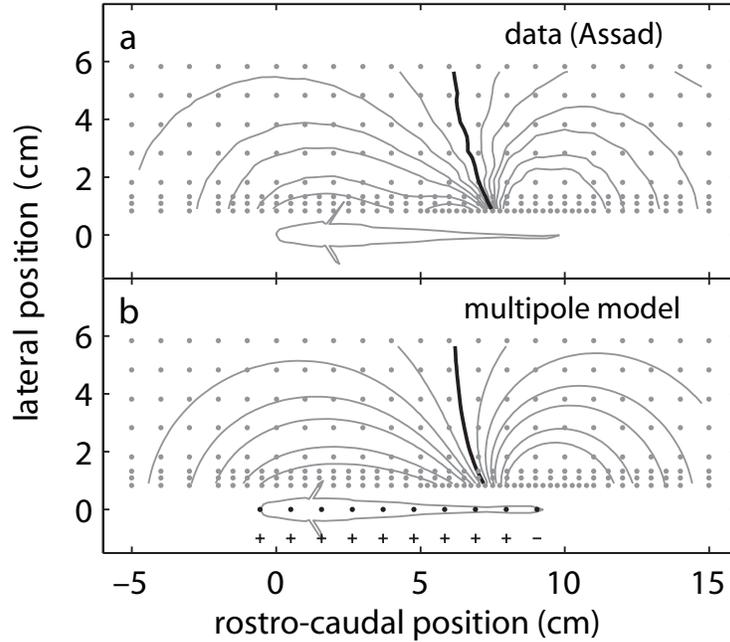


Figure 2.8. Comparison of multipole model with measured data from an actual fish. This figure is taken from [71].

$\text{mV}\cdot\text{cm}$ (equivalent to $q = 10^{-4} \text{ N}\cdot\text{m}^2/\text{C}$) will be used, which was empirically estimated from [71]. In this chapter only a 2-D slice will be considered.

Figure 2.9 depicts the 2-D slice of the electric field and iso-potential contours of the ellipse model. Figure 2.10 shows the EEV for four different sensors located on the ellipse. The yellow contour is the localization subspace for each sensor if a 10 mm conducting sphere is located at $(x = 10, y = 15, z = 0)$ (this position is shown as an orange “x”). Note that a target near the tail would create a larger voltage perturbation at a sensor near the tail than a target near the head for a sensor near the head (the white region in Figure 2.10(C) is larger than the black region in Figure 2.10(A)). This is because the charge at the tail is 99 times greater in magnitude than the other 99 charges, thus the

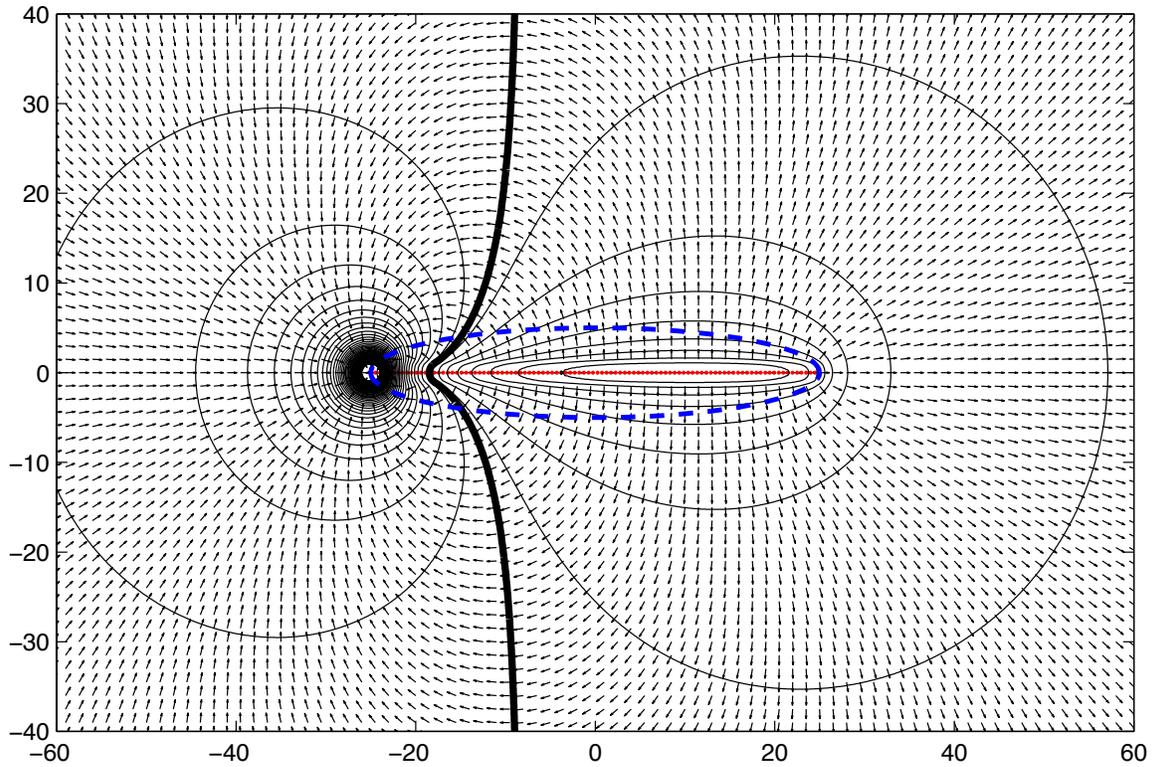


Figure 2.9. The electric field generated from a line of charges. Arrows show the direction of the electric field. Black curves are iso-potential contours (lines of constant voltage); the thick line is $V = 0$ volts. Red dots are the locations of the 100 poles. The blue dashed line is the edge of an ellipse that is 50 mm long and 10 mm tall. This is the electric field model used in the ellipse model.

electric field is greater at the tail. According the Equation (2.13), larger electric fields create larger voltage perturbations at the sensor.

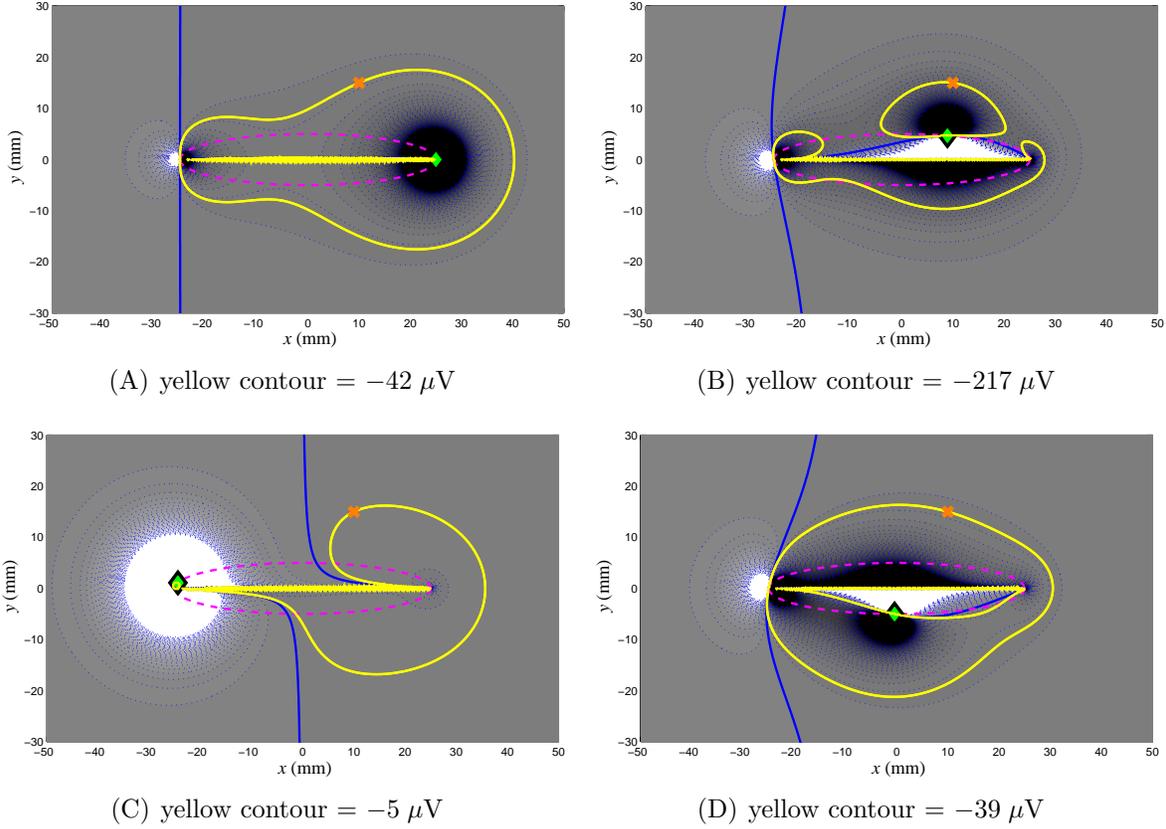


Figure 2.10. An example of four EEVs. The target is a 10 mm diameter spherical ideal conductor. The $\text{EEV}_e(\mathbf{x})$ is computed from Equation (2.13). White represents target locations that would result in an observation of greater than $+1,000 \mu\text{V}$, and black represents target locations that would result in an observation of less than $-1,000 \mu\text{V}$. Intermediate grays are linear interpolations of these extreme values. The solid blue contour is $w = 0 \text{ V}$. Blue dashed contours are at $25 \mu\text{V}$ increments. The yellow contour is the localization subspace for the single observation of a target centered at $(x = 10, y = 15, z = 0)$ (shown as an orange “x”). The center of the green diamond is location of the sensor. The pink dashed line is the edge of an ellipse that is 50 mm long and 10 mm tall.

CHAPTER 3

Belief Maintenance During Electrolocation

In the previous chapter models of idealized electrosensory observations were presented without any discussion of sensor noise. Any real measurement system has noise. In this chapter we show data fusion and belief maintenance techniques that account for sensor noise.

Two classes of beliefs are considered in this chapter—possibilistic and probabilistic. If the sensor noise is bounded, and if statistics about the observations are not available (or, if it is known that all possible observations are equally likely), then a *possibilistic* belief maintenance scheme should be used. Possibilistic belief maintenance can also be used for a “worst-case” analysis. Instead, if statistics about the observations are known, then a *probabilistic* belief maintenance scheme should be used. Two types of probabilistic belief maintenance schemes are presented in the chapter—the histogram filter and the particle filter.

At this point we have abandoned the numerical models of the electric field and electrosensory observations that were discussed in the first part of the previous chapter. Instead, the computationally-efficient analytical models discussed in Section 2.2 are used for the remainder of this document. Thus, targets are spheres and electric fields are generated by point charges. The two example systems used in this chapter are the dipole model (described in Section 2.3) and ellipse model (described in Section 2.4)

3.1. Possibilistic Belief Maintenance

First, consider a *possibilistic* noise model (also called *nondeterministic*; see [142]). If the noise is bounded, then possibilistic sensor fusion can be used. Possibilistic models assume all possible observations are equally likely. A possibilistic model simply prescribes whether or not something is *possible* without addressing the probability of the possibility.

3.1.1. Single Observation Possibilistic Electrolocation

A possibilistic noise model is now imposed on the electrolocation problem. Recall, Figure 2.7 shows the theoretical limit on how well a single voltage sensor observation can localize a target. Once noise is introduced into the picture, localization will become more uncertain; the *localization subspace* will change dimension, and the *localization quality* is reduced.

3.1.1.1. Single Observation with Dipole Model. Figure 3.1 demonstrates localization with a possibilistic model. The electric field, sensor locations, and conductivities of the target and medium are identical to those used in Figure 2.7. An observation of $w = -350 \mu\text{V}$ has been received (just like in Figure 2.7), and the noise of the sensor is known to be $\pm 200 \mu\text{V}$ and is assumed to be constant (not state or sensor dependent). Knowing only this sensor reading and noise level, we can conclude that the *noise-free* sensor reading could have been anywhere from -150 to $-550 \mu\text{V}$. The blue band represents all the locations a target could be that correspond to a -150 to $-550 \mu\text{V}$ observation. Thus, it is known with certainty that the target must lie within this band.

Figure 3.1 is constructed by first discretizing the 2-D world into a finite number of cells. Each cell maintains a binary variable, which defines whether or not it is possible the target is in this cell (1 for possible and 0 for not possible). Initially, when it is completely

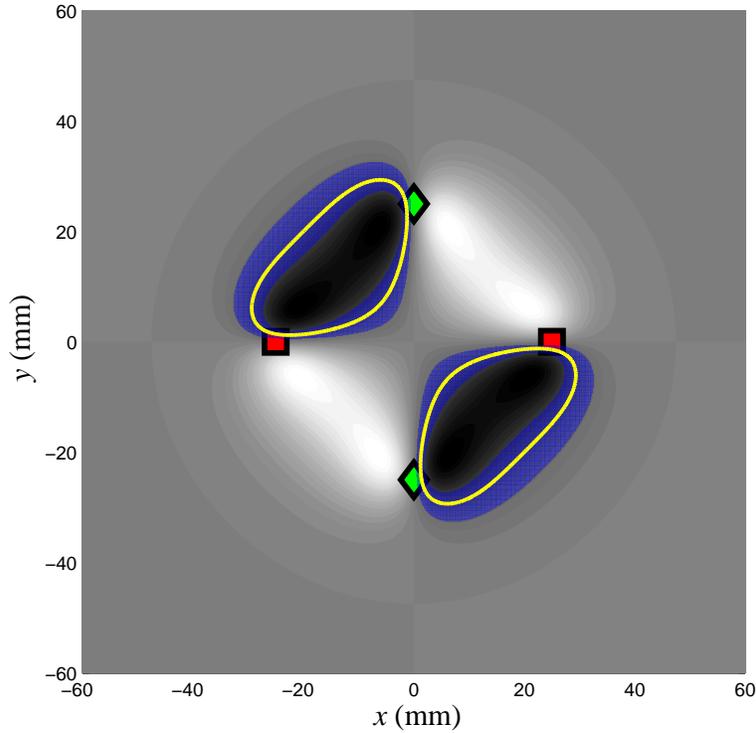


Figure 3.1. Possibilistic noise model with the dipole model. The target is a 10 mm diameter spherical ideal conductor. The $EEV_e(\mathbf{x})$ is shown in the background and is computed from Equation (2.13). White represents target locations that would result in a large positive observation, black are locations that result in largest negative observations, and grays are linear interpolations of these extreme values (this is the same as in Figure 2.7). A single observation of $w = -350 \mu V$ is received, thus if sensors were noise-free, then the belief of the target would be the yellow lines. But, the noise on the observation is $\pm 200 \mu V$, which results in a belief of the blue colored area. The blue area contains all possible target locations that would result in an observation of $w = -350 \pm 200 \mu V$.

uncertainty where the target is, all cells are set to 1. When new sensor data are received, all cells that are inconsistent with the new data are set to 0. The blue in Figure 3.1 represents all cell set to 1 and represents the belief of the target.

Formally, let \mathbf{x} be the world-frame position of the target, and \mathcal{X} and \mathcal{W} be the complete configuration space of the target and the space of all possible observations, respectively. Let $h : \mathcal{X} \rightarrow \mathcal{W}$ be a function mapping target position to (noise-free) observation. Thus, for our electrolocation scenario $h(\mathbf{x}) = \text{EEV}_{\mathbf{e}}(\mathbf{x}) = w$. The set $\mathcal{X}_1 \subseteq \mathcal{X}$ is a valid belief of the target position for a single observation, w_1 if

$$(3.1) \quad \mathcal{X}_1 = \{\mathbf{x} \in \mathcal{X} | (w_1 - n) \leq h(\mathbf{x}) \leq (w_1 + n)\}$$

where, $\pm n$ is the noise.

With a single observation of the noisy voltage sensors the localization subspace forms a 2-D area (in this 2-D slice); the localization subspace has gained an entire dimension from the noise-free case! Thus, the localization quality has also worsened.

3.1.1.2. Single Observation with Ellipse Model. The blue regions in Figure 3.2 represent the belief of the target location for each sensor making a single observation (no sensor fusion yet). In this particular example it is assumed that the noisy sensors did, in fact, observe the correct observation (i.e., the sensors observed the noise-free measurement). The belief of this noise-free observation is signified by the yellow contour in Figure 3.2. The blue shaded areas represent the belief assuming a possibilistic sensor model with $\pm 20 \mu\text{V}$ noise.

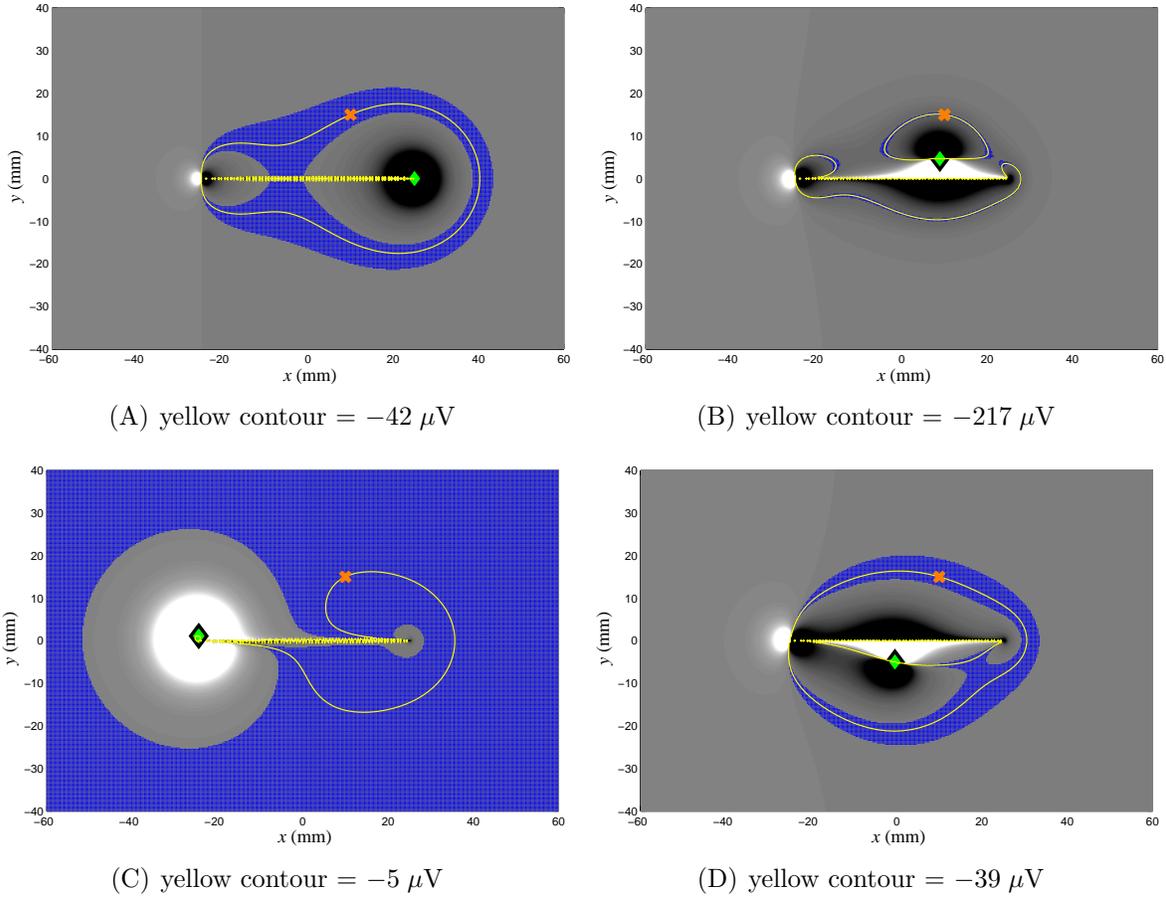


Figure 3.2. Possibilistic noise model with the ellipse model. Figures (A) through (D) show four separate beliefs and EEVs for the four different sensors. The blue shaded regions in the foreground represent the possibilistic beliefs. The target is a 10 mm diameter spherical ideal conductor. The $\text{EEV}_e(\mathbf{x})$ is computed from Equation (2.13). The EEV is shown behind the belief, where white represents target locations that would result in an observation of greater than $+1,000 \mu\text{V}$, and black represents target locations that would result in an observation of less than $-1,000 \mu\text{V}$. Intermediate grays are linear interpolations of these extreme values. The yellow contour is the noise-free localization subspace for the single observation of a target centered at $(x = 10, y = 15, z = 0)$ (shown as an orange “x”), and the value of the contour is stated in the caption of each figure. The center of the green diamond is the location of the sensor. The noise on each sensor observation is $\pm 20 \mu\text{V}$.

3.1.2. Fusion in Multi-Observation Possibilistic Electrolocation

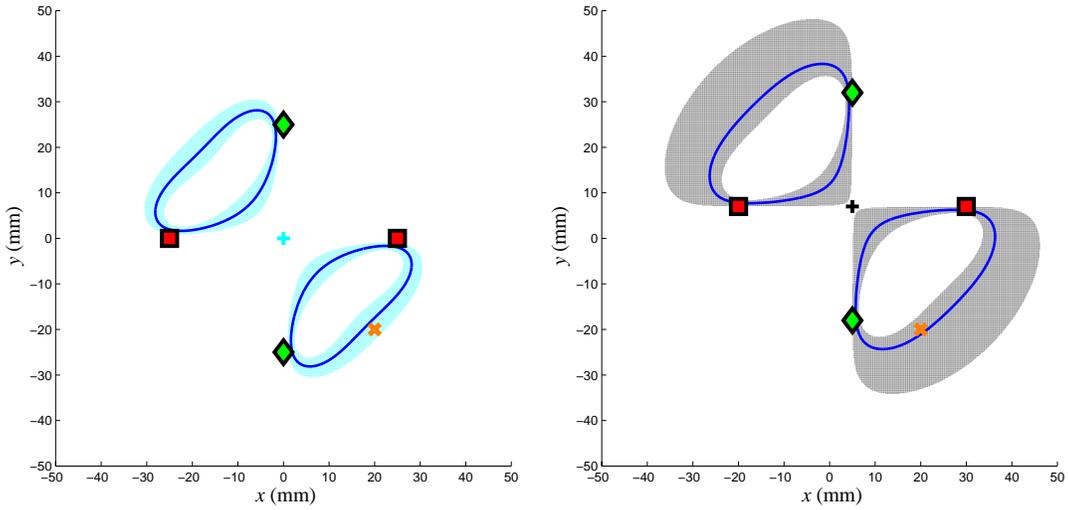
In Figure 3.2 beliefs were constructed for single observations, and possibly valuable information from other sensors were not utilized in the individual belief constructions. Sensor fusion is a crucial component of the belief construction process [104]. With the possibilistic sensor model that has been used, the data fusion process is merely taking the intersections of the beliefs in the configuration space of the target.

Formally, for all sensor observations, $w_i \in \mathcal{W}$, and beliefs, $\mathcal{X}_i \subseteq \mathcal{X}$ (constructed individually from Equation (3.1)), the minimal set $\mathcal{X}_{\text{fused}} \subseteq \mathcal{X}$ belief representation is:

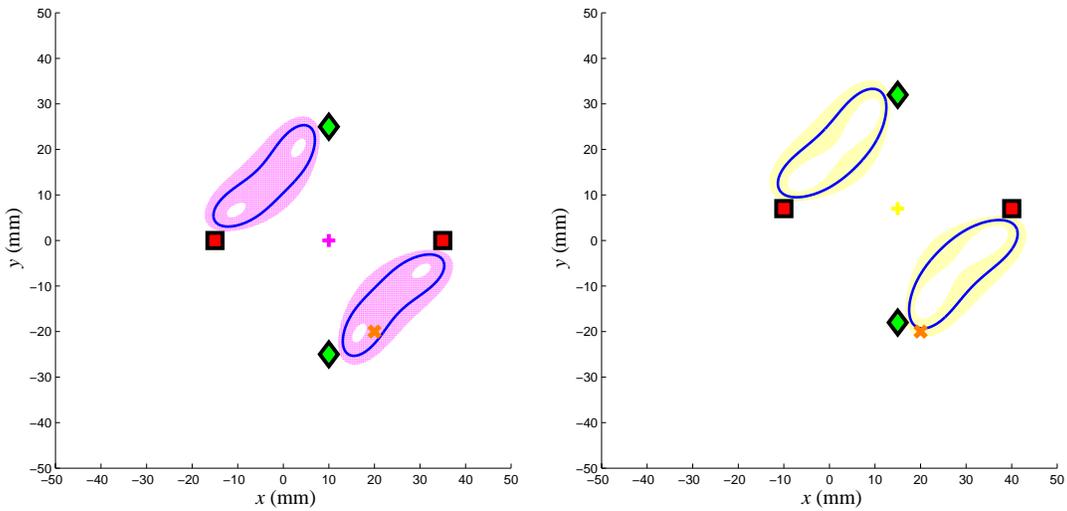
$$(3.2) \quad \mathcal{X}_{\text{fused}} = \bigcap_i \mathcal{X}_i.$$

3.1.2.1. Possibilistic Sensor Fusion with the Dipole Model. Figure 3.3 demonstrates possibilistic sensor fusion for the dipole model. In this example the noise on any observation is $\pm 200 \mu\text{V}$. Initially the robot is centered at location $(x = 0, y = 0)$ in the 2-D slice (the center of the robot is signified by a “+”). The robot makes a single observation, which results in the belief depicted in Figure 3.3(A). The robot then moves to location $(x = 5, y = 7)$. The belief for a single observation (no fusion) at that location is shown in Figure 3.3(B). Figures 3.3(C) and 3.3(D) show the single-observation beliefs when the robot is at positions $(x = 10, y = 0)$ and $(x = 15, y = 7)$, respectively.

When all four beliefs are superimposed on top of each other, the actual target location can only be in the regions where all the beliefs overlap. This union of the four beliefs is shown in Figure 3.4.



(A) $w_1 = -460 \mu\text{V}$ ($w_{1,\text{nf}} = -350 \mu\text{V}$ is the noise-free observation) (B) $w_2 = -210 \mu\text{V}$ ($w_{2,\text{nf}} = -250 \mu\text{V}$ is the noise-free observation)



(C) $w_3 = -750 \mu\text{V}$ ($w_{3,\text{nf}} = -870 \mu\text{V}$ is the noise-free observation) (D) $w_4 = -650 \mu\text{V}$ ($w_{4,\text{nf}} = -570 \mu\text{V}$ is the noise-free observation)

Figure 3.3. Single-observation possibilistic electrolocation for the dipole model. The shaded regions in Figures (A) through (D) represent the single-observation belief for each of the four emitter/detector locations. The actual location of the center of the 10 mm diameter spherical ideal conductor is $\mathbf{x}_{\text{target}} = (x = 20, y = -20, z = -8)$, and is signified by the orange “x.” The blue contour is the localization subspace if the actual observation, w_i , was considered noise-free. “+” is the center of the robot. $w_{i,\text{nf}} = h(\mathbf{x}_{\text{target}})$ would be the noise-free observation.

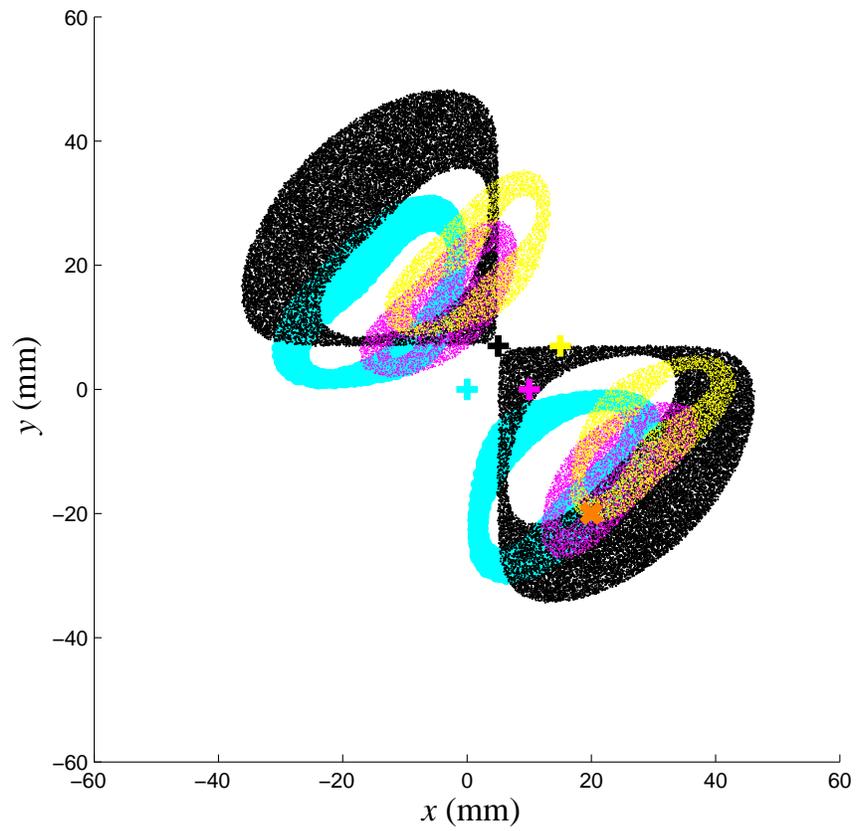


Figure 3.4. Four beliefs superimposed for dipole model. The shaded region is $\mathcal{X}_{\text{shaded}} = \bigcup_{i=1,2,3,4} \mathcal{X}_i$

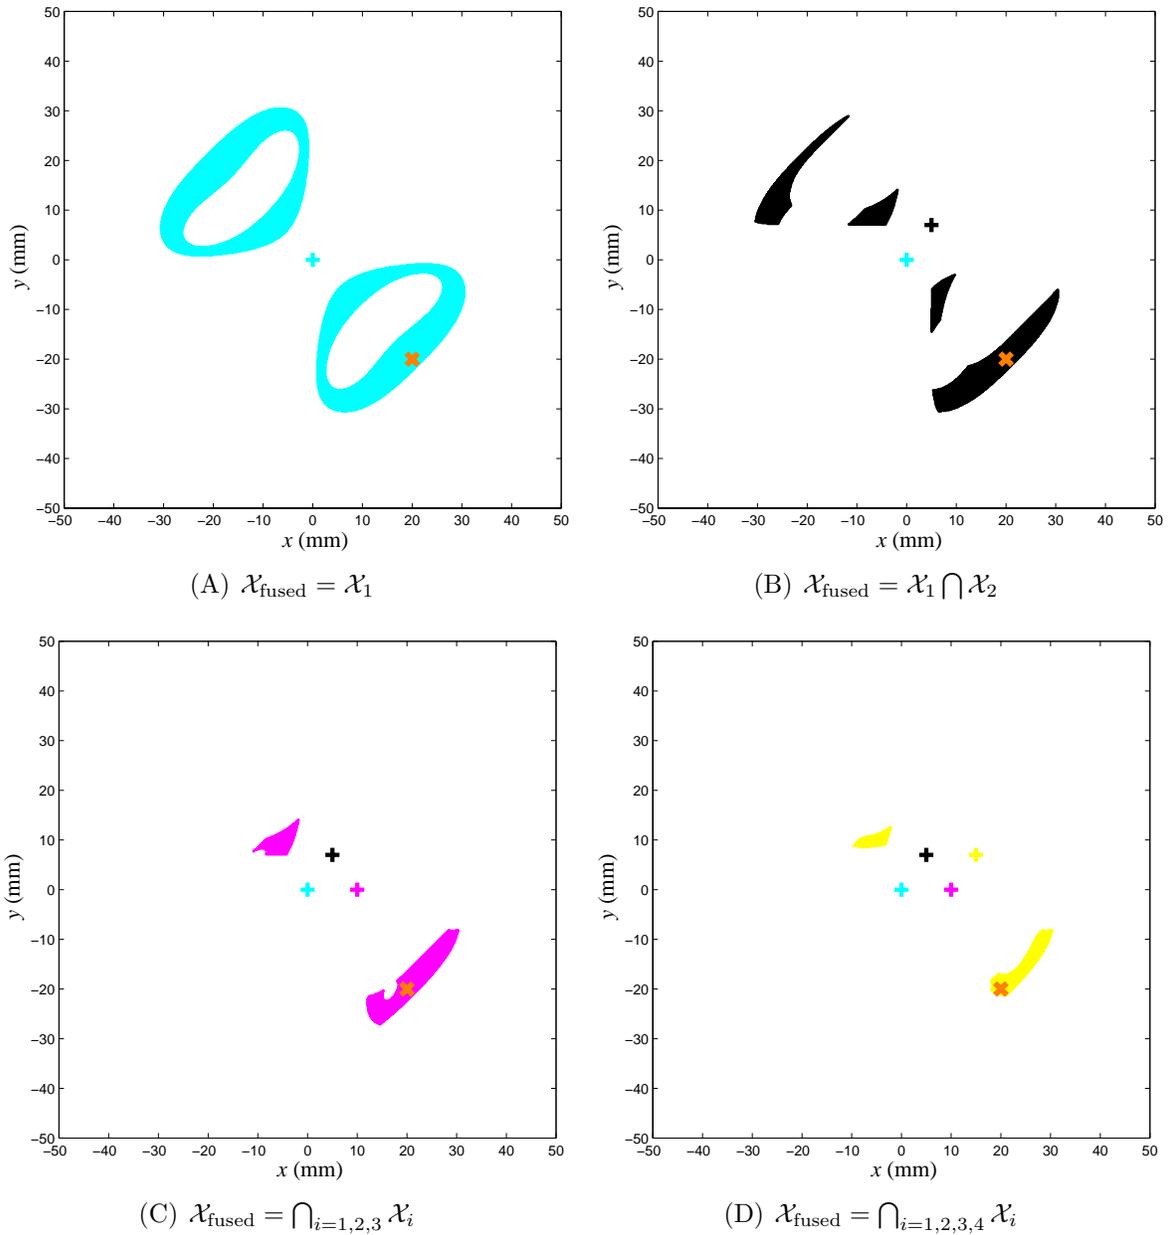


Figure 3.5. Fused belief after each observation in the dipole model. The shaded regions in (A) through (D) represent $\mathcal{X}_{\text{fused}}$, which is the fused belief after incorporating the current sensor observation into the prior belief. \mathcal{X}_i are the single-observation beliefs depicted in Figure 3.3

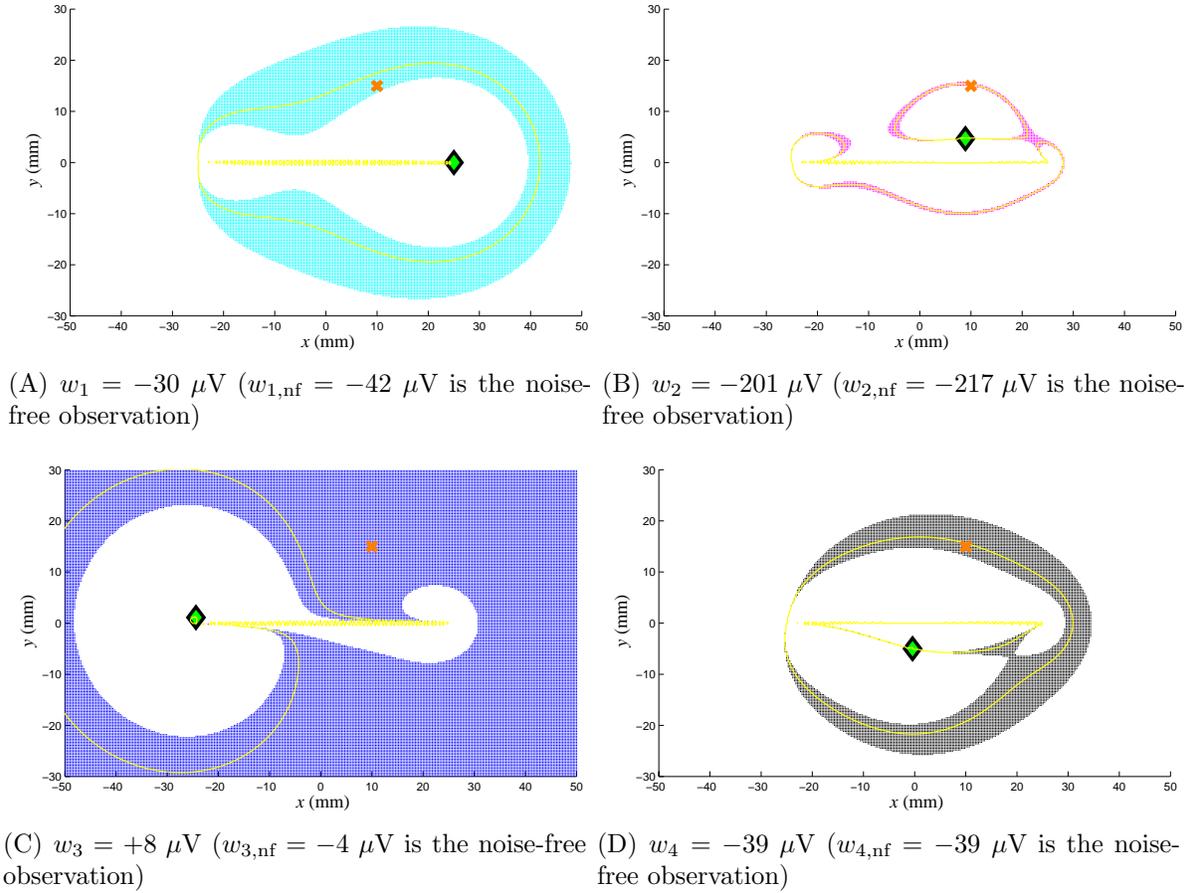
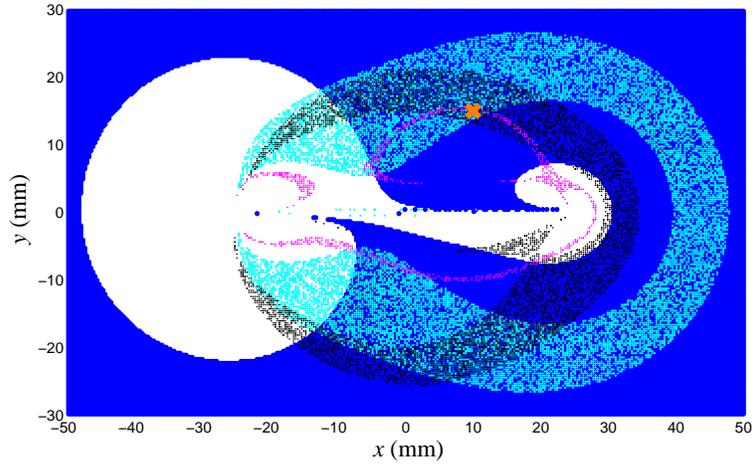
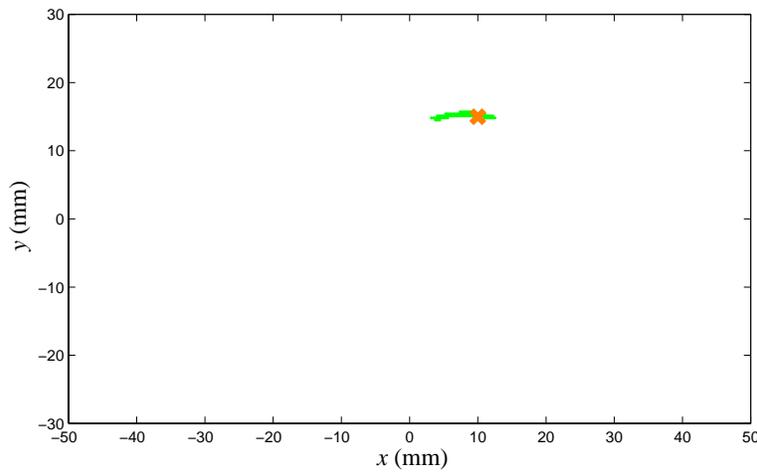


Figure 3.6. Possibilistic electrolocation. The actual location of the center of the 10 mm diameter spherical ideal conductor is $(x = 10, y = 15, z = 0)$, and is signified by the orange “x.” The yellow contour is the set of target locations that would render observations identical to the one actually received by the sensor. The shaded regions represent the individual beliefs (not fused) for each of the robot locations. The center of the green diamond is the location of the sensor. $w_{i,\text{nf}} = h(\mathbf{x}_{\text{target}})$ would be the noise-free observation.

To incorporate sensor information across observation, the intersection of the individual beliefs is used to create the minimal-set belief (see Equation (3.2)). The fused beliefs after each observation are shown in Figure 3.5; the target (orange “x” at location $(20, -20)$) does indeed always reside in the fused belief.



(A) Superposition of the four single-observation beliefs: $\mathcal{X}_{\text{shaded}} = \bigcup_{i=1,2,3,4} \mathcal{X}_i$



(B) The resultant belief is the intersection of the other four: $\mathcal{X}_{\text{fused}} = \bigcap_{i=1,2,3,4} \mathcal{X}_i$

Figure 3.7. Possibilistic sensor fusion for the ellipse model. The actual location of the center of the 10 mm diameter spherical ideal conductor is $(x = 10, y = 15, z = 0)$, and is signified by the orange “x.”

3.1.2.2. Possibilistic Sensor Fusion with the Ellipse Model. Figure 3.6 shows four separate observations of the same target. The green diamond shows the sensor location (different in each subfigure), and the orange “x” shows the actual location of the target (same position in each subfigure and unknown to the observer). The sensors are modeled as possibilistic with absolute noise of $\pm 20 \mu\text{V}$. For example, in Figure 3.6(A) the actual observation from the voltage sensor is $w_1 = -30 \mu\text{V}$. The sensor model says a noise-free observation could have been between -10 and $-50 \mu\text{V}$. The cyan colored region represents all the possible locations of the target that would give a voltage reading of -10 to $-50 \mu\text{V}$ for an ideal (i.e., noise-free) sensor.

As shown before, possibilistic beliefs have no statistics associated with them; all possible configurations of the target are counted as either “possible” or “not possible”. The subfigures of Figure 3.6 display target possibilities in target configuration space. So, if we superimpose all four maps on top of each other, then the actual target location can only be in the regions where all the beliefs overlap. The data fusion process is simply an intersection of the beliefs from each voltage sensor once the sensor model has been applied and then mapped to target configuration space.

This intersection of the four beliefs can be seen in Figure 3.7(B). In 3.7(A) The beliefs of the target position rendered from the single observation of each sensor are shown as shaded regions that correspond to the color of the sensor (this same information can be seen in Figure 3.6, but is shown together here to demonstrate the intersection data fusion process).

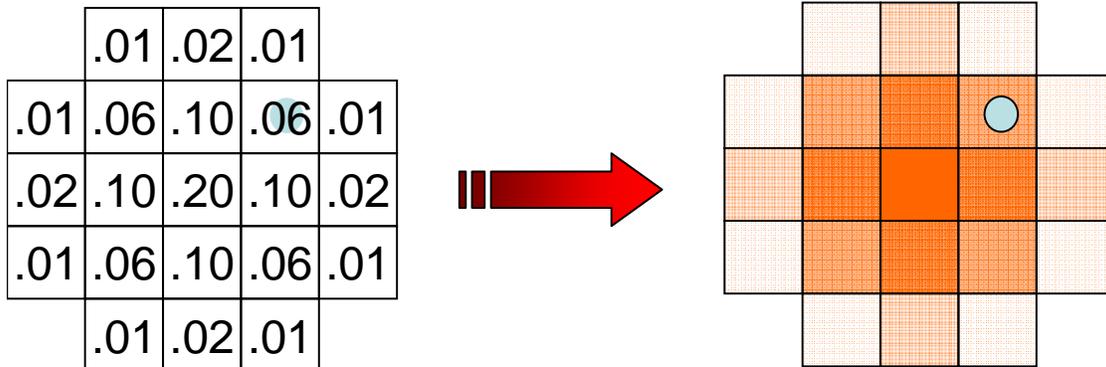


Figure 3.8. A simple example of a histogram belief scheme. The location of the target (light-blue circle) can be in one of the 21 regions (cells), and its true location is unknown. The figure on the left—minus the light-blue circle—represents the belief of the target. The number in each cell represents the probability of the target occupying the region (The probabilities of all 21 cells sum to 1). Alternatively, a color scheme could be used to represent the probabilities in each region; this is shown on the right.

3.2. Belief Maintenance via Histogram Filter

Possibilistic models are intuitive and provide absolute bounds for analyzing a worst-case scenario. But, it is often more advantageous to utilize the statistics of the uncertainties. A *probabilistic* model does exactly this. Probabilistic models assume enough information has been gathered to estimate the statistical properties of the induced uncertainties (these statistics could alternatively just be given).

In this section a type of probabilistic belief maintenance scheme called the *histogram filter* is used in the electrolocation task. The histogram filter is a type of Bayes filter where the state space is decomposed into a finite number of regions (cells). The cumulative probability of each cell is represented by a single probability value. Figure 3.8 depicts an example of utilizing a 2-D histogram for the belief of target.

3.2.1. Shannon Entropy as an Information Metric

When comparing uncertainties associated with different beliefs, it is convenient to establish a common measure. With a possibilistic belief from the previous section, one possible uncertainty metric is the area of the belief (i.e., the area of the 2-D localization subspace). But, this is not a useful method for probabilistic beliefs that take the form of probability distributions. In order to quantify the uncertainty for a given probability distribution, we need a mapping from the probability distribution to a scalar. Shannon Entropy [200] (also referred to as information entropy; for the remainder of this document it shall be referred to as just entropy) is usually accepted as the most generalized metric of uncertainty for probability densities (see [1, 55, 136, 184] for instances of entropy for probability distribution metrics). Entropy is given as

$$(3.3) \quad H = - \sum_{\text{cells}} p(\mathbf{x}) \cdot \log_2 p(\mathbf{x}),$$

where $p(\mathbf{x})$ is the probability of the target occupying that cell. Entropy is often called the “compactness” of a distribution. If a distribution is spread out, then its entropy is relatively high. It refers to how much “surprise” is in a distribution. An entropy of zero means there is no uncertainty. As the entropy increases so does the uncertainty. When \log_2 is used, entropy can be interpreted as the shortest average message length, in bits, that can be sent to communicate the true value of the random variable to a recipient. Figure 3.9 shows the entropy for two different beliefs; both have 21 regions, but they have different probabilities.

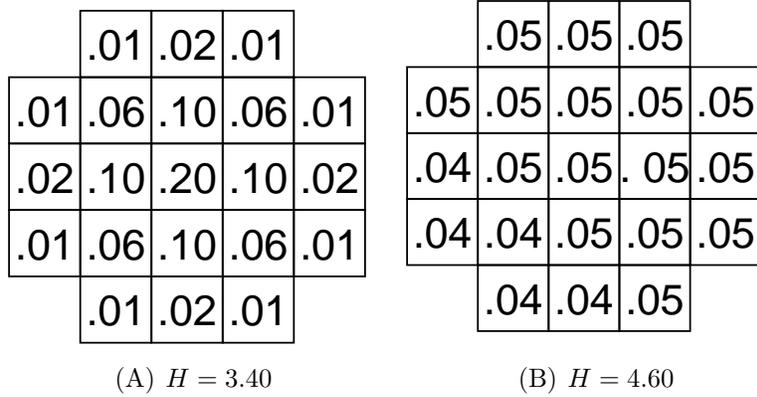


Figure 3.9. The entropy of two beliefs. Both beliefs have 21 possible regions where the target could be. But, each have different probability distributions. The belief in (B) is much more uniform than (A), thus, (B) has more uncertainty associated with it, and this is reflected in the higher entropy.

3.2.2. Sensor Model

The actual sensor observation, w , is the idealized (i.e., noise-free) observation, $\text{EEV}_{\mathbf{e}}(\mathbf{x}_{\text{target}})$, plus the noise, n . Thus, $w = \text{EEV}_{\mathbf{e}}(\mathbf{x}_{\text{target}}) + n$, where $p(n) \sim \mathcal{N}(0, \sigma_n^2)$, and $\text{EEV}_{\mathbf{e}}(\mathbf{x})$ is derived from Equation (2.13). The variance of the sensor noise is σ_n^2 .

The probabilistic sensor model can be written as $p(w|\mathbf{x}, \mathbf{e})$, which represents the likelihood of observing w conditioned on the position of the target in the sensor frame, \mathbf{x} . The information contained in \mathbf{e} defines all relevant unchanging aspects of the environment that the idealized observation depends on, which here includes the size, shape, and material properties of the target, the applied electric field, and the conductivity of the medium.

Since only Gaussian noise is considered, the sensor model can be parameterized as σ_n^2 (i.e., the variance of n) and the expected observation for each possible target location. Recall Equation (2.11) from Section 2.1.4. If only Gaussian sensor noise is considered, then the EEV is precisely the collection of expected observations as a function of the

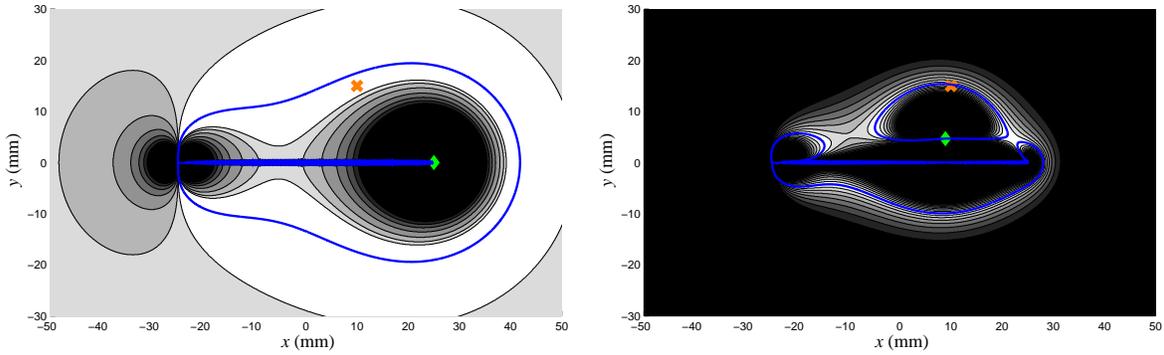
possible target positions. Thus, $E[w] = \text{EEV}_{\mathbf{e}}(\mathbf{x})$, where $E[\cdot]$ denotes the expected value function.

3.2.3. Single Observation Probabilistic Electrolocation

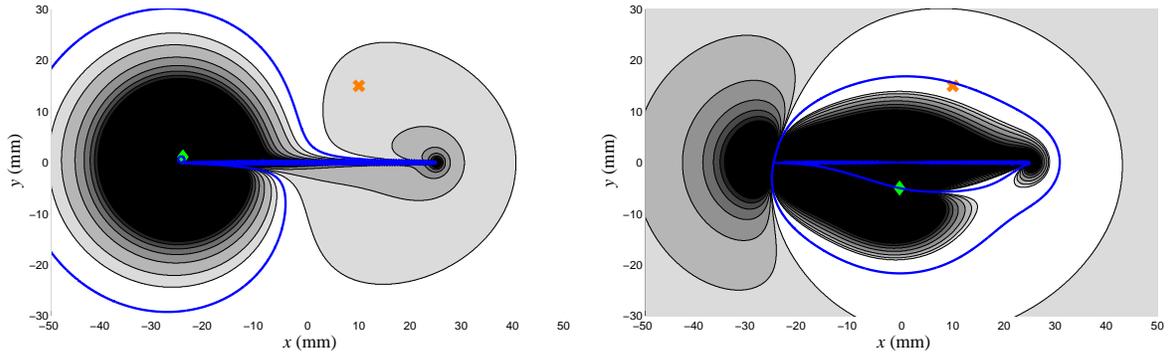
The key difference between possibilistic sensor models and probabilistic sensor models is that probabilistic models also designate the probability of possible states (not just if it is possible). All probabilistic noise models used in this document assume Gaussian noise of known standard deviation.

When utilizing the histogram belief scheme in the electrolocation task, the configuration space of possible target locations is discretized into a finite number of cells. The target occupies exactly one of these cells. A probability of occupancy is maintained for each cell.

Figure 3.10 (and all histogram beliefs used in this document) is constructed as a specific type of histogram called a *position probability grid* [54]. Position probability grids are originally derived from Elfes' *occupancy grids* [91], which are used to map an environment. With occupancy grids, the world is discretized into finite cells. Each cell contains the probability of being occupied by an obstacle. As the robot receives more sensory data, it appropriately updates the grid to reflect this new data. While Elfes' occupancy grid maps environments of various number and size of obstacles, a position probability grid only keeps track of a single object. The intended use of position probability grids is localization, which is determining one's position in the world. Here it is used for locating a single target.



(A) Sensor 1: $w_1 = -30 \mu\text{V}$ ($w_{1,\text{nf}} = -42 \mu\text{V}$); $H = 19.58$ (B) Sensor 2: $w_2 = -201 \mu\text{V}$ ($w_{2,\text{nf}} = -217 \mu\text{V}$); $H = 15.20$



(C) Sensor 3: $w_3 = +8 \mu\text{V}$ ($w_{3,\text{nf}} = -4 \mu\text{V}$); $H = 19.66$ (D) Sensor 4: $w_4 = -39 \mu\text{V}$ ($w_{4,\text{nf}} = -39 \mu\text{V}$); $H = 19.32$

Figure 3.10. Single-observation histogram beliefs for ellipse model. Black signifies cells outside the 7-sigma interval. The belief is represented as a grid of 750×1250 cells (937,500 total cells). The actual location of the center of the 10 mm diameter spherical ideal conductor is $(x = 10, y = 15, z = 0)$, and is signified by the orange “x.” The blue contour is the set of target locations that would render observations identical to the one actually received by the sensor. The center of the green diamond is the location of the sensor. Let, $w_{i,\text{nf}} = \text{EEV}_{\mathbf{e}}(\mathbf{x}_{\text{target}})$ be the noise-free observation (This is the reading the sensor would receive if there were no noise). The caption under each subfigure states the sensor number; actual (i.e., noisy) sensor observation (w_i); noise-free observation ($w_{i,\text{nf}}$); and entropy of the belief (H). Compare with 3.6

Figure 3.10(A) shows the belief of the target position after a single sensor observation from Sensor 1. The other three subfigures in 3.10 show the beliefs for the other three sensors in the ellipse model. Since Gaussian noise is assumed, there is no finite bound on the actual position of the target, but the probabilities do decay quickly outside of the expected value. Figure 3.10 uses a gray-scale filled contour plot to show the probabilities of the belief of the position of the target from a single observation. From highest to lowest probability the gray colormap scheme goes from white to black. The contour lines correspond to confidence intervals of standard deviations. For example, the confidence interval of one standard deviation is 0.683. This means that if a very large number of samples is drawn from a Gaussian distribution, 68.3% of the samples live within one standard deviation of the mean. Likewise, 95.5% of all the samples live within 2 standard deviations. Table 3.1 displays the confidence intervals for the first seven standard deviations. In Figure 3.10 the sum of all the probabilities in the white cells is approximately 0.683; and the sum of all the probabilities in the white and lightest gray together is approximately 0.955.

| std: | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|-------|--------|--------|--------|---------|-----------|-------------|---------------|
| C.I.: | 0.6827 | 0.9545 | 0.9973 | 0.99994 | 0.9999994 | 0.999999998 | 0.99999999997 |

Table 3.1. Confidence intervals of standard deviations.

Construction of Figure 3.10 begins by discretizing the 2-D world into cells (N_{cells} = number of cells). Once the sensor obtains an observation (let's call this observation w_1), it compares this observation to the expected observation at each ($i = 1 \dots N_{\text{cells}}$) cell in a precomputed EEV. With these two voltages and the noise on the sensor, a *likelihood* of target occupancy can be found as follows. Since the normally-distributed noise and the sensor model are not state dependent (i.e., always the same), computing the likelihood

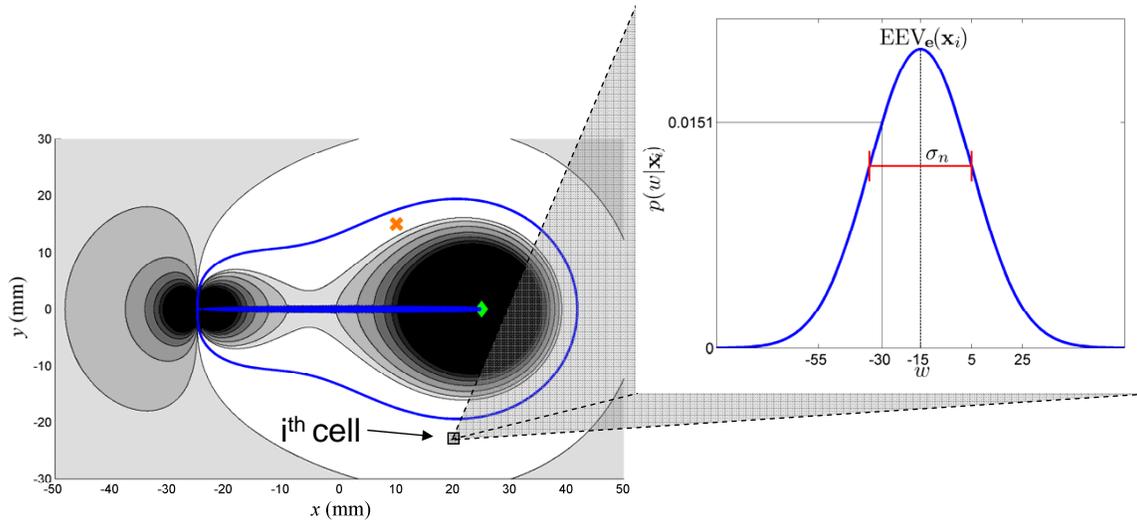


Figure 3.11. Determining the likelihood of w_1 for the i^{th} cell in the belief from Figure 3.10(A). The expected observation when the target is in this cell is $-15 \mu\text{V}$; i.e., $E[w_i] = EEV_e(\mathbf{x}_i) = -15 \mu\text{V}$. The standard deviation of the sensor noise is $20 \mu\text{V}$. Thus, the Gaussian curve associated with this cell has mean of $EEV_e(\mathbf{x}_i) = -15 \mu\text{V}$ (dashed, vertical line) and standard deviation of $\sigma_n = 20 \mu\text{V}$ (red, horizontal error bar). The likelihood of w_1 when the target is in this cell can be read from this Gaussian curve (for example, $p(w_1 = -30|\mathbf{x}_i) = 0.0151$).

can be done by reading off the probability of a Gaussian curve with mean $EEV_e(\mathbf{x}_i)$ and variance σ_n^2 . This Gaussian curve represents the distribution of observations *if the target were actually in the i^{th} cell* (i.e., $p(w|\mathbf{x}_i)$). From this curve we find the probability of observing the actual voltage, w_1 , assuming the target is actually in the i^{th} cell— $p(w_1|\mathbf{x}_i)$ —this is called the likelihood. These probabilities are found for every cell in the grid. Figure 3.11 depicts an example of determining the likelihood of an observation for a cell in the belief. Figure 3.10 shows these probabilities once they have been normalized (i.e., all probabilities sum to one).

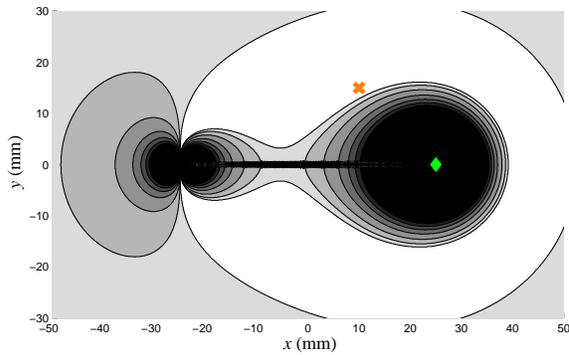
3.2.4. Fusion in Multi-Observation Probabilistic Electrolocation

Belief construction from the possibilistic beliefs was as simple as taking the intersection across all beliefs from each of the observations. A successful technique for integrating multiple beliefs from successive observations from probabilistic models must utilize probabilities of each candidate target location of the belief probability distribution. A *Bayes filter* is the tool of choice for such a task. A new belief probability distribution derived from an observation, w , is computed cell by cell with the prior belief probability distribution by applying Bayes rule to each cell. If \mathbf{x}_i is the candidate target location for the i^{th} cell in the belief, then Bayes rule is written:

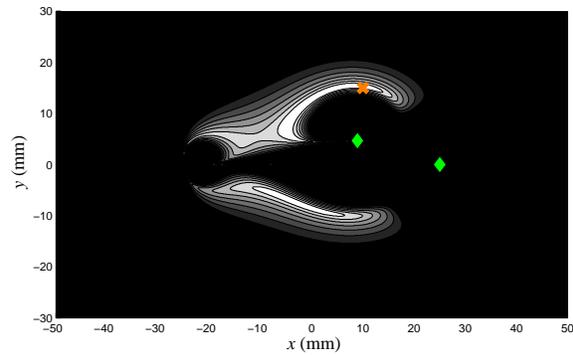
$$(3.4) \quad p(\mathbf{x}_i|w, \mathbf{e}) = \eta \cdot p(w|\mathbf{x}_i, \mathbf{e}) \cdot p(\mathbf{x}_i|\mathbf{e})$$

$$\text{posterior} = \text{normalizer} \cdot \text{likelihood} \cdot \text{prior}$$

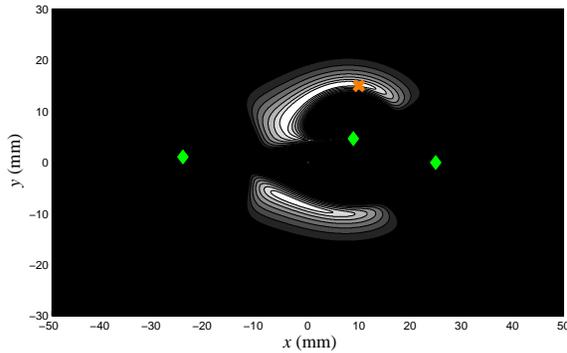
where $p(\mathbf{x}_i|w, \mathbf{e})$ is the *posterior* probability of the target occupying the i^{th} cell, conditioned on the observation w . We call $p(\mathbf{x}_i|\mathbf{e})$ the *prior* (sometimes called the *subjective belief*), and it represents the belief of the target (occupying the i^{th} cell) immediately before the observation. The observation along with the sensor model will yield the *likelihood*. Let η (sometimes called *evidence*) be a scaling factor that ensures all of the probabilities add up to 1.



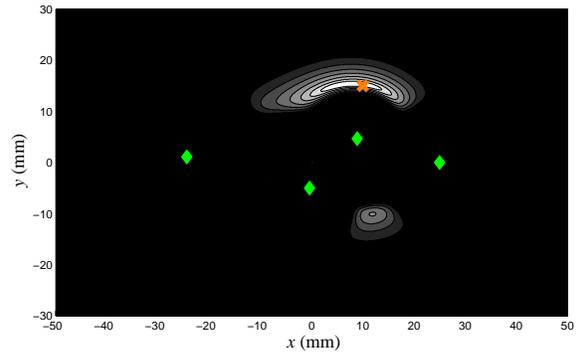
(A) Belief from sensor 1 (same as Figure 3.10(A)).
 $H = 19.58$



(B) Fusion of sensors 1 and 2. $H = 14.30$



(C) Fusion of sensors 1, 2, and 3. $H = 13.72$



(D) Fusion of sensors 1, 2, 3, and 4. $H = 11.83$

Figure 3.12. Fusion of probabilistic beliefs. The orange “x” is the actual location of the target. The gray-scale color map represents the current belief of the target location after the current sensor reading is fused with the prior. The green diamonds in each subfigure show which sensors have been used for the construction of the current belief. The caption under each subfigure displays the entropy to the belief.

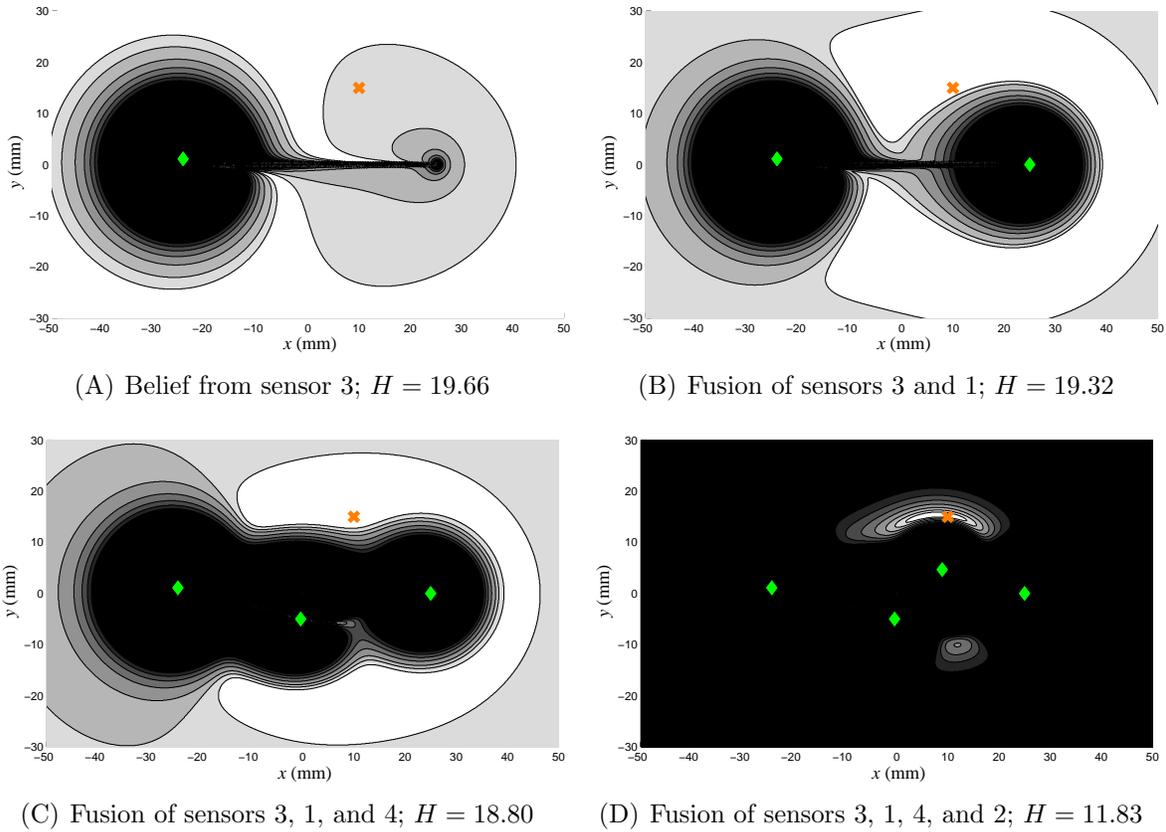
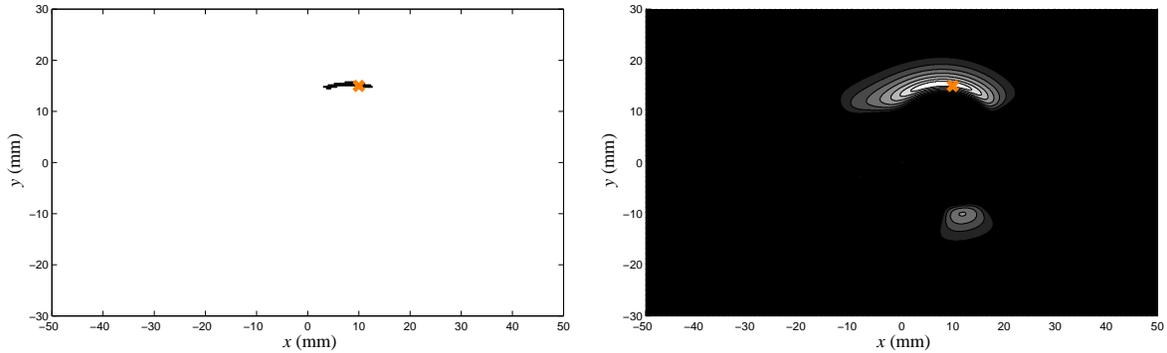


Figure 3.13. Fusion of probabilistic beliefs. Everything is the same as Figure 3.12 except the order of the sensors has changed.

3.2.5. Probabilistic Sensor Fusion with the Ellipse Model

Figure 3.12 demonstrates the successive application of Bayes' rule to each of the beliefs from each of the observations. Figure 3.12(A) is the belief after a single observation (same belief as Figure 3.10(A)). After the second observation is made from sensor 2, the belief from observation 1 is the *prior* and the belief from observation 2 (i.e., Figure 3.10(B)) is the *likelihood*. The resultant belief is the *posterior* and is shown in Figure 3.12(B). Likewise, Figure 3.12(B) is the prior and Figure 3.10(C) is the likelihood for the posterior belief shown in Figure 3.12(C).

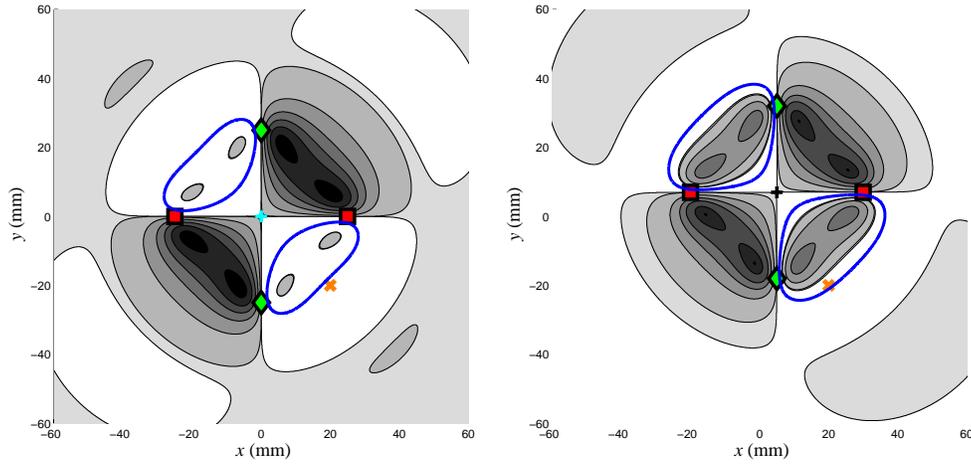


(A) possibilistic belief (same as Figure 3.7(B)) (B) probabilistic belief (same as Figure 3.12(D) and 3.13(D))

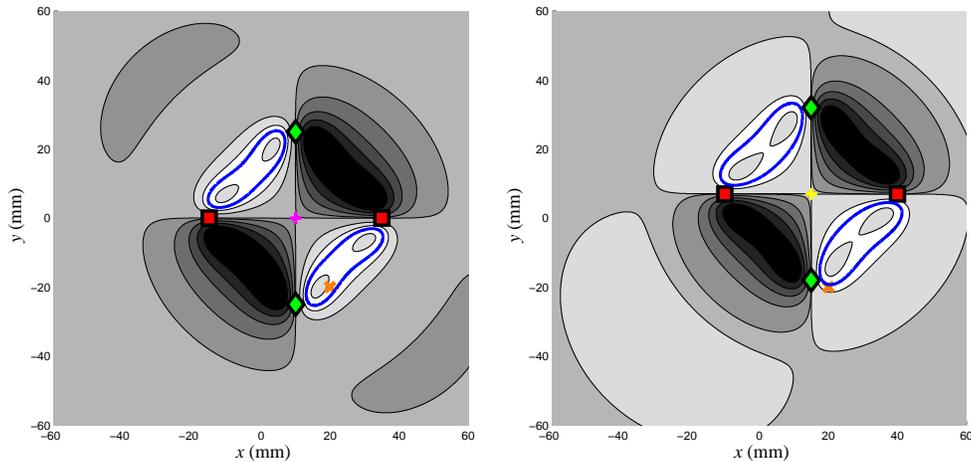
Figure 3.14. Comparing beliefs from possibilistic and probabilistic belief schemes for the ellipse model.

Because of the commutative property of Bayes filters, the order of the sensors in the sensor-fusion process does not matter. Figure 3.13 depicts successive Bayes filter steps for the same conditions as Figure 3.12, but only the order of the sensors has been altered. Note that the final beliefs in each figure (i.e., subfigure (D)) are identical.

Figure 3.14 compares the final beliefs from both the possibilistic and probabilistic models. The noise of both models are characterized by $\pm 20 \mu\text{V}$. With the possibilistic case this number defines absolute bounds of the noise, but with the probabilistic case $\pm 20 \mu\text{V}$ defines the standard deviation of the normal distribution. The possibilistic model states that the noise-free observation must always be within $\pm 20 \mu\text{V}$ from the actual observation, while the probabilistic model states that there is a 68.3% chance that the noise-free observation is within $\pm 20 \mu\text{V}$ from the actual observation. Thus, the possibilistic version is much more restrictive; this is reflected in the belief.



(A) $w_1 = -460 \mu\text{V}$ ($w_{1,\text{nf}} = -350 \mu\text{V}$); (B) $w_2 = -210 \mu\text{V}$ ($w_{2,\text{nf}} = -250 \mu\text{V}$);
 $H = 20.46$; robot location: $(0, 0, 0)$ $H = 21.14$; robot location: $(5, 7, 0)$



(C) $w_3 = -750 \mu\text{V}$ ($w_{3,\text{nf}} = -870 \mu\text{V}$); (D) $w_4 = -650 \mu\text{V}$ ($w_{4,\text{nf}} = -570 \mu\text{V}$);
 $H = 17.53$; robot location: $(10, 0, 0)$ $H = 18.28$; robot location: $(15, 7, 0)$

Figure 3.15. Single-observation histogram beliefs for dipole model. Black signifies cells outside the 7-sigma interval. The actual location of the center of the 10 mm diameter spherical ideal conductor is $(x = 20, y = -20, z = -8)$, and is signified by the orange “x” (recall, only a 2-D slice at $z = -8$ is considered). The blue contour is the set of target locations that would render observations identical to the one actually received by the sensor. The colored “+” is the center of the robot. The center of the green diamond is the location of the sensor. Let, $w_{i,\text{nf}} = \text{EEV}_{\mathbf{e}}(\mathbf{x}_{\text{target}})$ be the noise-free observation (This is the reading the sensor would receive if there were no noise). The caption under each subfigure states the sensor number; actual (i.e., noisy) sensor observation (w_i); noise-free observation ($w_{i,\text{nf}}$); and entropy of the belief (H). Compare with Figure 3.3.

3.2.6. Probabilistic Sensor Fusion with the Dipole Model

Figure 3.15 depicts the single-observation beliefs for each of the four sensors (no sensor fusion). The target is stationary, but the robot moves three times from its initial position. The sensor noise is normally distributed with a standard deviation of $200 \mu\text{V}$. The four robot positions and observations are the same as used with the possibilistic model. Figure 3.16 depicts the sensor fusion process for the four observations from Figure 3.15. Figure 3.17 compares the final beliefs from the possibilistic and probabilistic schemes.

3.3. Belief Maintenance via Particle Filter

3.3.1. An Introduction to the Particle Filter

A particle filter tracks a finite number of candidate target positions based on a Monte Carlo simulation forward simulation [161]. Belief of the target location is represented by a large number of hypotheses (or “particles”) in the target configuration space, allowing approximate representations of complex multi-modal beliefs. The particle filter uses Monte Carlo simulation [161] to approximate a Bayes filter, and the particle representation is updated as new information comes in [232].

One of the many applications of particle filters is localization [95] and mapping in mobile robotics [85, 96, 232]. The electrolocation problem is a variant of the localization problem, where the task is to determine the location of an external target as opposed to the location of the robot relative to some external coordinate frame. The particle filter algorithm consists of recursive implementations of both a predictive and a measurement update step, which are described below.

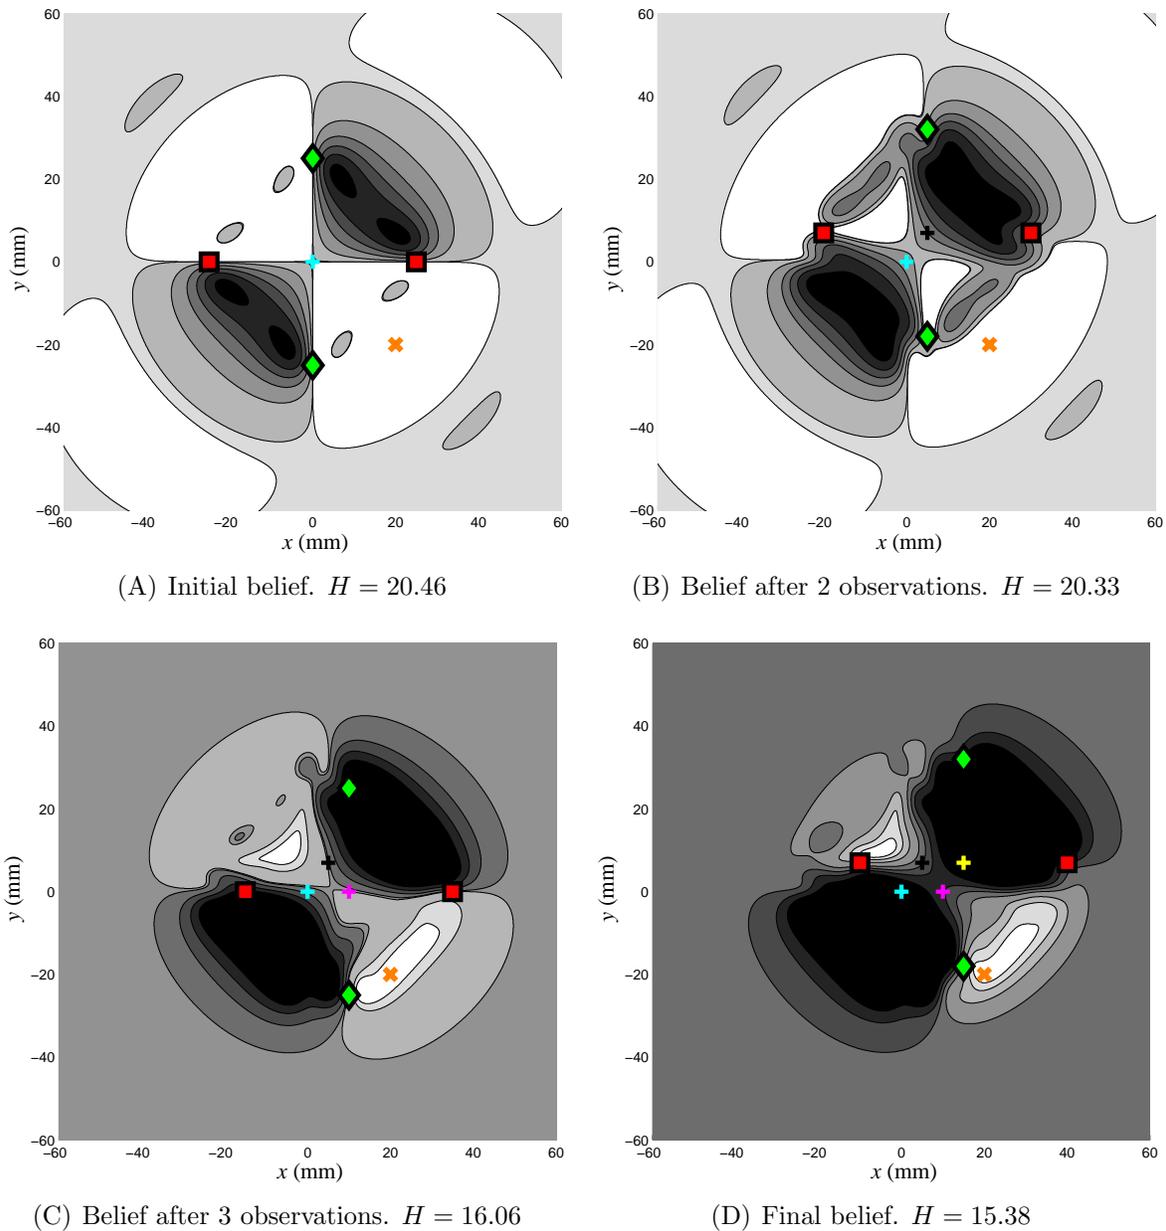
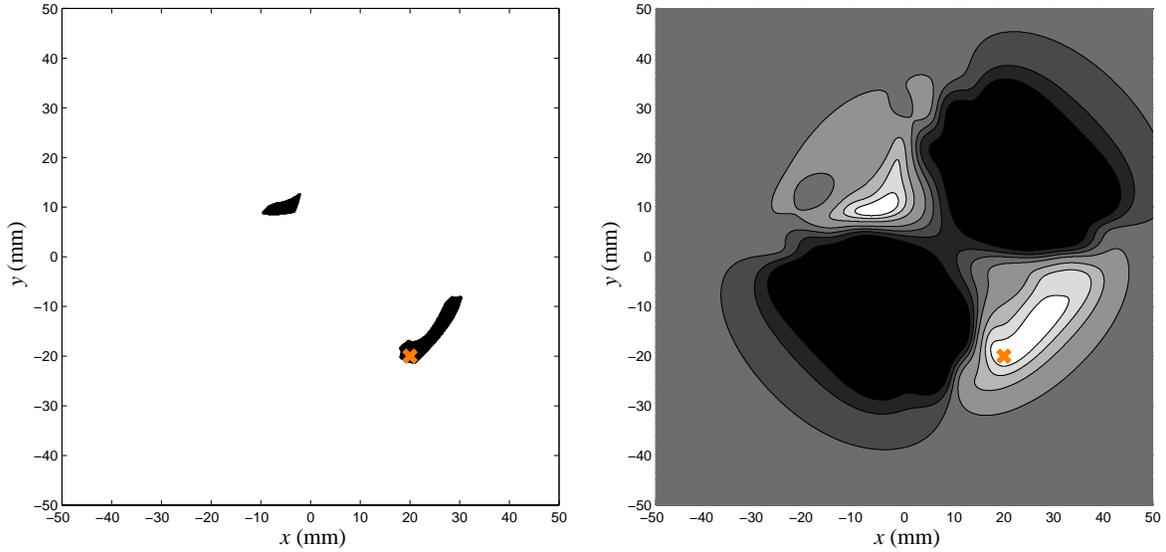


Figure 3.16. Fusion of probabilistic beliefs. The orange “x” is the actual location of the target. The gray-scale color map represents the current belief of the target location after the current sensor reading is fused with the prior. The “+” in each subfigure show which robot positions have been used for the construction of the current belief. The caption under each subfigure displays the entropy to the belief. Compare with Figure 3.5.



(A) possibilistic belief (same as Figure 3.5(D)) (B) probabilistic belief (same as Figure 3.16(D))

Figure 3.17. Comparing beliefs from possibilistic and probabilistic belief schemes for the dipole model.

3.3.1.1. Predictive step. The predictive step takes the motion control and the current belief of the target location relative to the sensor and creates a new belief. This step requires a motion model $p(\mathbf{x}'|\mathbf{x}, \mathbf{u})$, a probability distribution on the new position \mathbf{x}' of the target in the sensor frame as a function of the current position \mathbf{x} and the control \mathbf{u} . Each of the M particles \mathbf{x}_i in the current belief is mapped to a new particle by sampling from the distribution $p(\mathbf{x}'_i|\mathbf{x}_i, \mathbf{u})$.

The two electrolocation systems examined in this document are the Electrosensster (similar to dipole model) and the weakly electric fish (similar to the ellipse model). The the Electrosensster has high-precision encoders and controllers to achieve accurate motion. Thus, the motion model introduced very little uncertainty into the belief about the target position. For this reason, its motion model is deterministic (i.e., $\mathbf{x}' = \mathbf{x} + \mathbf{u}$). Similarly,

with the weakly electric fish model, it is assumed that motion uncertainties are much smaller than sensing uncertainties. Therefore, a probabilistic motion model is not used.

3.3.1.2. Measurement Update Step. Once sensor data are recorded, the measurement update step weights the likelihood of each observation conditioned on the predicted state \mathbf{x}'_i of each particle. The likelihood weight for particle i is written $\lambda_i = p(w|\mathbf{x}'_i, \mathbf{e})$, where \mathbf{e} is the vector of sphere conductivity, sphere diameter, water conductivity, and the value of the applied field. *Importance resampling* then randomly chooses M times from a roulette wheel of the particles $\mathbf{x}'_i, i = 1 \dots M$, where particle i occupies a slice of the wheel proportional to its weight λ_i [230]. The chosen particles will include duplicates at the same location. To introduce diversity into the particle set, a small amount of normally-distributed noise is added to the position of each particle. This new set of particles comprises the new belief.

3.3.1.3. The Particle Set as Another Type of Probabilistic Belief. Each particle in the belief represents a candidate location of the target. Regions of high-density particles represents regions of higher probably of target location. The collection of particles is simply an alternative to the histogram for representing the belief of the target location. A particle belief can be translated into a histogram belief, as shown in Figure 3.18. Appendix A on page 257 gives a simple example of locating a target using a range, bearing, and GPS sensor via particle filtering.

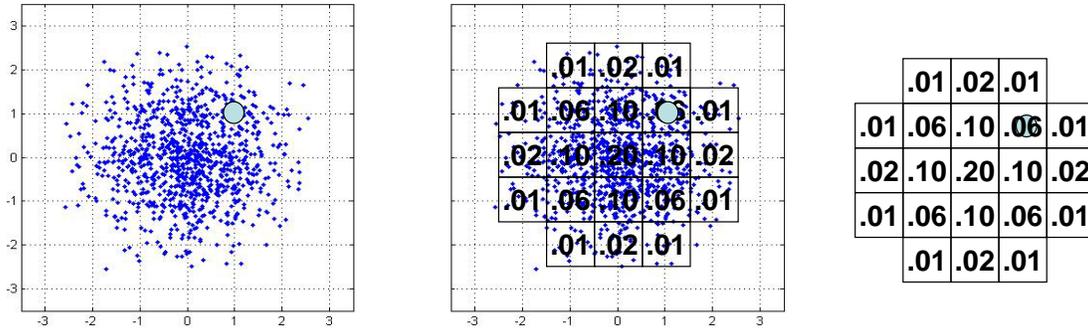


Figure 3.18. Converting a particle belief representation to a histogram belief representation. The left subfigure depicts the particle belief representation of the target. The target-configuration space is then decomposed into 21 cells. The percentage of total particles in each cell is computed. This results in the histogram belief representation.

3.3.2. An Uncertainty Metric for Particle Filters

With the histogram filter the entropy of the belief was used as the uncertainty metric. To compute the entropy of a belief represented by particles, one typically uses a multi-dimensional grid representing the possible (discretized) states. Each cell, i , in this grid stores a probability which is given by the sum of the normalized weights of the samples corresponding to that cell. The entropy is then computed by summing up $p(i) \cdot \log p(i)$ of each cell in that grid. In the case of multi-modal distributions, however, the entropy does not consider the distance between the different modes. For example, consider the two multi-modal beliefs, A and B , which both have identical entropies and the same number of modes. In belief A the modes are located close to each other (e.g. in the same quadrant of the workspace), but the modes in B are dispersed randomly throughout the workspace. Intuitively, belief A is better since the target appears to be localized to a quadrant. But, this is not reflected in the entropy metric.

A metric that captures the spatial distribution of particles will be used instead of entropy. The square root of the trace of the covariance matrix of the particles will be used as the uncertainty metric. Formally, we define P as the spatial covariance matrix of the particle belief, with elements

$$P = \begin{bmatrix} p_{xx} & p_{xy} \\ p_{yx} & p_{yy} \end{bmatrix},$$

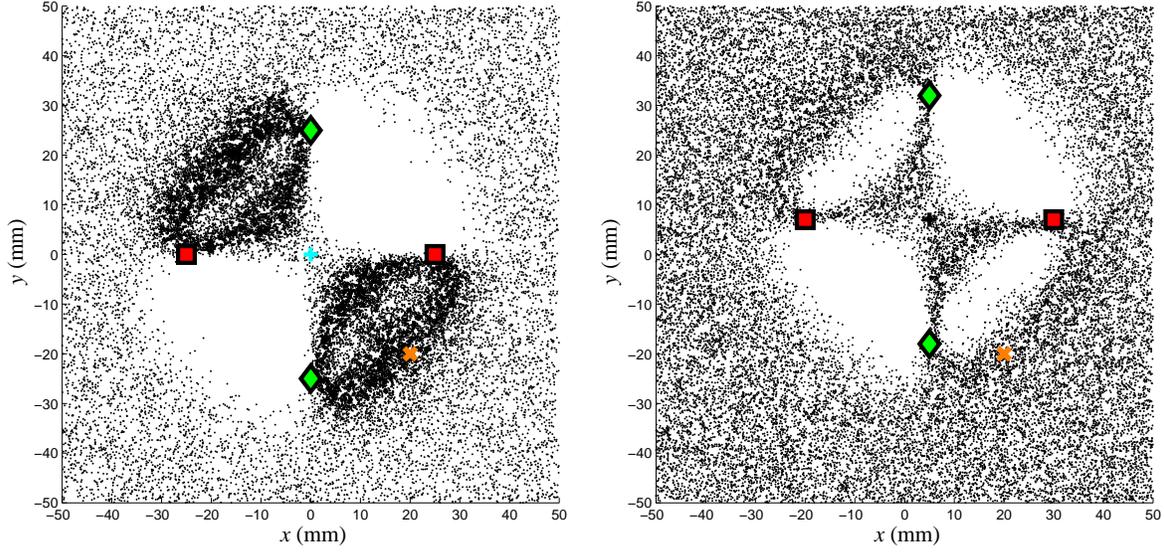
thus, the uncertainty metric is $\text{parts-std} = \sqrt{\text{trace}(P)}$, where $\text{trace}(P) = p_{xx} + p_{yy}$. If we let $b \in \mathcal{B}$ represent the set of particles that represents the belief of the target, then we can define the following function:

$$(3.5) \quad \text{parts-std}(b) = \sqrt{\text{trace}(P_b)}$$

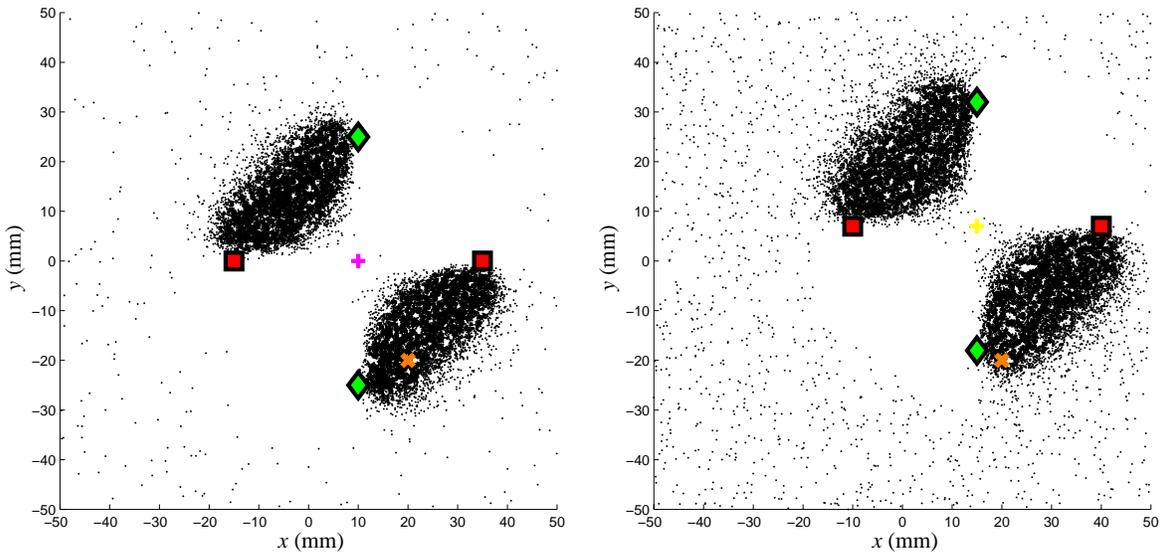
where P_b is the covariance matrix of the particle set b .

3.3.3. Particle Filter with the Dipole Model

Figure 3.19 depicts particle-belief representations for each of the four robot positions after a single observation (no sensor fusion across steps). Figure 3.20 depicts the sensor fusion process—via particle filtering—across the four robot positions. Note parts-std is greater in Figure 3.20(B) than in Figure 3.20(A). Thus, even though more information has been incorporated into the belief the uncertainty metric increased. Due to the non-linear properties of the system and multi-modal beliefs, there is no single metric on the belief space that is able to reduce the high-dimensional belief representation to a one-dimensional measure of uncertainty that completely encapsulates the quality of the history



(A) step 1: $w_1 = -460 \mu\text{V}$; $w_{1,\text{nf}} = -350 \mu\text{V}$; (B) step 2: $w_2 = -210 \mu\text{V}$; $w_{2,\text{nf}} = -250 \mu\text{V}$; parts-std = 33.3; position: (0, 0, 0) parts-std = 43.3; position: (5, 7, 0)



(C) step 3: $w_3 = -750 \mu\text{V}$; $w_{3,\text{nf}} = -870 \mu\text{V}$; parts-std = 21.8; position: (10, 0, 0) parts-std = 23.9; position: (15, 7, 0)

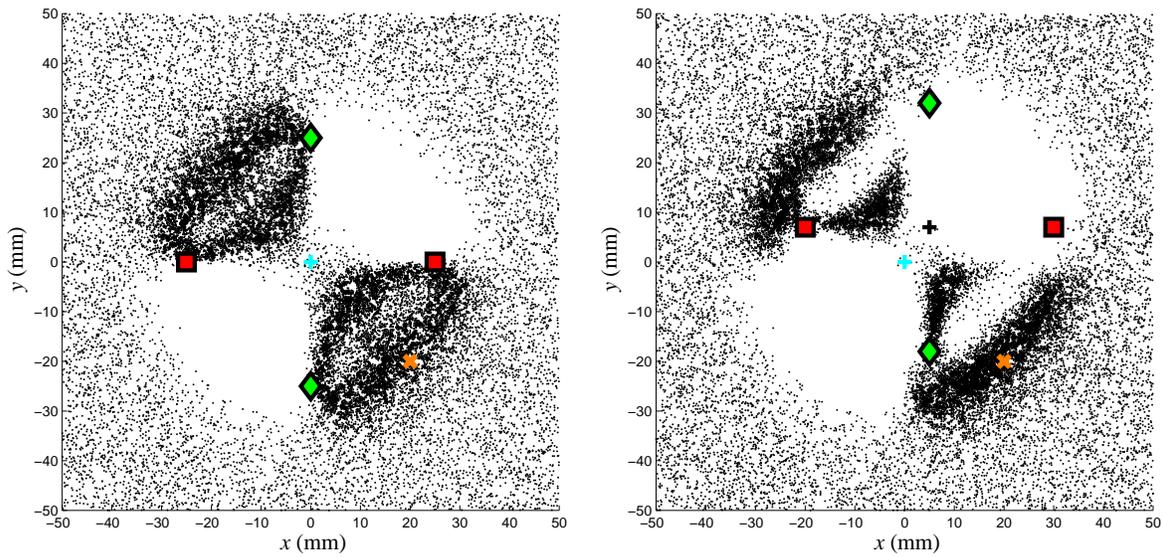
Figure 3.19. Single-observation particle beliefs for dipole model (no sensor fusion across steps). The green diamonds and red squares represent the locations of the sensors and emitters, respectively. The target (orange “x”) is stationary, while the robot (centered at the “+”) moves in each subfigure. The caption in each subfigure states the actual observation, noise-free observation, the uncertainty metric (defined in Equation 3.5), and robot position (in that order). Compare with Figures 3.3 and 3.15.

of observations. Intuitively, the addition of sensor observations should not increase the uncertainty of the belief.

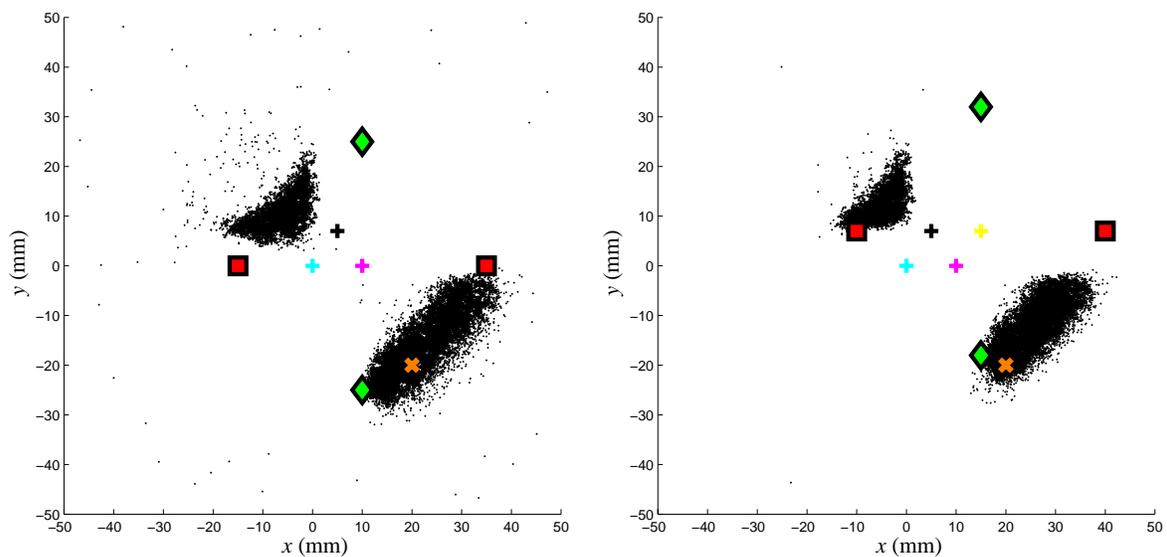
Figure 3.21 shows the final beliefs for the three different types of belief representations (possibilistic, histogram, and particles) for the dipole model. Both the histogram and particle schemes are based on a probabilistic sensor model and use variants of the Bayes filter to update beliefs. Because the belief in Figure 3.21(A) uses a possibilistic sensor model, the target must be contained in the locus of points that is represented by the belief. A probabilistic scheme, on the other hand, prescribes the likelihood of target occupancy. And, because we have assumed Gaussian noise, the likelihood of every possible target location is non-zero. For example, there is less than once chance in a trillion that the target is located somewhere in the black region in Figure 3.21(B) (i.e., the probability mass of the black region is less than 10^{-12}).

3.3.4. Particle Filter with the Ellipse Model

Figure 3.22 depicts the particle-belief representation for each of the four sensors in the ellipse model (no sensor fusion). Figure 3.23 shows the fusion of the four individual beliefs, where the order of the sensors is 3, 1, 4, and finally 2. Note `parts-std= 34.7` in Figure 3.22(C), but 34.8 in Figure 3.23(A). The slight discrepancy is due to the stochastic nature of resampling the particles. Thus, even though both beliefs received identical observations, the particle filter rendered slightly different beliefs. Figure 3.24 compares the final beliefs from the three different sensor fusion techniques used for the ellipse model.

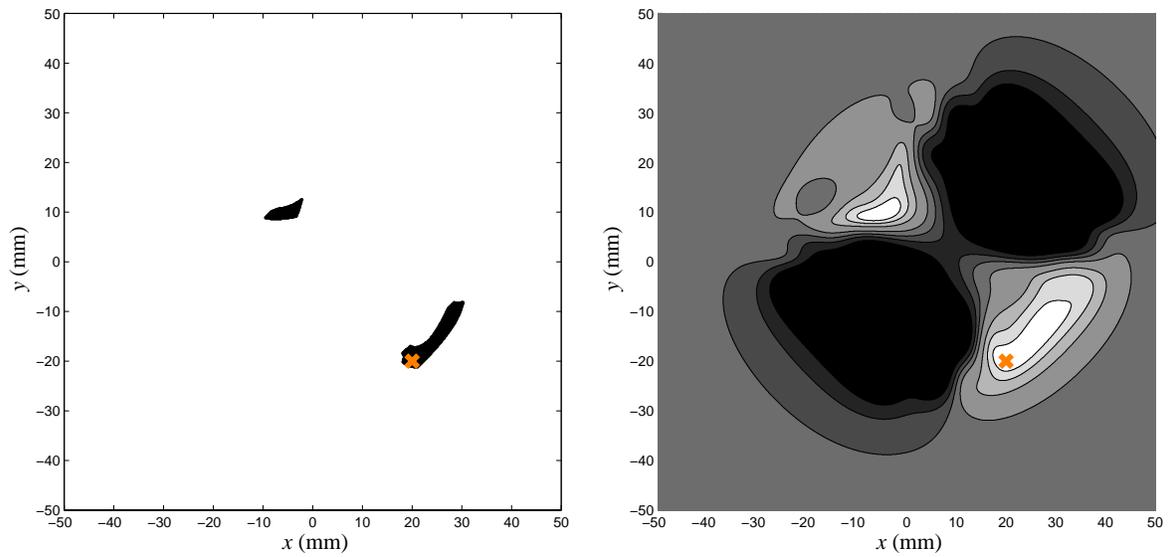


(A) belief after first observation. $\text{parts-std} = 33.3$ (B) belief after fusion from steps 1 and 2. $\text{parts-std} = 35.6$



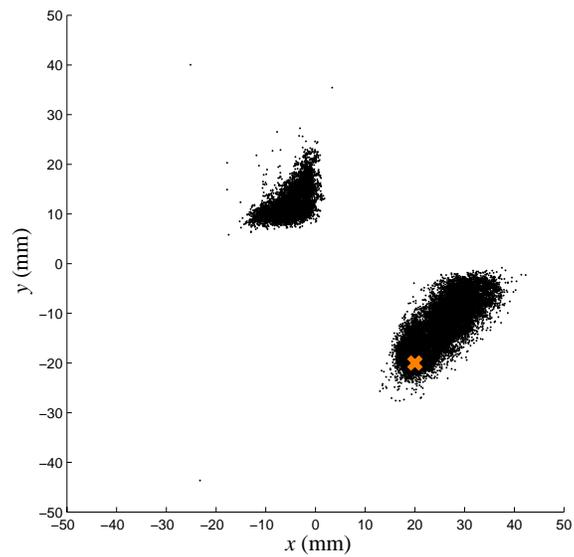
(C) belief after fusion from steps 1, 2, and 3. $\text{parts-std} = 19.6$ (D) final belief after fusion from steps 1, 2, 3, and 4. $\text{parts-std} = 18.7$

Figure 3.20. Fusion of probabilistic beliefs via particle filter. The orange “x” is the actual location of the target. The particles (black dots) represent the current belief of the target location after the current sensor reading is fused with the prior. The “+” in each subfigure show which robot positions have been used for the construction of the current belief. The caption under each subfigure displays the uncertainty metric of the particles (defined in Equation 3.5). Compare with Figures 3.5 and 3.16.



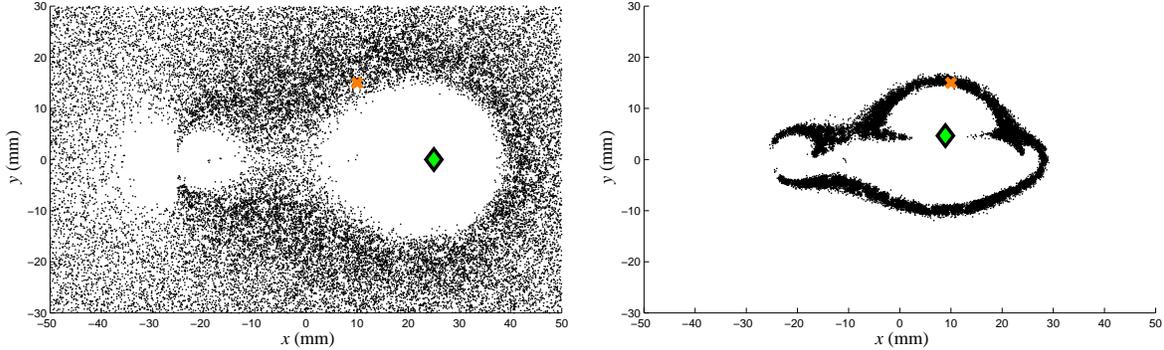
(A) possibilistic belief (same as Figure 3.5(D))

(B) histogram belief (same as Figure 3.16(D))

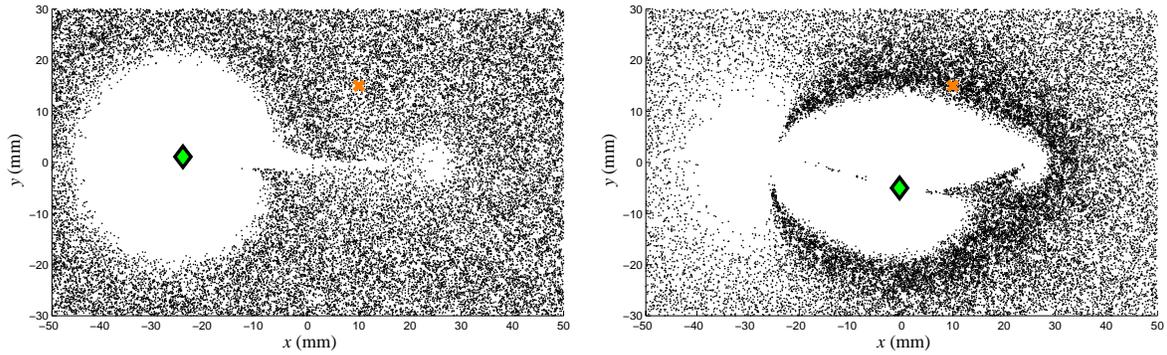


(C) particle belief (same as Figure 3.20(D))

Figure 3.21. Comparing beliefs from possibilistic and probabilistic belief schemes (both histogram and particle) for the dipole model.

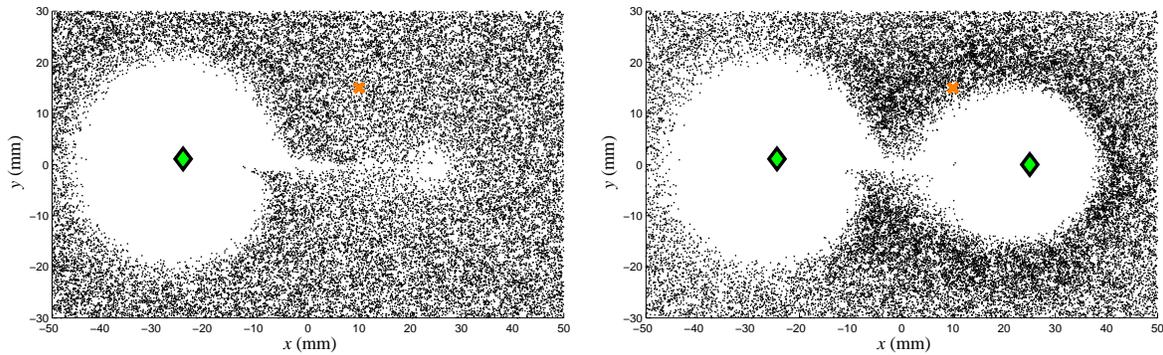


(A) Sensor 1: $w_1 = -30 \mu\text{V}$ ($w_{1,\text{nf}} = -42 \mu\text{V}$); (B) Sensor 2: $w_2 = -201 \mu\text{V}$ ($w_{2,\text{nf}} = -217 \mu\text{V}$);
parts-std = 32.5 **parts-std** = 16.9

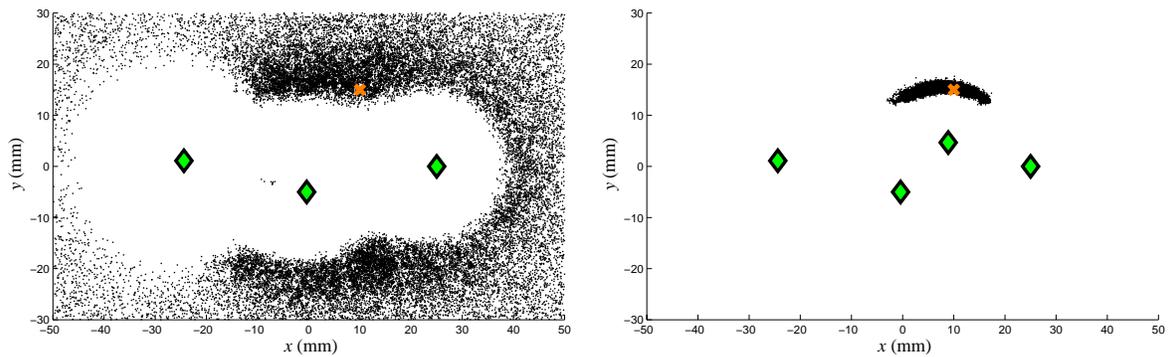


(C) Sensor 3: $w_3 = +8 \mu\text{V}$ ($w_{3,\text{nf}} = -4 \mu\text{V}$); (D) Sensor 4: $w_4 = -39 \mu\text{V}$ ($w_{4,\text{nf}} = -39 \mu\text{V}$);
parts-std = 34.7 **parts-std** = 30.5

Figure 3.22. Single-observation particle beliefs for ellipse model. The particles (black dots) represent the belief of the target. The actual location of the center of the 10 mm diameter spherical ideal conductor is $(x = 10, y = 15, z = 0)$, and is signified by the orange “x.” The center of the green diamond is the location of the sensor. Let, $w_{i,\text{nf}} = \text{EEV}_{\mathbf{e}}(\mathbf{x}_{\text{target}})$ be the noise-free observation (This is the reading the sensor would receive if there were no noise). The caption under each subfigure states the sensor number; actual (i.e., noisy) sensor observation (w_i); noise-free observation ($w_{i,\text{nf}}$); and the uncertainty metric (**parts-std**). Compare with Figures 3.6 and 3.10.

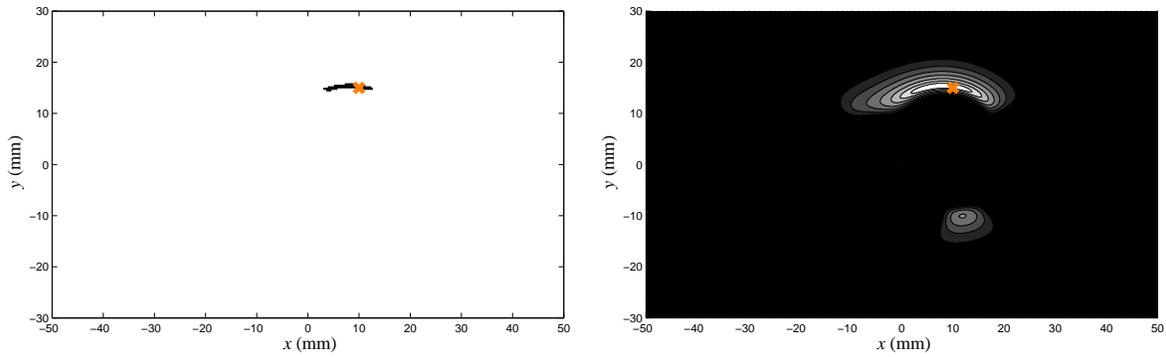


(A) belief after fusion of sensor 3; $\text{parts-std} = 34.8$ (B) belief after fusion of sensors 3 and 1; $\text{parts-std} = 32.6$



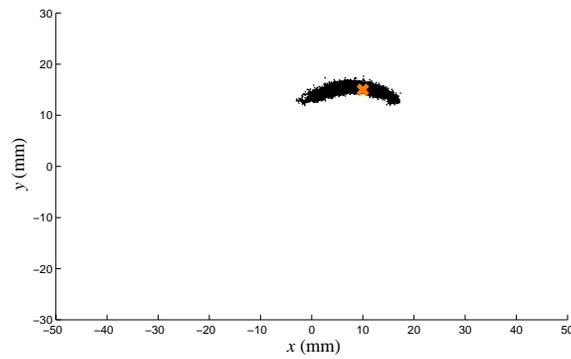
(C) belief after fusion of sensors 3, 1, and 4; $\text{parts-std} = 29.3$ (D) belief after fusion of sensors 3, 1, 4, and 2; $\text{parts-std} = 3.7$

Figure 3.23. Fusion of probabilistic beliefs via particle filtering. The orange “x” is the actual location of the target. The particles (black dots) represent the current belief of the target location after the current sensor reading is fused with the prior. The green diamonds in each subfigure show which sensors have been used for the construction of the current belief. The caption under each subfigure displays the uncertainty metric of the belief. Compare with Figure 3.13



(A) possibilistic belief (same as Figure 3.7(B))

(B) histogram belief (same as Figure 3.13(D))



(C) particle belief (same as Figure 3.23(D))

Figure 3.24. Comparing beliefs from possibilistic and probabilistic belief schemes (both histogram and particle) for the ellipse model.

3.4. Uncertainty maps

For a single-sensor electrosensory system the EEV can give some intuition about the expected quality of sensor information as a function of the relative target location (The dipole model is effectively a single-sensor system since an observation is defined as the difference between the two sensors). For example, the white and black regions of the dipole EEV (see Figure 2.7 on page 53) represent the target locations that generate the greatest expected magnitude of observation. We would expect these regions to be the best places for the target to be for reducing the uncertainty of the belief. For multi-sensor systems—such as the ellipse model—it is more difficult to predict the good target locations since each sensor has its own EEV. It would be convenient to have a tool for estimating the resulting belief uncertainty after each sensor has executed one observation and one belief update. Such a tool would reveal how good or bad a particular emitter/detector layout is.

We have developed such a tool, which is called the *uncertainty map*. An uncertainty map displays the belief uncertainty (e.g., `parts-std` or entropy) as a function of target position. To construct the map, the target is placed on a grid near the emitters and detectors. The belief is initially uncertain (e.g., a set of randomly dispersed particles) and each sensor updates the belief once. The uncertainty of the belief is then assigned to that point on the grid.

In this section we develop uncertainty maps for the histogram belief scheme (i.e., position probability grid). Figure 3.25 shows the uncertainty maps for 4 different levels of noise for the dipole model. Sensor noise corresponding to $\sigma_n = 200 \mu\text{V}$ (Figure 3.25(C)) has been used in the previous examples in this chapter. White represents target locations

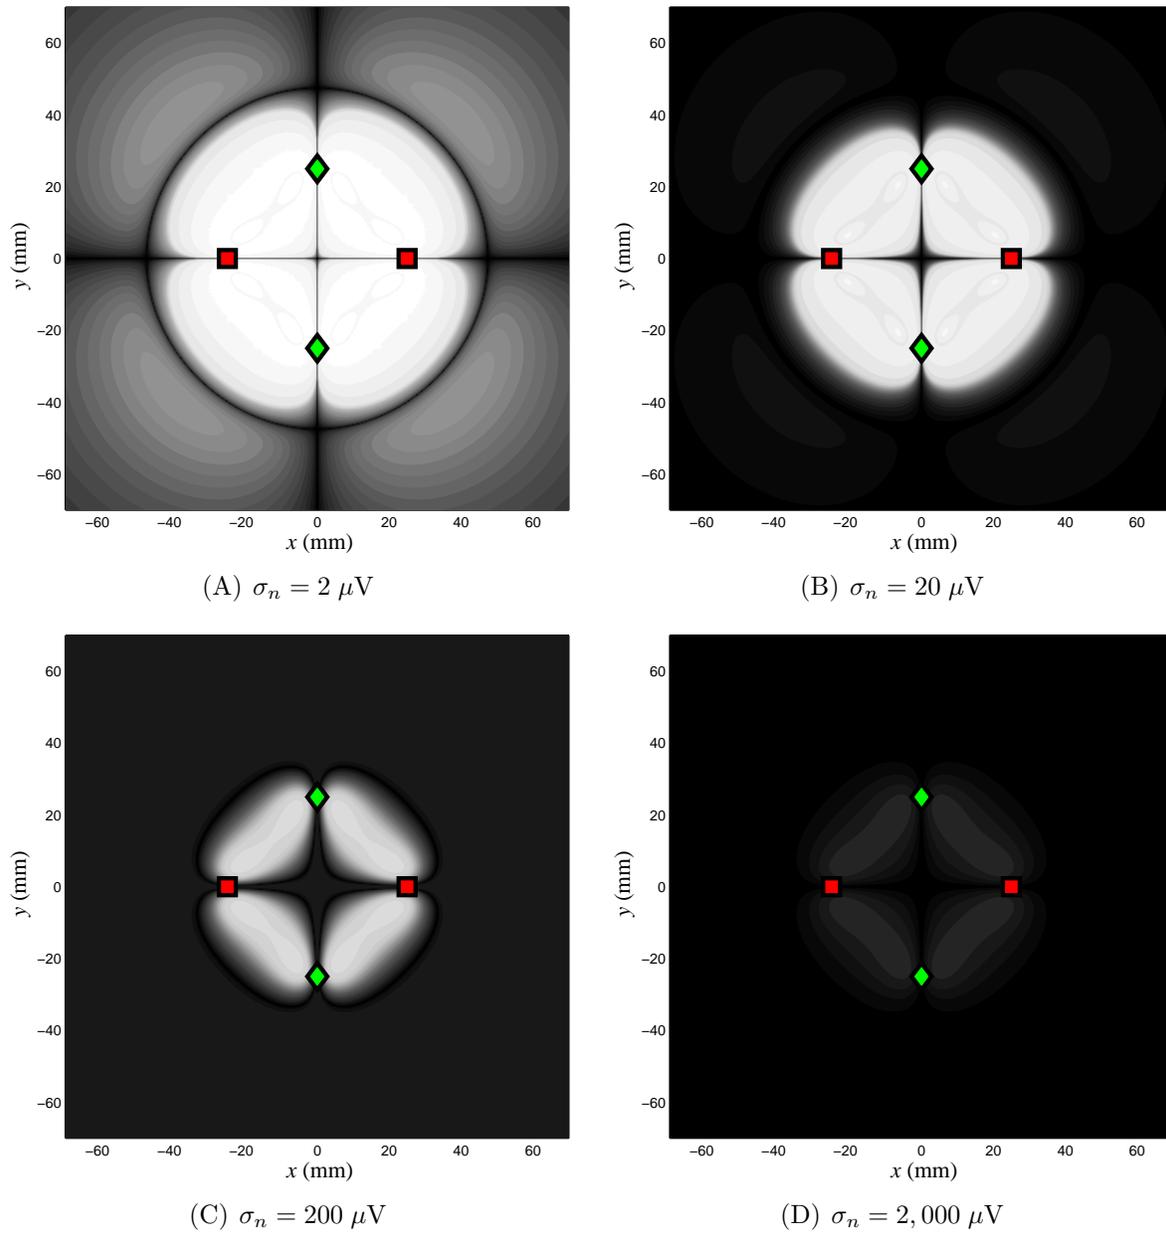


Figure 3.25. Uncertainty maps of the dipole model for 4 different noise levels; all other parameters (e.g. target and water) are identical to those used for the dipole model previously. White corresponds to almost no uncertainty in the belief (entropy is near zero). Black corresponds to a completely uncertain belief (maximum entropy). Intermediate grays are on a logarithmic scale. $200 \mu\text{V}$ was used in the previous examples in this chapter.

that would be effectively localized with a single observations. The resulting entropy of a belief from a target located in the white is near zero. Black regions represent target locations where an observation would effectively do nothing to reduce the uncertainty of a belief. In these regions the resulting entropy of the belief is near the maximum.

One might be tempted to use an uncertainty map as a potential field for a controller. In general it may be good to move towards the white regions in the uncertainty maps, but arriving at the lowest potential doesn't guarantee an optimal sensing trajectory. Uncertainty maps (as described here) only predict belief uncertainty for the specific case of beginning with a flat prior (completely uncertain belief) and then having each sensor update the belief exactly once. Thus, for anything but a flat prior an uncertainty map cannot precisely predict the resulting belief uncertainty.

An uncertainty map is a good tool for evaluating the localization effectiveness of a particular combination of sensor/emitter layout, noise level, target properties, and other environmental parameters (e.g., water conductivity). In Figure 3.25 four different levels of noise are considered; all other properties are identical to those used for the dipole model previously. A noise level of $\sigma_n = 2,000 \mu\text{V}$ (Figure 3.25(D)) appears to be an unacceptable level of noise for localizing this target under these conditions. With $\sigma_n = 200 \mu\text{V}$ (the noise level used in previous examples with the dipole model) the uncertainty map still has a lot of black, but it is possible to localize the target if the target can be positioned in the lighter regions. When the noise level is reduced to $\sigma_n = 2 \mu\text{V}$ the range of the electrolocation system is significantly increased.

In Figure 3.26 all parameters are identical to those used for the ellipse model previously. Two additional cases of 10 and 100 sensors are considered. Figure 3.26(A) uses

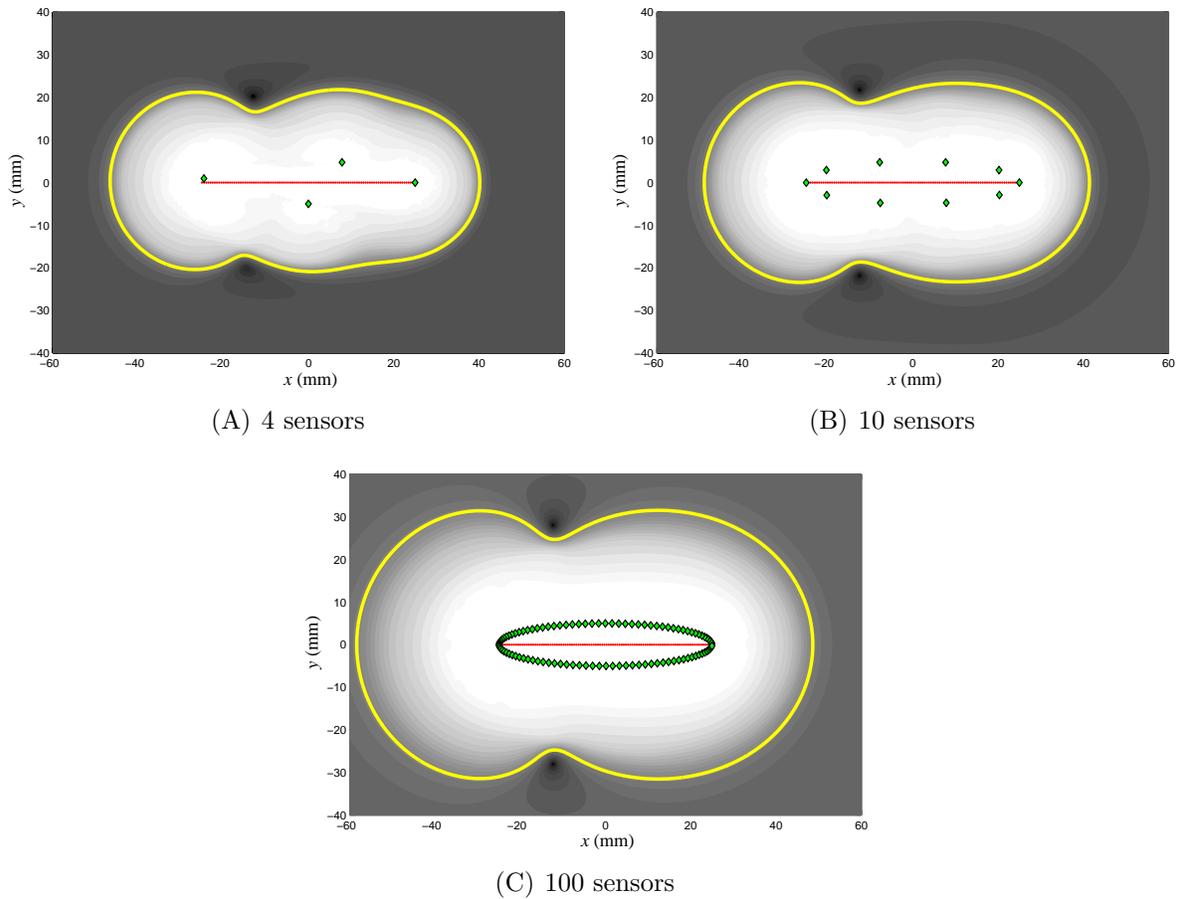


Figure 3.26. Uncertainty maps of the ellipse model for 3 different number of sensors; all other parameters (e.g. target and noise) are identical to those used for the ellipse model previously. White corresponds to almost no uncertainty in the belief (entropy is near zero). Black corresponds to a completely uncertain belief (maximum entropy). Intermediate grays are on a logarithmic scale. Number of sensors: 4, 10, and 100. Four sensors were used in the previous examples in this chapter. Green diamonds are the locations of the sensors, and red dots are the locations of the charges that generate the electric field. The yellow contour is at the same entropy level for all three subfigures.

the same sensor locations as used in previous examples of the ellipse model. To a first approximation, the four-sensor and ten-sensor layouts appear to yield similar uncertainty maps. When the number of sensors is increased to 100, the performance of the electrosensory system increases. Increasing the number of sensors has the same effect as decreasing the sensor noise (This statement is formally investigated later in Section 5.1). Thus, increasing the sensor density from 10 to 100 increases the range of the electrosensory system.

Part 2

**Electrosense Optimization for Prey Capture
in Weakly Electric Fish**

Sensing is a fundamental operation for almost any motion-based system—biological or artificial. We develop a general framework to investigate optimal sensing in the setting of optimal motion control and optimization of the sensory system. We use the prey-capture task of the weakly electric fish as a model system to explore optimal sensing. We develop a method for simulating sensory acquisition and probabilistic belief maintenance for the prey-capture task. We investigate two aspects of sensory system optimization. First, we examine the structure of the fish’s self-generated electric field. We speculate that the electric field is stronger near the tail to complement the fish’s omnidirectional sensing and motor system. Second, we show that the weakly-electric fish has evolved an efficient sensor layout for the prey-capture task. Within the optimal control framework we investigate co-optimization between the sensory and motor systems. We show that if motions are limited to those similar to the fish then mechanically-optimal trajectories are also good for sensing.

CHAPTER 4

Electrosense Optimization for Prey Capture:

Introduction and Formulation

In this part of the thesis the tools and techniques discussed in the Part 1 are applied to the prey-capture task of the weakly electric fish. This chapter provides an introduction and background for the results and discussion presented in the next chapter. First, the stage is set with a brief overview of the problem, solution methods, and results. Then, a formal mathematical framework of the optimal sensing problem is discussed. Finally, the methods for the simulations are presented in the final section of this chapter.

4.1. Introduction

Since the very first motor systems evolved in ancient organisms, a primary purpose of motion has been to control sensors to acquire better information (i.e., active sensing). In almost any system with both sensing and motion control, it is necessary to address their interdependencies. For example, design engineers are often faced with the daunting task of integrating sensory, motion, and control systems into a single robot. In general each of these components request conflicting designs. For example, a large sensory array may induce drag-based friction, which would cause the motion system to be inefficient; or, a particular actuation system may be incapable of meeting the objectives of the control task. Biological systems face the same co-optimization dilemma. We investigate the

fundamental problem of optimal sensing in motor systems, which includes both active sensing and optimization of the sensory system.

We use the prey-capture task of the weakly electric fish as a model system to investigate optimal sensing. The weakly electric fish is an excellent model system since much is known about its sensorimotor system. In general it is virtually impossible to simulate all the components of a biological sensory system, and we make no claim to model all sensory pathways in the weakly electric fish. Instead, we use simple, established models to estimate sensor information. Even with these simple models we show that the sensing and motor systems of the weakly electric fish are well-suited for the prey-capture task.

We begin by formally stating the problem of optimal control under uncertainty, which is to choose a set of future control actions that will minimize an appropriately defined expected future cost [52]. This lays the theoretical foundation for optimal sensing. In general there may be uncertainty in sensing and control, which make exact solutions intractable for almost any practical system. Nonetheless, good sensing strategies can still be investigated under the proper heuristics and assumptions. We then describe our methods for modeling and simulating a probabilistic belief-maintenance scheme for the electrolocation task. To make the simulations computationally tractable, we prescribe some simplifying assumptions such as a simplified electric field model [71], a simple model for sensing electric field perturbations induced by a spherical target [187], and a Bayes filter to maintain the belief of the target. Our model only uses a fraction of the full $\sim 14,000$ sensors found on the surface of a typical weakly electric fish. In the next chapter we show this simplification is justified since there is a tradeoff between the number of

sensors and the level of noise of the sensors. Thus, the *sensing effectiveness* of many high-noise sensors can be equivalent to fewer lower-noise sensors through a simple equivalence relationship.

The first hypothesis we test is that the weakly electric fish has evolved an effective electric field and sensor distribution for the prey-capture task. We first examine the self-generated electric field. We suggest that while the fish’s electric field is stronger near the tail, their anisotropic sensor distribution facilitates a near-isotropic sensing system that complements the omnidirectional motor system. We then investigate the utility of the fish’s distribution of electrosensors for the electrolocation task. The general formulation of the sensor layout optimization problem is conceptually the same as the optimal control problem. But, we abandon notion of “optimal” for merely finding the best among a sample of all possible sensor layouts. We sample 16 sensor layouts from a parameterized space of sensor layouts, and one of these layouts is very similar to the fish’s layout. For each combination of sensor layout and trajectory, the uncertainty of the prey position during the prey-capture task was quantified. Of the 16 sensor layouts and across the 25 trajectories tested, the sensor layout most similar to the fish was the most effective at locating the target (i.e., the lowest average uncertainty of the prey position during the task).

The other hypothesis tested with our simulator is that the sensing and motor system have been co-optimized for the prey-capture task. This would mean that the trajectories executed by the fish are optimal for both sensing and motor efficiency. This hypothesis was tested by comparing the sensing utility of the actual fish trajectory to a mechanically optimal trajectory performing the same prey-capture task. If we consider a fully actuated

ellipsoid, then mechanically-optimal trajectories are much worse for sensing than the trajectories of the fish. But, the mechanically-optimal trajectories are as good as the fish trajectories for sensing when only allowed to perform maneuvers similar to what the fish can do.

4.1.1. Relevant Work

There has been quite a bit of work on the notion of optimal sensing in both biological and engineering systems, much of which has been devoted to active sensing. In the field of artificial intelligence, there are primarily two areas of application: active vision for optical sensors and exploration for robots. In active vision a common objective is to orient a camera to obtain the best view [5, 14, 59, 58, 193, 251]. The details of the cost functional are dependent on the particular application, but the general formulation is identical across all application areas.

Since the mid-1990s active sensing—also referred to as exploration—has become a popular area in the robotics community [1, 226]. Active Markov Localization [94] is a framework for actively localizing a robot under uncertainty, but can be easily generalized to the active sensing problem. Others have used active sensing techniques in coastal navigation [192], exploration for a mobile rover [116, 152], sensor scheduling [136], tracking a dynamic target [77], robot localization [258], tactile sensing [127], and controlling the head of a robot [3].

Optimal sensing also includes optimization of the sensory system itself. Since sensory systems have such a dominant role in food acquisition, predator avoidance, and reproductive success, they have a profound influence on evolutionary fitness, hence are subject to

strong evolutionary pressures. In vision systems higher resolution comes with the penalty of higher metabolic expense from both the eye and the added brain size. In birds it has been shown that the size of the eye is positively correlated with resolution [49]. Birds in particular need to keep body weight as low as possible, thus, evolution is responsible for establishing a suitable compromise between visual resolution and body weight/metabolic expense. Similarly, in the blowfly (*Calliphora vicina*) at rest, visual acquisition and processing represents 13% of its metabolic rate [243]. When the blowfly takes off this is reduced to less than 0.4%, but the cost of carrying the eye is $\sim 3\%$ of its total metabolic rate. Thus, signal acquisition dominates at rest, while eye carriage dominates during flight [141]. Once again evolution must determine a suitable solution that simultaneously considers cost during rest and flight.

More closely related to the system examined here is the work in [50], where they modeled electrosensory landscapes in elasmobranchs (sharks, skates, and rays). They investigated various morphologies to see which is most advantageous for the electrosensory system.

4.1.2. Bayes Filter for Belief Maintenance in the Fish

It is obvious that biological systems deal with uncertainties in sensing and motion, but it is unclear what algorithms the nervous system uses to maintain a belief of the state of the world [242]. Here, we assume the fish employs a Bayesian data fusion scheme to update the belief of the position of the target (i.e., prey). Without a complete neuro-biological model of the weakly electric fish, there is no way to definitively determine if this biological system truly operates under Bayes' rule. But, many researchers have suggested that

biological systems perform Bayes' optimal [128] computations within their sensorimotor systems [133, 132, 180]. We assume that the weakly electric fish does indeed use a Bayes filter to update its belief of the prey. A Bayes filter is the optimal method for data fusion and thus is the limiting case for data fusion in the weakly electric fish.

4.2. Optimal Sensing Formulation

Before we invoke simplifications and heuristics on the optimal sensing problem, it is important to first understand the complete optimal sensing formulation. We begin with developing the general framework for state estimation under uncertainty. We then apply the formulation to active sensing and sensory system design. Finally, we discuss a few possible solution methods.

4.2.1. Motion Optimization Under Uncertainty Formulation

We first formulate the general problem of optimal motion under uncertainty. Let $x \in \mathcal{X}$, $w \in \mathcal{W}$, and $u \in \mathcal{U}$ represent the state of the system, an observation, and control action, respectively. Let $\mathcal{T} = \{k | k \in (0, 1, 2, \dots, F)\}$ be a set of discrete time steps, where F is the final time. Let $\tilde{w}_k = (w_1, w_2, \dots, w_k) \in \tilde{\mathcal{W}}_k$ be the *observation history* up to time step k . Similarly, let $\tilde{u}_k = (u_1, u_2, \dots, u_k) \in \tilde{\mathcal{U}}_k$ be the *action history* [142].

If we let \mathcal{I}_0 be the space of all the known initial conditions at $k = 0$, then $\mathcal{I}_{\text{hist}} = \mathcal{I}_0 \times \tilde{\mathcal{U}}_{k-1} \times \tilde{\mathcal{W}}_k$ represents the complete *history information space* [142] at time k . All the “raw” information known about the system is contained in $\mathcal{I}_{\text{hist}}$. Instead of operating in $\mathcal{I}_{\text{hist}}$ we choose a more convenient *derived information space*, where $b \in \mathcal{B}$ is a probability density function (PDF) over the state space and \mathcal{B} is a derived information space (i.e., \mathcal{B}

is a space of PDFs). We refer to this derived information space as *belief space*. To operate in belief space, there must exist an *information mapping* that maps elements in $\mathcal{I}_{\text{hist}}$ to elements in \mathcal{B} .

The belief is responsible for maintaining a probabilistic representation of the system state and is conditioned on all available information, thus $b_k = p(x_k | \tilde{w}_k, \tilde{u}_{k-1}, b_0)$, where $b_0 \in \mathcal{B}$ is the belief at $k = 0$. Here, we assume the current belief, b_k , is not conditioned on the current control action, u_k , since the current control action will not affect the state until the next time step. Thus, while observations and belief updates occur instantaneously, control actions are manifested over a time step.

Both observations and control actions influence the evolution of the belief through the sensor model and motion model, respectively. The term $p(w|x)$ defines the sensor model, which is conditioned on the state of the system. The term $p(x'|u, x)$ defines the motion model, which is conditioned on the selected control action and previous state, where the prime (') denotes the next time step. By imposing the Markov assumption and utilizing Bayes rule, the belief of the system state after executing control action, u but before the observation (often referred to as the *prediction* step), is [230]

$$(4.1) \quad b'_- = \int p(x'|u, x) b \, dx$$

where the “-” subscript denotes the belief before the observation, and b is the prior belief. After an observation is made, the belief is updated via the sensor model (often referred to as the *measurement update* step):

$$(4.2) \quad b' = \eta p(w'|x') b'_-$$

where η is a normalization constant. Equations (4.1) and (4.2) can be combined into a single “move-then-sense” step [227]:

$$(4.3) \quad b' = \eta p(w'|x') \int p(x'|u, x) p(x) dx$$

Equation (4.3) is the fundamental equation for all Bayes filters including Kalman filters and particle filters.

Figure 4.1 demonstrates the implementation of Equations (4.1) and (4.2) (i.e., Bayes filter) for a simple 1-D prey localization task. Equation (4.2) uses the prior in Figure 4.1(A) along with the likelihood depicted in Figure 4.1(B) (which is derived from the sensor model and the observation) to update the belief, which is shown in Figure 4.1(C). The observation, w , used to construct the likelihood in Figure 4.1(B) suggests it is highly likely the prey is located near $x = 15$ or $x = -15$. Between Figures 4.1(C) and 4.1(D) the fish moves +5 units to the right (thus, moving the prey belief 5 units to the left in this fish-centered coordinate system). Equation (4.1) is used to update 4.1(C) to the belief shown in 4.1(D). Due to imperfect actuation, the belief becomes more uncertain as shown by the spread of each of the two modes in Figure 4.1(D). The process begins again by combining 4.1(D) and 4.1(E) via Equation (4.2) to create the final belief shown in 4.1(F). The observation, w' , used to construct the likelihood in Figure 4.1(E) suggests it is highly likely the prey is located near $x = -20$.

Ultimately, we wish to find the best trajectory that minimizes some predefined cost. Formally, if J is the cost functional, then the optimal control problem is to find the

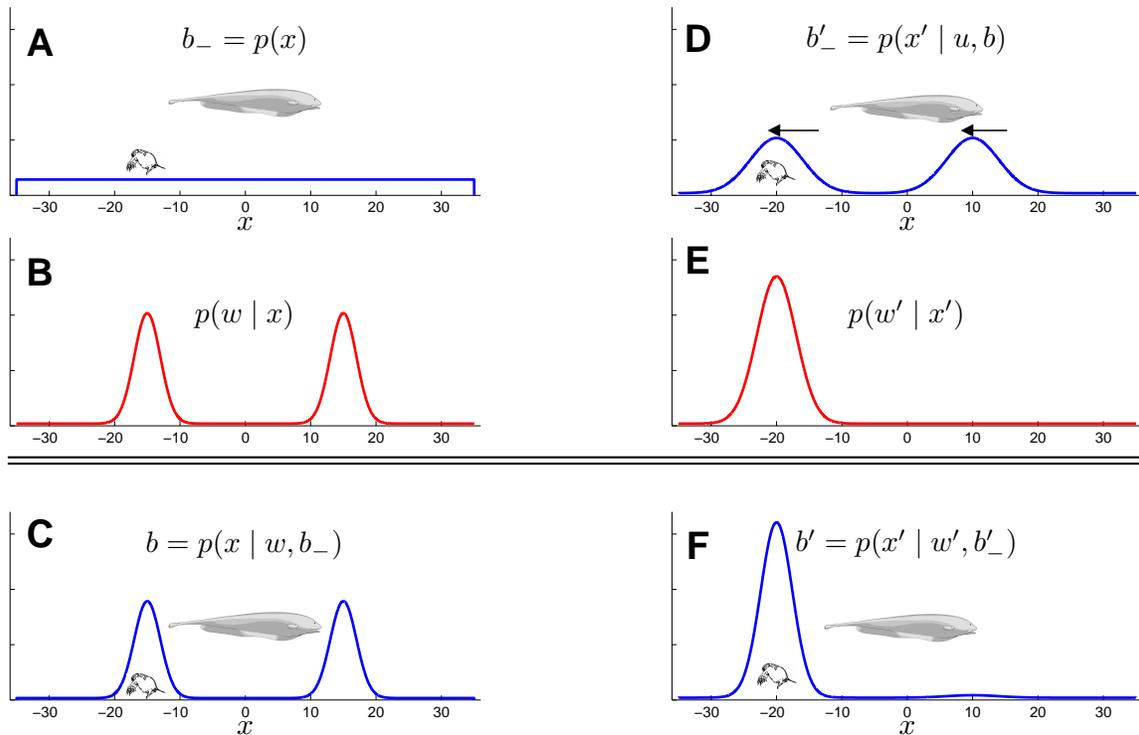


Figure 4.1. A 1-D example of prey localization using a Bayes filter. The x -coordinate is the position of the prey relative to the fish. The actual location of the prey—shown by the picture of the water flea—is initially at $x = -15$ relative to the fish and later moves to $x = -20$ in (D). (A) Initially the fish is completely uncertain about the position of the prey (i.e., *flat prior*). (B) The fish receives an observation that suggests the prey is either near $x = -15$ or $x = 15$. (C) The observation updates the belief (i.e., measurement update step). (D) The fish moves +5 to the right, and the motion model updates the belief (i.e., prediction step). (E) The second observation suggests the prey is near $x = -20$ relative to the fish. (F) The observation updates the belief (another measurement update step). The mode at $x = 10$ in (D) is inconsistent with the likelihood depicted in (E), thus this mode is almost completely eliminated in final belief shown in (F).

trajectory that minimizes J , where

$$(4.4) \quad J[\tilde{u}_F, b_0] = \phi[\tilde{u}_F, b_0] + \sum_{k=1}^F \mathfrak{L}[\tilde{u}_k, b_0, k]$$

The function $\phi[\cdot]$ is the terminal cost, and $\mathfrak{L}[\cdot]$ is the accumulated cost along the way. In general the cost function can be composed of one or both of these components.

Once the cost functional is established, the optimization problem is to find the best trajectory that will minimize the cost function, or

$$(4.5) \quad \tilde{u}_F^* = \underset{\tilde{u}_F}{\operatorname{argmin}} J[\tilde{u}_F, b_0].$$

In general, a new \tilde{u}_F^* must be recomputed every time b_0 is updated via a prediction step (i.e., Equation (4.1)) or a measurement update step (i.e., Equation (4.2)). Alternatively, an off-line method could calculate a control policy for every possible belief, creating a feedback-control law, which maps belief states (i.e., $b \in \mathcal{B}$) into control actions (The most popular method of calculating such a control policy is via *value iteration*, which uses dynamic programming). Such a feedback-control law operates over belief space (in general) since the true state is uncertain. Thus, it may be convenient to determine the transition of the belief to the next time step conditioned on the current belief and the control choice (This is useful, for example, when defining the control policy via value iteration once a cost is defined for each state). The result is a probability distribution over belief space (i.e., a PDF of a PDF) for all possible next beliefs and is a function of the control choice, u . The distribution of next beliefs, $p(b')$ is determined by

$$(4.6) \quad p(b'|u, b) = \int p(b'|w', u, b) p(w'|u, b) dw'$$

where $p(b'|w, u, b)$ is computed through Equation (4.3) and

$$(4.7) \quad p(w'|u, b) = \int p(w'|x') \int p(x'|u, x) p(x) dx dx'$$

where $p(w'|x')$, $p(x'|u, x)$, and $p(x)$ are the sensor model, motion model, and prior belief, respectively. If the motion model is deterministic, then in Equation (4.7) the $p(w'|u, b)$ becomes $p(w'|u, b, x')$, and the integral of $\int p(x'|u, x) p(x) dx$ can be removed.

Equation (4.6) defines a PDF over all possible next beliefs, $b' \in \mathcal{B}$, conditioned on the control inputs and current belief. Since the left side of Equation (4.6) is not conditioned on w' , the integral over all possible observations is needed to determine the expected value. It is interesting to note that Equation (4.6) is infinite dimensional for continuous states (e.g., $x \in \mathbb{R}$) since it defines a PDF for all possible states.

Because future observations are uncertain, the cost functional needs to be defined as an expected cost over all possible observations conditioned on the set of control actions, \tilde{u}_F and the belief at $k = 0$, b_0 . If the cost is a function of the belief at $k > 0$, then recursive implementation of Equation (4.6) will be needed to compute b_k as a function of b_0 , \tilde{u}_k , and \tilde{w}_k . In general, since \tilde{w}_k is unknown, an expected value over $\tilde{\mathcal{W}}_k$ will need to be computed.

4.2.2. Active Sensing

Here, we define the active sensing problem as finding the set of controls that minimize the expected final uncertainty. Let $\text{Uncert} : \mathcal{B} \rightarrow \mathbb{R}$ denote a function for computing an uncertainty norm on a belief. For example, Uncert could be the entropy of a belief.

Let $\text{Update} : \mathcal{B} \times \widetilde{\mathcal{W}}_k \times \widetilde{\mathcal{U}}_{k-1} \rightarrow \mathcal{B}$ denote a function that produces an updated belief (posterior) from the initial belief, the set of observations, and the control history. Update uses recursive implementation of Equation (4.6).

The Lagrangian \mathcal{L} of a cost functional that only accounts for the final uncertainty from $k = 0$ through $k = F$ (i.e., active sensing) is

$$(4.8) \quad \mathcal{L}_{\text{AS}}[\widetilde{u}_{k-1}, b_0, k] = \int_{\widetilde{\mathcal{W}}_k} \text{Uncert}[b_k = \text{Update}(b_0, \widetilde{u}_{k-1}, \widetilde{w}_k)] p(\widetilde{w}_k | \widetilde{u}_{k-1}, b_0) d\widetilde{w}_k$$

This function describes how the expected uncertainty varies with the initial belief and the control inputs up to time $k - 1$. Note the integral over $\widetilde{\mathcal{W}}_k$ is needed to find the expected value over all possible observations histories.

A cost functional can also account for costs of controls as well as the uncertainty of the belief. Let $\text{CCost} : \mathcal{U} \rightarrow \mathbb{R}^1$ denote a function that determines the cost associated with a set of control actions. Thus, $\mathcal{L}_{\text{CCost}}[u_k] = \text{CCost}(u_k)$. An example of a specific class of active sensing cost functionals that also accounts for the cost of control actions is:

$$(4.9) \quad J[\widetilde{u}_F, b_0] = \phi[\widetilde{u}_F, b_0] + \sum_{k=1}^F [\mathcal{L}_{\text{AS}}[\widetilde{u}_{k-1}, b_0, k] + \alpha \mathcal{L}_{\text{CCost}}[u_k]]$$

where α is a weighting factor, which controls the relative importance between uncertainty and control costs.

Equation (4.9) demonstrates a specific class of cost functionals where the uncertainty costs are decoupled from the control costs, and the control cost is only a function of the immediate control action. In general the total cost need not be a weighted sum. Also, in

general the control cost may be a function of the state and previous controls. Since the state is unknown, an expected value over the possible states is used.

4.2.3. Sensory System Design Optimization

Optimal sensing is not restricted to optimal control and can include optimization of parameters of the sensing system. In biological systems the optimization process could take place over evolutionary timescales through natural selection. For example, the self-generated electric field can be optimized to render the best sensory signal for a given allowance of energy to produce it. With parameter optimization, the equations are the same as presented above, except the set of parameters, p , shows up in the arguments of the cost functional and in the statement of the optimization problem. Thus,

$$(4.10) \quad J[\tilde{u}_F, b_0, p] = \phi[\tilde{u}_F, b_0, p] + \sum_{k=1}^F \mathfrak{L}[\tilde{u}_k, b_0, k, p]$$

The optimization problem is finding the best parameters that will optimize the cost function, or,

$$(4.11) \quad p^* = \underset{p}{\operatorname{argmin}} J[\tilde{u}_F, b_0, p]$$

Equation (4.11) finds the optimal parameters only for the trajectory \tilde{u}_F and the initial belief b_0 , which is, most likely, a single instance in the space of all possible combinations of trajectories and initial beliefs. In general, the parameters would need to be optimized across all possible conditions.

4.2.4. Solution Methods

Fundamental optimal control theory says the optimal set of controls from Equation (4.5), \tilde{u}_F^* can be derived using Pontryagin’s maximum principle (a necessary condition), or by solving the Hamilton-Jacobi-Bellman equation (a sufficient condition). In practice exact solutions via literal application of these techniques is computationally infeasible. Thus, approximations and heuristics must be used.

Perhaps the least restrictive formulation is the partially-observable Markov decision process (POMDP). In fact we only need complete discretization (i.e., discrete states, observations, actions, and time) to use this formulation [122]. Dynamic programming, which is an application of the Bellman equation, is the most popular technique for solving a POMDP. In addition to a sensing and motion model, a reward function is needed to map a state and action to a scalar reward. In general this is a computationally expensive approach, but in some cases heuristics are acceptable [64].

If it can be assumed that the dynamics are linear and the disturbances on the sensor and process models are normally distributed, and the cost functional is quadratic, then the Linear Quadratic Gaussian (LQG) control can be used. LQG dramatically decreases the computational cost by decoupling the control and estimation problems (the *separation principle*). The Kalman filter is an optimal state estimator for the LQG case that parameterizes the belief with only a vector of means and the covariance matrix. The optimal control policy can be determined in closed-form via the Ricatti equations.

The solution methods used in this paper take a more pragmatic route than the ones described above. We evaluate a small sample from the space of possible solutions. Even

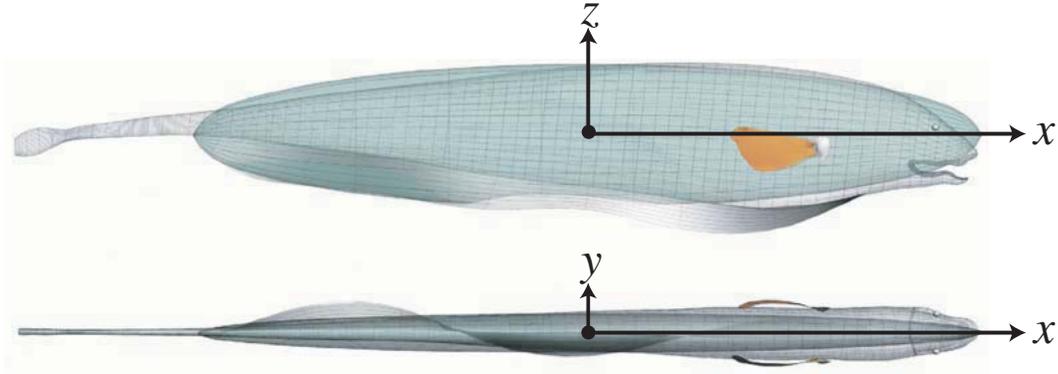


Figure 4.2. The black ghost knifefish fit to an ellipsoid. This figure is taken from [155].

though we cannot claim to find an optimal solution, we can still gain some insight into optimal sensing.

4.3. Methods

We use the prey-capture task of the weakly electric fish to investigate several aspects of optimal sensing. In this section we explain the methods for simulating the belief-maintenance process during this task. The morphology of the fish's body is modeled as a 6-DOF rigid ellipsoid of similar length (see Figure 4.2). In addition, other simplifications of the electric field, sensor layout parametrization, and electrosensory observations are needed for computational tractability.

4.3.1. Electric Field Model

The model of the electric field is described in [71], which built off the model developed in [12]. It is almost identical to the model used for the ellipse model (Section 2.4, page 54), where the unperturbed electric field is modeled as a line of point charges in a homogeneous

medium. The model of the electric field with n total poles, one negative pole, and $m = n - 1$ positive poles is [211]:

$$(4.12) \quad \mathbf{E}_{\text{fish}}(\mathbf{x}) = \left[-\frac{q}{|\mathbf{x} - \mathbf{x}_p^n|^3}(\mathbf{x} - \mathbf{x}_p^n) + \sum_{i=1}^m \frac{q/m}{|\mathbf{x} - \mathbf{x}_p^i|^3}(\mathbf{x} - \mathbf{x}_p^i) \right] \frac{\sigma_{\text{mes}}}{\sigma_{\text{mod}}}$$

The vector \mathbf{x}_p^i is the position of the i th pole. The quantity q is analogous to electric charge in an electrostatic model and is distributed such that the first m poles have a “charge” of q/m and the remaining pole has a charge of $-q$, resulting in a total net charge of zero. The poles are uniformly distributed along the midline from head to tail. We use $q = 10 \text{ mV}\cdot\text{cm}$ (equivalent to $q = 10^{-4} \text{ N}\cdot\text{m}^2/\text{C}$), and the pole locations run from the nose to the tail of the fish along the central axis of the fish body with equidistant spacing. These values result in field values within 10% of measurements of the electric field vector $\mathbf{E}_{\text{fish}}(\mathbf{x})$ of *A. albifrons* obtained by other researchers (B. Rasnow, C. Assad, P. Stoddard, unpublished data) at $\sigma_{\text{mes}} = 210 \text{ }\mu\text{S}$ using a multi-axis electrode array [8, 7, 187]. The $\sigma_{\text{mes}}/\sigma_{\text{mod}}$ term scales the field strength to the water conductivity used in simulation. This scaling is based on empirical measurements of the knifefish field at different water conductivities [129] which suggest the electric organ can be idealized as constant current source. We selected $35 \text{ }\mu\text{S}$ because an earlier study [154] found detection range was highest for trials at this conductivity, and this conductivity is most similar to rivers of the fish’s native habitat.

4.3.2. Idealized Electrosensory Observation

Electrosensors (i.e., voltage sensors) on the surface of the ellipsoid measure changes in the self-generated electric field. These voltage observations—along with a sensor model—are

used to update the belief of the target. At each time step, each sensor performs exactly one observation. The observations are simulated with Rasnow’s model for spherical objects in a uniform electric field [187], which is Equation (2.13) on page 48. In our simulations we used a target radius of $r_{\text{target}} = 0.083$ cm, which is a typical size of a water flea. The conductivities of the target and water are $\sigma_{\text{target}} = 300 \mu\text{S} / \text{cm}$; and $\sigma_{\text{water}} = 35 \mu\text{S} / \text{cm}$, which results in a contrast factor of approximately $\chi = +0.72$.

4.3.3. Probabilistic Sensor Model

Of course, any real sensor has non-zero noise. Therefore, we employ a probabilistic sensor model to update the belief. The noise on the sensors is assumed to be normally distributed. The actual sensor observation, w , is the idealized observation, δV , plus the noise, n . Thus, $w = \delta V(\mathbf{x}_t) + n$, where $p(n) \sim \mathcal{N}(0, \sigma_n^2)$, and δV is computed from Equation (2.13). We define σ_n^2 as the variance of the sensor noise.

Our probabilistic sensor model can be written as $p(w|\mathbf{x}, \mathbf{e})$, which represents the likelihood of observing w conditioned on the position of the target in the sensor frame, \mathbf{x} . The vector \mathbf{e} defines all relevant unchanging aspects of the environment that the idealized observation depends on, which here includes the size, shape, and material properties of the target, the applied electric field, and the conductivity of the medium.

We are assuming all sensors have an equal level of noise for any level of idealized observation. Because we are assuming normally distributed noise, our sensor model can be parameterized as σ_n^2 (the variance of n) and the expected observation for each possible target location. Recall, the ensemble of electrosensory viewpoints (EEV) is a mapping from the target position to the expected observation for each possible target location

[213], and can be written as:

$$(4.13) \quad \text{EEV}_{\mathbf{e}}(\mathbf{x}) = \text{E}[w]$$

$\text{E}[w]$ is the expected (i.e., noise-free) observation for a target at location \mathbf{x} . The EEV is constructed directly from Equation (2.13).

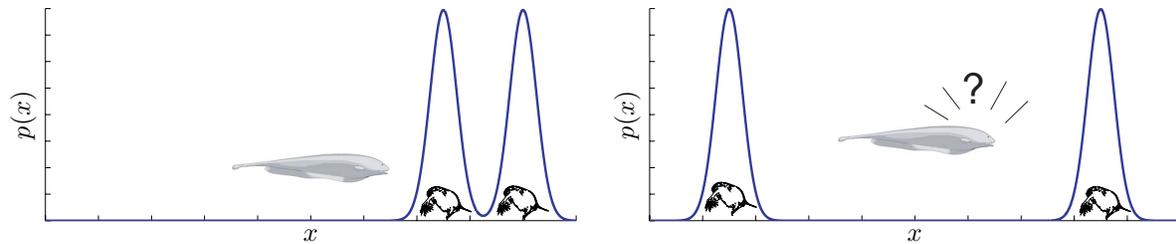
4.3.4. Uncertainty Norm

For the prey-capture simulations, the belief of the prey location is maintained via a particle filter (See Section 3.3 on page 83), allowing approximate representations of complex multi-modal beliefs. A norm on the belief space (i.e., the space of all possible sets of particles) is needed to quantify the uncertainty associated with the belief.

Shannon Entropy [200] (see Section 3.2.1 on page 72) is usually accepted as the most generalized norm of uncertainty for probability densities [55, 136]. Entropy can be computed for a set of particles by discretizing the state space into cells. However, in the case of multi-modal distributions, the entropy does not consider the distance between the different modes.

An uncertainty norm that captures the spatial distribution of particles is the 3-D version of the `parts-std` norm originally discussed in Section 3.3.2. Thus, the uncertainty norm is `parts-std` = $\sqrt{\text{trace}(P)}$, where $\text{trace}(P) = p_{xx} + p_{yy} + p_{zz}$ (Recall, P is the spatial covariance matrix of the particle belief).

For the prey-capture task it is important to penalize larger distances between modes within a belief to avoid having the controller to choose which mode to act on. For example, let us consider the 1-D prey capture task, where at each time step the fish must decide



(A) Clustered-modes belief has a low standard deviation (B) Unclustered-modes belief has a high standard deviation

Figure 4.3. The motivation for using the `parts-std` uncertainty norm rather than entropy. In this 1-D example of prey-capture, the fish must decide whether to move to the left or right. The plots show the belief of the x -coordinate (the 1-D position of the prey). Both beliefs are bi-modal, thus the fish roughly believes the prey could be in one of two locations (signified by the pictures of the water fleas). Both PDFs have the same utility by the Shannon entropy norm, but the standard deviation of the PDF in (B) is approximately 4.5 times greater than the PDF in (A).

to move left or right until it captures the prey. Figure 4.3 depicts two scenarios with bi-modal beliefs of the prey location, thus the fish roughly believes the prey could be in one of two locations (signified by the pictures of the water fleas). It is unambiguous for the fish choose its control option in the case of the PDF with clustered modes (Figure 4.3(A)). In the case of the spread apart modes (Figure 4.3(B)), the fish is just as likely to move further away from the prey than move towards it. For the prey-capture task, the clustered-modes belief (Figure 4.3(A)) is clearly better than the unclustered-modes belief (Figure 4.3(A)). The entropy for both beliefs are identical, but the standard deviation is approximately 4.5 times greater in Figure 4.3(B) than in 4.3(A) (In 1-D the `parts-std` norm is equivalent to the standard deviation of the belief). Similar situations of multi-modal beliefs were observed in our prey-capture simulations in higher-dimension space. While both entropy and `parts-std` approximate the uncertainty of a belief, only `parts-std` accounts for the

spatial distribution of the belief. Thus, a controller with `parts-std` in its cost function—as opposed to entropy—favors clustered-modes beliefs over nonclustered-modes beliefs, while the entropy norm is indifferent to the distance between modes.

4.3.5. Sensor Layout Parametrization

We use a compact parametrization to approximate the sensor densities along the antero-posterior axis (the AP dimension) and along dorsal-ventral-right-left (DVRL) dimension. Each sensor layout corresponds to a sensor PDF for each of the two dimensions. The density of sensors along each dimension is characterized by a region of low density and a region of high density. The density function is parameterized by three numbers: 1) center of high-density region, c_{hd} ; 2) width of high-density region, w_{hd} ; and 3) percentage of sensors in high-density region, p_{hd} .

The range of the AP is from -1 (for the extreme anterior point) to $+1$ (for the extreme posterior point). The range for the DVRL dimension is $(-\pi, +\pi]$, where 0 is the most dorsal part, $-\pi/2$ is the fish’s right, $+\pi/2$ is the fish’s left, and $\pm\pi$ is the most ventral part of the fish. An example of a sensor density distribution using this parametrization is shown in Figure 4.4.

The sensor density that most closely approximates the distribution of sensors in *A. albifrons* as measured from [61]—called the “fish-like” layout—is $\{c_{\text{hd}}, w_{\text{hd}}, p_{\text{hd}}\} = \{-0.84, 0.32, 0.62\}$ for the AP dimension and $\{c_{\text{hd}}, w_{\text{hd}}, p_{\text{hd}}\} = \{0, \pi/2, 0.35\}$ for the DVRL dimension. 11 other layouts were examined that had identical w_{hd} and p_{hd} as the fish-like layout for both the DVRL and AP dimension, but differed in the c_{hd} parameter for each dimension. Three different c_{hd} s for the AP dimension ($-0.84, 0$, and $+0.84$) and

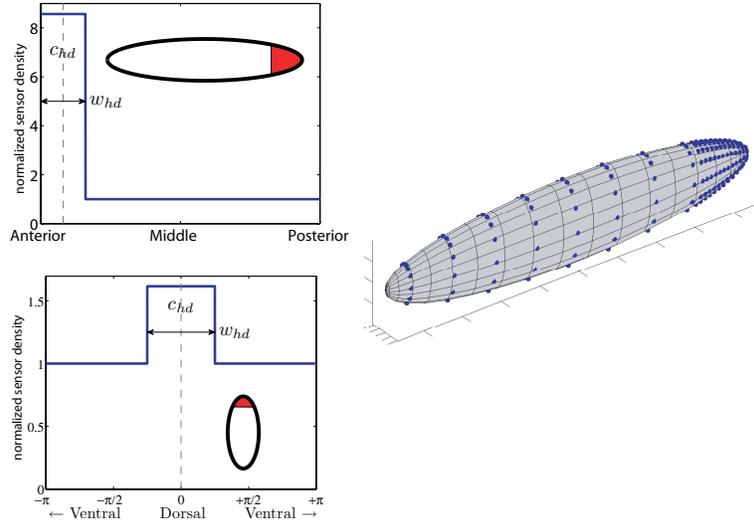


Figure 4.4. An example of a sensor layout using our parametrization. This is the sensor layout closest to the fish (i.e., the fish-like layout). The top left shows the sensor density in the AP dimension, and the bottom left shows the density in the DVRL dimension. The gray dashed line signifies c_{hd} , and w_{hd} is shown by the double-headed arrow. The figure on the right shows the locations of the sensors on the surface of the ellipsoid that corresponds to these two density functions.

four different c_{hd} s for the DVRL dimension ($-\pi/2$, 0 , $+\pi/2$, and $+\pi$) were considered for a total of twelve permutations (including the fish layout).

In addition, three more distributions had a uniform distribution in the DVRL dimension, and in the AP dimension had $w_{hd} = 0.32$ (just like the fish-like layout) and $p_{hd} = 1$ (i.e., all sensors in the high-density region). The three c_{hd} s used are -0.84 , 0 , and 0.84 , which place all the sensors near either the head, middle (in the A-P direction), or tail, respectively. Lastly, a uniform sensor distribution was examined for a grand total of 16 sensor distributions. The sensor distributions are summarized in Table 4.1.

| sensor concentration | | layout number | c_{hd} AP | c_{hd} DVRL |
|----------------------|---------|------------------|--------------------|----------------------|
| AP | DVRL | | | |
| head | right | 1 | -0.84 | $-\pi/2$ |
| | dorsal | 2, FL | -0.84 | 0 |
| | left | 3 | -0.84 | $\pi/2$ |
| | ventral | 4 | -0.84 | π |
| middle | right | 5 | 0 | $-\pi/2$ |
| | dorsal | 6 | 0 | 0 |
| | left | 7 | 0 | $\pi/2$ |
| | ventral | 8 | 0 | π |
| tail | right | 9 | 0.84 | $-\pi/2$ |
| | dorsal | 10 | 0.84 | 0 |
| | left | 11 | 0.84 | $\pi/2$ |
| | ventral | 12 | 0.84 | π |
| uniform | uniform | 13 | uniform | uniform |
| head only | uniform | 14 | -0.84 | uniform |
| middle only | uniform | 15 | 0 | uniform |
| tail only | uniform | 16 | 0.84 | uniform |

Table 4.1. The sixteen different sensor layouts. For sensor layouts 1 through 12 $\{w_{\text{hd}}, p_{\text{hd}}\} = \{0.32, 0.62\}$ for the AP dimension and $\{w_{\text{hd}}, p_{\text{hd}}\} = \{\pi/2, 0.35\}$ for the DVRL dimension. In layout number 13 the sensors are uniformly dispersed over the surface. For sensor layouts 14 through 16 $\{w_{\text{hd}}, p_{\text{hd}}\} = \{0.32, 1\}$ for the AP dimension and $\{w_{\text{hd}}, p_{\text{hd}}\} = \{2\pi, 1\}$ for the DVRL dimension. FL = fish-like sensor layout.

4.3.6. Three Types of Trajectories

Three different types of prey-capture trajectories are considered. One trajectory is based on the actual trajectory of a fish during the prey-capture task (mo-cap). Another is the mechanically optimal trajectory of the ellipsoid (mech-opt). The third type has neither sensing-optimal nor motion-optimal significance—it is merely another type of trajectory the fish could have executed (poly-fit). Twenty-five of each type of prey-capture trajectory were simulated.

4.3.6.1. Motion Caption (mo-cap). In a previously published study [154] motion capture data from weakly electric fish (*Apteronotus albifrons*) was recorded while the fish were performing the prey capture task. Individual water fleas (*Daphnia magna*) were introduced, and video recorded the fish’s motion as it captured the prey. Prey-capture events were digitized, and 3-D trajectories of the fish and prey were rendered using a model-based tracking system with spatial resolution of 0.5 mm and a temporal resolution of 1/60 second. Once the fish detects the prey it performs an abrupt deceleration that begins the prey-capture motion. We define *detection time* to coincide with the beginning of this abrupt deceleration.

4.3.6.2. Polynomial Fit of 6-DOFs (poly-fit). To see if the fish’s actual prey-capture trajectory is a good choice for efficient sensing we generated other trajectories that were not actual fish motion. A fourth-order polynomial was fit to each of the 6 degrees of freedom (DOFs) of the ellipsoid from the time of prey detection to the time of prey capture (the pre-detection trajectory remained identical to the motion capture trajectory). Four constraints were imposed to each of the 6 polynomials to match the configuration and

velocity at the time of detection and at the time of prey capture. The remaining degree-of-freedom of each polynomial was used to ensure the ellipsoid did not intersect the sphere (the prey is modeled as a sphere of radius $r_{\text{target}} = 0.083$ mm).

For example, if q_1 is a coordinate defining the configuration of the ellipsoid (e.g., the x position of the center of mass of the ellipsoid), then the time history of q_1 is defined by $q_1(t) = c_0 + c_1t + c_2t^2 + c_3t^3 + c_4t^4$. The coefficients $c_0, c_1, c_2,$ and c_3 are determined by imposing the position and velocity boundary conditions at detection time and prey capture time. The smallest c_4 coefficient that ensures that the surface of the ellipsoid does not intersect the surface of the prey is used as the final coefficient.

4.3.6.3. Mechanically Optimal Trajectories (mech-opt). We also investigate the sensing efficiency of mechanically optimal trajectories. The fish is modeled as a neutrally buoyant solid rigid ellipsoid immersed in an inviscid fluid moving under the influence of applied forces and torques F and T . The motion of the ellipsoid can be described using the Kirchhoff equations [139]. Kirchhoff's equations of motion are based on potential flow theory, and provide a dynamic model of the motion of the body [112].

$$(4.14) \quad M\dot{v} = Mv \times \omega + F,$$

$$(4.15) \quad I\dot{\omega} = I\omega \times \omega + Mv \times v + T.$$

where, v is the velocity of the center of mass of the ellipsoid in a frame rotating with the body, ω is the angular velocity of the body. M and I are the mass and inertia matrices, respectively. We allow the body to be fully actuated, that is, F and T can be any vectors in \mathbb{R}^3 .

The details of the optimization procedure are presented in [185]. The objective function to be minimized can be thought of as analogous to the metabolic cost of muscle activation for the fish to complete the trajectory. If $q(t)$ defines the time-history of the six degrees of freedom of the ellipsoid, then the objective function takes the following form:

$$(4.16) \quad J(q(t), \dot{q}(t), F(t), T(t)) = \sum_{i=1}^3 F_i(t)^2 + (\alpha_i T_i(t))^2.$$

The scaling factors α_i are included so that the terms in the sum have the same dimension. For the α_i they use the reciprocals of the radii of gyration of the ellipsoid, which turn out to be in the range from 0.26 to 1.93 cm⁻¹. The optimization is performed using an implementation of sequential quadratic programming [185]. The optimization routine finds a local minima of the objective function.

4.3.7. Computing environment

Computations were performed on a 54 CPU (2 GHz G5, 1 GB RAM) cluster of Xserves (Apple Computer Inc., Cupertino CA, USA) running OS X. An open source distributed computing engine (Grid Engine, Sun Microsystems, Santa Clara CA, USA) was utilized to manage the computation across the nodes. Simulation and analysis programs were written in MATLAB (The Mathworks, Natick MA, USA) and compiled to portable executables for execution on the cluster.

For each of the 75 different trajectories, the 16 different sensor layouts were run 20 times for a total of 24,000 prey capture simulations, taking more than a year of CPU time (approximately one week on the 54 CPU cluster).

CHAPTER 5

Electrosense Optimization for Prey Capture: Results and Discussion

In this chapter we present results and discussion from simulating the sensory acquisition and belief maintenance of the weakly electric fish. The fish's body is modeled as a 6-DOF rigid ellipsoid of similar dimensions as the fish (see Figure 4.2). The modeled electric field is generated by a line of point charges running from the head to the tail with a single negative charge at the tail and the remaining charges positive (see Section 4.3.1 and [71]). Electrosensors (i.e., voltage sensors) on the surface of the ellipsoid measure changes in the self-generated electric field. The observations are simulated with Rasnow's model for spherical objects in a uniform electric field [187]. The belief of the position of the target is maintained via a particle filter, which is a type of Monte Carlo Bayes filter. All of these methods have been thoroughly described in Section 4.3

5.1. The Tradeoff Between Number of Sensors and Sensor Noise

Apteronotus albifrons has approximately 14,000 probability type (P type) tuberous electroreceptor organs on its surface, yet we base our conclusions on simulations using much fewer than 14,000 sensors. To justify this simplification we demonstrate the existence of a sensing equivalence between an array of many noisy sensors and an array of fewer, less noisy sensors.

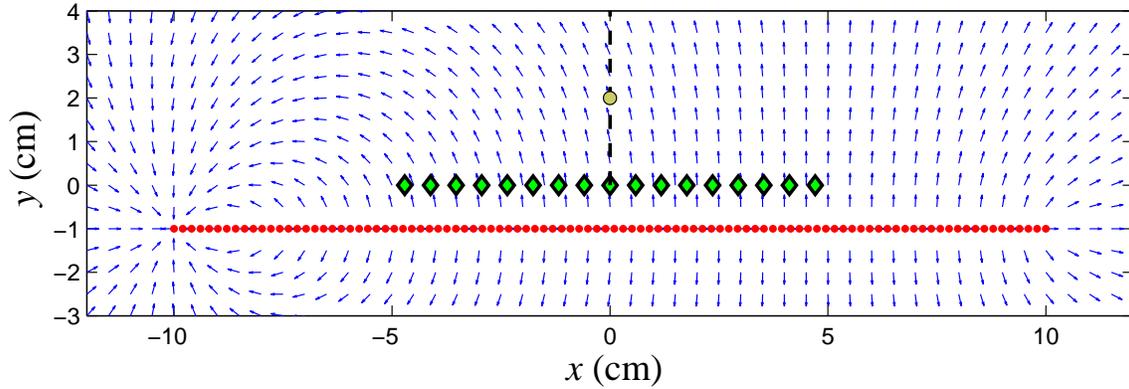


Figure 5.1. Setup for toy problem. Green diamonds and red dots are the sensors and emitter positions, respectively. Blue arrows show the direction of the electric field. The target location is confined to the black dashed line at $x = 0$.

5.1.1. Toy Problem

A simple experiment was devised to investigate the tradeoff between number of sensors and noise level of the sensors. A 20 cm line of 100 charges (99 positive and 1 negative, see Section 4.3.1 for details) is used to generate the electric field as depicted in Figure 5.1. The negative charge is at $x = -10$ cm. A 10 cm line of evenly spaced sensors is positioned 1 cm above the charges (see Figure 5.1). A single target is constrained to $x = 0$.

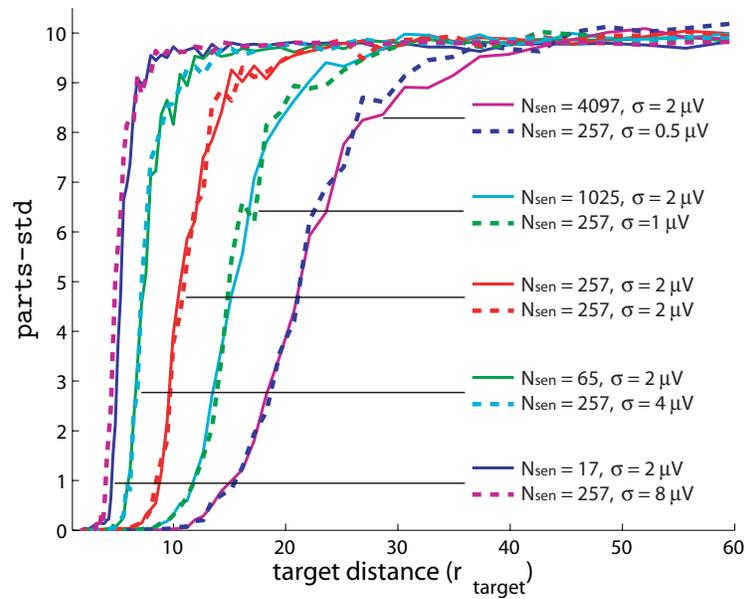
In the experiments a single target (same properties as the modeled water flea, see Section 4.3.2) begins at $(x = 0, y = 4)$ and moves to $(x = 0, y = r_{\text{target}})$, where r_{target} is the radius of the target in cm. We looked at two different types of sensor fusion. In *multi-step fusion* sensory information across all steps is used to construct the belief. In *single-step fusion* only information from the current time step is used to construct the belief. With single-step fusion at each time step the belief is initially completely uncertain.

Two sets of simulations were performed. In the first set the sensor noise was fixed at $\sigma_n = 2 \mu\text{V}$ and the number of sensors varied across the 13 trials, where $N_{\text{sen}} = (2^i + 1)$

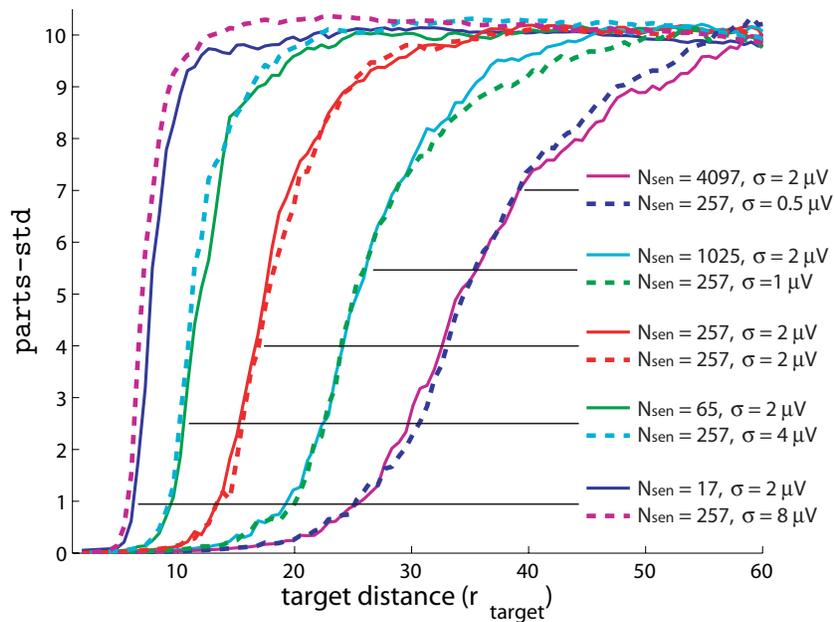
for $i = 1, 2, \dots, 12$ with an additional trial of $N_{\text{sen}} = 1$. The odd number of sensors ensured that there was always a sensor directly below (same x value) the target. Initially, the location of the target is completely uncertain and the belief is initialized as a set of particles randomly dispersed across feasible target locations. At each time step each sensor updates the belief exactly once. We use `parts-std` (See Section 4.3.4) as the measure of uncertainty of the belief. Simulations were performed for both the single- and multi-step fusion case. In the second set of simulations, the number of sensors was fixed at $N_{\text{sen}} = 257$. Sensors are assumed to have Gaussian noise of standard deviation σ_n . We considered ten levels of sensor noise, $\sigma_n = 2^i \mu\text{V}$, for $i = -4, -3, \dots, 5$.

As expected, as the number of sensors in the array increases, so does the effectiveness of the sensor array, and the sensing effectiveness also improves as the sensor noise decreases. This is true for the single- and multi-step fusion cases. In the multi-step sensor fusion scenario the `parts-std` norm decreased quicker than in the single-step sensor fusion case because the belief in the multi-sensor case has more updates.

From these experiments we found different combinations of σ_n and N_{sen} rendered almost identical histories of `parts-std`. Let $\xi = (\text{number of sensors})/(\text{variance of sensor noise}) = N_{\text{sen}}/\sigma_n^2$ denote the *sensing effectiveness* of particular combination of total number of sensors and the variance of the noise on the sensors. Figure 5.2 shows pairs of curves with similar ξ . This figure suggests there is a predictable tradeoff between the number of sensors and level of sensor noise. If this relationship holds for the conditions of the 3-D ellipsoid trajectories—which we show next to be true—then a smaller number of sensors can be offset by using a lower level of noise.



(A) single-step fusion



(B) multi-step fusion

Figure 5.2. Sensing equivalence. The solid lines have $2 \mu\text{V}$ noise on the sensors. The dashed lines all have 257 sensors. Thus, the four red curves (2 solid and 2 dashed) are $2 \mu\text{V}$ and 257 sensors. For each curve in a pair, the sensing effectiveness, $\xi = N_{\text{sen}}/\sigma_n^2$, is almost identical.

5.1.2. Number of Sensors vs. Sensor Noise in 3-D Simulation

To demonstrate this relationship between number of sensors and level of sensor noise is not just an artifact of the simple toy problem, we investigated this relationship in the trajectories of the 3-D ellipsoid using the simulation methods described in the previous chapter. In these simulations a prey-capture trajectory with the sensor layout most similar to the fish (sensor layout 2; see Section 4.3.5) was used. We examined three different numbers of sensors—64, 256, and 1024—and six different levels of sensor noise—0.5, 1, 2, 4, 8, and 16 μV . Each combination of number of sensors and level of sensor noise were tested for a total of 18 combinations. To ensure this relationship is valid in the range of the number of sensors on the fish ($\sim 14,000$), we also tested 16,384 sensors with 16 μV noise.

Figure 5.3 depicts the sensing effectiveness for different combinations of number of sensors and sensor noise. As with the toy example, curves of constant sensing effectiveness are very similar. It is apparent that the relationship suggested from the toy problem also holds for a much more complex system. The heavy dashed line represents the trial with 16,384 sensors.

The gray vertical line in Figure 5.3 is at the time of behavior prey detection (as defined in Section 4.3.6.1). The detection point is defined as the time when the fish just begins to decelerate, thus the information to make the decision to decelerate must have been present prior to the behavior detection point due to sensorimotor delays. Notice that the knees of curves of $\xi = 64 N_{sen}/\mu V^2$ are just before the time of detection. From this data it can be concluded that the fish probably had a sensing effectiveness somewhere between $\xi = 64$ and $\xi = 1024$ based on this detection point and a sensorimotor delay

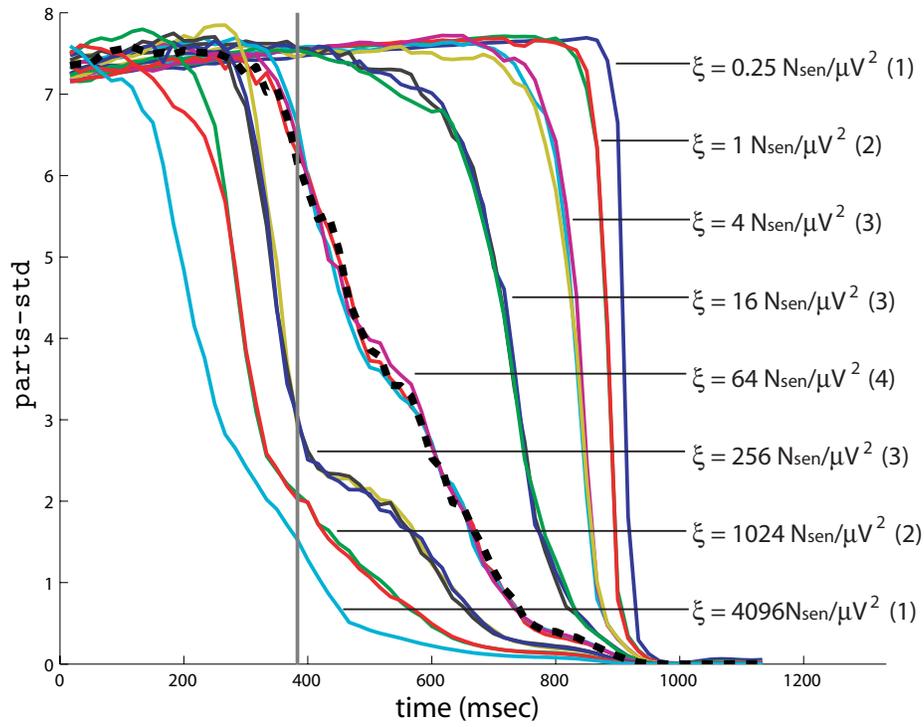


Figure 5.3. Sensing equivalence (ξ) for the 3-D ellipsoid. Each of the 19 curves corresponds to a different combination of number of sensors and sensor noise, where each combination corresponds to one of the eight ξ 's shown by the annotations. The numbers in parentheses signify the number of curves corresponding to the sensing effectiveness. As predicted, curves with equal sensing effectiveness are very similar. The heavy dashed line signifies the use of 16,384 sensors in the simulation. The gray vertical line denotes the time of behavioral prey detection. For each combination of sensor noise and number a sensors, 40 simulations were run, and each curve represents the mean of the 40 simulations.

of ~ 110 ms. With approximately 14,000 sensors on the fish's body, Gaussian noise with standard deviation $14.8 \mu\text{V}$ would equate to $\xi = 64$, and noise with standard deviation $3.7 \mu\text{V}$ would equate to $\xi = 1024$. Thus, if we are able to lump all the sensor-noise-related effects into a single parameter, we could say that the “effective noise” of the sensors is probably between $3.7 \mu\text{V}$ and $14.8 \mu\text{V}$ for this trial.

5.1.3. Why Does the Fish Have a Large Number of Sensors?

The number of sensors and the level of sensor noise of the weakly electric fish is the result of millions of years of evolution. This sensing optimization process can be stated using the framework developed in Section 4.2.3 and Equations (4.10) and (4.11), where $p = \xi(N_{\text{sen}}, \sigma_n^2)$ is being optimized as a function of number of sensors and level of sensor noise. If the fish has truly found an optimal solution for ξ (we do not really know if it has), then Pontryagin's principle states that $\partial\xi/\partial N_{\text{sen}} = 0$ and $\partial\xi/\partial\sigma_n^2 = 0$, and nearby (for local optimality) ξ 's are lower than ξ^* (the * denotes an optimal solution).

The lumped parameter σ_n includes endogenous sources such as transduction imperfections at the afferent level as well as exogenous sources such as clutter and external electric fields. The fish has no control over the exogenous components of σ_n and it is very difficult to evolve better physiology that would allow improvement of the endogenous components of σ_n . Thus, the fish has evolved a large number of sensors to improve its ξ .

5.2. Why is the Electric Field Stronger Near the Tail?

The electric field of the weakly electric fish is well-modeled as a line of point charges with a single, negative charge at the tail [71]. The direction of the electric field using

this simple model can be seen in Figure 5.1, where the tail is at $x = -10$ and the head is at $x = +10$. In this section we investigate why the fish would evolve such a structure for the electric field. The electric field is stronger near the tail than near the head, yet intuitively, one would think that the stronger part of the electric field should be located towards the head since this is where the fish localizes the prey. As in the previous section, this can be stated as a parameter optimization problem, where p defines the set of fish parameters that generate the electric field and are subject to optimization via evolution.

To help answer this question, we have constructed an *uncertainty map* of nearby prey positions. This uncertainty map displays the belief uncertainty (i.e., `parts-std`) as a function of prey position. To construct the map, the prey is placed on a grid near the fish. The belief is initially uncertain (i.e., a set of randomly dispersed particles) and each sensor updates the belief once via particle filtering. The uncertainty of the belief is then assigned to that point on the grid. (In the future, something akin to an uncertainty map could be used by a planner in an active sensing paradigm, similar to what was done in [222].)

Figure 5.4 shows two different slices (mid-sagittal plane and mid-coronal plane) for both the fish-like electric field and for an electric field when the large negative pole has been placed at the head (i.e., flipped). The fish-like electric field creates a very regular uncertainty map, while the flipped electric field is anisotropic. While the flipped electric field is better for localizing prey near the head, the fish-like electric field is a better complement to the fish's omnidirectional sensing and motor system [10].

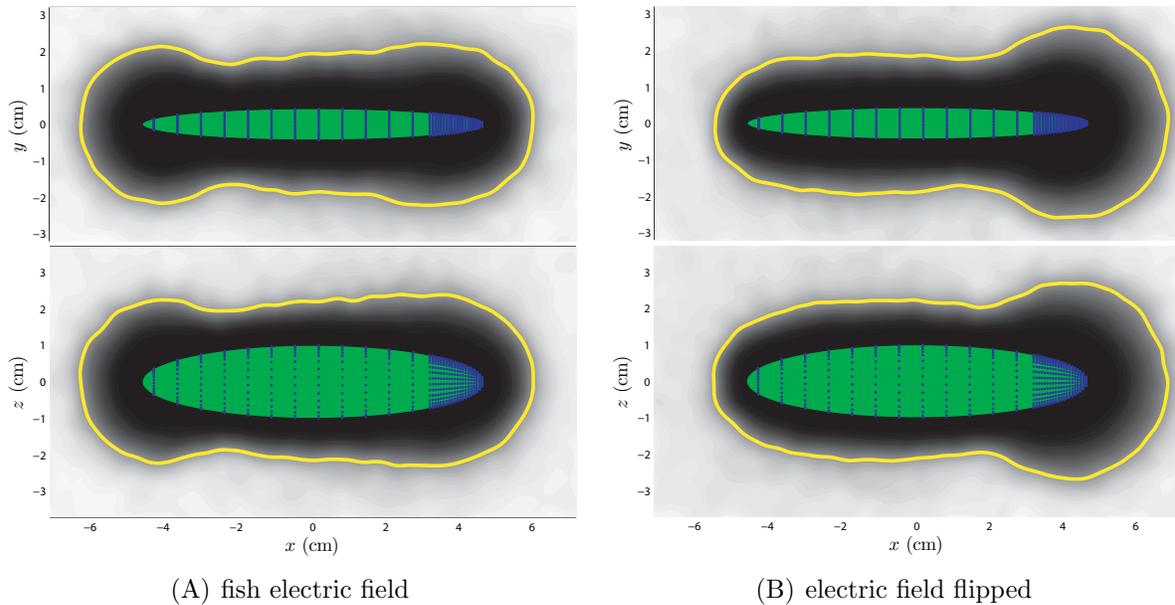


Figure 5.4. Uncertainty map of prey position for the normal electric field of the fish and when the electric field is flipped. We consider two cross-sections here—the mid-sagittal plane (bottom in each figure) and the mid-coronal plane (top in each figure). The green ellipses are the cross sections of the ellipsoids, and the blue dots are the 2-D projections of the sensors (the fish-like layout—layout 2—is used here). The gray-scale filled contours correspond to the resulting `parts-std` for a target in that location after each sensor has updated an initially uncertain prior. The yellow contour is at `parts-std` = 4, which is roughly half the value of `parts-std` with a completely uncertain belief (for example see Figure 5.3). We use $N_{\text{sen}} = 1024$ sensors and a sensor-noise level of $\sigma_n = 4 \mu\text{V}$, which corresponds to the sensing effectiveness used in the other simulations.

One method of quantifying the effectiveness of the electric field is by estimating the area that will render a belief with an uncertainty below a predefined threshold. For example, if $\text{parts-std}_{\text{max}}$ represents a completely uncertain belief, then the yellow contours in Figure 5.4 is at $\text{parts-std} = 0.5\text{parts-std}_{\text{max}}$. The area inside this contour, but not including the ellipse, of the flipped electric field case is 4% greater in the mid-sagittal plane and 3% greater in the mid-coronal plane than the normal electric field case. Across all parts-std from $0.01\text{parts-std}_{\text{max}}$ to $0.9\text{parts-std}_{\text{max}}$ the flipped electric field always creates a larger area but only by between one and five percent for both planes.

To quantify the level of isotropism of the sensing system we first determined the shortest distance from each point on the perimeter of the ellipse to a specified parts-std contour. The standard deviation of these distances is an indicator of the uniformness of the sensing system. For both the mid-sagittal and mid-coronal planes this standard deviation was approximately twice as large in the flipped electric field than in the normal electric field across all parts-std from $0.01\text{parts-std}_{\text{max}}$ to $0.9\text{parts-std}_{\text{max}}$. Also, the ratio of largest distance to the smallest distance was almost twice as great in the flipped electric than in the normal electric field case for the same range of parts-std and for both planes. Both of these measures of isotropism suggest the normal electric field facilitates a much more isotropic sensing system than the flipped electric field. Thus, a small increase in sensing volume may be less important than a large difference in the isotropism of the sensing system.

5.3. Determining the Best Sensor Layout

In this section we address the question of “Is the sensor layout of the *Apteronotous albifrons* a good choice for the prey-capture task?” An efficient distribution of sensors can positively affect the fitness (in evolutionary terms) of a fish. For example, a good sensor layout may require less expended mechanical energy to obtain the same level of prey uncertainty (i.e., a good sensor layout can offset motion).

Finding the best sensor layout is yet another case of parameter optimization. For our particular parametrization scheme, p defines the sensor densities on the two-dimensional surface of the ellipsoid (see Section 4.3.5 for details). For computational considerations we have reduced the set of possible p 's to 16 (listed in Table 4.1). Because the space of possible parameters is so small, we do not need to resort to any of the solution methods described in Section 4.2.4; we can try all 16 and see which is best.

In order to determine if the *A. albifrons*' sensor layout is a good choice for the electrolocation task, one the sensor layouts is similar to that of the actual fish (sensor layout 2). In the *A. albifrons* there are higher densities in the dorsal and anterior (head) regions (approximate densities were recorded in [61]).

16 total layouts are considered. Sensor layouts 1 through 12 are characterized by a region of high sensor density and the remainder of the ellipsoid surface has a lower density of sensors. These layouts were constructed by shifting the region of higher sensor density to three locations in the anterior-posterior direction (head, tail, and middle) and four locations in the dorsal-ventral-right-left direction for a total of 12 permutations (Figure 5.5(A) shows the one of these 12 that is most similar to the fish). Layout number 13 is a uniform distribution of sensors (See Figure 5.5(B)). Layouts 14, 15, and 16 only have

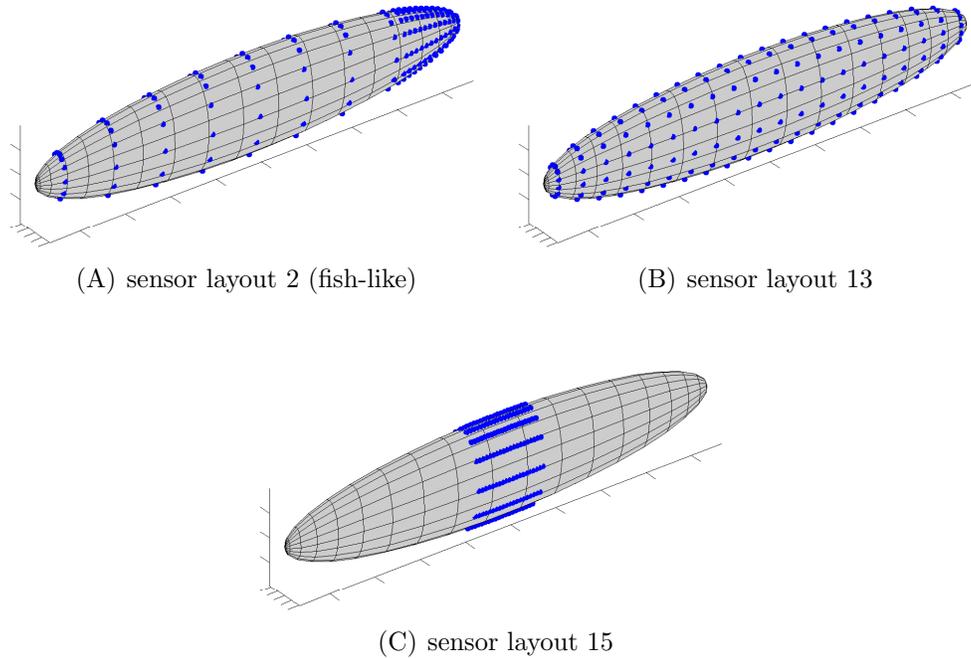


Figure 5.5. Three examples of sensor layouts with 256 sensors.

sensors near the head, middle, and tail, respectively (Figure 5.5(C) shows sensor layout 15). Section 4.3.5 describes the 16 layouts in more detail.

Three different types of prey-capture trajectories are considered. One trajectory is based on the actual trajectory of a fish during the prey-capture task (mo-cap). Another is the mechanically optimal trajectory of the ellipsoid (see mech-opt in Section 4.3.6.3). The third type fits polynomials to the six degrees-of-freedom of the ellipsoid, where the pose and velocities at detection time and prey-capture time are constrained to match the corresponding mo-cap trajectory (see poly-fit in Section 4.3.6.2). The time-history of each of the six degrees of freedom of the ellipsoid are defined by a polynomial. Twenty-five of each type of prey-capture trajectory were simulated. Both the mech-opt and poly-fit

trajectories are constrained to match the configuration and velocity of the corresponding mo-cap trajectory at the time of prey detection and at the time of prey capture.

5.3.1. Sensory Performance of a Trajectory

The `parts-std` norm estimates the uncertainty of a belief at a single point in time. For trajectory-comparison purposes we need to quantify the sensory performance of an entire trajectory. The uncertainty of the prey position at the time of prey capture is irrelevant since the shortest sensorimotor delay of weakly electric fish is about 100 ms, thus, the fish’s current actions are based on information that is at least 100 ms old. If we wish to examine how accurate the fish’s final action in the prey-capture sequence could be, then we can look at the average uncertainty metric over a time window ending at $t_{\text{pc}} - 100$ ms, where t_{pc} is the time of prey capture. It is within this window that the fish will decide on its final action. We consider `parts-std` over a 100 ms window from $t_{\text{pc}} - 200$ ms to $t_{\text{pc}} - 100$ ms. If we define `sense-utility` to be this sensing utility for the trajectory, then,

$$(5.1) \quad \text{sense-utility}[b(t)] = \frac{1}{100 \text{ ms}} \int_{t_{\text{pc}}-200}^{t_{\text{pc}}-100} \text{parts-std}[b(\tau)] d\tau$$

where $b(t)$ is the time history of the belief. Even though this measure only explicitly considers the belief near the end of the prey-capture task, the utility of earlier observations is still reflected in this measure since beliefs take into account all prior information.

5.3.2. Data for all Sensor Layouts and Trajectories

All 16 sensor layouts were tried on the 25 different trials and for each of the 3 types of trajectories. This was repeated (for statistical purposes) 20 times for a total of 24,960 prey-capture simulations. The average across all $N = 20$ repeats is shown in Table 5.1. In these simulations we used $N_{\text{sen}} = 64$ and $\sigma_n = 1 \mu\text{V}$, which corresponds to $\xi = 64 N_{\text{sen}}/\sigma_n^2$.

The **sense-utility** values in columns 4 through 6 in Table 5.1 (the numbers not in parentheses) are normalized by the best (out of 36) of the combination of layout number and trajectory type. The best **sense-utility** out of the 36 possibilities is denoted in boxed text. The numbers in parentheses are another method of evaluating sensor-layout performance, which is described below. A lower **sense-utility** is better, thus layout number 2 with the mo-cap trajectory is the best of 36 possibilities. Two claims can be made from the **sense-utility** data. First, for the given sample of 25 motion capture prey-capture trajectories, sensor layout 2 is the best, since the lowest **sense-utility** in column 4 is for the fish-like layout (FL = fish-like sensor layout). Secondly, given the fish has sensor layout 2, the mo-cap trajectory is better than both mech-opt and poly-fit for sensing, since the lowest **sense-utility** in the layout 2 row is in the mo-cap column. Thus, it appears that the fish has made a good choice—at least in the interest of sensing—of sensor layout and trajectory type for the prey-capture task.

In order to verify the implications of the **sense-utility** data in Table 5.1 we perform pairwise t-test for all combinations of sensor layouts for each trajectory. For a given pair of sensor layouts, if one pair is significantly better ($\alpha = 0.05$), then 1 point is added to the score of the superior sensor layout, and 1 point is subtracted from the inferior sensor layout (zero-sum game). With 16 different sensor layouts there are 120

| sensor layout description | | | normalized sensor-utility | | | sensor rank |
|--------------------------------|---------|---------------|---------------------------|--------------------|--------------------|-------------|
| sensor concentration head-tail | D-V | layout number | mo-cap | poly-fit | mech-opt | |
| head | right | 1 | 1.1 (92) | 1.7 (81) | 1.7 (108) | 2 (2) |
| | dorsal | 2, FL | 1 (108) | 1.5 (102) | 1.7 (111) | 1 (1) |
| | left | 3 | 1.1 (91) | 1.7 (78) | 1.8 (107) | 3 (3) |
| | ventral | 4 | 1.1 (92) | 1.7 (76) | 1.8 (105) | 4 (4) |
| middle | right | 5 | 1.5 (48) | 1.9 (71) | 2 (55) | 7 (7) |
| | dorsal | 6 | 1.4 (69) | 1.7 (84) | 1.9 (65) | 5 (6) |
| | left | 7 | 1.5 (52) | 1.9 (69) | 2.1 (47) | 8 (8) |
| | ventral | 8 | 1.6 (47) | 1.9 (68) | 2.1 (44) | 9 (9) |
| tail | right | 9 | 2.1 (-12) | 2.7 (-34) | 2.8 (-24) | 12 (12) |
| | dorsal | 10 | 1.9 (-2) | 2.5 (-24) | 2.7 (-32) | 11 (11) |
| | left | 11 | 2.1 (-21) | 2.6 (-46) | 2.9 (-42) | 13 (13) |
| | ventral | 12 | 2.2 (-45) | 2.8 (-42) | 3.1 (-64) | 14 (14) |
| uniform | uniform | 13 | 1.4 (79) | 1.8 (70) | 2.1 (70) | 6 (5) |
| head only | uniform | 14 | 1.7 (7) | 2.1 (20) | 2.1 (36) | 10 (10) |
| middle only | uniform | 15 | 3.7 (-231) | 4.8 (-210) | 5.3 (-213) | 15 (15) |
| tail only | uniform | 16 | 10.7 (-374) | 11.6 (-363) | 11.7 (-373) | 16 (16) |
| trajectory total | | | 36.3 | 44.8 | 47.9 | |

Table 5.1. The sensing performance for each combination of sensory layout and trajectory type. The numbers in parenthesis are the scores from performing the pair-wise t-tests between alternative layouts within the same trajectory type. The first three columns describe the type of sensor layout (see Table 4.1 for details; sensor layout 2 is most like the fish, which is denoted by **FL**). D-V is short for dorsal-ventral. The next three columns show the average normalized **sense-utility** for each combination of layout number of trajectory type. The last column ranks the sensor layout according to sum of columns 4 through 6 (lower rank is always better). The bottom row sums the normalized **sense-utility** for each trajectory type across all sensor layouts. The three rows in blue text (sensor layouts 2, 6, and 10) correspond to sensor layouts with higher densities on the dorsal edge. The best **sense-utility** is denoted by boxed text, which is the mo-cap trajectory with layout number 2. The highest pair-wise t-tests within each trajectory type are also denoted by boxed text. All of these are from sensor layout number 2.

possible pairs for each trajectory type ($16 \text{ choose } 2 = 120$). Performances of the sensor layouts are only compared within the same type of trajectory (zero-sum game for each type of trajectory). Therefore, the maximum possible score for any sensor layout is $(15 \text{ other layouts}) \times (25 \text{ trajectories}) = 375$. Similarly, if a sensor layout was significantly worse than all others across all trajectories it would receive a score of -375 . These results are summarized as the values in parentheses in Table 5.1. Evidence supports the hypothesis that the sensor layout closest to the fish is the best. In both analyses sensor layouts with highest concentrations near the head (layouts 1 to 4) are the four best for all three trajectory types. Also, within each group of 4 with the same head-to-tail density (sets $\{1, 2, 3, 4\}$, $\{5, 6, 7, 8\}$, and $\{9, 10, 11, 12\}$) the best sensor layout is the one with the highest density at the dorsal edge (i.e., layouts 2, 6, and 10). In Table 5.1 these three layouts are signified with blue text.

It is important to note that with the significance scores from Table 5.1 we cannot draw any conclusions across trajectory types, since the t-tests were performed across the 16 layouts but within the same trajectory type. Thus, we can determine the relative ordering of the sensor layouts only within each trajectory. Comparisons across trajectory types using the significance scores in Table 5.1 are meaningless.

5.3.3. Is There Sensor Specialization?

The four best sensor layouts from Table 5.1 all had the highest sensor density near the head. But, sensor layout 14, which only has sensors near the head, performed relatively poorly. Upon inspection of the time history of `parts-std`, in most trials it appears that sensors near the middle (along the head-to-tail direction) are responsible for gross

refinement of the belief, while sensors near the head are used for fine localization. Early in the time history of the uncertainty of the prey, there is little belief refinement as the signal-to-noise (SNR) ratio is well below unity. As the SNR approaches unity `parts-std` rapidly decreases. Often, the prey is nearest the middle of the body during this phase. This may explain why the uniform sensory layout (number 13) performs much better than layout 14 (all sensors near head). Is the location of sensors on the fish’s body specialized for different levels of localization accuracy? For example, sensors further from the fish’s mouth could be responsible for gross detection due to possible high convergence of afferent signals downstream, and sensors closer to the mouth could have better spatial accuracy due to low spatial convergence of the signal downstream.

5.4. Interdependencies Between Sensing and Motion

Ideally, acquiring information effectively and moving in a mechanically efficient manner should be done simultaneously. This requires a co-optimization of the sensing and mechanical systems. In this section we show the fish’s sensing and motor systems have co-evolved for the prey-capture task

The active sensing problem, which was formally defined in Section 4.2.2, is moving to acquire the best information possible. The Lagrangian for a pure active sensing cost is defined in Equation (4.8), which can be used as the cost functional J in Equation (4.5) to define the active sensing optimal control problem. As suggested in Section 4.2.4, we make no claim of finding the global optimal trajectory. Instead, we examine 75 trajectories generated from either motion-capture data (`mo-cap`), a mechanical optimization solver (`mech-opt`), or fitting polynomials to the degrees-of-freedom (`poly-fit`). For each

25 mo-cap trials, both poly-fit and mech-opt have configuration and velocity boundary conditions that match the corresponding mo-cap trial. Our conclusions are based on the sense-utility metric of the 75 trajectories.

5.4.1. Two Types of Trials

In many of the mechanically optimal trajectories, just after detection time the ellipsoid will roll to a configuration that allows its dorsal edge to slice through the fluid. Then, just prior to prey capture, it must roll back so its pose and velocity match the mo-cap trajectory at the time of prey capture. When the ellipsoid leads with its dorsal edge, the amount of fluid it is accelerating is minimized, which is good for minimizing mechanical cost. Unfortunately for sensing, this maneuver often increases the distance between the target and its sensors and emitters. Roughly speaking, the sensor signal decreases with the 4th power of target distance for small targets over the behaviorally relevant range of target distances [71].

For each trajectory we determined if the ellipsoid performed a “big-roll” maneuver or not for both the mech-opt and mo-cap trajectories. The presence of a big-roll is determined by examining the angular velocity about the major axis (the axis running head to tail) of the ellipsoid at a 150 ms window after detection and a 150 ms window prior to prey capture. If the angular velocity in both windows reach a max (roll) velocity of 0.9 rad/sec, then it is concluded that this trajectory has performed a big-roll maneuver.

The optimization routine used to find the mech-opt trials is not aware of the constraints of the real fish and searches across a fully-actuated movement repertoire for the optimal trajectory. In many situations the real fish is unable to perform a big-roll maneuver, but

in some situations it can [185]. This motivates a distinction between the fish-like mech-opt trials and the non-fish-like mech-opt trials. In the *fish-like* trials at least one of the two following conditions is met: 1) The mech-opt trajectory did not perform a big-roll; or 2) The mo-cap trajectory did perform a big-roll. Ten of the twenty-five trials fit this criteria. The remaining 15 trials are considered *non-fish-like*, since these are the trials where the mech-opt did perform a big-roll and the mo-cap did not.

5.4.2. Sensing Performance of Trajectory Types

To compare sensing performance across trajectory types we use the normalized **sense-utility**, where the **sense-utility** is divided by the average **sense-utility** for all three trajectory types. Thus, if the normalized **sense-utility** is much less than unity for a particular trajectory type, then it is pretty good relative to the other two. Columns (2) through (4) of Table 5.2 show the sensing performance (normalized **sense-utility**) for all 25 trials as well as just the fish-like and non-fish-like trials. The table only considers trajectories using sensor layout 2, since this is closest to the fish.

5.4.3. The Sensor-Distance Metric as a Predictor of Sensing Performance

A possible predictor of sensing performance is the distance from the closest sensor on the ellipsoid to the prey averaged along the trajectory, which we call **sensor-dist**. To compute **sensor-dist** for a trajectory, at each time step we determine the distance from the center of the prey to the nearest sensor on the ellipsoid. The average of these distances across all time steps from detection to prey capture is the **sensor-dist** metric. Column (5) of Table 5.2 is the normalized **sensor-dist** metric for the mech-opt trials, which

| (1) | sense-utility | | | sensor-dist |
|-----------------------------------|--|-----------------|--|-----------------|
| | (2) mo-cap | (3) poly-fit | (4) mech-opt | (5) mech-opt |
| all trials ($N = 25$) | 0.89 | 1.02 | 1.09 | 1.05 |
| fish-like trials ($N = 10$) | 0.97 | 1.06 | 0.97 | 0.99 |
| non-fish-like trials ($N = 15$) | 0.83 | 1.00 | 1.17 | 1.10 |

Table 5.2. Performance of fish-like and non-fish-like trials. Columns (2) through (4) compare the trajectory types for each of the three sets of trials. The **sense-utility** is normalized by the average **sense-utility** (of the three trajectory type) of the trial. Thus, for each row, columns (2) through (4) sum to 3. The best **sense-utility** for each set of trials is shown in boxed text. Column (5) is the normalized **sensor-dist** metric for the mech-opt trials. There is a significant different ($\alpha = 0.05$) in **sense-utility** between the mo-cap and mech-opt trajectories for both the “all trials” set of data and the non-fish-like trials.

is computed by dividing the **sensor-dist** of the mech-opt trajectory by the average **sensor-dist** of all three trajectory types.

Figure 5.6 suggests a positive correlation between the normalized **sense-utility** of the mech-opt trajectories (column (4) in Table 5.2) and the normalized **sensor-dist** metric (column (5) in Table 5.2). Thus, for good sensing performance it is good to get as close to the target as possible as quickly as possible. Of course, the fish does not know the exact location of the prey, so it cannot use this metric in a real-time control scenario. Here, we use it to gain some insight as to why the mech-opt trials perform poorly.

5.4.4. Co-optimization in the Prey-Capture Task

If the ellipsoid is endowed with full actuation over its 6 degrees-of-freedom, then the mechanically-optimal trajectory usually performs poorly for sensing. The conclusion is drawn from the first row of Table 5.2 (“all trials”). The second row of Table 5.2 is the average across the 10 trials where the mech-opt trajectories performed maneuvers similar

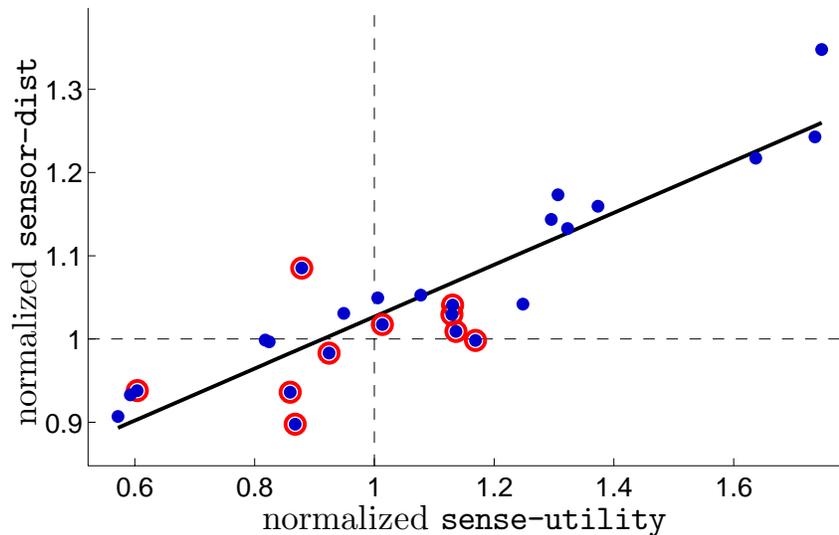


Figure 5.6. The normalized **sense-utility** of the mech-opt trajectories (column (4) in Table 5.2) versus the normalized **sensor-dist** metric (column (5) in Table 5.2). Each blue dot is a trial number. The solid line is the linear-least-squares fit ($R^2 = 0.82$). Red circles signify the 10 fish-like trials.

to what the actual fish could do. For these trials the mech-opt performed as well as the mo-cap trajectories. *Thus, the mechanically-optimal trajectories are also good for sensing when only allowed to perform maneuvers similar to what the fish can do.* This suggests that for sensing prey during the prey-capture task a fully-actuated fish possesses unnecessary actuation capabilities. These additional capabilities would come at the cost of additional tissue and bone mass, which could decrease the motion performance of the fish.

Table 5.2 shows the relative performance between trajectory types for the fish-like sensor layout (layout number 2). Thus, our conclusions are based on the fact that sensor layout is fixed and the actuation system has evolved for this layout. But, even if we fix the trajectories and are able to vary the sensor layout across all sixteen described earlier,

then it is still the case that the mech-opt trajectory type is the worst of the three for any of the sixteen layouts. Thus, even if the fish was fully-actuated, then none of the sixteen sensor layouts would allow the mechanically-optimal trajectory to be the best for sensing.

Part 3

The Electrosenster: A Robotic Electrosensory System

The electrosensory system presented in this part consists of a 2-degree-of-freedom (DOF) XY planar robot, which controls the position of the generated electric field and the sensors used to estimate the location of the target. Using various control algorithms, the robot can successfully locate nearby targets, as well as localize itself when placed in a pre-mapped environment, in both fresh and saltwater. Most of the controllers address the stochastic nature of the electrosensory observations by employing Bayesian fusion to calculate a belief of the position of the target. These controllers use empirical calibration of the sensor for the specific target and environment. Another controller, used solely for electrolocating individual targets, exploits invariant spatial features of the electric field signatures of a class of targets to achieve electrolocation without calibration. In most trials its precision was better than one millimeter.

CHAPTER 6

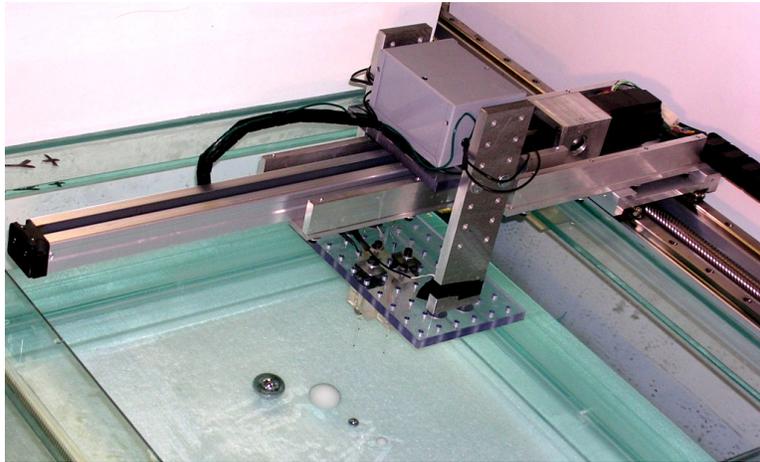
Experimental Setup

This chapter describes a planar Cartesian robot that uses voltage measurements from its self-generated electric field to estimate the state of its world. *The Electrosenster* is composed of the robot, its signal processing circuitry, and its supporting software [214]. The Electrosenster is shown in Figure 6.1. Primary, the robot’s task is determining the position of a single target with some assumed properties (electrolocation).

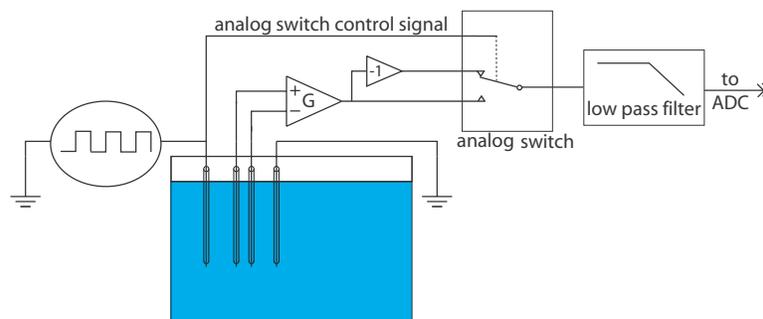
6.1. Related Systems

A survey of the literature to date reveals only two different physical systems that utilize measured voltage perturbations from a self-generated electric field to estimate the position of a target. One group built an artificial electrosensory system to investigate the possibility of using it for obstacle avoidance in underwater robots [73] (referred to here as the “Chetty system”). Their system could detect either a conducting or insulating sphere as small as 25 mm in diameter as far away as 5 mm. In their preliminary experiments their underwater robot could perform obstacle avoidance using this electrosensory system.

A few years later another study built an artificial electrosensory array to study the feasibility of using such a system in underwater robots [156] (referred to here as the “MacIver system”). They were able to determine the distance of a submerged 1 cm diameter plastic sphere up to a distance of about 12 mm away from the sensor. Their distance estimation algorithm was based on the time history of the sensor measurements.



(A) Photo of the Electro-sensiter with four spherical targets at the bottom of the tank. The larger targets have diameters of $1\ 1/2''$, and the other two are $1/2''$ diameter.



(B) Simplified schematic of electronics. The electronics will be thoroughly described in Section 6.5.

Figure 6.1. The Electro-sensiter.

The Electrosenster is an improvement to both of these systems in four ways. First, the Electrosenster employs probabilistic belief maintenance techniques. As previously seen, such techniques utilize statistics of sensing and motion to maintain a belief of the position of the target. Equipped with this belief of the target position, the controller can make informed decisions. Secondly, neither of the other two systems actively seek to improve the belief of the target (i.e., exploration or active sensing). Such a control scheme allows for an efficient localization paradigm. Thirdly, only the Electrosenster can provide a 2-D spatial estimate of the target. The MacIver system only estimates a scalar distance, and the Chetty system makes no attempt to estimate quantitative distances. The Chetty system uses the measured electric field perturbation to steer a submerged vehicle away from the object. Lastly, the Electrosenster boasts superior detection range to either of the other two.

6.2. Aquatic Environment

Experiments were conducted in a 750 mm by 750 mm glass tank filled to a depth of approximately 160 mm. To minimize the effects of the tank walls on the electric field, experiments were conducted in a central region of 200 mm by 200 mm. We used two different types of water, the first to mimic conductivity conditions in the natural habitat of weakly electric fish (“fresh”) [157, 175], and the second to mimic ocean water conductivity. For each type of water sodium-chloride was added to de-ionized water until the desired salinity was obtained. The salinities of the two water types resulted in conductivities of 160 $\mu\text{S}/\text{cm}$ for fresh water, similar to the native habitat of weakly electric

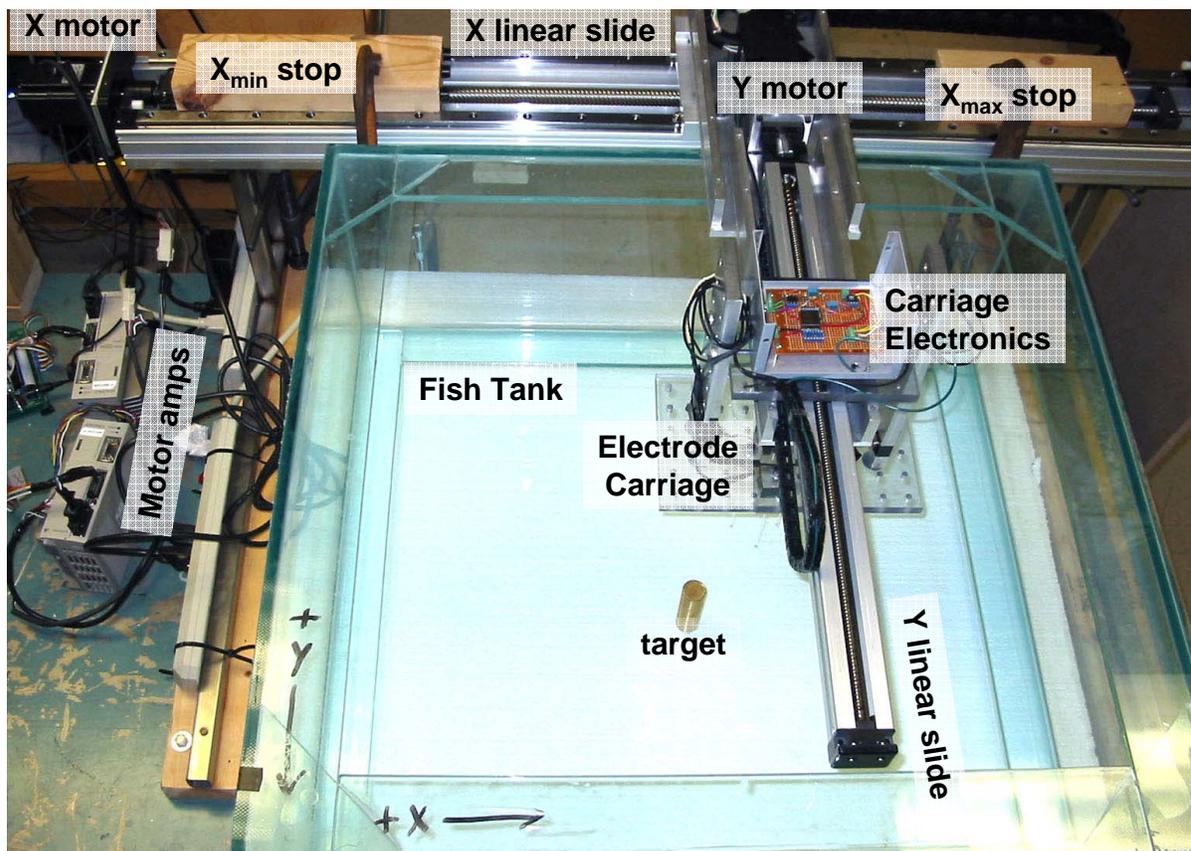


Figure 6.2. Top view photo of the hardware. A 1 inch diameter brass cylinder is in the fish tank. The electrodes are fastened to a translucent plastic breadboard.

fish [80], and $56,000 \mu\text{S}/\text{cm}$, an upper bound for marine water conductivity (typically $43,000 \mu\text{S}/\text{cm}$ [81]).

For many of the targets used, a small stand in the tank suspends them such that the tops of the targets cleared the bottom of the electrodes. The electrical signature of the target stand alone did not significantly interfere with any of the measurements.

6.3. Hardware

The hardware of the Electroster consists for the following components, which are shown in Figure 6.2:

- (1) two linear slides
- (2) two motors
- (3) two motor amps
- (4) a platform
- (5) the carriage
- (6) electrodes

The two linear slides provide the planar, two degrees-of-freedom of motion. The **X** linear slide (P/N GL20-S-40-1250L from THK Company, LTD) has a 1.25 meters ball screw stoke and 40 millimeter per revolution pitch. Physical stops were affixed to the track (labeled “ x_{\min} stop” and “ x_{\max} stop” in Figure 6.2) that allowed for only 37 cm of travel in the x direction. The **Y** linear slide (P/N VLA-ST-45-05-0500 also from THK Company, LTD) has 0.5 m ball screw stoke and 6 mm/rev.

Both of the motors used to power the linear slides are Yaskawa AC servomotors (P/N SGM-02B312). The motors are driven by the Yaskawa SGD-02BS Servo Pack that are able to provide 200 watts of power.

A welded steel platform raises the **X** linear slide 30 cm above the ground. This allows for the **Y** linear slide to hover above the fish tank and for the carriage to suspend its four electrodes into the water. The carriage contains some of the preliminary signal processing circuitry (see green shaded area in Figure 6.3). The carriage also suspends a

translucent plastic breadboard for configuring the positions of the electrodes. The holes on the breadboard are spaced at an inch (25.4 mm) on a 10 by 6 grid.

6.4. Electrodes

This electric field is transmitted between two silver electrodes submerged in water, which were plated with silver chloride to improve the metal-water electrical interface [191]. Refer to Appendix B for a description of the electroplating process. For most of the experiments the electrodes are 50 mm apart. The two sensing electrodes are identical to the electric field emission electrodes. Both the electric field emission electrodes and the sensing electrodes are made from 0.38 mm diameter silver wires and are stabilized by 0.5 mm borate silicate glass pipettes (See [188] for more details on hardware for measuring electric fields). The two sensing electrodes were placed in positions that nominally correspond to the same voltage in the electric field when no target is present. This allows us to input these signals to a high-gain differential amplifier.

6.5. Electrical Circuitry

The electrical circuitry described in this section is responsible for generating the electric field (stages (1) through (5) in Figure 6.3), measuring the voltage at the sensing electrodes, and then producing a low-noise, low-frequency signal that the analog-to-digital converter (ADC) can utilize. The twelve components of the circuit depicted in Figure 6.3 are:

- (1) **oscillator**: generate 2 MHz TTL-level pulses
- (2) **binary counter**: downsample to 1 and 2 kHz
- (3) **XOR gate**: produce 90° phase shift

- (4) **high-pass filter**: remove DC bias and set gain
- (5) **fish tank**: interface between generated electric field and sensing electrodes
- (6) **high-pass filter**: buffer and remove DC bias
- (7) **offset adjustment**: adjust input levels for the differential amplifier
- (8) **instrumentation amplifier**: produce difference of two signals
- (9) **high-pass filter**: clean up signal before processing
- (10) **signal inverter**: produce negative of signal
- (11) **analog switch**: keep portion of signal that is in phase with control signal
- (12) **low-pass filter**: smooth out for ADC

The electric organ discharge (EOD) of the weakly electric fish is a sinusoid-like emission of about 1 kHz. We use a biphasic square wave of 1 kHz. The choice of square wave is simply because the electronics are easier to implement. In the future if we would like to use a sinusoid, we could send the square wave through a multi-pole low-pass filter before transmitting through the water. In simulation an 8-pole low-pass Butterworth filter with cutoff frequency of 1.2 kHz produced a nice clean 1 kHz sinusoid from a 1 kHz square wave. It is also important to high-pass filter the signal before it is transmitted through the electrode to minimize corrosion of the electrode. The high-pass filter will subtract a DC offset and minimize charge buildup at the electrode/water interface.

Generation of the electric field signal begins at stage (1) in Figure 6.3, which is a crystal oscillator (P/N SG-531P) that puts out pulses at 2 MHz. The pulses are then put into a 14-stage binary ripple counter (Philips P/N 74HC4020N) that generates two square waves of interest; one at 1 kHz and the other at 2 kHz (stage (2)). We also need a 1 kHz square wave that is 90 degrees out of phase for gating one of the two analog switches (for

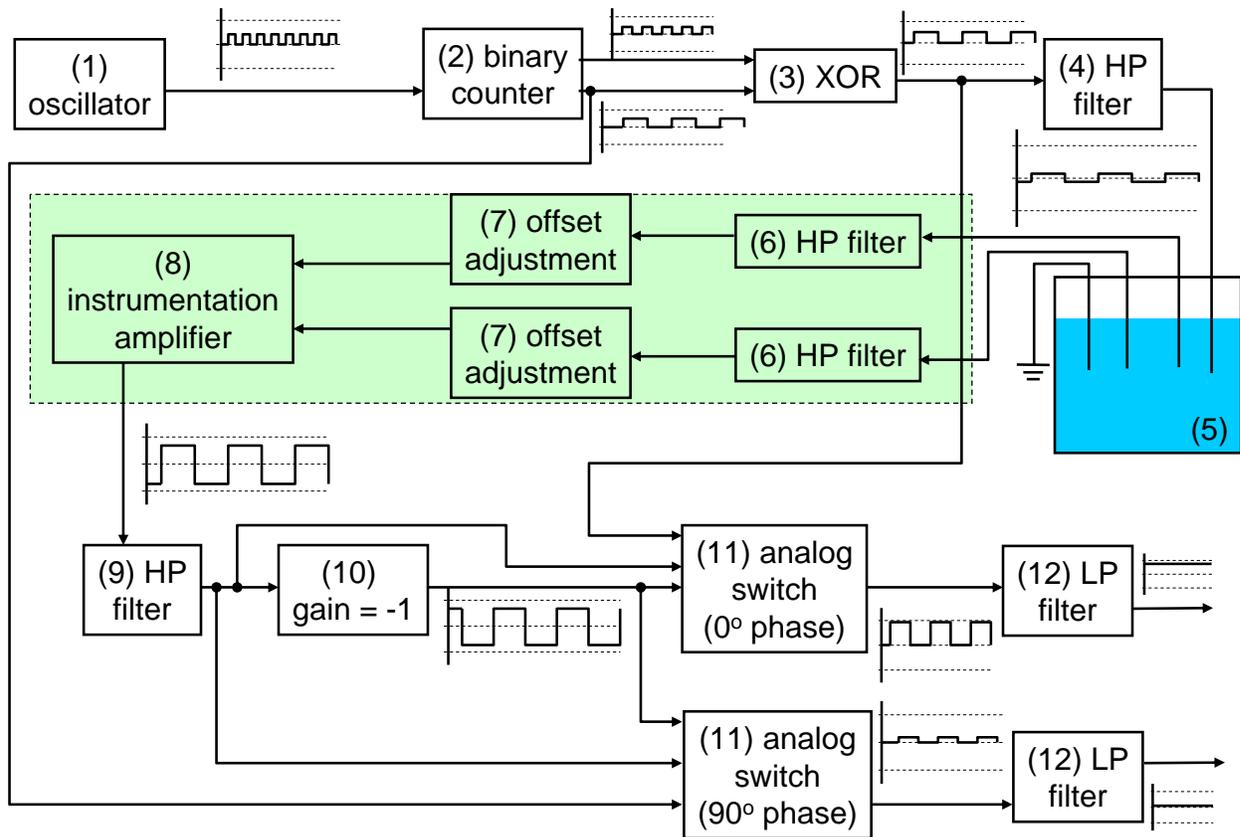


Figure 6.3. Block Diagram of the signal processing circuitry. The blocks represent the twelve stages of the circuit. Arrows indicate the flow of the signal. An illustration of the signals are shown at various points throughout the diagram; the three dashed lines in the insets represent -12 , 0 , and $+12$ volts. The green highlighted blocks (6 through 8) are implemented on the carriage of the Electroster.

later in the sensor circuitry). This 90 degree shifted wave can be created by sending the 1 kHz and 2 kHz square waves (of TTL levels) through an XOR logic gate (stage (3)). The 1 kHz square wave is sent through a high-pass filter (stage (4)) to remove DC offset and amplify the signal before being transmitted through the fish tank (stage (5)). Stage (6) is responsible for eliminating any induced DC bias on the sensing electrodes (minimize electrode corrosion).

Both the electric field emission electrodes and the sensing electrodes are simply silver wires that are stabilized by glass pipettes. The two sensing electrodes were placed in positions that nominally correspond to the same voltage in the electric field. This means that when there is no target in the tank, both electrodes will measure the same voltage. The circuit can be fine tuned with two potentiometers that act as voltage dividers (stage (7)). So, when we put these two signals through a differential amplifier (e.g. instrumentation amplifier, Burr-Brown INA 129) with nothing in the tank there is zero volts at the output of the amplifier (stage (8)). But, when these voltages at the electrodes deviate from their nominal values, these small differences can be greatly amplified. The differential amplifier serves two purposes: 1) common mode rejection: any voltage experienced by both electrodes (e.g. external noise) is canceled out. 2) increased resolution: We really only care about changes in voltages, and these changes are usually on the order of mV to 10's of μV , which may be difficult to detect. But, these voltage changes are amplified to levels that can be more accurately read by our electronics.

After leaving the differential amplifier the signal is put through another high-pass filter to once again ensure there is no DC bias in the signal (stage (9)). Stage (11) of the circuit involves two analog switches. An analog switch passes one of the two input

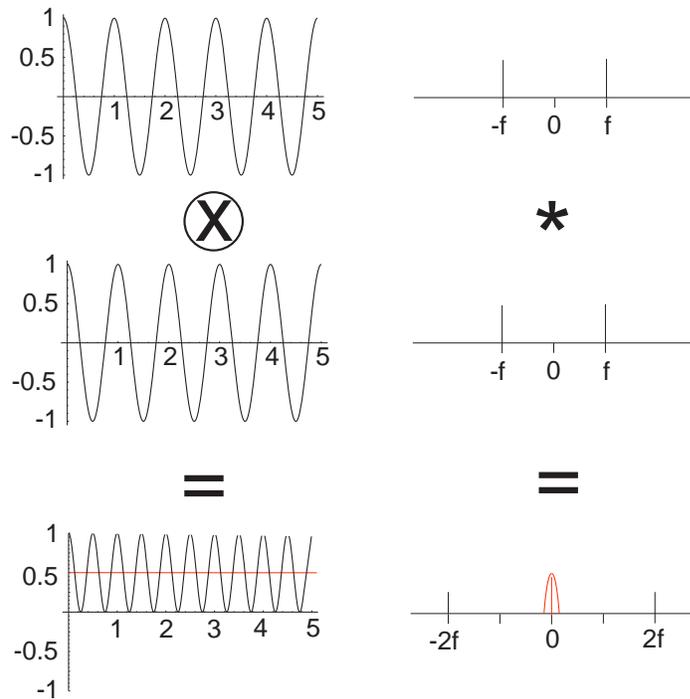


Figure 6.4. Synchronous detection. The left shows the process in the time domain, while the right column occurs in the frequency domain. This figure is taken from [208]

signals depending on the state of the switch. In our circuit the two signals are the signal from the differential amplifier and the negative of that signal, which was generated from stage (10). The square wave used for the electric field is used as the controller for the switch. The key point of this stage of the circuit is that the analog switch only allows portions of the signal that have the exact same frequency as the switch controller signal (electric field signal) to have a non-zero bias at the output of the analog switch. While all the signal is passed through the analog switch, half of the signal is the original signal's negative. Recall that we had a high-pass filter stage before the analog switch. This ensures a zero-mean signal before entering the analog switch. But, the component of the signal that exactly matches the frequency of the switching frequency will have a

non-zero-mean at the output of the analog switch and all other frequency components will keep their zero-mean characteristic. A low-pass filter will essentially average out all frequency components that are significantly higher than the filter's cutoff frequency. Thus, all frequency components of the signal leaving the analog switch that are higher than the low-pass filter's cutoff frequency and do not match the switching frequency of the analog switch are averaged out to zero. This process is very similar to its analog counterpart, synchronous detection [2], which is demonstrated in Figure 6.4. The signal leaving the low-pass filter is read in by the analog-to-digital-converter (ADC) on the PC104 stack. The analog switch/low-pass filter stages are duplicated, but the switching control signal is now phase shifted by 90 degrees. So, in the end, the Electrosenster produces two analog channels.

6.6. Computer Hardware Structure

The XY table is controlled by a PC-104 form-factor computer running the xPC real-time operating system from Mathworks. This computer is referred to as the xPC target and runs a compiled Simulink model in real time. The xPC target has four functions:

- (1) home the encoders of the motors
- (2) position motors to set point
- (3) collect raw sensor data
- (4) perform preliminary signal processing on sensor data

High level control is handled by a local desktop PC, referred to as the xPC host. The xPC host is connected to the xPC target through a router via ethernet. The xPC host

simply passes goal positions to the xPC target and receives the conditioned signal data. The xPC target runs a PID controller at 1 kHz on the current goal position.

6.7. Electric Field

The electric field is generated by two exposed silver-chloride spheres in a fish tank. The qualitative structure of this electric field is well-approximated by the electric field generated from two point charges in an infinite medium (see dipole model in Section 2.3). The justification of this approximation is in Appendix C. A 2-D slice of the electric field using this approximation is depicted in Figure C.1 on page 264. Thus, the dipole model discussed in previous chapters serves as a good model for the Electrosenster.

The power consumption of a sensing system can be an important consideration. In fresh water conditions the electric field of the Electrosenster consumed 0.3 mW of power. In saltwater, the electric field consumed 40 mW of power, corresponding to the two orders of magnitude higher conductivity of the saltwater over the fresh water.

CHAPTER 7

Sensor Characterization

In this chapter the electrosensory system of the Electrosenster is characterized. We first show how to construct an empirical sensor model using data collected from the Electrosenster. Then, the detection distance is investigated. We use simple models to develop an analytical solution of the detection distance. This theoretical model is compared with measured detection distances from the Electrosenster.

7.1. Empirical Construction of Sensor Model

The “dipole model” (see Section 2.3, page 50) discussed in previous chapters is based on the Electrosenster. Thus, essentially the same type of sensor model used for the dipole model is used for analyzing the Electrosenster. The only difference is that the sensor model for the Electrosenster will be constructed from measurements of the target rather than from a theoretical model. As with the dipole model, the position of the target in the z -direction (vertical) is set such that the bottom of the electrodes clear the top of the target by approximately 2 mm (the emitters and detectors are constrained to move in the $z = 0$ plane). When reference is made to the location of the electrodes or target, the geometric center is used to define its location. For example, with spherical targets of radius r_{target} mm if the emitter has radius r_{emitter} mm, then the z -distance between the emitter and target when they both are located at the same x and y coordinates is $r_{\text{target}} + r_{\text{emitter}} + 2$ mm.

7.1.1. The Sensor Model

The probabilistic sensor model used for the Electrosenster can be written as $p(w|\mathbf{x}, \mathbf{e})$, which represents the likelihood of observing w conditioned on the position of the target in the sensor frame, \mathbf{x} , and under the environmental conditions defined by \mathbf{e} . For the Electrosenster $w = V_1 - V_2$, and experimental data indicates that sensor noise is approximately zero-mean and normally distributed. Thus, the sensor model can be parameterized as a mean and variance for each possible target location in the predefined workspace.

The sensor model used to construct the beliefs discussed in Sections 3.2 (histogram) and 3.3 (particles) used Equation (2.13) on page 48 (This is Rasnow's equation for the voltage perturbation in a uniform electric field induced by a sphere) to estimate the expected value of $p(w|\mathbf{x}, \mathbf{e})$ for each \mathbf{x} . These expected values are precisely the EEVs depicted in Figures 2.7 (dipole model on page 53) and 2.10 (ellipse model on page 58).

7.1.2. Data Collection for the Sensor Model

For the Electrosenster, rather than derive the mean and standard deviation of $p(w|\mathbf{x}, \mathbf{e})$ from an analytical model, we measure these two constituents of the sensor model by a calibration step. For a given target at a known location, we collect sensor samples at a grid of points near the target. By taking a number of readings with the target located at each grid point, the sensor model can be parameterized by two grids of values: a grid of average sensor readings and a grid of sensor reading variance. Letting $\mathbf{x}_s = (x_s, y_s)$ be the known target location in the sensor frame, we define the mean and variance of ten measurements at this location as $\text{ave}(\mathbf{x}_s)$ and $\text{var}(\mathbf{x}_s)$. Both $\text{ave}(\cdot)$ and $\text{var}(\cdot)$ can be viewed as mappings from \mathbb{R}^2 to \mathbb{R} . Off-grid values of these functions are obtained by

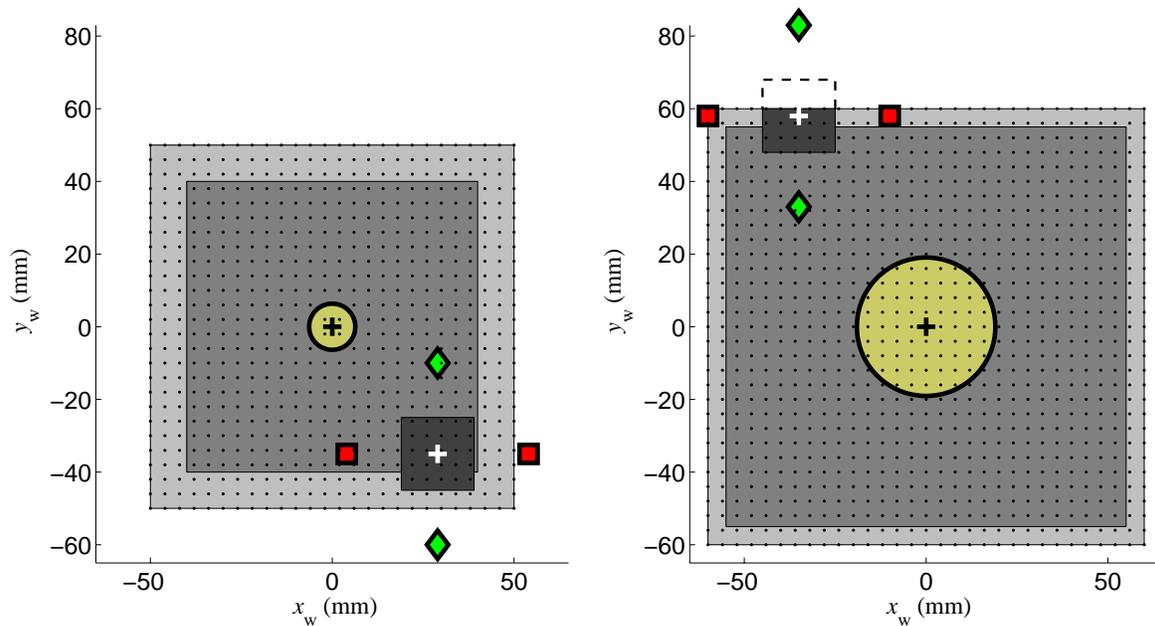
interpolation. Qualitatively, experimental data was consistent with our analytical model (e.g., compare the theoretical EEV in Figure 2.7 with the measured EEV in Figure 7.2).

The Electrosenster was tested with different target geometries—the most common being cylinders and spheres—and sizes ranging from 1/4" diameter to several inches wide. We also tested targets of different materials. Plastics and glass targets were classified as electrical insulators and metal targets were classified as electrical conductors. In general, smaller targets perturbed the electric field less than larger targets. Thus, the workspace was adjusted depending on the target. The workspace for a large sphere (1 1/2" diameter) and small sphere (1/2" diameter) are shown in Figure 7.1. This figure also depicts the local (i.e., permissible control choices at the current time step) and global control limits for this example.

7.2. Empirical EEVs

As previously mentioned, various sets of experiments were performed using many different types of targets in different water conductivities. In the interest of brevity, only a few of these EEVs will be examined in this section. More EEVs are shown in Appendix D.

Here, we examine the EEVs for eight experimental conditions corresponding to all permutations of two target diameters (12.7 and 38.1 mm), two target conductivities (conductor and insulator), and two water conductivities (fresh and salt). For each target/water combination the target was placed at the center of the tank. For the 38.1 mm diameter spheres the robot moved the sensor carriage on a 120 mm by 120 mm grid, stopping every 4 mm for a total of $31 \times 31 = 961$ positions (dots in Figure 7.1(B)). For the 12.7 mm



(A) workspace for 12.7 mm ($1/2''$) diameter sphere (B) workspace for 38.1 mm ($1\ 1/2''$) diameter sphere

Figure 7.1. A 2-D projection of the robot workspace for two different size spheres. The white cross is the center of the robot's carriage; this is defined as the position of the robot. The black cross is the center of the 2-D projection of the spherical target; in the plane of the sensor and emitter electrodes, this is defined as the position of the target. The green diamonds and red squares are the positions of the sensing electrodes and field emitter electrodes, respectively for the given robot position. The darkest gray area represents the area of possible locations the robot can travel to at the next step. The next darker area contains all the possible starting positions for the robot. The largest square represents all possible locations the robot is allowed to travel after the starting position. The black dots are all the locations the robot visited while constructing the associated EEV. (A) Workspace when locating the 12.7 mm sphere. The robot is at position $(x_w = 29, y_w = -35)$. (B) Workspace when locating the 38.1 mm sphere. The robot is at position $(x_w = -35, y_w = 58)$.

diameter spheres the robot moved the sensor carriage on a 100 mm by 100 mm grid, stopping every 4 mm for a total of $26 \times 26 = 676$ positions (dots in Figure 7.1(A)). The $\text{ave}(\cdot)$ and $\text{var}(\cdot)$ functions were obtained by taking the average and variance of 10 time-averaged sensor readings at each grid point.

For each EEV in Figure 7.2, solid red contours indicate the detection range of the sensor: if a target is placed inside the contours, the expected observation will correctly detect its presence with at least a 95% likelihood. This is determined by collecting a number of measurements \mathcal{W}_{nt} when no target is present and calculating the mean $w_{\text{nt}} = E[\mathcal{W}_{\text{nt}}]$ and standard deviation $\sigma_{\text{nt}} = \sqrt{\text{variance}[\mathcal{W}_{\text{nt}}]}$. The no-target observation, w_{nt} , is defined as the expected observation when nothing perturbs the electric field (i.e., when no target is nearby), and the no-target noise, σ_{nt} , is defined as the standard deviation of the observations when no target is nearby. The detection range is defined here by the EEV voltage contours at $w_{\text{nt}} \pm 2\sigma_{\text{nt}}$. For the empirical sensor models σ_{nt} is between 40 and 100 mV for the 6 EEVs where $G < 200$ in Figure 7.2, and around 600 mV for the two EEVs where $G = 1977$ (See Figure 6.1(B) on page 153 for a simplified schematic of the Electroster's circuitry). To determine the noise at the sensor (before being amplified), divide σ_{nt} by differential amplifier gain, G . The noise was between 0.2 mV and 1 mV for all experimental conditions.

Each of the two EEVs in each pair of constant gain (e.g., the pair in Figures 7.2(A) and 7.2(B)) in Figure 7.2 only differ by the material of the target. If we use the simple model from Equation (2.13), the two EEVs in each pair only differ by the contrast factor, χ , which takes the value of +1 for an ideal conducting sphere and $-1/2$ for an ideal insulator. Let us look at Figures 7.2(A) and 7.2(B). Because metal has a high electrical conductivity,

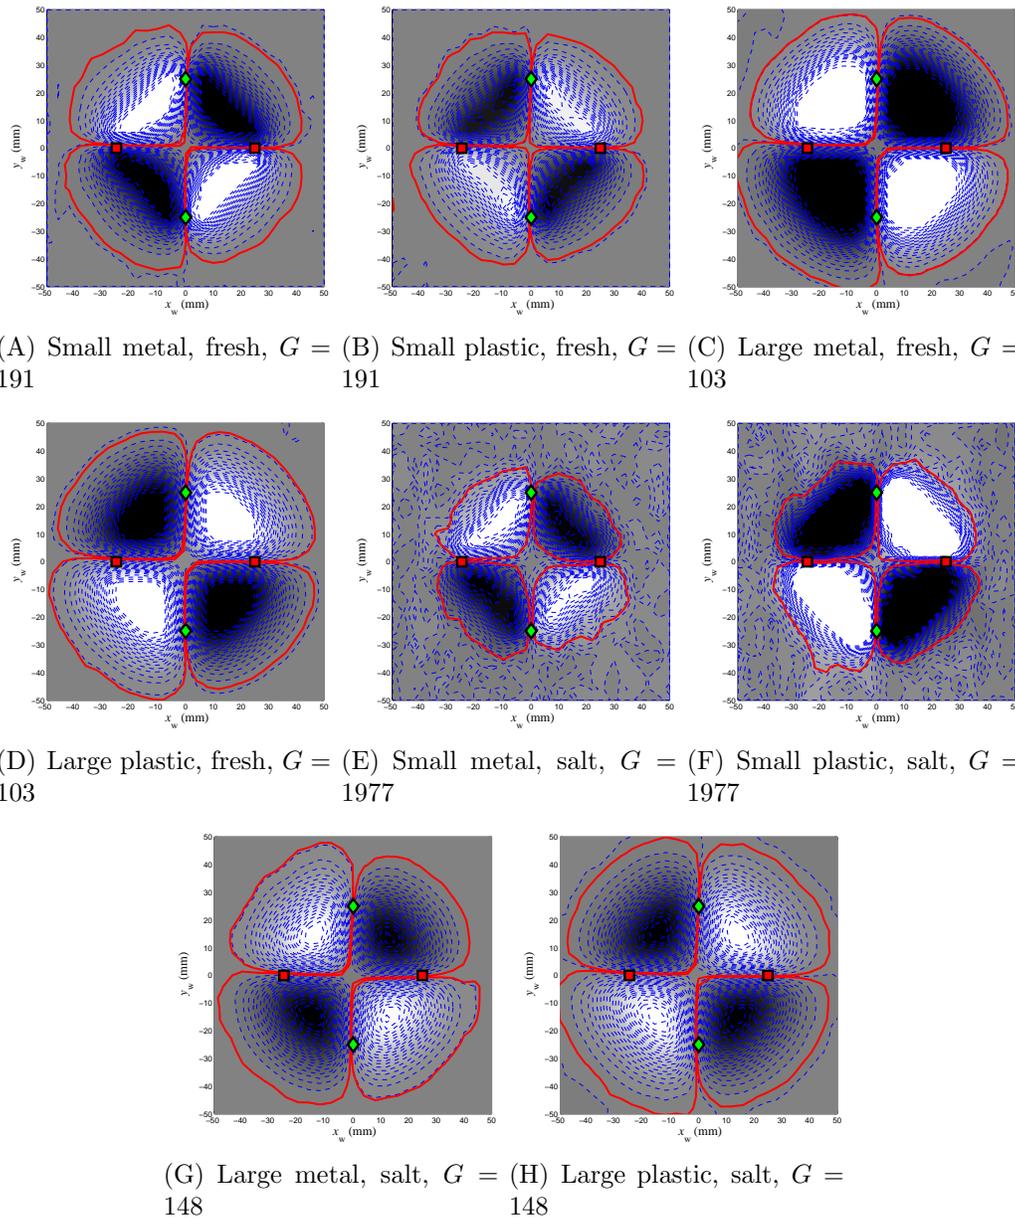


Figure 7.2. EEVs for eight different experimental conditions. Caption legend: Small is a 12.7 mm diameter sphere, Large is a 38.1 mm diameter sphere, fresh is $160 \mu\text{S}/\text{cm}$ water, salt is $56,000 \mu\text{S}/\text{cm}$ water, and G is the differential amplifier gain. Metal is type 440-C stainless steel and plastic is Delrin plastic. Red squares indicate the location of the electric field dipoles. Green diamonds indicate the location of the voltage sensors. The color map goes from white for $+10 \text{ V}$ to gray for intermediate values to black for -10 V . Blue dashed lines are at 500 mV increments. The solid red contours represent an observation that results in detection of the target with 95% confidence.

in fresh (low conductivity) water the target has high electrical contrast magnitude, $|\chi|$. By comparison, Delrin plastic has low electrical conductivity magnitude; thus in fresh water this target has low contrast. Consequently, the magnitudes of the sensor readings are greater for the same relative robot positions, which is evident by the larger saturated lobes in Figure 7.2(A) (our data acquisition card saturates at ± 10 V). But, as will be seen in the analysis from Section 7.3, a relatively large difference in $|\chi|$ translates to a small change in target position that would correspond to the same sensor observation. This results in qualitatively similar detection contours for the metal and plastic spheres, as seen in Figure 7.2.

Disregarding the effects of noise, a high gain or large contrast factor favors longer range target detection, but these same factors can lead to adverse effects. Most of the EEVs in Figure 7.2 exhibit significant regions that correspond to $+10$ V or -10 V due to saturation of the data acquisition hardware. Referring back to the EEVs in Figures 7.2(A) and 7.2(B) (same gain, but different $|\chi|$), a sensor reading of $+10$ V would localize the target to within the two relatively large white lobes in Figure 7.2(A), but the same sensor reading for the EEV in Figure 7.2(B) would localize the target to a much smaller region.

7.3. Detection Range Analysis for a Single Sensor

In the previous section empirical measurements were used to approximate the detection distance of the Electrosenster. In this section we revert back to theoretical models to estimate the detection distance of a target. The framework presented here can be generalized beyond the Electrosenster to predict detection distances for any electrosensory system that fit the models discussed here. We will first consider the simplest emitter/detector

setup: a single sensor in an electric field created by an electric dipole. This layout is a departure from the Electrosenster, but this is one of the few systems that are simple enough to write out closed-form solutions. In the next section a similar analysis will be attempted on the Electrosenster setup but due to the geometry of the emitter/detector layout, closed-form solutions become too lengthy to be displayed on a reasonable number of pages. Also in the next section, the theoretical detection distances will be compared to the measured detection distances from the previous section.

Once the specific sensor model is known, it is possible to determine the distance at which a target can be reliably detected [213]. For this analysis we first consider an emitter/detector layout which consists of a single voltage detector located at the origin and the two emitters at locations $(x = \pm L/2, y = 0, z = 0)$. We wish to determine the dimensionless detection distance, r_d , of a sphere of radius A along a 45° line passing through the origin. Figure 7.3 shows the system examined here. Note that these contours of the EEV are equally spaced on a logarithmic scale.

Our detection distance model begins with the closed-form solution of the dimensionless voltage perturbation $\delta v = \delta V/V_e$ (derived from Equation (2.13)) where δV is the perturbation and V_e is half of the difference between the emitter electrodes (e.g., one emitter is at $+V_e$ and the other is at $-V_e$). The dimensionless voltage perturbation is given by:

$$(7.1) \quad \delta v(a, \chi, e, r) = -2 a^3 \chi e \frac{N_1(r) + N_2(r)}{D(r)}$$

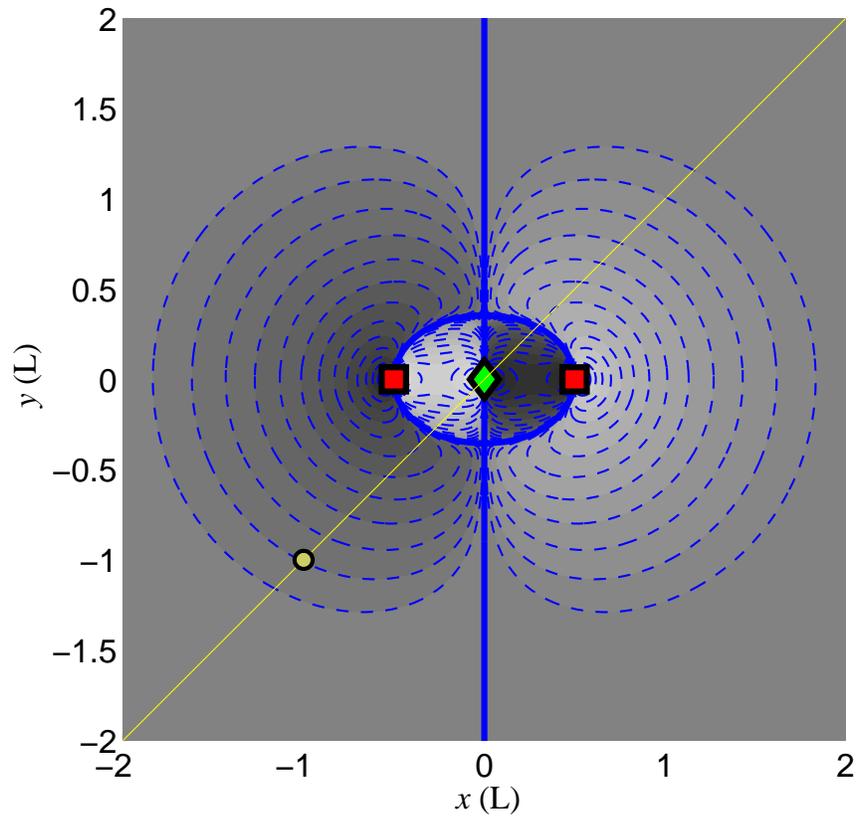


Figure 7.3. Sensor observation as a function of target position. This is precisely the EEV for the single-sensor setup constructed using Equation (2.13) (Rasnow's model from page 48). The emitters are red squares and the detector is a green diamond. Only target positions along the yellow line will be considered (e.g., the one depicted at coordinates $(-1, -1)$). Contours represent sensor observation for the target position. The color map uses a logarithmic scale. The solid blue contour represents an observation of zero volts. Lengths are normalized to L , which is the distance between the emitters.

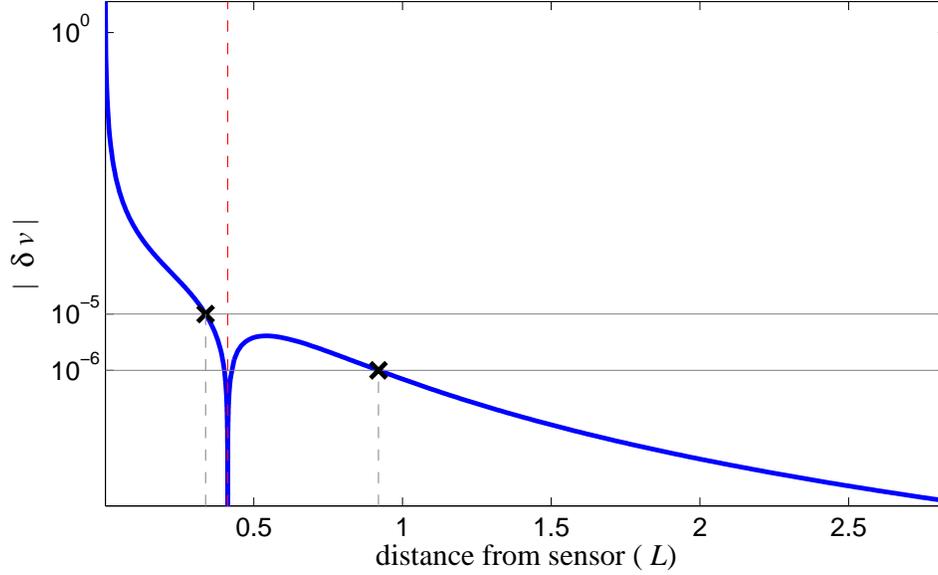


Figure 7.4. The absolute value of the dimensionless sensor signal as a function of the dimensionless distance from the sensor along the slice depicted in Figure 7.3. The red dashed vertical line represents the singularity when the sensor observes zero volts. Horizontal lines are at constant sensor observations, and the ‘x’ on each line represents the maximum distance for the observation. This maximum distance is used in the definition of the detection distance in Equation (7.2). δv is computed using Equation (7.1) with $a = 0.05$, $e = 0.005$, and $\chi = +1$.

where

$$\begin{aligned}
 N_1(r) &= \left(\sqrt{2} - 4\sqrt{2}r^2 + 16r^3 \right) \sqrt{1 - 2\sqrt{2}r + 4r^2} \\
 N_2(r) &= \left(\sqrt{2} - 4\sqrt{2}r^2 - 16r^3 \right) \sqrt{1 + 2\sqrt{2}r + 4r^2} \\
 D(r) &= r^2 \left(1 - 2\sqrt{2}r + 4r^2 \right)^{\frac{3}{2}} \left(1 + 2\sqrt{2}r + 4r^2 \right)^{\frac{3}{2}}
 \end{aligned}$$

where a is the dimensionless target radius given by A/L where A is the radius of the target and L is the distance between the two emitters, χ is the contrast factor, e is the dimensionless emitter radius given by E/L where E is the emitter radius, and r is the

dimensionless distance to the center of target from the detector given by R/L , where R is the distance to the center of the target. Equation (7.1) is derived in Appendix E.

The no-target observation, w_{nt} , is defined as the expected observation when nothing perturbs the electric field (i.e., when no target is nearby). In this particular example $w_{\text{nt}} = 0$. The no-target noise, σ_{nt} , is defined as the standard deviation of the observations when no target is nearby. For a two-sigma ($\sim 95\%$) confidence interval, when there is no target nearby, there is only a $\sim 5\%$ chance that any given observation will fall outside the interval $[w_{\text{nt}} - 2\sigma_{\text{nt}}, w_{\text{nt}} + 2\sigma_{\text{nt}}]$. Thus, if an observation does fall outside the interval of $[w_{\text{nt}} - 2\sigma_{\text{nt}}, w_{\text{nt}} + 2\sigma_{\text{nt}}]$, then it can be concluded with at least 95% confidence that a target has perturbed the electric field. We define the dimensionless expected detection distance, $r_{\text{d}} = R_{\text{d}}/L$ (hereafter simply detection distance), as the distance when the expected observation is equal to either $w_{\text{nt}} + 2\sigma_{\text{nt}}$ or $w_{\text{nt}} - 2\sigma_{\text{nt}}$. If we define the dimensionless no-target noise as $n = \sigma_{\text{nt}}/V_e$, then the detection distance is

$$(7.2) \quad r_{\text{d}}(a, \chi, e, n) = \max\{r \mid \delta v(a, \chi, e, r) = \pm 2n\}.$$

The max function is used to select from multiple r 's when there is more than one solution r to $\delta v(a, \chi, e, r) = \pm 2n$. For example, Figure 7.4 shows the sensor observation along the yellow line in Figure 7.3. For an observation of $|\delta v| = 10^{-5}$ (in Figure 7.4 this is signified by the upper horizontal line) there is only one possible r (represented as the 'x'). But, for $|\delta v| = 10^{-6}$ (lower horizontal line in Figure 7.4) there are three possible values of r . In this case the max of these is chosen (shown as the 'x').

Fixing a at 0.05 and e at 0.005, Figure 7.5 shows the relationship between the dimensionless no-target noise, $n = \sigma_{\text{nt}}/V_e$, and r_{d} . The black points were computed using

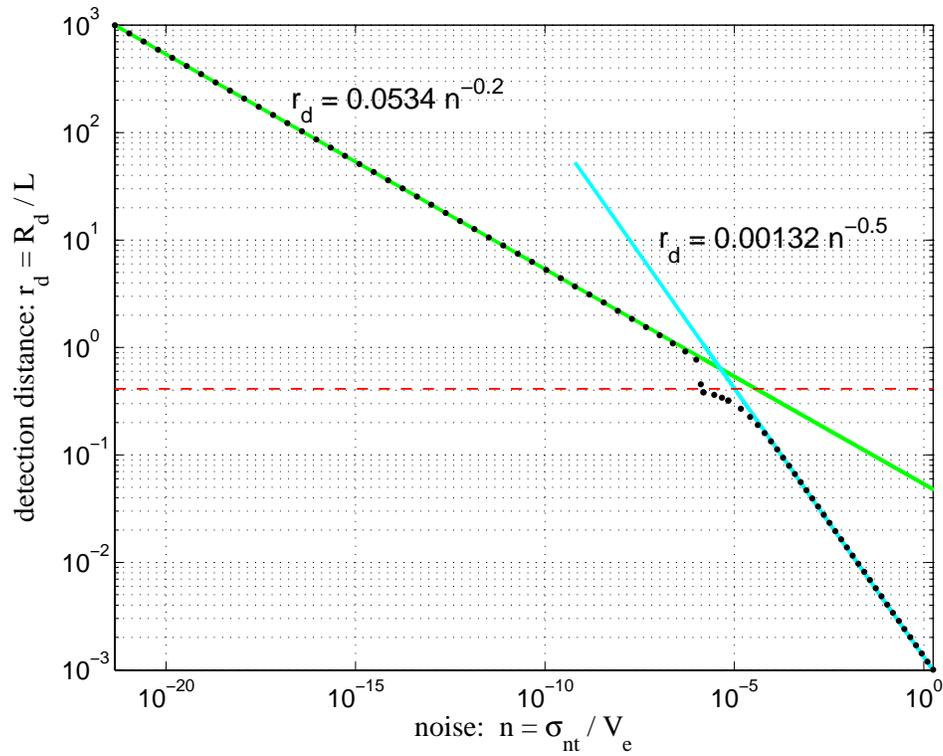


Figure 7.5. Dimensionless detection distance ($r_d = R_d/L$) as a function of dimensionless noise ($n = \sigma_{nt}/V_e$). The normalized radius of the target is $a = A/L = 0.05$, and the radius of the emitting electrode is $e = E/L = 0.005$. The red dashed line is the singularity detection distance where the target can never be detected. Data points very near this singularity are not shown. Away from the singularity detection distance can be well approximated by one of the two curves shown as green and blue solid lines, which take the form of $r_d = k n^c$.

the method just described using Equation (7.1). When these points are plotted on a log-log scale, most of them fall very near one of two lines, where each line takes the form $r_d = k n^c$.

Figure 7.5 reveals two distinct regimes for dimensionless detection distance with normalized noise. The *near field* regime is locations inside the $w = 0$ oval (thick blue contour) in Figure 7.3. In this regime detection distance falls off as $1/\sqrt{n}$. The *far field* regime is locations outside the $w = 0$ oval. Here, detection distance falls off as $1/\sqrt[5]{n}$. Thus, in the far field with a 32-fold reduction of noise extends the detection distance by a factor of two. Both of the fitted curves shown in Figure 7.5 fail as the detection distance approaches the singularity point (shown as a red dashed line). This singularity happens when the observation is exactly zero. This is when the yellow line crosses the solid blue line in Figure 7.3. The singularity is also represented by the red dashed line in Figure 7.4.

The choice of a , e , and χ only affect the constant of proportionality, k , and not the exponent, c , of the fitted power-law equations, $r_d = k n^c$, used to approximate the near and far field relationships. In fact, a , e , and χ do not influence the location of the transition from the near and far fields (this transition is signified by the solid blue oval contour in Figure 7.3 and the red dashed line in Figure 7.5). Thus, the r_d corresponding to this transition is fixed at approximately 41% of L for any choice of a , e , and χ . As shown by Equation (7.1), the choice of a , e , and χ does, however, shift the curves in Figure 7.5 to the right or left.

| experimental conditions | sensor noise (mV RMS) | detection distance, R (mm) | |
|---------------------------|--------------------------|------------------------------|-------------|
| | | measured | theoretical |
| 12.7 mm, steel in fresh | 0.25 | 32.7 | 32.7 |
| 12.7 mm, plastic in fresh | 0.23 | 31.3 | 29.5 |
| 38.1 mm, steel in fresh | 0.92 | 41.1 | 43.2 |
| 38.1 mm, plastic in fresh | 0.86 | 41.0 | 39.3 |
| 12.7 mm, steel in salt | 0.24 | 25.3 | 32.9 |
| 12.7 mm, plastic in salt | 0.35 | 27.4 | 26.6 |
| 38.1 mm, steel in salt | 0.49 | 39.6 | 46.5 |
| 38.1 mm, plastic in salt | 0.31 | 41.3 | 45.3 |

Table 7.1. Empirically measured and theoretical detection distance.

7.4. Computed vs Measured Detection Distances

A similar analysis as used in Section 7.3 was performed for the emitter and detector layout used in the Electrosenster. The two emitters are again located at positions $(x = \pm L/2, y = 0, z = 0)$, and the detectors are located at positions $(x = 0, y = \pm L/2, z = 0)$. The z coordinate of the center of the target is fixed such that there is approximately 2 mm clearance between the sensing/emitting electrodes at the top of the target when directly above the center of the target. If $L = 50$ mm, then this is precisely the setup for the Electrosenster. The detection distance is the distance from the center of the target to the origin, i.e., $(x = 0, y = 0, z = 0)$. As before, we wish to find the detection distance along a 45° line passing through the origin, but, as with the experimental setup, the center of the target is not coplanar with the emitters and detectors.

Table 7.1 shows the measured and computed detection distance for the eight experimental conditions. In this table the sensor noise is defined as σ_{nt}/G (G is the gain on the differential amplifier and σ_{nt} is determined experimentally). The x and y coordinates of detection are determined by the intersection of the red contours in Figure 7.2 with a 45°

line passing through the origin. Then, the constant offset z -coordinate of the center of the target is taken into account. The equations used to compute the theoretical values in Table 7.1 are the same as those discussed in Appendix E and Section 7.3 except for the location of the sensors and the fact that the observation is the difference of the two sensor voltages. When computing the detection distance, the following parameters were used: $L = 50$ mm, $E = 1$ mm, $A = \{12.7 \text{ or } 38.1\}$ mm, $V_e = 2$ volts, and $\chi = \{+1 \text{ or } -1/2\}$.

The average magnitude of the difference between the measured and theoretical values in Table 7.1 is about 3 mm. These relatively small discrepancies suggest that the theoretical analysis presented in Section 7.3 and Appendix E is able to predict performance in real electrosensory systems.

The average difference between the measured and theoretical values is about -2 mm, which means the theory is more likely to overestimate the measured detection distance. The model appears to predict detection distance much better with the plastic spheres (where the expected error is ~ 0 mm and error magnitude is ~ 2 mm) and worse with the metal spheres (expected error is ~ -4 mm and error magnitude is ~ 4 mm). Perhaps oxidation, or other electrically impeding films, on the surface of the metal sphere rendered the sphere far from an ideal conductor, thus significantly reducing $|\chi|$, and decreasing the actual detection distance.

CHAPTER 8

Probabilistic Target Electrolocation

In this chapter the Electrosenster performs the task of electrolocating a single target using a predetermined EEV and a probabilistic belief maintenance scheme. The performance of four different types of greedy active-sensing controllers is investigated and compared to two different open-loop controllers. Three different sets of experiments are examined, which cover five different target geometries, two target conductivities, three water conductivities, and the six different controllers.

In this Chapter the belief of the target is maintained via a particle filter. With a particle filter the belief of the target location is represented by a large number of hypotheses (or “particles”) in the target configuration space, allowing approximate representations of complex multi-modal beliefs [6]. The particle filter uses Monte Carlo simulation [161] to approximate a Bayes filter, and the particle representation is updated as new information comes in [228]. Refer to Section 3.3 (beginning on page 83) for a more detailed discussion on particle filters and their application to electrolocation.

8.1. Probabilistic-Based Control Algorithms

The complete active sensing problem was formally introduced in Section 4.2.2. Essentially, the goal of an active sensing controller is to move its sensors to gather the “best” information, where “best” depends on the specific task. Given an infinite amount of computing power, the ideal active sensing control scheme for the Electrosenster would

be defined by Equation (4.5) with the cost functional resembling Equation (4.9), where, for example, the POMDP formulation would be solved via dynamic programming. The control choice at each step would be conditioned on all current information and optimized over all possible sequences of future actions. In reality, real-time systems don't have the luxury of waiting for optimal trajectories to be computed. Instead, heuristic controllers based on the expected information at the next time step (i.e., greedy controllers) are used in the “active” controllers for the Electrosenster. In the next subsections, four different greedy controllers are presented.

The four greedy controllers presented below attempt to maximize the expected information gain on the next step. Similar techniques have been successfully used in [42, 94, 165], where controllers minimize some scalar function of the covariance matrix (in the case of a Gaussian belief) or the entropy (in the case of nonparametric belief) [13].

8.1.1. Minimize Expected Variance of Particles, min-var

Recall Equation (3.5) (on page 88) defined an uncertainty metric, `parts-std`, on particle belief representations that is the square root of the trace of the covariance matrix of the particles. The `min-var` controller chooses the control option that minimizes the expected `parts-std` at the next time step.

A mathematical formulation of `min-var`'s objective will now be presented. Let $w \in \mathcal{W}$ and $u \in \mathcal{U}$, where \mathcal{W} is the space of all possible observations and \mathcal{U} is the space of possible controls. The belief of the system state (i.e., the set of particles), b , is an element of the belief space, \mathcal{B} . Let `Update` : $\mathcal{B} \times \mathcal{W} \times \mathcal{U} \rightarrow \mathcal{B}$ denote a function that produces an updated

belief (posterior) from the prior belief, the observation, and the control option. `Update` uses a single instance of the update step in a Bayes filter. The expected uncertainty at the next time step after choosing control option u_i is found by

$$(8.1) \quad E[\text{Uncert}(u_i, b)] = \int_{\mathcal{W}} \text{Uncert}[\text{Update}(b, u_i, w)] p(w|u_i, b) dw.$$

This equation describes how the expected uncertainty varies with the choice of control input. If `Uncert(b)` is used to compute `parts-std(b)`, then objective of the `min-var` controller is to choose the best control option, u^* such that

$$(8.2) \quad u^* = \underset{u_i}{\text{argmin}} E[\text{Uncert}(b)|u_i].$$

Finding the control option, u^* through solving Equations (8.1) and (8.2) requires integrating over all possible observations. Instead, the `min-var` controller takes a shortcut by employing *particle voting* (see Algorithm 1 and Section 8.1.1.1) to choose the control option [213], where each particle is responsible for simulating a possible next observation and casts a vote based on this observation. For the `min-var` controller replace `Uncert` in Algorithm 1 with `parts-std`.

8.1.1.1. Particle Voting for Approximating the Uncertainty of the Next Belief. The particle voting algorithm is formally stated in Algorithm 1, which will now be explained in words. In particle voting each particle's vote for the best control option is determined by evaluating the `Uncert` (`parts-std` in the case of the `min-var` controller) of the particle cloud after executing each of the possible control actions in simulation. Particle voting looks to each particle in the belief representation and chooses the best

control action under the assumption that the target is actually at that particle's location. Each particle can then be thought of as voting for a particular control action that works best for that particle. Once all of the votes are tallied, the control action with the most votes is chosen. Each particle's vote for the best control option is determined by evaluating the expected belief after executing the control action (in simulation) and taking an observation (in simulation). This observation is assumed to be the expected observation for this particular particle and control option combination. In the case of the `min-var` controller, after each simulated control action and observation, the current particle cloud is resampled based on this simulated observation. The `Uncert` metric (`parts-std` in the case of the `min-var` controller) of this new (simulated) particle cloud serves as the cost for this control action. The control action that yields the lowest expected cost gets the particles vote. Three of the Electrosenster's four active controllers use particle voting.

Algorithm 1 Particle voting for approximating the uncertainty of the next belief

Given the set of particles $\mathcal{X} = \{\mathbf{x}_m | m = 1 \dots M\}$.

for each particle, $m = 1 \dots M$, **do**

\mathbf{x}_m is a target location hypothesis (particle).

for each control option, \mathbf{u}_c , $c = 1 \dots C$, **do**

(1) Simulate movement of particles in sensor frame: $\mathcal{X}' = \{\mathbf{x}'_i | \mathbf{x}'_i = \mathbf{x}_i + \mathbf{u}_c, i = 1 \dots M\}$.

(2) Simulate measurement: $w_{m,c} = E[w | \mathbf{x}'_m, \mathbf{e}]$, where $E[\cdot]$ denotes expected value.

(3) Generate corresponding weights, λ , based on $w_{m,c}$ for each particle in \mathcal{X}' : $\mathcal{L}' = \{\lambda'_i | i = 1 \dots M\}$.

(4) Generate \mathcal{X}'' via importance resampling based on \mathcal{X}' and \mathcal{L}' .

(5) Compute `Uncert` for particle set \mathcal{X}'' .

end for

Particle m votes for the control option that resulted in the lowest predicted `Uncert`.

end for

The control option with the most votes is chosen.

8.1.2. Minimize Distance to Centroid of Particles, min-mean

The `min-mean` controller chooses the control option that moves the center of the robot closest to the centroid of the belief. This controller is motivated by the gross generalization that moving as close to the target as possible is the best thing to do (We will see that this generalization often fails). The centroid of the particles (geometric mean of the particles) is the estimate of the target location. This controller bears the least computational load since it does not attempt to estimate future expected beliefs. It is the only active controller that does not use particle voting.

8.1.3. Maximize Expected Change in Sensor Reading, max-grad

The `max-grad` controller chooses the control option that maximizes the magnitude the difference between the current sensor reading and the expected next sensor reading. This controller is motivated by another gross generalization that it is best to position the target in regions of high gradient in the EEV. For example, look at Figures 2.10(B), 2.10(C), 3.22(B), and 3.22(C), which have been preprinted as Figures 8.1(A), 8.1(B), 8.1(C), and 8.1(D), respectively, for convenience. The target is in a much steeper gradient in Sensor 2 than with Sensor 3; i.e., $dE[w_2]/dx > dE[w_3]/dx$. As a result Sensor 2 yields a belief with lower uncertainty than the belief from Sensor 3.

Let, $\text{EEV-grad}(x, u_i) = |\text{EEV}_e(x) - \text{EEV}_e(x')|$ be the absolute value of the difference between the expected observation when the target is at x and the expected observation when the target is at x' , where x' is the target location after invoking action u_i . The `max-grad` algorithm uses particle voting, but because `EEV-grad` is only a function of x

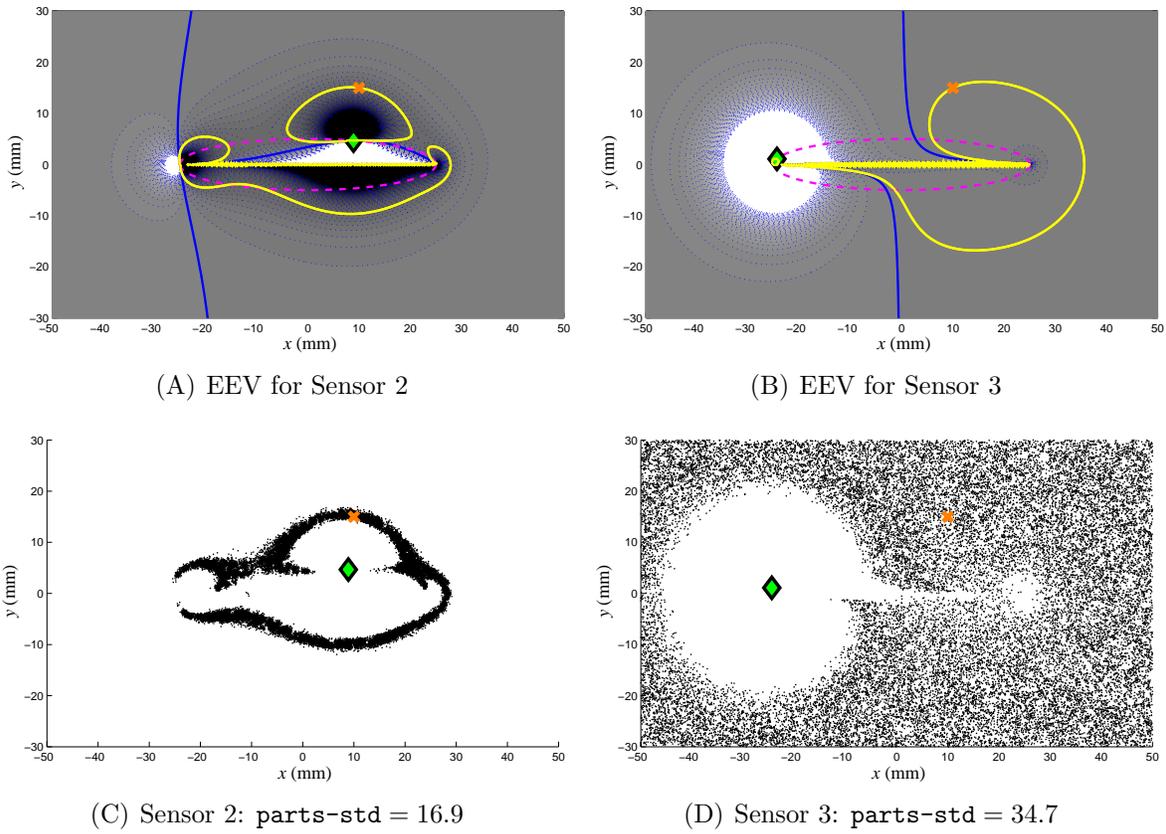


Figure 8.1. Motivation for max-grad controller. Subfigures (A) and (B) are the EEVs for Sensors 2 and 3 in the ellipse model, and subfigures (C) and (D) are the beliefs after a single observation. In (A) the target is located in steep gradient on the EEV, and the resulting belief—shown in (C)—has relatively low uncertainty. In (B) the target is in a very flat region of the EEV, and the resulting belief—shown in (D)—has high uncertainty.

and u , the resulting belief does not have to be computed for each combination of particle and control option, thus, it has less computational load than `min-var`.

Algorithm 2 Particle voting for `max-grad` controller

Given the set of particles $\mathcal{X} = \{\mathbf{x}_m | m = 1 \dots M\}$.
for each particle, $m = 1 \dots M$, **do**
 \mathbf{x}_m is a target location hypothesis (particle).
for each control option, \mathbf{u}_c , $c = 1 \dots C$, **do**
(1) Simulate movement of particles in sensor frame: $\mathbf{x}'_m = \mathbf{x}_m + \mathbf{u}_c$.
(2) Compute `EEV-grad` based on \mathbf{x}'_m , \mathbf{x}_m , and \mathbf{u}_c
end for
Particle m votes for the control option that resulted in the highest `EEV-grad`.
end for
The control option with the most votes is chosen.

Algorithm 2 defines the particle voting algorithm for the `max-grad` controller. Algorithm 2 is similar to Algorithm 1 except steps (2) through (5) in the inner loop in Algorithm 1 have been replaced by step (2) in Algorithm 2, and step (1) in Algorithm 2 invokes the motion model for only the single particle instead of the entire particle set as in Algorithm 1.

8.1.4. Minimize Expected Number of Particles, `parts-eater`

The `parts-eater` controller chooses the control option that minimizes the expected number of unique particles at the next time step. Each particle eliminated is a possible state that is no longer considered a candidate target location.

To gain a better understanding on the motivation behind this controller, we need to revisit the update step of the particle filter. The resampling process stochastically chooses the most likely particles based on the observation (see Section 3.3.1.2 on page 86 for details). This process does not generate any unique particles. More likely particles

are duplicated, and less likely particles are eliminated. The intuition is that the control action that results in the fewest number of unique particles is the best option.

Let $\text{Unique}(b) = \text{count}(\bigcup b)$, where b is the set of particles, and the function *count* returns the number of particles in the set. The union operator, \bigcup , eliminates duplicates. For example, if control option u_1 results in the next particle set (after simulated resampling) to be M (M is the total number of particles) copies of the same particle, then there is no uncertainty and the target has been perfectly localized— u_1 is the best possible control option. At the other extreme, if control option u_2 results in the identical posterior particle set as the prior, then nothing is gained by this control action— u_2 is the worst possible control option.

Since the `parts-eater` algorithm must predict a belief (i.e., set of particles) based on a control action, particle voting is used. Algorithm 1 describes the `parts-eater` algorithm if `Unique` is replaced by `Uncert`.

8.1.5. Open-Loop Controllers

The purpose of the feedback controllers presented above is to minimize uncertainty as quickly as possible. The advantage of feedback over open-loop controllers is that feedback controllers can adapt to unpredictable environments. But, if the feedback controllers do no better than a pre-computed open loop control plan, then there is no advantage to the feedback controller, and the CPU-friendly open-loop controller should be used. We present two open-loop controllers that will be compared to the feedback controllers.

The open-loop controllers still use the same belief update mechanism (i.e., a particle filter). However, their control decisions are not based on meeting some objective, but instead from a pre-defined plan.

8.1.5.1. Random Walk, `rnd-wlk`. The `rnd-wlk` controller randomly chooses one of the possible control actions. This controller serves as the baseline to compare the performance of the other controllers.

8.1.5.2. Move to Way Points Sampled from a Spiral, `spiral`. The `spiral` controller moves towards way points sampled from a spiral. The spiral is centered at the initial location. A way point sampled from the spiral is assigned to each time step. The controller chooses the control option that moves it closest to the current way point. The `spiral` controller represents a systematic method of exploring the workspace. Figure 8.2 shows the way points.

8.1.6. Completion of Electrolocation Sequence

For each trial the controller terminates when either a maximum run time is exceeded (failure) or the target is considered sufficiently localized (success). We use the `parts-std` uncertainty metric defined in Section 3.3.2 on page 87 to determine if the target is sufficiently localized. We set the “success” termination condition as $\text{parts-std}(b) = \sqrt{\text{trace}(P_b)} \leq \beta$ for an appropriately chosen β , where P_b is the covariance matrix of the particle set b .

8.1.7. Preview of Next Three Sections

In the next three sections, three experimental sets are examined. In all sets the task is the same: electrolocation of a single target with a predetermined sensor model. The

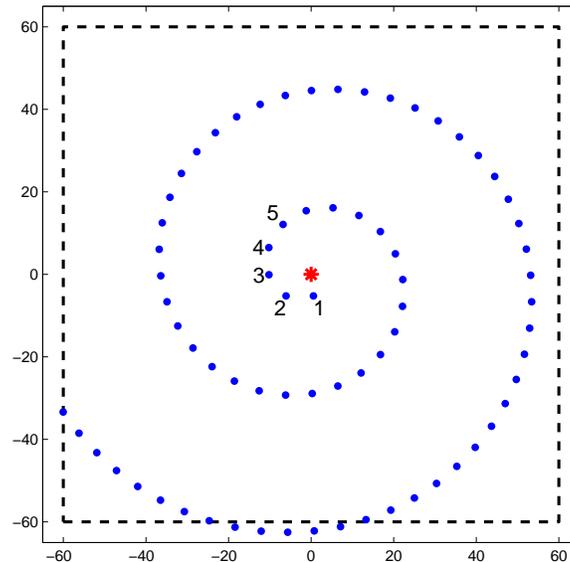


Figure 8.2. Way points for `spiral` controller. The initial location of the robot is signified by a red asterisk. At the first time step the controller chooses the control action that brings it closest to the blue dot labeled “1”. The process is repeated for the remainder of the time steps. A typical workspace of the robot is usually no bigger than a square of sides 120 mm, which is shown as black dashed lines.

experimental conditions (e.g., target type, water conductivity, and workspace size) and which controllers are considered are varied across the three experimental sets. The targets are cylinders and spheres made of either metal (stainless steel or aluminum) or plastic (PVC or Delrin) in either fresh or salt water.

8.2. Experimental Set 1: Cylinders

8.2.1. Methods for Experimental Set 1

In this set of experiments the robot performs electrolocation of cylinders for three experimental conditions:

- (1) 12.7 mm (1/2”) diameter aluminum (electrical conductor) cylinder in fresh water.

- (2) 25.4 mm (1”) diameter aluminum (electrical conductor) cylinder in fresh water.
- (3) 25.4 mm (1”) diameter plastic (electrical insulator) cylinder in salt water.

The corresponding EEVs are depicted in Figure D.1 on page 267. The 80 mm heights of the cylinders were chosen such that the bottom of the sensors were just able to pass over them without touching.

We test four different controllers in these conditions: `rnd-wlk`, `min-mean`, `min-var`, and `max-grad`. For these experiments we choose a discrete-time formulation with a discrete control set consisting of 17 different control options. One of these is to remain stationary, and the other 16 are 5 mm moves at 22.5° intervals on the unit circle.

8.2.1.1. Sensor Model Construction. Empirical sensor models were constructed for each of the three experimental conditions in Table 8.1. The protocol is almost identical to that used for sensor characterization (See Section 7.1.2, page 165). For each target/water combination the target was placed at the center of the tank. The robot moved the sensor carriage on a 160 mm by 160 mm grid, stopping every 5 mm to take a data point, for a total of $33 \times 33 = 1089$ positions. The `ave()` and `var()` maps were obtained by taking the average and variance of 10 sensor readings at each grid point, where each sensor reading was the average of 100 readings taken at 1 kHz.

8.2.1.2. Electrolocation Protocol. Automatic electrolocation was performed for the three experimental conditions shown in Table 8.1 using each of the four different types of controllers (described in Section 8.1). For each condition, each of the four controllers started from ~ 50 different positions for a total of ~ 200 trials for each condition. Fifty starting positions were chosen randomly and this same set of 50 positions was used for all controllers.

The robot keeps moving until $\text{parts-std}(b) = \sqrt{p_{xx} + p_{yy}} \leq 10$ mm. If this stopping condition is not satisfied within 35 steps, the trial is labeled a failure.

Each complete step of each trial consisted of the following operations:

- (1) **observation:** Acquire sensor data from sensor circuit
- (2) **weighting each particle:** Apply sensor model (i.e., $p(w|x, e)$) to each particle
- (3) **importance resampling of particles:** Create new particle population based on particle weights.
- (4) **check trial termination conditions:** If uncertainty of the belief is low enough or the number of steps is too many, then the trial is over
- (5) **choose next position:** Use one of the four controllers to pick the next position
- (6) **artificially add noise to particles:** This is to add some diversity into the particle population.

A completed trial is labeled a success only if the uncertainty threshold defined above is met in fewer than 35 moves, and failure otherwise.

8.2.2. Description of Performance Statistics for All Data Sets

For a given starting position, if either controller yielded a failure trial, then data from these trials was omitted in calculating the mean error and the mean number of steps statistics. This was done because failure trials often resulted in errors much larger than successful trials. All data (including failure trials) were used in calculating the medians.

Tables 8.1, 8.2, 8.3 show the performance of the electrolocation controllers (for all three data sets) using three statistical measures:

- (1) **success trials:** This is the number of trials that were labeled a success. The number next to it in parentheses is the number of trials that were labeled failure.
- (2) **error:** This is the distance between the actual target location and the estimated target location at the end of the trials. The *mean* error only considers the trials that were a success for both controllers for that starting position, while the *median* error considers all trials.
- (3) **steps:** This is the number of steps needed to complete the trials. The mean calculation only considers successful trials while the median calculation accounts for all trials.

This “error” statistic (the final error between the actual location of the target and the centroid of the particles) was not given heavy importance for two reasons: 1) For general beliefs (e.g., multi-modal) the centroid is a somewhat arbitrary method of state estimation; and 2) The actual value of the error is dependant on the threshold of `parts-std` stopping condition. For example if the threshold is low, then more steps and observations will be taken, which will create a more accurate estimate.

8.2.3. Results from Experimental Set 1

The target was always placed at the center (coordinates (0,0)) of the workspace. The robot could move anywhere within the 160 mm by 160 mm (± 80 in both x and y coordinates), but the starting positions were restricted to ± 50 mm for the 12.7 mm aluminum cylinder in freshwater and the 25.4 mm PVC cylinder in salt water, and ± 60 mm for the 25.4 mm aluminum cylinder in freshwater.

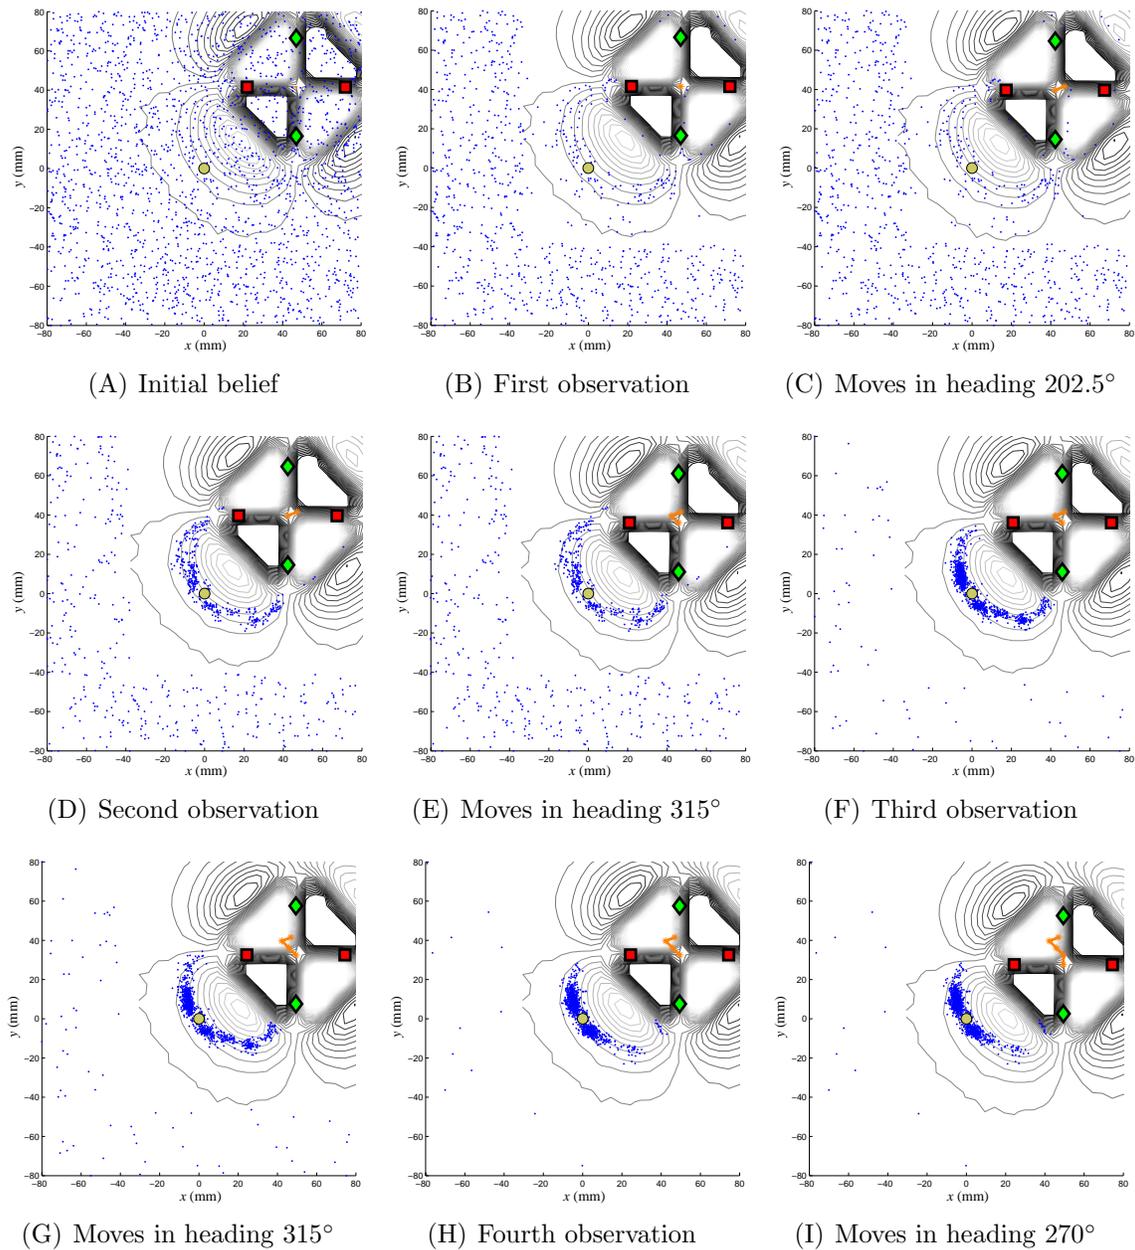


Figure 8.3. Frames from a typical electrolocation sequence. The particles that make up the belief of the target position are blue dots. The actual target location is a yellow circle, in this case the 25.4 mm Al cylinder. The position of the electrodes on the robot are the green diamonds (sensors) and red asterisks (electric field emitter). The orange asterisks represent the path of the robot. The contours are iso-electric-field-perturbation gradations at 650 mV steps. The target is at coordinate $x = 0$, $y = 0$, and the axes range from ± 80 mm in both dimensions.

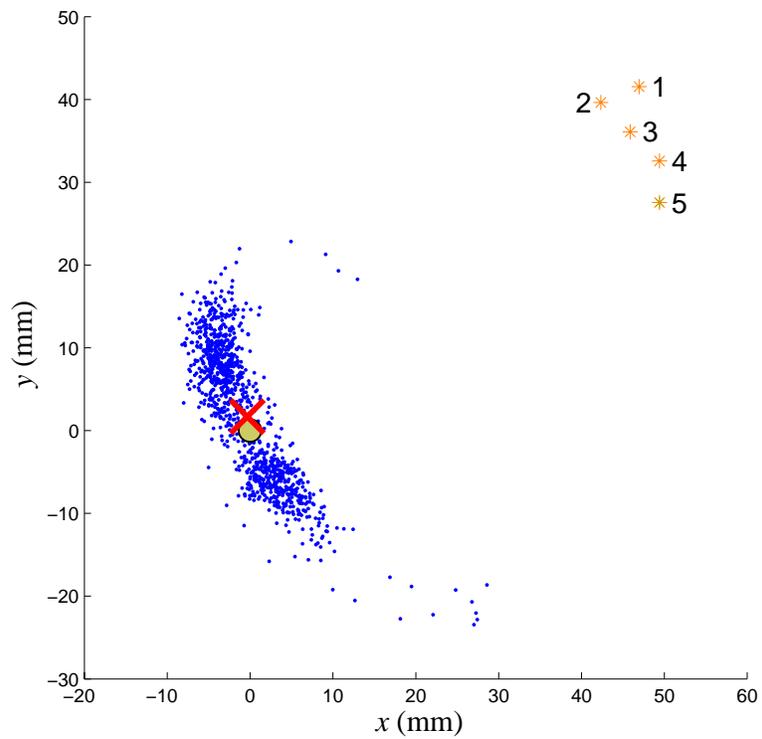


Figure 8.4. Final frame of the same trial. The red “x” located at coordinate $x = -0.34$ and $y = 1.68$ is the centroid of the particles and is the final estimate of the location of the target (shown in yellow). The numbered orange asterisks are the positions visited by the robot. Note that the axes have changed from the previous figure.

Figure 8.3 shows a typical electrolocation sequence. In Figure 8.3(B) the robot has updated its belief using a single observation and the particle filter algorithm. Since the sensor measurement is near the RMS noise assumed for the sensor, the particle filter kept most of the particles. This particular trial uses the controller that minimizes the expected spatial variance of the particles at the next step, i.e., `min-var`. Based on the particles shown in Figure 8.3(B) the robot moves at a heading of 202.5° . This is repeated until the uncertainty is below the threshold described above. Figure 8.4 shows the final belief (dots). The centroid of these particles is taken as the estimate of the target location (“x”), and has coordinates $x = -0.34$ mm and $y = 1.68$ mm (total error = 1.72 mm). The numbered asterisks are the positions of the robot for each of the five observations.

8.2.3.1. 12.7 mm aluminum in freshwater. The top part of Table 8.1 summarizes the statistics from the 196 trials under these conditions. Of the 49 unique starting positions 39 had successful trials from all four controllers. Thus, the error and steps statistics in Table 8.1 were averaged over those 39 trials for each controller. The starting positions were restricted ± 50 in both x and y coordinates.

8.2.3.2. 25.4 mm aluminum in freshwater. The middle part of Table 8.1 summarizes the statistics from the 228 trials used for these experimental conditions. Of the 57 unique starting positions 32 had successful trials from all four controllers. Since the diameter of the target cylinder for these trials was large, the starting positions were expanded to ± 60 in both x and y coordinates.

8.2.3.3. 25.4 mm PVC in salt water. The bottom part of Table 8.1 summarizes the statistics from the 204 trials used for these experimental conditions. Of the 51 unique starting positions 18 had successful trials from all four controllers. Since the contrast

12.7 mm diameter metal (aluminum) cylinder in freshwater

| | success trials (failure trials) | error, mm mean / median | steps mean / median |
|----------|------------------------------------|----------------------------|------------------------|
| rnd-wlk | 46 (3) | 4.3 / 4.3 | 7.4 / 7 |
| min-var | 46 (3) | 3.8 / 3.8 | 6.8 / 6 |
| min-mean | 45 (4) | 3.7 / 3.3 | 8.1 / 7 |
| max-grad | 47 (2) | 4.3 / 4.3 | 6.6 / 6 |

25.4 mm diameter metal (aluminum) cylinder in freshwater

| | success trials (failure trials) | error, mm mean / median | steps mean / median |
|----------|------------------------------------|----------------------------|------------------------|
| rnd-wlk | 46 (11) | 3.7 / 4.1 | 10.2 / 8 |
| min-var | 50 (7) | 2.6 / 2.6 | 7.2 / 6 |
| min-mean | 46 (11) | 5.1 / 5.1 | 7.3 / 7 |
| max-grad | 51 (6) | 4.9 / 4.3 | 9.6 / 7 |

25.4 mm diameter plastic (PVC) in salt water

| | success trials (failure trials) | error, mm mean / median | steps mean / median |
|----------|------------------------------------|----------------------------|------------------------|
| rnd-wlk | 38 (13) | 5.1 / 5.2 | 9.4 / 12 |
| min-var | 40 (11) | 4.3 / 5.5 | 12.4 / 9 |
| min-mean | 37 (14) | 5.6 / 6.4 | 7.2 / 7 |
| max-grad | 45 (6) | 6.4 / 5.7 | 9.1 / 6 |

Table 8.1. Electrolocation statistics for the three experimental conditions tested.

factor (see Equation (2.13)) of a plastic target is about half that of an aluminum one, the starting positions were restricted back to ± 50 in both x and y coordinates.

8.2.4. Discussion

The performance of the four controllers did not vary greatly. This may be a result of the small step size (5 mm) of the robot. A second sensor measurement taken nearby the first measurement often yields less additional information than if the second measurement were taken further away from the first.

Nonetheless, there are some trends that can be seen in the data. It appears that the different controllers perform better for different statistics. For example the ‘maximize expected change in sensor reading’ controller seemed to consistently have the fewest failed trials, while the `min-var` controller (minimize the expected variance of the particles) controller tended to have the lowest error across different experimental conditions.

A failed trial would almost always be the result of one of the following causes:

- (1) robot is too far away from the target
- (2) robot gets stuck in a local minimum
- (3) the target falls on the boundary between the saturated lobes

If the robot gets lost from being too far away from the target, then really the only remedy is to increase the range of the sensor. This could be accomplished by increasing the gain on the differential amplifier, decreasing the level of noise in the system, and increasing the repeatability of the sensor. With the current setup the electrical circuitry can only sustain a finite gain in the differential amplifier before it went unstable. We usually operated near that instability point in order to maximize sensitivity.

A failed trial would often be the result of the controller getting stuck in a local minimum. The most common scenario would be when the target is in one of the saturated lobes of the sensor map, and the particles form a bimodal belief (a group of particles in the two lobes of the same sign). If the gain on the differential amplifier were decreased, then the voltages in these saturated lobes could be brought down within the range of the ADC at the price of decreasing the range of the sensory system. In future iterations we could actively change the gain of the differential amplifier to establish sensitivity where needed.

Another failure mode occurs when the target falls on the boundary between the lobes. When this happens the robot receives a measurement of near zero volts. Unless there happens to be a particle right at this boundary, the robot interprets the measurement as the target being far away and eliminates any nearby particles. A remedy for this problem would be to have multiple sets of sensors. This redundancy would make this failure mode highly unlikely.

8.3. Experimental Set 2: Eight Different Conditions for Spheres

8.3.1. Methods for Experimental Set 2

Electrolocation of spheres was performed for eight experimental conditions corresponding to all permutations of two target diameters (12.7 and 38.1 mm), two target conductivities (conductor and insulator), and two water conductivities (fresh and salt). A small stand in the tank suspends the spherical targets such that the tops of the targets cleared the bottom of the electrodes. The electrical signature of the target stand alone did not significantly interfere with any of the measurements.

In these experiments we only considered the `rnd-wlk` and `min-var` controllers. Because of the fewer controllers, we were able to explore more experimental conditions for the same total number of trials.

There are 16 possible control options at each time step. One of these options is to remain stationary, and the other 15 are randomly sampled from the interior points of a square centered at the current robot position and with sides of length 20 mm. Thus, the robot can never move more than 10 mm in any dimension in a single time step. At each time step, a new set of 15 control options are randomly generated.

8.3.1.1. Sensor Model Construction. EEVs were recorded for eight experimental conditions. For each target/water combination the target was placed at the center of the tank. For the 38.1 mm diameter spheres the robot moved the sensor carriage on a 120 mm by 120 mm grid, stopping every 4 mm for a total of $31 \times 31 = 961$ positions (dots in Figure 7.1(B)). For the 12.7 mm diameter spheres the robot moved the sensor carriage on a 100 mm by 100 mm grid, stopping every 4 mm for a total of $26 \times 26 = 676$ positions (dots in Figure 7.1(A)). The $\text{ave}(\cdot)$ and $\text{var}(\cdot)$ functions were obtained by taking the average and variance of 10 time-averaged sensor readings at each grid point.

8.3.1.2. Electrolocation Protocol. Electrolocation was performed 50 times for each of the eight experimental conditions using the two controllers, for a total of $50 \times 8 \times 2 = 800$ trials. The robot was initiated at the same 50 random positions for each experimental condition. The target was always placed at the center of the tank (coordinates $(x_w = 0, y_w = 0)$). For the larger targets the robot was confined to start within coordinates $\{\pm 55, \pm 55\}$ (darker gray area in Figure 7.1(B)) and was permitted to move anywhere

with coordinates of $\{\pm 60, \pm 60\}$. For the smaller targets the robot started at $\{\pm 40, \pm 40\}$ and stayed within $\{\pm 50, \pm 50\}$ (darker gray area in Figure 7.1(A)).

The two probabilistic controllers were terminated either after 35 steps (failure) or when the square root of the trace of the particle covariance matrix dropped below 10 mm (success).

8.3.2. Experimental Results from Experimental Set 2

Figure 8.5 shows a typical electrolocation sequence. Figure 8.5(A) shows the particles randomly distributed over the workspace since the robot initially only knows that the target is somewhere within the bounds of the workspace. In Figure 8.5(B) the robot has updated its belief using a single observation and the particle filter algorithm. Since the sensor measurement is near the RMS noise assumed for the sensor, the particle filter kept most of the particles that correspond to a sensor reading of near 0 V. This trial uses the `min-var` controller. Based on the particles shown in Figure 8.5(B) the robot moves up 9.5 mm and right 4.5 mm. This “sense and move” sequence is repeated until the uncertainty is below the threshold (square root of the trace of the covariance is less than or equal to 10 mm). Figure 8.6 shows the final belief (blue dots).

The random walk controller (`rnd-wlk`) was compared to the greedy minimum-variance controller (`min-var`). First, we look at two statistics to evaluate the performance of the controllers. The first statistic is the number of failed trials (i.e., trials that need more than 35 steps) out of the 50 total trials for each controller for each condition. The other statistic is the median number of steps for completion. Figure 8.7 displays these statistics for the 800 target electrolocation trials (broken down by controller type and experimental

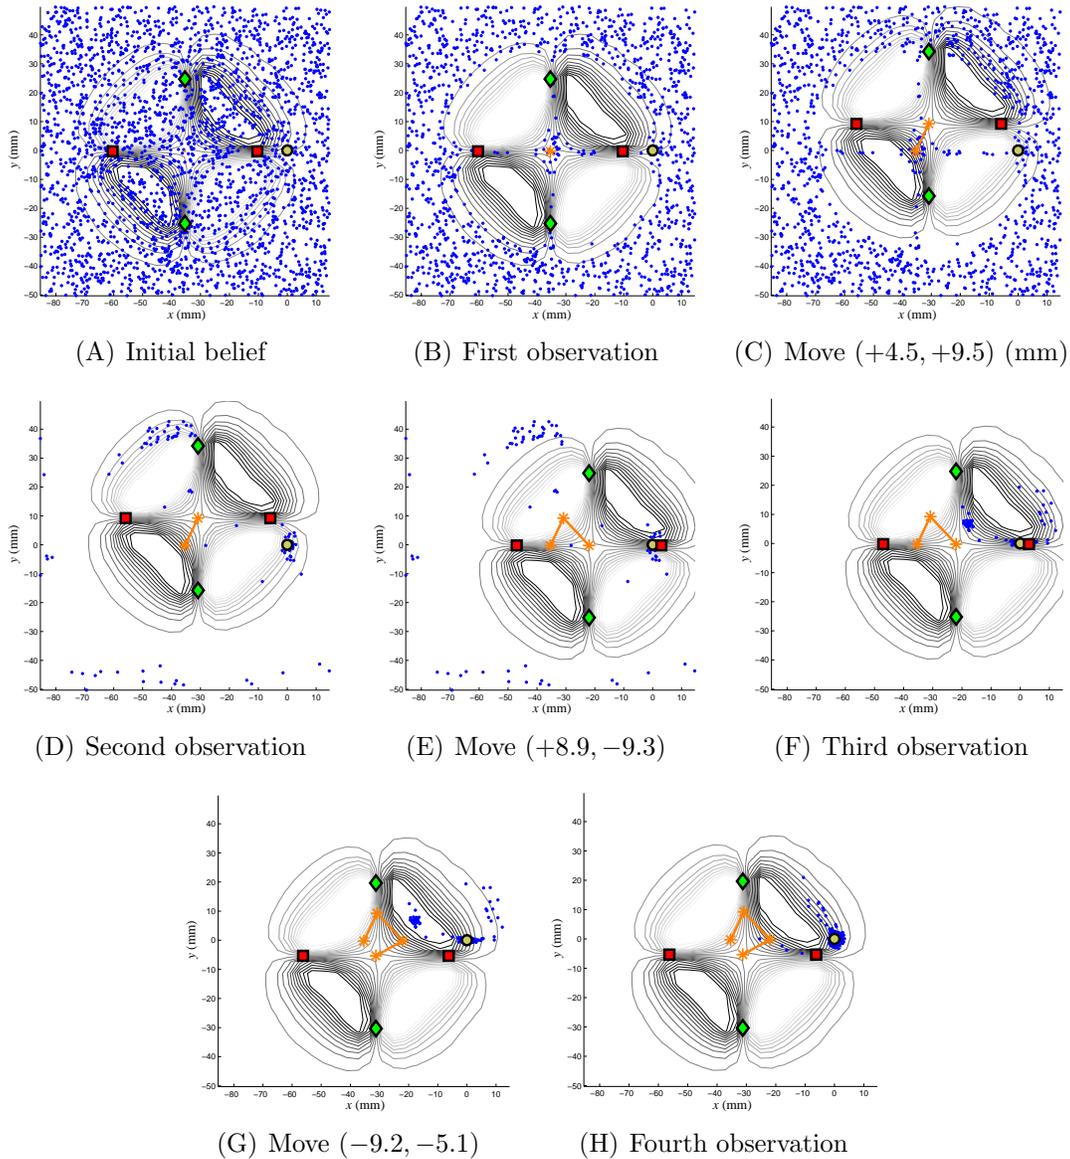


Figure 8.5. Snapshots of a typical electrolocation sequence. The EEV corresponds to Figure 7.2(A). The 2500 particles are shown as blue dots. The actual target location is a yellow circle, in this case a 12.7 mm diameter stainless steel sphere. The position of the electrodes on the robot are the green diamonds (sensors) and red squares (electric field emitter). The orange asterisks represent the path of the robot. The target is at coordinate $x = 0$, $y = 0$. The `min-var` controller was used to choose the action at each time step. The “move” frames are the beliefs immediately before the observation, and the “observation” frames are the beliefs after the observation taken from this position is fused.

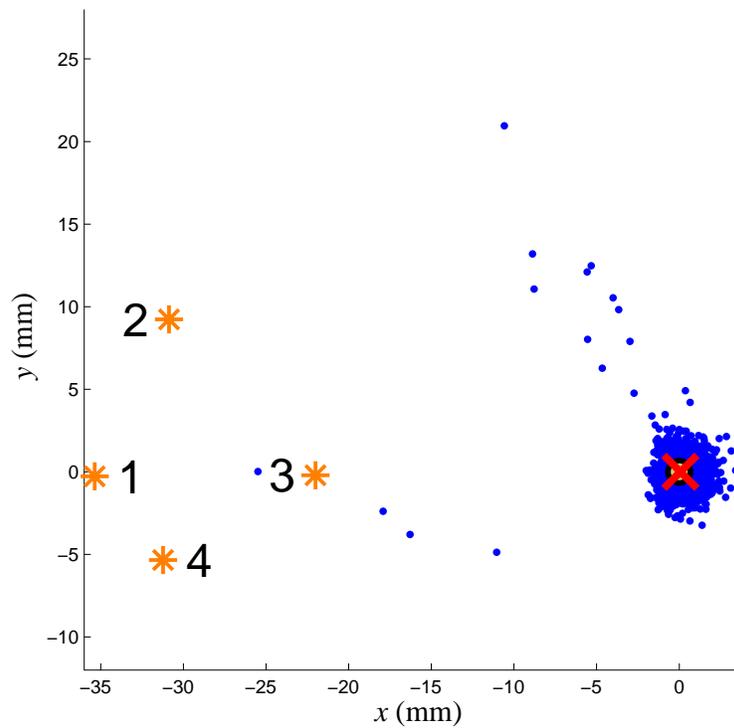


Figure 8.6. Final snapshot of the same trial. The red “x” located at $(x = 0.074, y = 0.001)$ is the centroid of the particles and is the final position estimate for the target located at $(x = 0, y = 0)$. The numbered orange asterisks are the positions visited by the robot. Note that the axes have changed from the previous figure.

conditions) and suggests that the `min-var` controller is better at reducing both the number of failures and the median number of steps to complete the task. The total number of failures for the `rnd-wlk` is 113 out of the 400 total, which is almost six times as many as the 19 failures for the `min-var` controller. For all eight of the experimental conditions the median number of steps for completion is lower for the `min-var` controller with the average of those medians being 12.4 steps for `rnd-wlk` and 6.7 for the `min-var` controller.

Table 8.2 shows the performance of the electrolocation controllers using the same three statistical measures as before (see Section 8.2.2 for explanation). If we examine the success trial statistics only, the mean error and the mean number of steps for completion do not make a compelling argument for choosing the minimum variance over the random walk controller. However, if we examine all trials including the failures, the median error and the median number of steps is significantly better for the active controller over the random walk (See Figure 8.7).

8.3.2.1. Pre-Detection Search. In 22% of the electrolocation trials the robot’s starting position was outside the detection range of the target (i.e., the initial position of the target fell outside the red contours in Figure 7.2). Thus, for these trials the initial electro-sensory feedback does not provide any reliable estimate to the location of the target. However, the resulting belief is still improved since significant portions of the target configuration space are inconsistent with the observation (e.g., Figure 8.5(B)). The `min-var` controller will move to minimize the expected particle variance, which will usually drive the robot towards exploring novel regions. Prior to initial target detection the `min-var` controller performs better than the `rnd-wlk` controller since it systematically eliminates regions of the target configuration space. In addition `rnd-wlk` performs poorly in this case

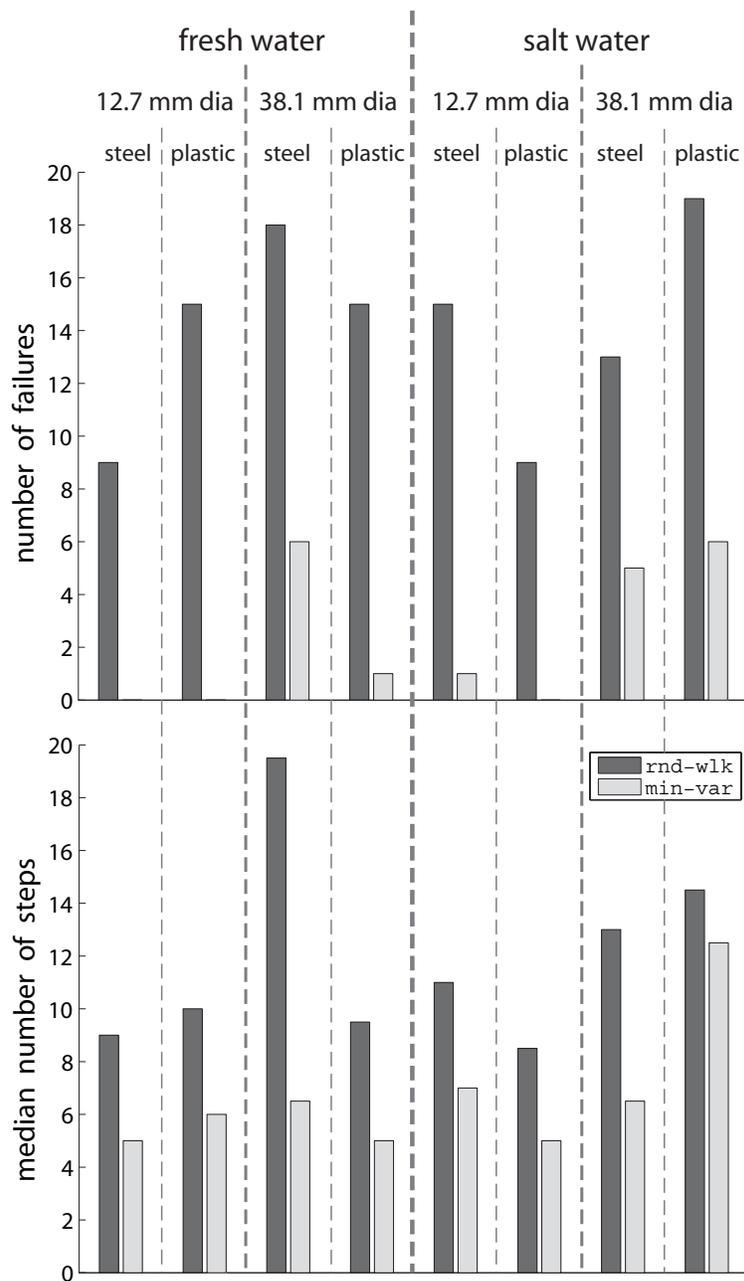


Figure 8.7. Performance of the `rnd-wlk` versus the `min-var` controllers. The height of the bars are determined for the 50 trials for each of the controllers for each of the experimental conditions. The top bar graph displays the number of failed trials. The bottom bar graph displays the median number of steps needed to complete the electrolocation task.

| water salinity | sphere dia. [mm] / material | control type | success (failure) | error [mm] mean / median | # steps mean / median |
|-----------------------|------------------------------------|---------------------|--------------------------|---------------------------------|------------------------------|
| fresh | 12.7 / S | rnd-wlk | 41 (9) | 2.0 / 1.5 | 10 / 9 |
| fresh | 12.7 / S | min-var | 50 (0) | 1.7 / 1.1 | 10 / 5 |
| fresh | 12.7 / P | rnd-wlk | 35 (15) | 2.1 / 2.0 | 9 / 10 |
| fresh | 12.7 / P | min-var | 50 (0) | 1.5 / 1.1 | 10 / 6 |
| fresh | 38.1 / S | rnd-wlk | 32 (18) | 2.2 / 2.8 | 15 / 23.5 |
| fresh | 38.1 / S | min-var | 44 (6) | 1.4 / 1.5 | 9 / 6.5 |
| fresh | 38.1 / P | rnd-wlk | 35 (15) | 2.4 / 2.8 | 10 / 9.5 |
| fresh | 38.1 / P | min-var | 49 (1) | 2.2 / 2.4 | 9 / 5 |
| salt | 12.7 / S | rnd-wlk | 35 (15) | 4.3 / 4.9 | 11 / 11 |
| salt | 12.7 / S | min-var | 49 (1) | 3.9 / 3.7 | 10 / 7 |
| salt | 12.7 / P | rnd-wlk | 41 (9) | 2.3 / 2.2 | 10 / 8.5 |
| salt | 12.7 / P | min-var | 50 (0) | 2.7 / 1.9 | 7 / 5 |
| salt | 38.1 / S | rnd-wlk | 37 (13) | 1.1 / 1.0 | 10 / 13 |
| salt | 38.1 / S | min-var | 45 (5) | 1.3 / 1.0 | 11 / 6.5 |
| salt | 38.1 / P | rnd-wlk | 31 (19) | 3.1 / 4.9 | 11 / 14.5 |
| salt | 38.1 / P | min-var | 44 (6) | 2.9 / 2.8 | 11 / 12.5 |
| all conditions | | rnd-wlk | 287 (113) | 2.4 / 2.8 | 10.8 / 12.4 |
| all conditions | | min-var | 381 (19) | 2.2 / 1.9 | 9.7 / 6.7 |

Table 8.2. Electrolocation statistics for the eight experimental conditions tested for the particle filter-based control scheme. P=Plastic; S=Steel. Water salinity: fresh is 160 $\mu\text{S}/\text{cm}$; salt is 56,000 $\mu\text{S}/\text{cm}$. Refer to Section 8.2.2 for a definition of each of the statistics displayed here.)

because it will often revisit locations, which often does very little to improve the belief. This hypothesis is consistent with the fact that when electrolocation statistics (like those shown in Figure 8.7 and Table 8.2) were reanalyzed using only the 22% of the trials that began outside the detection range, the `min-var` controller still outperformed the `rnd-wlk` controller. In future versions we could use an efficient pre-detection search strategy like the one in [31]

8.4. Experimental Set 3: Comparison of Six Controllers

In the final set of data we compared all six controllers electrolocating a one inch (25.4 mm) metal sphere in 350 μS water.

8.4.1. Methods for Experimental Set 3

The workspace was a 100 mm by 100 mm square at the center of the tank. The EEV was constructed by sampling every 4 mm on a 100 mm by 100 mm grid (sampling points seen in Figure 7.1(A) on page 167). The corresponding EEV is depicted in Figure 8.8. The gain on the differential amplifier was $G = 107$.

Each of the six controllers started from the same 64 starting positions. The starting positions were generated by first dividing the 100 by 100 mm workspace into an 8 by 8 grid of cells. Exactly one starting position was chosen randomly from each set of interior points that define each of the cells.

The remainder of the experimental methods for this set of data were identical as those used for data set 2.

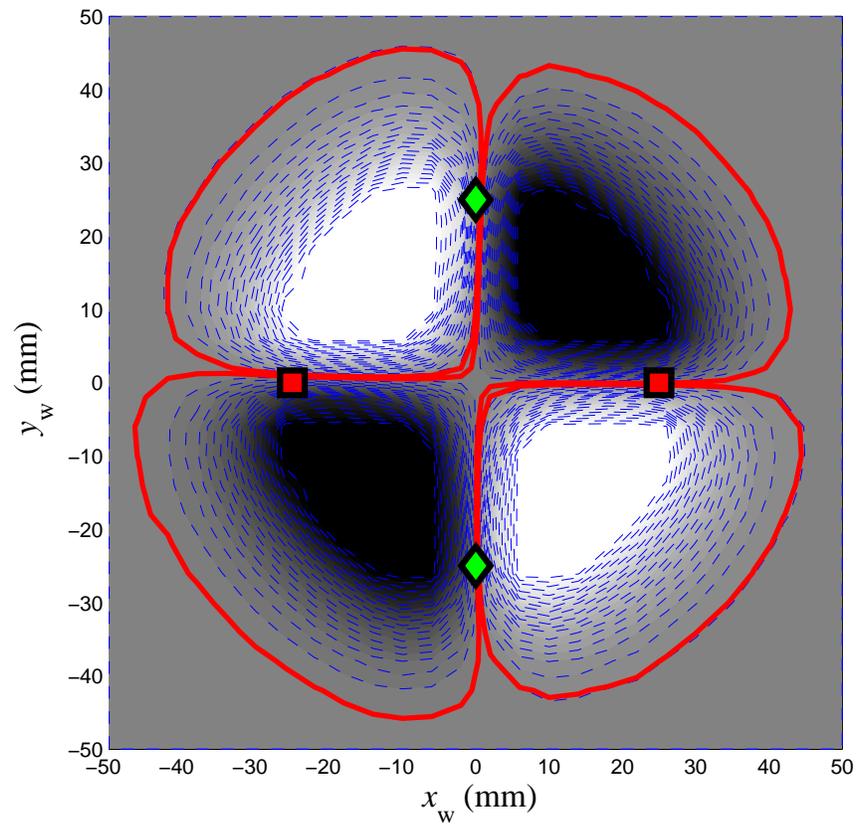


Figure 8.8. EEV for experimental set 3. Red squares indicate the location of the electric field dipoles. Green diamonds indicate the location of the voltage sensors. The color map goes from white for +10 V to gray for intermediate values to black for -10 V. Blue dashed lines are at 500 mV increments. The solid red contours represent an observation that results in detection of the target with 95% confidence.

8.4.2. Results for Experimental Set 3

Figure 8.9 shows the median number of steps for completion and the number of failures (out of 64) for each of the six controllers. Table 8.3 shows all 5 statistical measures for each of the six controllers (see Section 8.2.2 for explanation of the statistical measures). Both the `max-grad` and `parts-eater` had no failed trials. The `min-var` controller had the lowest mean and median number of steps for completion. As before, only the trials where all controllers successfully located the target were used in the mean calculations. In this data set there were 19 such trials.

It is not obvious which of the six controllers performed the best if we consider all statistical measures. Both the `max-grad` and `parts-eater` successfully located the target within the allowed number of steps all 64 times. The `min-var` controller appeared to perform the electrolocation task the quickest, although there is no significant difference between `min-var`, `max-grad`, and `parts-eater` in the number of steps to complete the task in the 19 trials. Due to the limited number of trials, a definitive choice of superior controller cannot be established, but it is clear that `min-var`, `max-grad`, and `parts-eater` performed better than the other three (under our chosen measures). Thus, we conclude that the `min-var`, `max-grad`, and `parts-eater` controllers all performed roughly the same and share a three-way tie for “best” controller.

The two open-loop controllers—`rnd-wlk` and `spiral`—performed worse across the board compared with the three best active controllers. These differences are significant ($\alpha = 0.05$) for the number of steps for completion in the 19 trials. Between the two open-loop controllers, the `spiral` controller appears to edge out `rnd-wlk`, mostly due to its fewer failed trials. This was an expected result since `spiral` systematically explores

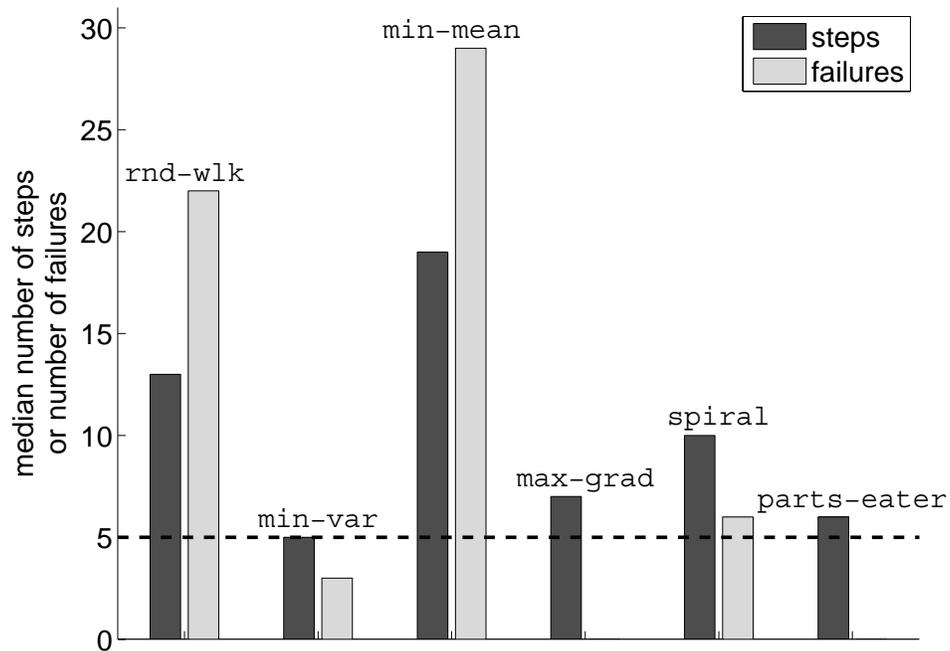


Figure 8.9. Performance of the six controllers for data set 3. The height of the bars are determined for the 64 trials for each of the controllers for each of the two performance measures, which are the number of failed trials (out of 64) and the median number of steps needed to complete the electrolocation task. Shorter bars indicate better controller performance. The horizontal dashed line indicates the minimum (across all six controllers) median number of steps for any single controller. The minimum number of failures for any controller is zero (both `max-grad` and `parts-eater`).

| | controller | success trials (failure trials) | error, mm mean / median | steps mean / median |
|-----------|--------------------------|------------------------------------|----------------------------|------------------------|
| open-loop | <code>rnd-wlk</code> | 42 (22) | 2.0 / 3.8 | 10.7 / 13 |
| | <code>spiral</code> | 58 (6) | 2.2 / 2.0 | 9.1 / 10 |
| active | <code>min-mean</code> | 35 (29) | 2.1 / 7.0 | 10.8 / 19 |
| | <code>min-var</code> | 61 (3) | 1.3 / 1.4 | 6.2* / 5* |
| | <code>max-grad</code> | 64* (0*) | 1.9 / 1.7 | 6.3 / 7 |
| | <code>parts-eater</code> | 64* (0*) | 1.1* / 1.1* | 6.4 / 6 |
| | average | 54 (10) | 1.8 / 2.9 | 8.3 / 10 |

Table 8.3. Statistics for electrolocation data set 3. Asterisks (*) indicate the best value for each statistic. The last row displays the average (mean) for each statistic across all six controllers. The “mean” statistics are computed across the 19 trials where all six controllers successfully completed the electrolocation task. The “median” statistics considered all 64 trials.

the workspace, where `rnd-wlk` may often wander off to a corner of the workspace without ever exploring the rest of the workspace.

The active controller `min-mean` appears to have performed the worst across all six controllers. This is mainly due to its 45% failure rate. In many of its trials the center of robot arrives at the centroid of the particles before the target is successfully located. If the target is located in one of the saturated lobes (see black and white lobes in Figure 8.8), and if the particles are clustered in these lobes, then observations will do little to refine the belief. The controller will keep the robot stationary since it is already at the centroid of the particles. When this happens the robot will remain there until its 35 time steps have expired, which will result in a failed trial. See Figure 8.10 for an illustration of the “stuck” scenario.

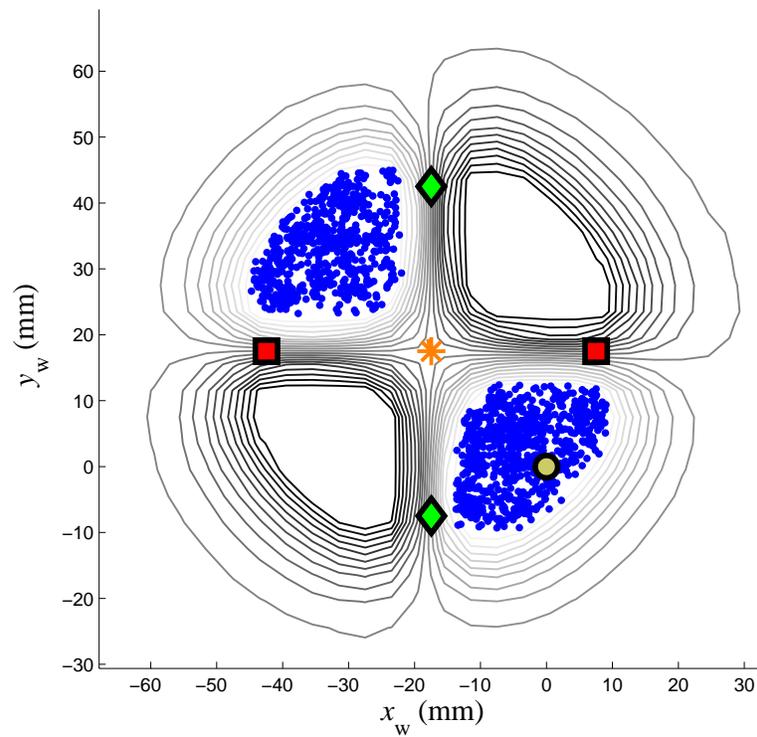


Figure 8.10. The mean-min controller gets stuck. The blue dots are the particles in the belief, and the yellow disk is the actual location of the target. The center of the robot (orange asterisk) has arrived at the centroid of the particles. Since the target is located in a saturated lobe, the belief will not change. Thus, the robot will not move.

8.5. Conclusions Across Experiment Sets

We now draw generalizations across all twelve experimental conditions (three experimental conditions from set 1, eight from set 2, and one from set 3). In all twelve experimental conditions, `rnd-wlk` was compared with `min-var`, thus the strongest generalization can be made comparing these two controllers. In all twelve conditions, the `min-var` performed better with both fewer failed trials and a lower median number of steps to complete the task. Thus, the `min-var` active controller does better than merely wandering around aimlessly.

Across all six controllers, in 11 of the 12 conditions, `min-var` had the lowest median number of steps for task completion. The only exception was with the one inch diameter plastic cylinder in salt in experimental set 1. In this case both `min-mean` and `max-grad` usually needed fewer steps than `min-var`.

In the 4 experimental conditions where the `max-grad` controller was used, it always had the fewest number of failed trials. It only failed 14 times out of a total of 221 attempts.

Not enough data was taken with the `parts-eater` controller to make any strong generalization. In the final set of experiments it did have an impressive estimation error (mean and median = 1.1). And, it never did fail to locate the target in its 64 tries.

CHAPTER 9

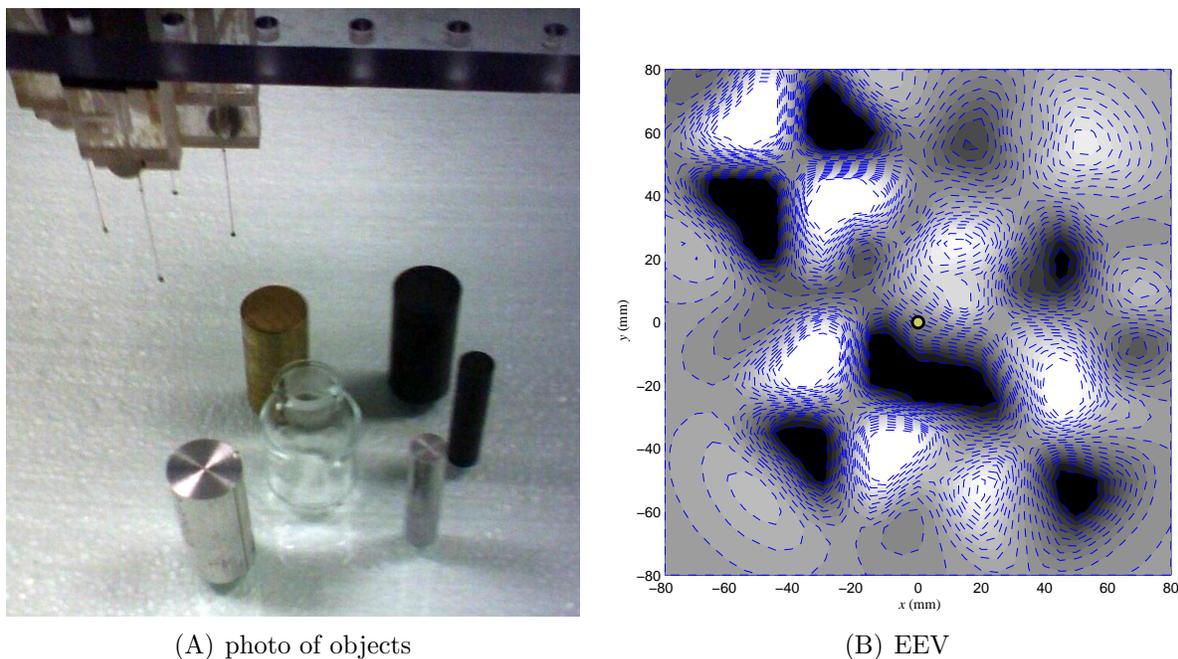
More Capabilities of the Electrosenster

Thus far we have only utilized the Electrosenster to localize single targets with a predefined sensor model (i.e., $\text{ave}(\cdot)$ and $\text{var}(\cdot)$). In this chapter we investigate additional capabilities of the Electrosenster. First, the tank is populated with many objects. The task of the Electrosenster is to locate itself amid these objects (i.e., global localization). Next, a controller is presented that does not need a predefined sensor model. It exploits invariant spatial features of the electric field signatures of a class of targets to achieve electrolocation without calibration. Finally, extensions to the current Electrosenster are proposed.

9.1. A Global Localization Scenario

We now demonstrate the robot's ability to localize itself in a much more complex environment—global localization. From the standpoint of the particle filter-based electrolocation algorithm, it does not matter if the measured EEV is of a single target or of a complex layout of diverse objects. Both the `min-var` and `rnd-wlk` controllers were used.

For this set of experiments six objects of different diameters and conductivities were dispersed randomly with x_w and y_w coordinates within ($x_w = \pm 80, y_w = \pm 80$). Figure 9.1 shows the environment. The goal of the robot was to localize itself relative to the origin of the map. Table 9.1 summarizes the performance of the global localization task for all 100 trials (50 trials for each controller). These global localization trials used the same



(A) photo of objects

(B) EEV

Figure 9.1. Global localization environment. (A): A photo of the six objects. There are two aluminum cylinders (silver color), one brass cylinder (gold color), two PVC plastic cylinders (black), and a glass vial. The four electrodes of the Electroster can be seen in the upper left corner. (B): EEV used for global localization. This EEV serves as the prior knowledge for the global localization experiments. The origin of the image is shown as a yellow circle. The color map goes from white for +10 V to gray for intermediate values to black for -10 V. Blue dotted lines are at 1 V increments.

| controller type | success trials (failure trials) | error, mm mean / median | steps mean / median |
|----------------------|------------------------------------|----------------------------|------------------------|
| <code>rnd-wlk</code> | 48 (2) | 1.3 / 1.0 | 11.8 / 11 |
| <code>min-var</code> | 49 (1) | 1.4 / 0.8 | 6.9 / 6 |

Table 9.1. Statistics for global localization. The “mean” statistics are computed across the 48 trials where both controllers successfully completed the electrolocation task. The “median” statistics considered all 50 trials. Refer to Section 8.2.2 for complete definitions of each of the statistics displayed here.

protocol and algorithms as the probabilistic-based single-target trials. Figure 9.2 depicts a self-localization sequence.

The results from the global localization trials show little difference in the success and error statistics between the two controllers tested. During many of the global localization trials it was common for multiple clusters of particles to form in the belief representation. As seen in Figure 9.1, multiple locations in the EEV look locally similar. As a result, the minimum-variance controller, which benefits from a local gradient, had trouble disambiguating the multiple hypotheses, but on average the minimum-variance controller was still able to complete the task faster than the random walk controller.

As seen from Table 9.1 the `min-var`—on average—performs the self-localization task much faster than the `rnd-wlk` controller. This difference is even more noticeable here than in the single-target electrolocation task that was investigated in the previous chapter. This may be because this is a much richer electrosensory environment. For example compare the EEV in Figure 9.1(B) with that of any of the single-target EEVs. The symmetry and regularity in the single-target EEVs actually hinders the electrolocation task when using these controllers since it is more difficult to disambiguate local electrosensory scenes. Notice there are steeper gradients in the multi-target EEV than in a single-target EEV.

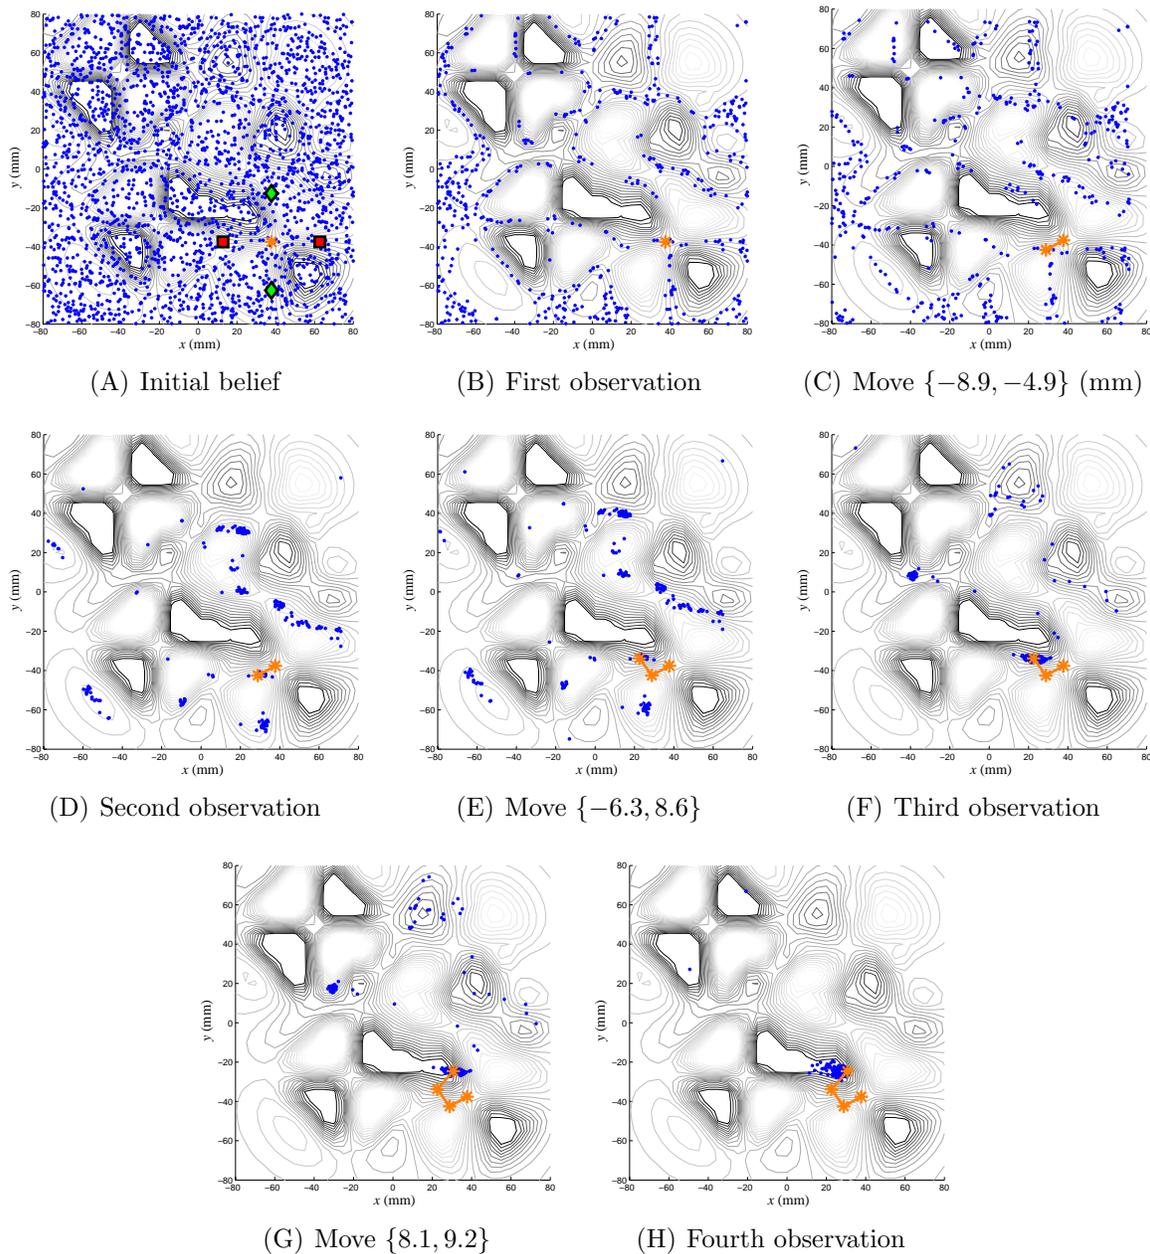


Figure 9.2. Snapshots of a typical global localization sequence. The 2500 particles are shown as blue dots and represent belief of the current location of the center of the robot. The position of the electrodes on the robot are the green diamonds (sensors) and red squares (electric field emitter). The electrodes are only shown in the first frame. The orange asterisks represent the actual path of the center of the robot. The EEV corresponds to Figure 9.1. The *min-var* controller was used to choose the action at each time step.

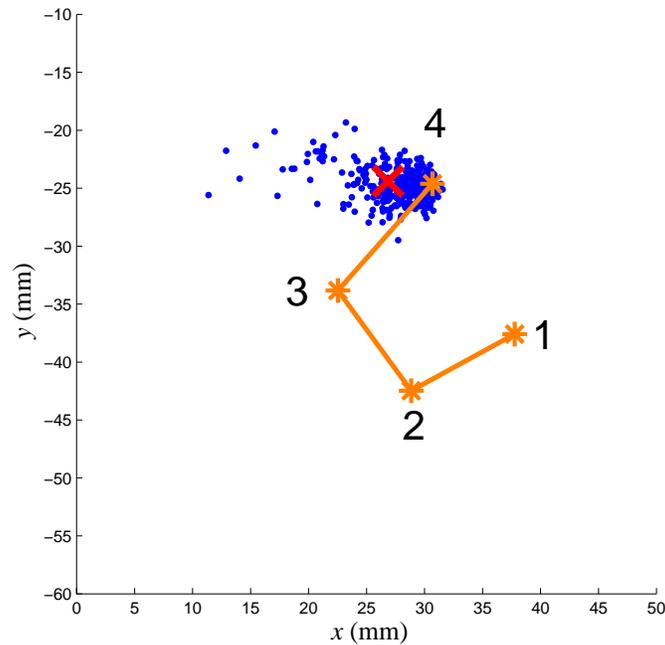


Figure 9.3. Final snapshot of the same trial. The red “x” located at $(x = 26.9, y = -24.4)$ is the centroid of the particles and is the final position estimate for the robot position, which is actually located at $(x = 30.7, y = -24.6)$. The numbered orange asterisks are the positions visited by the robot. Note that the axes have changed from the previous figure.

Thus, there is more reliable information to base control decisions on. This is why the feedback controller performs much better than the open-loop controller.

9.2. A Non-Probabilistic-Based Controller

One of the major disadvantages of the probabilistic-based controllers that has been used thus far is that the sensor model for the specific set of experimental conditions must be calibrated in advance. Empirically determining the sensor model can be time-consuming or impractical. Instead, we can use features of the EEV common to all isolated targets that are invariant to the salinity of the water and conductivity and size of the target. The feature that we use is the two orthogonal lines of 0 V that pass through

the center of the sensor (see, e.g., Figure 8.8). This feature was found for all isolated, spherical targets we observed, and it suggests the use of sensor voltage zero crossings for electrolocation. There is also typically a circular contour of 0 V (See Figure 2.7 on page 53), but the gradient of the EEV is small at this distance, and we require a steep gradient of the EEV for robust identification of zero crossings.

We tested a single policy for the deterministic controller, **zero-cross**, which searches for polarity change in sensor readings. Defining the coordinates of a world frame as (x_w, y_w) , the robot performs a binary search on the region of possible x_w -coordinates of the target. For each x_w -coordinate tested, the robot moves in the y_w -direction until a robust change in signal polarity is identified. Once this is found, the binary search ends, and the target's y_w coordinate is recorded as the y_w position of the sensor where the zero crossing occurs. The robot then scans along the x_w -direction until another robust zero crossing is found.

Figure 9.4 depicts an example of the **zero-cross** controller locating a metal sphere. The y -coordinate of the target is searched for first. In Figure 9.4(A), the Electrosenser divides the workspace in half (in the x -direction) and moves down the middle with a constant x -coordinate and moving in 10 mm increments. It did not find a zero crossing in its first pass (first pass shown in Figure 9.4(A)). It then continues to systematically search. In Figure 9.4(B) it does find a zero-crossing. In Figure 9.4(C) it moves in smaller increments to achieve an accurate estimate. In Figure 9.4(D) it begins to search for the x -coordinate of the target by moving along a line 15 mm above the estimated target y -coordinate, which is where the EEV shows its steepest gradient near the zero crossing.

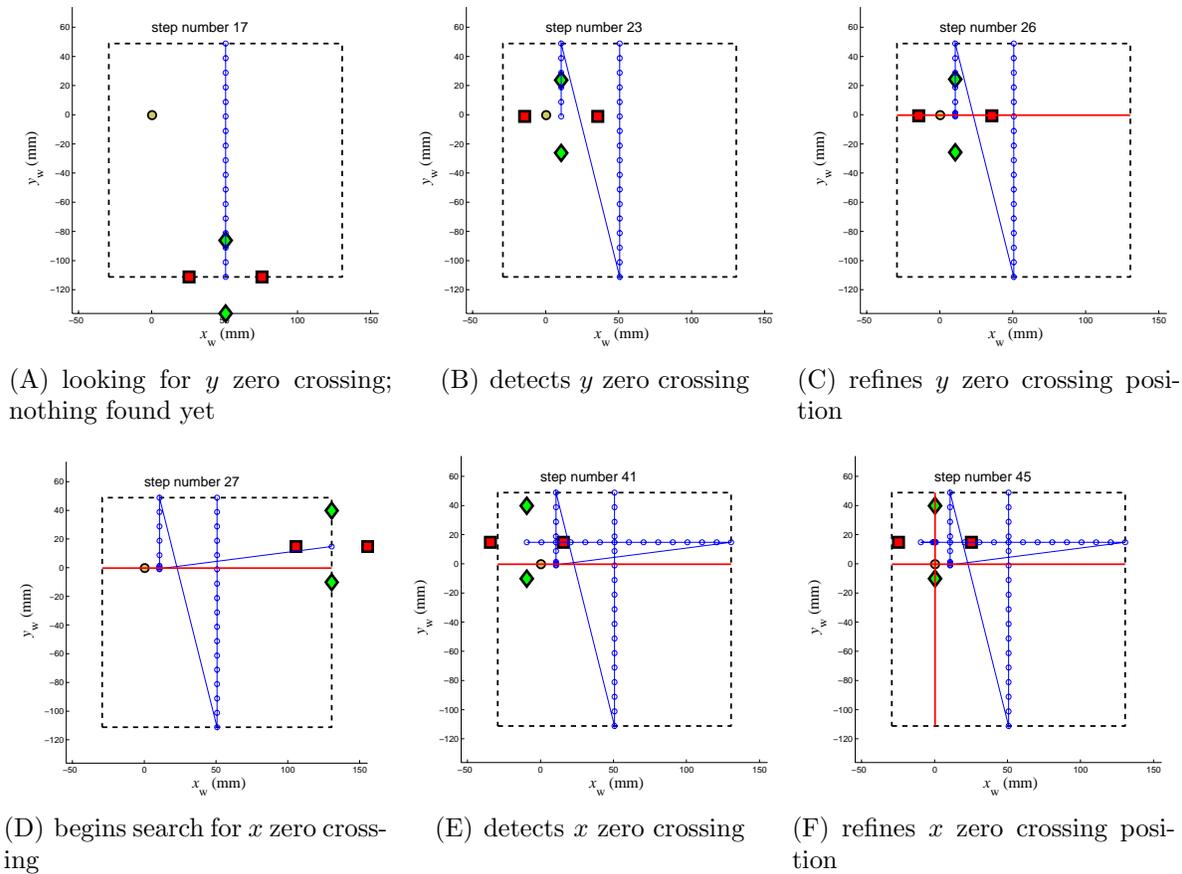


Figure 9.4. Electrolocation using zero-cross controller. The black dashed lines represent the predefined bounds of possible target locations; this defines the workspace of the robot. The blue line is the path of robot, and the blue circles are locations of the center of robot where it stopped to make an observation. The yellow disk is the actual location of the target. Red squares and green diamonds are the locations of the emitter and detectors for the current robot position, respectively. In (C) the red horizontal line signifies the y estimate of the target location. Similarly, the red vertical line in (F) is the x estimate of the target location. Thus, the intersection of the two red lines is the estimate of the target.

In Figures 9.4(E) and 9.4(F) the robot detects zero crossing and refines its x -coordinate estimate.

9.3. Performance of Non-Probabilistic-Based Controller

Two sets of experiments were performed using the `zero-cross` electrolocation algorithm. In the first set it was noticed that the estimate of the y -coordinate of the target position was much worse than the estimate of the x -coordinate. In the second set of experiments the algorithm was adjusted to include an additional pass to improve precision. This additional pass greatly improved the estimate of the target location at the expense of taking longer to perform the task. The two sets of experiments are now described in detail.

9.3.1. Experimental Set 1: Eight Different Conditions

Electrolocation of spheres was performed for eight experimental conditions corresponding to all permutations of two target diameters (12.7 and 38.1 mm), two target conductivities (conductor and insulator), and two water conductivities (fresh and salt). These are the identical target and water properties as used in Experimental Set 2 of the probabilistic-based controller (See Section 8.4 on page 205). A small stand in the tank suspends the spherical targets such that the tops of the targets cleared the bottom of the electrodes. The electrical signature of the target stand alone did not significantly interfere with any of the measurements.

Table 9.2 shows the average electrolocation error for the zero-crossing controller over the 50 trials for each of the eight conditions. The target always remains at coordinate

$(x_w = 0, y_w = 0)$, but the center of the robot's 160 mm by 160 mm workspace to scan is generated randomly from within coordinates $(x_w = \pm 75, y_w = \pm 75)$. This allows the relative starting position of the robot to change every trial without having to move and calibrate the position of the target at every trial. For each set of experimental conditions the estimated position of the target is calculated as the mean x_w coordinate and mean y_w coordinate of all 50 final positions for that set of conditions.

Since the estimated position is calculated after the electrolocation task, the numbers in Table 9.2 demonstrate the repeatability of the electrolocation system and not necessarily its ability to locate a predefined reference point on the target. However, theory and empirical observations both suggest that the zero-crossing reference point that we are localizing is the actual center of the 2D projection of the sphere. Thus, the data in Table 9.2 is a measure of the precision of the electrolocation system. A consistent offset bias in the location estimate would not be reflected in the statistics in Table 9.2. In other words, Table 9.2 does not capture the statistical first-order moment, but it does reflect the second-order moment of the estimates.

Table 9.2 reveals that the average magnitude of error in the x_w direction is less than 1 mm for each of the conditions and less than $65 \mu\text{m}$ in 6 out of the 8 conditions. The large discrepancy in precision between the x_w and y_w errors is because during the search for the y -coordinate (which is looked for first), the Electrostenster has no clue how far away the target is, thus, the zero-crossing could be a very shallow gradient. Once the y coordinate is known, the Electrostenster takes a trajectory during the x -coordinate search that will allow the target to pass approximately 15 mm from the center of the robot. Figure 9.5(A) shows a typical EEV. The blue solid line shows a possible relative trajectory the

| experimental conditions | mean abs error | |
|---------------------------|----------------|------------|
| | x_w (mm) | y_w (mm) |
| 12.7 mm, steel in fresh | 0.06 | 0.39 |
| 12.7 mm, plastic in fresh | 0.05 | 0.64 |
| 38.1 mm, steel in fresh | 0.05 | 0.93 |
| 38.1 mm, plastic in fresh | 0.06 | 0.89 |
| 12.7 mm, steel in salt | 0.23 | 0.47 |
| 12.7 mm, plastic in salt | 0.14 | 0.59 |
| 38.1 mm, steel in salt | 0.06 | 0.78 |
| 38.1 mm, plastic in salt | 0.05 | 0.69 |

Table 9.2. Electrolocation statistics for the eight experimental conditions tested for the zero-crossing-based control scheme.

target could make during the y -coordinate search. The red dashed line shows the relative trajectory the target would make during the x -coordinate search ($y = -15$ relative the estimated target y -coordinate from the previous pass). Figure 9.5(B) shows the expected history of observations for the two trajectories shown in Figure 9.5(A). Because of factors such as noise and a finite number of observations, an estimate of the zero crossing of the blue solid line would be less precise. Thus, we see larger variability in the y estimate than the x estimate. Next, we will add another pass to the `zero-cross` controller that will alleviate this problem.

9.3.2. Experimental Set 2: Extra Pass

In this set of experiments an additional pass is added to the `zero-cross` controller in an attempt to increase its precision at locating the spherical target. Here, a one inch

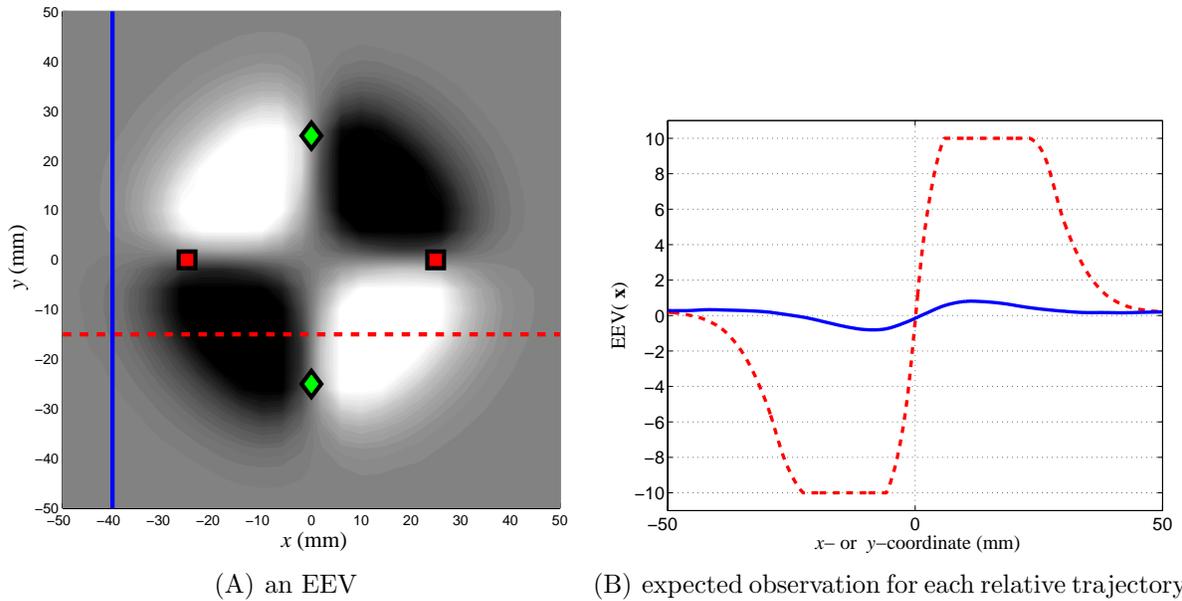


Figure 9.5. Two target trajectories relative to the robot. The solid blue line is a relative trajectory the target could take during the search for the x -coordinate of the target. The red dashed line corresponds to relative trajectory of the target during the x -coordinate search of the **zero-cross** controller. The target moves at a constant $y = 15$ relative to the center of the robot (based on the y estimate found in the previous pass). (B): The history of the expected observations of the two trajectories from (A).

(25.4 mm) metal sphere is electrolocated in 350 μS water (This is the same as Experimental Set 3 for the probabilistic-based controllers, see Section 8.4). The electrolocation task is repeated 100 times.

Recall from Table 9.2 that the **zero-cross** algorithm did a relatively poor job of estimating the y -coordinate of the target location compared with the x -coordinate. It was suggested from Figure 9.5 that certain relative trajectories could render poor estimates. In this final set of experiments an additional pass is executed by the ElectroSenser to obtain a more precise y -coordinate estimate. So, the new **zero-cross** protocol is 1) approximate y -coordinate, 2) estimate x -coordinate, and 3) re-estimate y -coordinate.

| experimental conditions | mean abs error | |
|-------------------------|----------------|------------|
| | x_w (mm) | y_w (mm) |
| 25.4 mm, steel in fresh | 0.11 | 0.02 |

Table 9.3. Electrolocation statistics for the zero-crossing-based control scheme using an additional pass.

Table 9.3 shows the average error estimate for both coordinates over the 100 trials. As before, these numbers reflect the precision of the electrolocation system. The average total error is $112 \mu\text{m}$, which is a great improvement from the previous set of data. Thus, the additional pass to re-estimate the y -coordinate will improve the target estimate.

9.4. Extensions of the Current Setup

The Electrosenster is the first actively-controlled robot to use electrosense to locate targets. Its simplicity did not limit us from investigating the basic questions of artificial electrosense. Nonetheless, there are many improvements we could make to the system. A few are discussed below.

The only difference between the EEVs shown in Figures 7.2(E) and 7.2(F) is that the magnitude of the electrical contrast factor (χ in Equation (2.13)) is greater with the plastic target than with the metal target. An alternative to increasing the magnitude of χ is to increase the gain of the differential amplifier. If we disregard the effects of noise, the sensor map shown in Figure 7.2(E) could be made to look identical to the one shown in Figure 7.2(F) by only changing the gain on the differential amplifier (negative gains are permissible). Armed with the controller's proposed ability to dynamically change G , an electrolocation algorithm could take advantage of having both a high-gain EEV (good for long-range detection) and a lower-gain EEV (good for precise localization when

near the target). Interestingly, there is good evidence for dynamic gain control in the electrosensory lateral line lobe, the structure in the hindbrain of weakly electric fish where all electrosensory data are initially processed [21, 146].

Because of the commutative property of Bayes filters, the final belief shown in Figure 8.4 could have been achieved by four stationary sensor pairs positioned at the four orange asterisks. As seen by the EEVs shown in Figure 7.2, a single observation is almost never enough to establish a satisfactory belief of the target location. But, often with just a few movements the particles quickly converge to an acceptable uncertainty level. A robot equipped with an array of sensors could potentially locate the target without any motion between the sensor array and object. This demonstrates how motion and sensing can be traded off in active sensing systems. A parallel example from the realm of biological active localization systems is provided by the dolphin, which uses a narrow (10°) sonar beam with head scanning motions, possibly for increased sensing range over what would be possible with a broader beam [174].

All control algorithms employed by the electrosensory robot never look ahead more than one time step. Greedy controllers are vulnerable to converging to unacceptable local minima. Ideally, at each time step the controller would evaluate every possible path and proceed to execute the first action in the trajectory that yields the best expected sensing information. To make this strategy computationally practical, we must limit the number of lookahead steps and perhaps use heuristics to select which paths to evaluate.

One of the major drawbacks of the particle-filter-based controller utilized here (all controllers except **zero-cross**) is that it requires the EEV of the target. It would be

convenient if this time-consuming step of EEV construction was eliminated. The similarities between the modeled (Figure 2.7 on page 53) and measured EEVs (Figure 7.2 on page 169) suggests that a theoretical EEV for a known target can replace a pre-measured EEV. Alternatively, an EEV for an unknown target could be estimated from EEVs that are expected to be similar to that of the unknown target. For example, if the EEV for 10 mm, 15 mm, and 20 mm diameter metal sphere in fresh water is known, and the robot is searching for a 14.5 mm diameter metal sphere in fresh water, the EEV for this novel target could be approximated from the known EEVs.

Currently the probabilistic-based controller searches for the target in a two-dimensional place once it is given EEV that depends on a particular water conductivity, target conductivity, and target size. Alternatively, the robot could determine the conductivity of the water by measuring electrical current and voltage drop across the emitting electrodes. The conductivity and diameter of the target could be two more dimensions in the search space. This would dramatically increase the computational costs associated with the Monte Carlo approach. Prior knowledge could lessen this additional burden. For example, if only pure conductors or insulators are possible, then the belief space would be reduced to two 3-D spaces, which is, of course, more tractable than the full 4-D space.

We do not solve the full EIT problem with our robot. EIT usually implies measuring electrical properties at fixed electrodes—which can act as both emitters and detectors—at the periphery to estimate an impedance map of the interior [248]. Our robot employs an active sensing strategy to move towards positions of maximal information gain. Rather than construct an impedance map, this movement generates a sequence of observations serving electrolocation.

For other studies using a similar setup as the Electrosenster, it was found that by shielding the experimental tank from electromagnetic interference from surrounding lights and equipment, ambient noise levels could easily be reduced by more than a factor of 32. Using the same technique as discussed in Section 7.4 to compute the theoretical values in Table 7.1, this noise reduction would result in roughly a doubling in detection distance.

CHAPTER 10

Conclusions and Future Directions**10.1. The Driving Objectives of this Thesis**

The two original objectives that have driven the work presented in this thesis are 1) implementation of an active artificial electrosensory system; and 2) obtaining a better understanding of sensing and motor control in weakly electric fish.

10.1.1. Objective #1: Implementation of an Active Artificial Electrosensory System

Motivating objective #1 is the development of a sensing modality well suited for low speed, highly maneuverable underwater vehicles operating in cluttered environments. Such vehicles require sensors with several specific properties [155]. First, the sensors must work well at short range. Second, rather than being concentrated on one portion of the hull of the vehicle (e.g., cameras at the front of a remotely operated vehicle), the sensors should be distributed over the entire surface of the vehicle. This is required to support highly maneuverable, possibly omnidirectional movement in tight spaces. Third, to more robustly control sensory acquisition in geometrically complex spaces, such sensors should generate the energy that they subsequently transduce, as with radar and sonar systems [174]. Finally, as a very inexpensive approach for acquiring a unique form of sensory information—as the sensors are simply exposed conductors—active electrosense may also

provide a useful complement to the usual suite of sensors provisioned on standard underwater vehicles.

Chapters 6 through 9 demonstrate objective #1 was successfully accomplished. The Electrosenster is a complete electrosensory system. Chapter 6 describes the hardware of this 2-DOF electrosensory-based robot. In Chapter 7 we demonstrate the ability of the Electrosenster to self-calibrate its sensor model. This sensor model is used by the probabilistic-based controllers during the target localization and (self) global localization tasks. Also in this chapter we develop an analytical solution to the detection range for a simple electrosensory system, and we compare this model to empirical target detection results from the Electrosenster.

Once the Electrosenster calibrates its sensor model it can use a probabilistic-based controller to electrolocate the target. Four such active controllers and two such open-loop controllers presented in Chapter 8 establish that the Electrosenster is able to locate targets under various experimental conditions. Three out of the four active controllers performed significantly better than the two open-loop controllers.

The abilities of the Electrosenster go beyond locating an isolated target when given a sensor model. Chapter 9 demonstrates the Electrosenster's ability to localize itself in a complex environment (i.e., amid many objects). Also in this chapter we develop a controller (called **zero-cross**) that is able to localize isolated spheres *without* a predetermined sensor model. Instead, the algorithm exploits invariant spatial features of the electric field signatures induced by spherical targets.

10.1.2. Objective #2: Understanding Sensing and Motion in Weakly Electric Fish

One of our major focuses within objective #2 is how weakly electric fish—and other biological systems—could perform electrolocation. While there are a number of hypotheses in the literature regarding the mechanisms of active electrolocation [246], these are still largely not well understood. While the Electrosenster utilizes the same physics as biological active electrolocation, their belief maintenance techniques almost certainly differ. The Electrosenster can therefore merely provide a useful baseline for comparative analysis of biological active electrolocation.

While this document has made steps towards understanding biological electrolocation, we have only begun to understand this complex process. Part 1 establishes the fundamentals of electrolocation. We consider a Bayes filter for belief maintenance. While it has been proposed by many researchers that biological systems perform Bayes' optimal [128] computations within their sensorimotor systems [133, 180], without a complete neurobiological model of the weakly-electric fish, there is no way to definitively determine if this biological system truly operates under Bayes' rule.

In Part 2 we use the models and methods developed in Part 1 to simulate sensory acquisition and probabilistic belief maintenance for the prey-capture task of the weakly-electric fish (We assume that the weakly-electric fish does indeed use a Bayes filter to update its belief of the prey). Using these simple models we investigate optimal sensing in weakly electric fish and are able to make the following statements:

- (1) We suggest that the weakly electric fish has evolved a large number of sensors to offset the noise the electrosensory system is subjected to.

- (2) We speculate that the electric field is stronger near the tail to complement the fish's omnidirectional sensing and motor system.
- (3) We show that the weakly-electric fish has evolved an efficient sensor layout for the prey-capture task.
- (4) We show that if motions are limited to those similar to the fish, then mechanically-optimal trajectories are also good for sensing.

In the grand scheme, this thesis was only able to provide a few pieces to the solution of objective #2. Most of the progress was made in the sensing realm, but sensing cannot be studied in isolation without considering motion and the mechanics of the sensorimotor pathways. Thus, in future work it is crucial the sensing, motion, and neural mechanisms be studied as a complete system. The next section begins a discussion on some possible approaches.

10.2. Future Directions: Modeling the Sensing and Motion Systems of the Weakly Electric Fish

Objective #2 has been the motivation of much research across many decades and across many labs. Knowing this objective could not be fully fulfilled in a single thesis, we selectively chose a few simple aspects of the fish sensorimotor system to investigate. In this section possible future approaches of acquiring an even better understanding of sensorimotor transformations in the weakly electric fish is discussed.

Throughout this document very simplified and abstract models of the weakly electric fish have been used in the interest of computational efficiency. With these models we have obtained useful insight into sensing and motion in the weakly electric fish, but these models

fall short of accurately modeling the real sensorimotor components of their corresponding biological system. Perhaps the first step in constructing a complete model of the fish's sensorimotor transformations is the use of more realistic models. In Appendix F more complex models that are grounded in the morphology, anatomy, and neurobiology of the weakly electric fish are considered. The Appendix begins by providing some of the relevant biology (particularly in the sensing realm) of the fish (Section F.1). We then discuss two relevant models. The first is a model of the neural coding in the primary electrosensory afferents. Up until now, we have disregarded this component and have assumed perfect transmission of the voltage at the fish's sensor to the representation of that voltage in the central processing center. In reality, there are many intermediate signal transformations including the encoding of transdermal potentials (voltage at the sensor) into spike trains¹, which is the topic of Section F.2 and depicted in Figure F.1. The other component of the fish system that is touched on in Appendix F is transforming spike trains into prey location. This process is less understood, but there are some models of the electrosensory lateral line lobe (ELL). The ELL is the sole recipient of all the primary afferents of the electrosensory system. This means that all the spike trains that propagate from the electrosensory organs first pass through the ELL before anywhere else in the CNS. Thus, it is a crucial component, and is discussed in Section F.3.

¹The central nervous system (CNS) uses spike trains to encode information to be sent to neurons within the CNS. Nerves transmit voltages and the "spike" refers to the impulse of voltage in the nerve. Spike trains are not unlike binary voltages in computer circuitry (e.g., both represent information as streams of binary voltages).

Up until now, we have mostly focused on sensing and processing models of the fish². We now suggest some possible methods for modeling sensing, motion, and control components of the fish within a single framework. Appendix G discusses some simple methods for modeling sensorimotor transformation. This framework is only valid for linear systems (linear motion and sensing models) with Gaussian disturbances. While real biological systems do not fall in this realm, the methods in Appendix G can be used in situations where linear approximations are appropriate for modeling the relevant components. If it can be assumed that the dynamics are linear and the disturbances on the sensor and motion models are normally distributed, and the cost function is quadratic, then the Linear Quadratic Gaussian (LQG) can be used. LQG dramatically decreases the computational cost by decoupling the control and estimation problems (the *separation principle*). The Kalman filter is an optimal state estimator for the LQG case that parameterizes the belief with only a vector of means and the covariance matrix. The optimal control policy can be determined in closed-form via the Ricatti equations.

Unfortunately, if a “simple” model is not sufficient, then more complex (computationally expensive) methods must be employed to capture the relevant aspects of the sensorimotor transformations. Perhaps the least restrictive formulation is the partially-observable Markov decision process (POMDP). In fact we only need complete discretization (i.e., discrete states, observations, actions, and time) to use this formulation [121]. Dynamic programming, which is an application of the Bellman equation, is the most popular technique for solving a POMDP. In addition to a sensing and motion model, a reward function is needed to map a state and action to a scalar reward. In general

²In Section 4.2 (beginning on page 108) we did provide the general formulation for optimal motion control. But, we quickly abandoned analytical solutions for computationally tractable ones.

this is a computationally expensive approach, but in some cases heuristics are acceptable [64]. POMDPs and similar methods (including reinforcement learning) are discussed in Appendix H.

Appendices G and H discuss general methods for modeling systems with sensing, motion, and control. These methods were not specifically developed for biological systems. In particular biological systems must compensate for significant sensorimotor delays. Such long delays make pure feedback control impractical for many situations. In Appendices I and J we investigate some critical components of fish (and most biological system) sensorimotor transformations with a focus on compensating for sensorimotor delays. Appendix I discusses the importance of both feedback and feedforward control in biological systems. Feedback control is slow, but can compensate for unpredictable disturbances. Feedforward control circumvents sensorimotor delays, but an accurate predictive model of the world is needed to successfully use it. In Appendix J we discuss feedforward (predictive) controllers in more detail.

It is now obvious that models of biological sensorimotor transformations can be very complex. In this section we have given just a glimpse of how one may approach such a daunting endeavor. Each biological system is different, and we have focused our discussion on the weakly electric fish. Thus, other biological systems may use methods not mentioned here.

References

- [1] B. Abidi. Automatic sensor placement. In *SPIE Conference on Intelligent Robots and Computer Vision XIV*, volume 2588, pages 387–398, Philadelphia, PA, 1995.
- [2] D. Allport, T. G. Zimmerman, and J. A. Paradiso. Electric field sensing and the "flying fish". Technical report, MIT Media Laboratory - Physics and Media Group, June 1995.
- [3] S. Ando. Intelligent three-dimensional vision sensor with ears. *Sensors and Materials*, 7(3):213–231, 1995.
- [4] J. Andrade-Cetto and A. Sanfeliu. The effects of partial observability when building fully correlated maps. *IEEE Transactions on Robotics*, 21(4):771–777, 2005.
- [5] T. Arbel. On the sequential accumulation of evidence. *International journal of computer vision*, 43(3):205–230, 2001.
- [6] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-gaussian Bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, 2002.
- [7] C. Assad. *Electric field maps and boundary element simulations of electrolocation in weakly electric fish*. PhD thesis, California Institute of Technology, 1997.
- [8] C. Assad, B. Rasnow, and P. K. Stoddard. Electric organ discharges and electric images during electrolocation. *Journal of Experimental Biology*, 202(10):1185–1193, 1999.
- [9] J. J. Atick. Could information-theory provide an ecological theory of sensory processing. *Network-Computation in Neural Systems*, 3(2):213–251, 1992.
- [10] D. Babineau, J. E. Lewis, and A. Longtin. Spatial acuity and prey detection in weakly electric fish. *PLoS Computational Biology*, 3(3):402–411, 2007.

- [11] D. Babineau, A. Longtin, and J. E. Lewis. Modeling the electric field of weakly electric fish. *Journal of Experimental Biology*, 209(18):3636–3651, 2006.
- [12] M. Bacher. A new method for the simulation of electric-fields, generated by electric fish, and their distortions by objects. *Biological Cybernetics*, 47(1):51–58, 1983.
- [13] M. Baglietto, M. Paolucci, L. Scardovi, and R. Zoppoli. Entropy-based environment exploration and stochastic optimal control. In *Decision and Control, 2003. Proceedings. 42nd IEEE Conference on*, volume 3, pages 2938–2941, 2003.
- [14] R. Bajcsy. Active perception. *Proceedings of the IEEE*, 76(8):996–1005, 1988.
- [15] D. C. Barber and B. H. Brown. Applied potential tomography. *Journal of Physics E-Scientific Instruments*, 17(9):723–733, 1984.
- [16] H. Barlow. Redundancy reduction revisited. *Network-Computation in Neural Systems*, 12(3):241–253, 2001.
- [17] A. G. Barto, A. H. Fagg, N. Sitkoff, and J. C. Houk. A cerebellar model of timing and prediction in the control of reaching. *Neural Comput*, 11(3):565–94, 1999.
- [18] A. G. Barto and S. Mahadevan. Recent advances in hierarchical reinforcement learning. *Discrete Event Dynamic Systems-Theory and Applications*, 13(1-2):41–77, 2003.
- [19] G. Barto, Andrew, S. J. Bradtke, and S. P. Singh. Learning to act using real-time dynamic programming. Technical Report UM-CS-1993-002, University of Massachusetts, 1993.
- [20] J. Bastian. Electrolocation I. How the electroreceptors of *Apteronotus albifrons* code for moving-objects and other electrical stimuli. *Journal of Comparative Physiology*, 144(4):465–479, 1981.
- [21] J. Bastian. Gain control in the electrosensory system mediated by descending inputs to the electrosensory lateral line lobe. *J. Neurosci.*, 6(2):553–562, 1986.
- [22] J. Bastian. Plasticity of feedback inputs in the apteronotid electrosensory system. *Journal of Experimental Biology*, 202(10):1327–1337, 1999.
- [23] J. Bastian and W. Heiligenberg. Neural correlates of the jamming avoidance-response of *Eigenmannia*. *Journal of Comparative Physiology*, 136(2):135–152, 1980.

- [24] K. Basye, T. Dean, J. Kirman, and M. Lejter. A decision-theoretic approach to planning, perception, and control. *IEEE Expert-Intelligent Systems and Their Applications*, 7(4):58–65, 1992.
- [25] J. S. Bay. Tactile shape sensing via single- and multi-fingered hands. In *IEEE international conference on robotics and automation*, pages 290–295, Scottsdale, AZ, 1990.
- [26] J. S. Bay. A fully autonomous active sensor-based exploration concept for shape-sensing robots. *IEEE Transactions on Systems Man and Cybernetics*, 21(4):850–860, 1991.
- [27] J. S. Bay and H. Hemami. Dynamics of a learning controller for surface tracking robots on unknown surfaces. *IEEE Transactions on Automatic Control*, 35(9):1051–1054, 1990.
- [28] C. Bell, D. Bodznick, J. Montgomery, and J. Bastian. The generation and subtraction of sensory expectations within cerebellum-like structures. *Brain Behavior and Evolution*, 50:17–31, 1997.
- [29] C. C. Bell. Memory-based expectations in electrosensory systems. *Current Opinion in Neurobiology*, 11(4):481–487, 2001.
- [30] V. E. Benes, K. L. Helmes, and R. W. Rishel. Pursuing a maneuvering target which uses a random process for its control. *IEEE Transactions on Automatic Control*, 40(2):307–311, 1995.
- [31] O. Benichou, M. Coppey, M. Moreau, P.-H. Suet, and R. Voiturez. Optimal search strategies for hidden targets. *Physical Review Letters*, 94(19):198101, 2005.
- [32] A. T. D. Bennett. Do animals have cognitive maps? *Journal of Experimental Biology*, 199(1):219–224, 1996.
- [33] J. O. Berger. *Statistical Decision Theory*. Springer-Verlag, Berlin, 1980.
- [34] N. J. Berman and L. Maler. Neural architecture of the electrosensory lateral line lobe: Adaptations for coincidence detection, a sensory searchlight and frequency-dependent adaptive filtering. *Journal of Experimental Biology*, 202(10):1243–1253, 1999.
- [35] N. Bhushan and R. Shadmehr. Computational nature of human adaptive control during learning of reaching movements in force fields. *Biol Cybern*, 81(1):39–60, 1999.

- [36] S. J. Blakemore. Deluding the motor system. *Conscious Cogn*, 12(4):647–55, 2003.
- [37] S. J. Blakemore, D. Wolpert, and C. Frith. Why can't you tickle yourself? *Neuroreport*, 11(11):R11–6, 2000.
- [38] S. J. Blakemore, D. M. Wolpert, and C. D. Frith. Central cancellation of self-produced tickle sensation. *Nat Neurosci*, 1(7):635–40, 1998.
- [39] H. Bleckmann, H. Schmitz, and G. von der Emde. Nature as a model for technical sensors. *Journal of Comparative Physiology A: Neuroethology Sensory Neural and Behavioral Physiology*, 190(12):971–981, 2004.
- [40] B. Bonet and H. Geffner. Planning with incomplete information as heuristic search in belief space. In S. C. Knoblock, S. Kambhampati, and C., editors, *6th International Conf. on Artificial Intelligence Planning and Scheduling*, pages 52–61, Breckenridge, CO, 2000. AAAI Press.
- [41] A. Borst and F. E. Theunissen. Information theory and neural coding. *Nature Neuroscience*, 2(11):947–957, 1999.
- [42] F. Bourgault, A. A. Makarenko, S. Williams, B. Grocholosky, and H. Durrant-Whyte. Information based adaptive robotic exploration. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2002)*, Lausanne, Switzerland, 2002.
- [43] C. Boutilier, T. Dean, and S. Hanks. Decision-theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, 11:1–94, 1999.
- [44] C. Boutilier, R. Reiter, M. Soutchanski, and S. Thrun. Decision-theoretic, high-level agent programming in the situation calculus. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*, pages 355 – 362, Austin, TX, 2000. AAAI Press / The MIT Press.
- [45] S. Bovet and R. Pfeifer. Emergence of delayed reward learning from sensorimotor coordinations. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Edmonton, Alberta, Canada, 2005.
- [46] K. Brammer and G. Siffing. *Kalman-Bucy Filters*. Artech House, Inc, Norwood, MA, 1989.

- [47] R. Brandman and M. E. Nelson. A simple model of long-term spike train regularization. *Neural Computation*, 14(7):1575–1597, 2002.
- [48] W. R. Breckon and M. K. Pidcock. Mathematical aspects of impedance imaging. *Clinical Physics and Physiological Measurement*, 8(4A):77, 1987.
- [49] M. D. Brooke, S. Hanley, and S. B. Laughlin. The scaling of eye size with body mass in birds. *Proceedings of the Royal Society of London Series B-Biological Sciences*, 266(1417):405–412, 1999.
- [50] B. R. Brown. Modeling an electrosensory landscape: behavioral and morphological optimization in elasmobranch prey capture. *Journal of Experimental Biology*, 205(7):999–1007, 2002.
- [51] H. Bruyninckx and J. de Schutter. The geometry of active sensing. Technical report, Katholieke Universiteit Leuven, Dept. Mechanical Engineering, 1999.
- [52] A. E. Bryson and Y.-C. Ho. *Applied optimal control : optimization, estimation, and control*. Blaisdell Pub. Co., Waltham, Mass, 1969.
- [53] T. H. Bullock. Electroreception. *Annual Review of Neuroscience*, 5:121–170, 1982.
- [54] W. Burgard, D. Fox, D. Hennig, and T. Schmidt. Estimating the absolute position of a mobile robot using position probability grids. In *Proc. of the Fourteenth National Conference on Artificial Intelligence*, pages 896–901, Portland, OR, 1996.
- [55] W. Burgard, D. Fox, and S. Thrun. Active mobile robot localization. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI)*, San Mateo, CA, 1997. Morgan Kaufmann.
- [56] P. Cain and S. Malwal. Landmark use and development of navigation behaviour in the weakly electric fish *Gnathonemus petersii* (Mormyridae; Teleostei). *Journal of Experimental Biology*, 205(24):3915–3923, 2002.
- [57] A. P. Calderon. On an inverse boundary value problem. In W. Meyer and M. Raupp, editors, *Seminar on Numerical Analysis and its Applications to Continuum Physics*, pages 65–73, Rio de Janeiro, 1980.
- [58] F. Callari. Active object recognition: Looking for differences. *International journal of computer vision*, 43(3):189–204, 2001.

- [59] F. G. Callari and F. P. Ferrie. Active recognition: Using uncertainty to reduce ambiguity. Technical Report TR-CIM-95-11, Centre for Intelligent Machines; McGill University, 1995.
- [60] A. Cameron and H. Durrantwhyte. A Bayesian-approach to optimal sensor placement. *International Journal of Robotics Research*, 9(5):70–88, 1990.
- [61] C. E. Carr, L. Maler, and E. Sas. Peripheral organization and central projections of the electrosensory nerves in gymnotiform fish. *Journal of Comparative Neurology*, 211(2):139–153, 1982.
- [62] A. R. Cassandra. Optimal policies for partially observable Markov decision processes. Technical Report Technical Report CS-94-14, Brown University, Department of Computer Science, 1994.
- [63] A. R. Cassandra. *Exact and approximate algorithms for partially observable Markov decision processes*. Phd, Brown University, 1998.
- [64] A. R. Cassandra, L. P. Kaelbling, and J. A. Kurien. Acting under uncertainty: Discrete Bayesian models for mobile robot navigation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1996.
- [65] A. R. Cassandra, L. P. Kaelbling, and M. L. Littman. Acting optimally in partially observable stochastic domains. In *Twelfth National Conference on Artificial Intelligence*, pages 1023–1028, Seattle, WA, 1994.
- [66] A. R. Cassandra, M. L. Littman, and N. L. Zhang. Incremental pruning: A simple, fast, exact algorithm for partially observable Markov decision processes. In D. G. Shenoy and P. Pundalik, editors, *Thirteenth Annual Conference on Uncertainty in Artificial Intelligence*, pages 54–61, 1997.
- [67] J. Castellanos. Multisensor fusion for simultaneous localization and map building. *IEEE Transactions on Robotics and Automation*, 17(6):908–914, 2001.
- [68] M. J. Chacron, A. Longtin, and L. Maler. Negative interspike interval correlations increase the neuronal capacity for encoding time-dependent stimuli. *Journal of Neuroscience*, 21(14):5328–5343, 2001.
- [69] M. J. Chacron, A. Longtin, M. St-Hilaire, and L. Maler. Suprathreshold stochastic firing dynamics with memory in p-type electroreceptors. *Physical Review Letters*, 85(7):1576–1579, 2000.

- [70] M. J. Chacron, L. Maler, and J. Bastian. Electroreceptor neuron dynamics shape information transmission. *Nature Neuroscience*, 8(5):673–678, 2005.
- [71] L. Chen, J. L. House, R. Krahe, and M. E. Nelson. Modeling signal and background components of electrosensory scenes. *Journal of Comparative Physiology A: Neuroethology Sensory Neural and Behavioral Physiology*, 191(4):331–345, 2005.
- [72] M. Cheney, D. Isaacson, and J. C. Newell. Electrical impedance tomography. *Siam Review*, 41(1):85–101, 1999.
- [73] G. Chetty and A. Russell. Electric field based obstacle avoidance. In *International Conference on Field and Service Robotics*, pages 378–382, Canberra, Australia, 1997.
- [74] G. Chetty and A. Russell. Electric field based sensing for underwater vehicle guidance. In *Proceedings of the 2nd International Conference on Bioelectromagnetism*, pages 139–140, 1998.
- [75] H. Choset and K. Nagatani. Topological simultaneous localization and mapping (SLAM): toward exact localization without explicit localization. *Robotics and Automation, IEEE Transactions on*, 17(2):125–137, 2001.
- [76] L. Chrisman. Planning for closed-loop execution using partially observed Markovian decision processes. In *AAAI Spring Symposium Series: Control of Selective Perception*, Stanford University, 1992.
- [77] T. H. Chung, V. Gupta, J. W. Burdick, and R. M. Murray. On a decentralized active sensing strategy using mobile sensor platforms in a network. In *Conference on Decision and Control*, 2004.
- [78] S. Collins, A. Ruina, R. Tedrake, and M. Wisse. Efficient bipedal robots based on passive-dynamic walkers. *Science*, 307(5712):1082–1085, 2005.
- [79] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley Series in Telecommunications. John Wiley and Sons, Inc, 1991.
- [80] W. G. R. Crampton. Electric signal design and habitat preferences in a species rich assemblage of gymnotiform fishes from the Upper Amazon Basin. *An. Acad. Bras. Ci.*, 70(4 Part 2):805–847, 1998.
- [81] F. Culkin and N. Smith. Determination of the concentration of potassium chloride solution having the same electrical conductivity, at 15°C and infinite frequency, as standard seawater of salinity 35.0000 ‰ (Chlorinity 19.37394 ‰). *IEEE J. Ocean. Eng.*, 5(1):22–23, 1980.

- [82] A. J. Davison. Modeling the world in real time: how robots engineer information. *Philosophical Transactions of the Royal Society of London Series A: Mathematical Physical and Engineering Sciences*, 361(1813):2875–2890, 2003.
- [83] J. de Geeter, J. de Schutter, H. Bruyninckx, H. van Brussel, and M. Decroton. Tolerance-weighted l-optimal experiment design for active sensing. In *1998 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Victoria, BC, Canada, 1998.
- [84] T. Dean and K. Kanazawa. Persistence and probabilistic projection. *IEEE Transactions on Systems Man and Cybernetics*, 19(3):574–585, 1989.
- [85] F. Dellaert, D. Fox, W. Burgard, and S. Thrun. Monte Carlo localization for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 1999.
- [86] F. Dickin and M. Wang. Electrical resistance tomography for process applications. *Measurement Science and Technology*, 7(3):247–260, 1996.
- [87] K. A. Dines and R. J. Lytle. Analysis of electrical-conductivity imaging. *Geophysics*, 46(7):1025–1036, 1981.
- [88] G. Dissanayake, S. B. Williams, H. Durrant-Whyte, and T. Bailey. Map management for efficient simultaneous localization and mapping (SLAM). *Autonomous Robots*, 12(3):267–286, 2002.
- [89] A. Doucet, S. Godsill, and C. Andrieu. On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, 10(3):197–208, 2000.
- [90] K. Doya. Complementary roles of basal ganglia and cerebellum in learning and motor control. *Curr Opin Neurobiol*, 10(6):732–9, 2000.
- [91] A. Elfes. Using occupancy grids for mobile robot perception and navigation. *Computer*, 22(6):46–57, 1989.
- [92] H. J. S. Feder, J. J. Leonard, and C. M. Smith. Adaptive mobile robot navigation and mapping. *International Journal of Robotics Research*, 18(7):650–668, 1999.
- [93] D. Fox, W. Burgard, F. Dellaert, and S. Thrun. Monte Carlo localization: efficient position estimation for mobile robots. In *Proc. of the Sixteenth National Conference on Artificial Intelligence*, Orlando, FL, 1999.

- [94] D. Fox, W. Burgard, and S. Thrun. Active Markov localization for mobile robots. *Robotics and Autonomous Systems*, 25(3-4):195–207, 1998.
- [95] D. Fox, J. Hightower, L. Liao, D. Schulz, and G. Borriello. Bayesian filtering for location estimation. *IEEE Pervasive Computing*, 2(3):24–33, 2003.
- [96] D. Fox, S. Thrun, W. Burgard, and F. Dellaert. Particle filters for mobile robot localization. In A. Doucet, N. de Freitas, and N. Gordon, editors, *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, 2000.
- [97] F. Gabbiani and W. Metzner. Encoding and processing of sensory information in neuronal spike trains. *Journal of Experimental Biology*, 202(10):1267–1279, 1999.
- [98] F. Gabbiani, W. Metzner, R. Wessel, and C. Koch. From stimulus encoding to feature extraction in weakly electric fish. *Nature*, 384(6609):564–567, 1996.
- [99] H. Geffner. Modelling intelligent behaviour: The Markov decision process approach. *Progress in Artificial Intelligence-Iberamia 98*, 1484:1–12, 1998.
- [100] J. P. Gonzalez and A. Stentz. Planning with uncertainty in position: an optimal and efficient planner. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Edmonton, Alberta, Canada, 2005.
- [101] M. S. Grewal and A. P. Andrews. Chapter 4: Linear optimal filters, predictors and smoothers. In *Kalman Filtering: Theory and Practice*. Prentice-Hall, Inc., Englewood Cliffs, NJ, 1993.
- [102] B. Grocholosky. *Information-theoretic control of multiple sensor platforms*. PhD thesis, The University of Sydney, 2002.
- [103] B. Grocholosky, H. Durrant-Whyte, and P. Gibbens. An information-theoretic approach to decentralized control of multiple autonomous flight vehicles. In *Sensor Fusion and Decentralized Control in Robotic Systems III*, Boston, MA, USA, 2000.
- [104] J. Gu, M. Meng, A. Cook, and P. X. Liu. Sensor fusion in mobile robot: some perspectives. In *Proceedings of the 4th World Congress on Intelligent Control and Automation, 2002.*, volume 2, pages 1194–1199, 2002.
- [105] S. Hagiwara and H. Morita. Coding mechanisms of electroreceptor fibers in some electric fish. *Journal of Neurophysiology*, 26(4):551, 1963.
- [106] M. J. Hartmann. Active sensing capabilities of the rat whisker system. *Autonomous Robots*, 11(3):249–254, 2001.

- [107] S. Haykin. Bayesian estimation and Kalman filtering: neural implications, 2004.
- [108] S. Healy. *Spatial Representation in Animals*. Oxford University Press, New York, 1998.
- [109] W. Heiligenberg and J. Bastian. The electric sense of weakly electric fish. *Annual Review of Physiology*, 46:561–583, 1984.
- [110] W. Heiligenberg and J. Dye. Labeling of electroreceptive afferents in a gymnotoid fish by intracellular injection of HRP - the mystery of multiple maps. *Journal of Comparative Physiology*, 148(3):287–296, 1982.
- [111] D. S. Holder. *Electrical Impedance Tomography: Methods, History and Applications*. Series in Medical Physics and Biomedical Engineering. Taylor and Francis, 2004.
- [112] P. Holmes, J. Jenkins, and N. E. Leonard. Dynamics of the Kirchhoff equations i: Coincident centers of gravity and buoyancy. *Physica D-Nonlinear Phenomena*, 118(3-4):311–342, 1998.
- [113] C. D. Hopkins. Stimulus filtering and electroreception - tuberous electroreceptors in 3 species of gymnotoid fish. *Journal of Comparative Physiology*, 111(2):171–207, 1976.
- [114] K. M. Horn, M. Pong, and A. R. Gibson. Discharge of inferior olive cells during reaching errors and perturbations. *Brain Research*, 996(2):148–158, 2004.
- [115] J. C. Houk and S. P. Wise. Distributed modular architectures linking basal ganglia, cerebellum, and cerebral cortex: their role in planning and controlling action. *Cereb Cortex*, 5(2):95–110, 1995.
- [116] G. Hovland. Control of sensory perception in a mobile navigation problem. *The International journal of robotics research*, 18(2):201–212, 1999.
- [117] X. M. Hu and T. Ersson. Active state estimation of nonlinear systems. *Automatica*, 40(12):2075–2082, 2004.
- [118] J. R. Inman, J. Ryu, and S. H. Ward. Resistivity inversion. *Geophysics*, 38(6):1088–1108, 1973.
- [119] M. Isard and A. Blake. Condensation - conditional density propagation for visual tracking. *International Journal of Computer Vision*, 29(1):5–28, 1998.

- [120] J. D. Jackson. *Classical Electrodynamics*. John Wiley and Sons, Inc, Hoboken, NJ, third edition, 1999.
- [121] L. P. Kaelbling, A. R. Cassandra, and J. A. Kurien. Acting under uncertainty: Discrete Bayesian models for mobile-robot navigation. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1996.
- [122] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1-2):99–134, 1998.
- [123] S. Kakei, D. S. Hoffman, and P. L. Strick. Sensorimotor transformations in cortical motor areas. *Neurosci Res*, 46(1):1–10, 2003.
- [124] M. Kalandros and L. Y. Pao. Controlling target estimate covariance in centralized multisensor systems. In *American Control Conference*, pages 2749–2753, Philadelphia, PA, 1998.
- [125] L. Kallenberg. Markov decision processes. Technical report, Leiden University, The Netherlands, May 28, 2000 2000.
- [126] R. E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME - Journal of Basic Engineering*, 82:35–45, 1960.
- [127] M. Kaneko, H. Maekawa, and K. Tanie. Active tactile sensing by robotic fingers based on minimum-external-sensor-realization. In *IEEE International Conference on Robotics and Automation*, Nice, France, 1992.
- [128] D. C. Knill and A. Pouget. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27(12):712–719, 2004.
- [129] E. I. Knudsen. Spatial aspects of the electric fields generated by weakly electric fish. *J. Comp. Physiol.*, 99(2):103–118, 1975.
- [130] S. Koenig and R. G. Simmons. Xavier: A robot navigation architecture based on partially observable Markov decision process models. In D. Kortenkamp, R. P. Bonasso, and R. Murphy, editors, *Artificial Intelligence and Mobile Robots*. American Association for Artificial Intelligence, Menlo Park, CA, 1998.
- [131] M. Konishi. Deciphering the brain’s codes. In *Neural codes and distributed representations: foundations of neural computation*. MIT Press, Cambridge, MA, USA, 1999.

- [132] K. P. Kording, S. P. Ku, and D. M. Wolpert. Bayesian integration in force estimation. *Journal of Neurophysiology*, 92(5):3161–3165, 2004.
- [133] K. P. Kording and D. M. Wolpert. Bayesian integration in sensorimotor learning. *Nature*, 427(6971):244–247, 2004.
- [134] R. Krahe, G. Kreiman, F. Gabbiani, C. Koch, and W. Metzner. Stimulus encoding and feature extraction by multiple sensory neurons. *Journal of Neuroscience*, 22(6):2374–2382, 2002.
- [135] C. Kreucher, K. Kastella, and A. O. Hero. Multi-target sensor management using alpha-divergence measures. *Information Processing in Sensor Networks, Proceedings*, 2634:209–222, 2003.
- [136] C. Kreucher, K. Kastella, and A. O. Hero. Sensor management using an active sensing approach. *Signal Processing*, 85(3):607–624, 2005.
- [137] S. Kristensen. Sensor planning with Bayesian decision theory. *Robotics and autonomous systems*, 19(3-4):273–286, 1997.
- [138] A. D. Kuo. The relative roles of feedforward and feedback in the control of rhythmic movements. *Motor Control*, 6(2):129–145, 2002.
- [139] H. Lamb. *Hydrodynamics*. Dover, New York, 6th edition, 1932.
- [140] M. J. Lannoo and S. J. Lannoo. Why do electric fishes swim backwards - an hypothesis based on gymnotiform foraging behavior interpreted through sensory constraints. *Environmental Biology of Fishes*, 36(2):157–165, 1993.
- [141] S. B. Laughlin. Towards the cost of seeing. In *Nervous Systems and Behaviour, 4th International Congress of Neuroethology*, Georg Thieme, Cambridge, 1995.
- [142] S. M. LaValle. *Planning Algorithms*. Cambridge University Press, <http://msl.cs.uiuc.edu/planning/>, 2006.
- [143] J. J. Leonard and H. F. Durrant-Whyte. Simultaneous map building and localization for an autonomous mobile robot. In *IEEE/RSJ International Workshop on Intelligent Robots and Systems*, pages 1442–1447, Osaka, Japan, 1991.
- [144] J. J. Leonard, H. F. Durrant-Whyte, and I. J. Cox. Dynamic map building for an autonomous mobile robot. *International Journal of Robotics Research*, 11(4):286–298, 1992.

- [145] J. J. Leonard, R. J. Rikoski, P. M. Newman, and M. Bosse. Mapping partially observable features from multiple uncertain vantage points. *International Journal of Robotics Research*, 21(10-11):943–975, 2002.
- [146] J. E. Lewis and L. Maler. Dynamics of electrosensory feedback: Short-term plasticity and inhibition in a parallel fiber pathway. *J. Neurophysiol.*, 88(4):1695–1706, 2002.
- [147] B. Lindner, M. J. Chacron, and A. Longtin. Integrate-and-fire neurons with threshold noise: A tractable model of how interspike interval correlations affect neuronal signal transmission. *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)*, 72(2):021911, 2005.
- [148] W. R. B. Lionheart. EIT reconstruction algorithms: pitfalls, challenges and recent developments. *Physiological Measurement*, 25(1):125–142, 2004.
- [149] H. W. Lissmann and K. E. Machin. The mechanism of object location in *Gymnarchus niloticus* and similar fish. *Journal of Experimental Biology*, 35(2):451–486, 1958.
- [150] M. L. Littman. *Algorithms for Sequential Decision Making*. PhD thesis, Brown University, 1996.
- [151] S. Liu and L. E. Holloway. Active sensing policies for stochastic systems. *IEEE Transactions on Automatic Control*, 47(2):373–377, 2002.
- [152] F. Lorussi, A. Marigo, and A. Bicchi. Optimal exploratory paths for a mobile rover. In *International Conference on Robotics and Automation*, pages 2078–2083, Seoul, Korea, 2001.
- [153] W. S. Lovejoy. Computationally feasible bounds for partially observed Markov decision-processes. *Operations Research*, 39(1):162–175, 1991.
- [154] M. A. MacIver. *The computational neuroethology of weakly electric fish: Body modeling, motion analysis, and sensory signal estimation*. PhD thesis, University of Illinois, Urbana-Champaign, 2001.
- [155] M. A. MacIver, E. Fontaine, and J. W. Burdick. Designing future underwater vehicles: Principles and mechanisms of the weakly electric fish. *IEEE Journal of Oceanic Engineering*, 29(3):651–659, 2004.
- [156] M. A. MacIver and M. E. Nelson. Towards a biorobotic electrosensory system. *Autonomous Robots*, 11(3):263–266, 2001.

- [157] M. A. MacIver, N. M. Sharabash, and M. E. Nelson. Prey-capture behavior in gymnotid electric fish: Motion analysis and effects of water conductivity. *Journal of Experimental Biology*, 204(3):543–557, 2001.
- [158] J. Manyika and H. Durrant-Whyte. *Data fusion and sensor management: A decentralized information-theoretic approach*. Ellis Horwood series in electrical and electronic engineering. Ellis Horwood, first edition, 1994.
- [159] P. S. Maybeck. Chapter 1; introduction. In *Stochastic models, estimation, and control; volume I*. Academic Press, New York, NY, 1979.
- [160] T. McGeer. Passive dynamic walking. *International Journal of Robotics Research*, 9(2):62–82, 1990.
- [161] N. Metropolis and S. Ulam. The Monte Carlo method. *Journal of the American Statistical Association*, 44:335–341, 1949.
- [162] W. Metzner, C. Koch, R. Wessel, and F. Gabbiani. Feature extraction by burst-like spike patterns in multiple sensory maps. *Journal of Neuroscience*, 18(6):2283–2300, 1998.
- [163] R. C. Miall, D. J. Weir, D. Wolpert, and J. Stein. Is the cerebellum a smith predictor? *J. of Motor Behavior*, 25(3):203–16, 1993.
- [164] F. A. Middleton and P. L. Strick. Basal ganglia and cerebellar loops: motor and cognitive circuits. *Brain Res Brain Res Rev*, 31(2-3):236–50, 2000.
- [165] L. Mihaylova, T. Lefebvre, H. Bruyninckx, K. Gadeyne, and J. de Schutter. Active sensing for robotics - a survey. Technical report, Katholieke Universiteit Leuven, Department of Mechanical Engineering, 2003.
- [166] L. E. Miller, R. N. Holdefer, and J. C. Houk. The role of the cerebellum in modulating voluntary limb movement commands. *Arch Ital Biol*, 140(3):175–83, 2002.
- [167] M. Miyazaki, D. Nozaki, and Y. Nakajima. Testing Bayesian models of human coincidence timing. *Journal of Neurophysiology*, 94(1):395–399, 2005.
- [168] G. E. Monahan. State-of-the-art - a survey of partially observable Markov decision-processes - theory, models, and algorithms. *Management Science*, 28(1):1–16, 1982.
- [169] M. Mundhenk, J. Goldsmith, C. Lusena, and E. Allender. Complexity of finite-horizon Markov decision process problems. *Journal of the Acm*, 47(4):681–720, 2000.

- [170] F. A. Mussa-Ivaldi and S. A. Solla. Neural primitives for motion control. *IEEE Journal of Oceanic Engineering*, 29(3):640–650, 2004.
- [171] M. E. Nelson. Electrophysiological models. In S. Koslow and S. Subramaniam, editors, *Databasing the Brain: From Data to Knowledge*, pages 285–301. Wiley, New York, 2005.
- [172] M. E. Nelson. Target detection, image analysis, and modeling. In T. Bullock, C. Hopkins, A. Popper, and R. Fay, editors, *Electroreception*, pages 290–317. Springer, New York, 2005.
- [173] M. E. Nelson and M. A. MacIver. Prey capture in the weakly electric fish *Apteronotus albifrons*: Sensory acquisition strategies and electrosensory consequences. *The journal of experimental biology*, 202(10):10, 1999.
- [174] M. E. Nelson and M. A. MacIver. Sensory acquisition in active sensing systems. *Journal of Comparative Physiology A: Neuroethology Sensory Neural and Behavioral Physiology*, 192(6):573–586, 2006.
- [175] M. E. Nelson, M. A. MacIver, and S. Coombs. Modeling electrosensory and mechanosensory images during the predatory behavior of weakly electric fish. *Brain Behavior and Evolution*, 59(4):199–210, 2002.
- [176] M. E. Nelson, Z. Xu, and J. R. Payne. Characterization and modeling of p-type electrosensory afferent responses to amplitude modulations in a wave-type electric fish. *Journal of Comparative Physiology A: Sensory Neural and Behavioral Physiology*, 181(5):532–544, 1997.
- [177] A. Y. Ng and M. I. Jordan. Pegasus: A policy search method for large MDPs and POMDPs. In *Uncertainty in Artificial Intelligence, Proceedings of the Sixteenth Conference*, 2000.
- [178] S. Park and A. D. Kuo. An optimal estimator model of multi-sensory processing in human postural control. *On the Convergence of Bio-Information-, Environmental-, Energy-, Space- and Nano-Technologies, Pts 1 and 2*, 277-279:148–154, 2005.
- [179] R. L. Parker. The inverse problem of resistivity sounding. *Geophysics*, 49(12):2143–2158, 1984.
- [180] M. G. Paulin. Evolution of the cerebellum as a neuronal machine for Bayesian state estimation. *Journal of Neural Engineering*, 2(3):S219–S234, 2005.

- [181] J. Pineau, G. Gordon, and S. Thrun. Point-based value iteration: An anytime algorithm for POMDPs. In *International Conference on Artificial Intelligence (IJCAI)*, 2003.
- [182] M. K. Pitt and N. Shephard. Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association*, 94(446):590–599, 1999.
- [183] B. Porat and A. Nehorai. Localizing vapor-emitting sources by moving sensors. *IEEE Transactions on Signal Processing*, 44(4):1018–1021, 1996.
- [184] J. Porta, B. Terwijn, and B. Krose. Efficient entropy-based action selection for appearance-based robot localization. In *IEEE International Conference on Robotics and Automation, 2003.*, volume 2, pages 2842– 2847, Taipei, Taiwan, 2003.
- [185] C. M. Postlethwaite, T. M. Psemeneke, J. Selimkhanov, M. Silber, and M. A. MacIver. Optimal movement in the prey capture behavior of weakly electric fish. *Submitted*, 2007.
- [186] B. Poucet. Spatial cognitive maps in animals - new hypotheses on their structure and neural mechanisms. *Psychological Review*, 100(2):163–182, 1993.
- [187] B. Rasnow. The effects of simple objects on the electric field of *Apteronotus*. *Journal of Comparative Physiology A: Sensory Neural and Behavioral Physiology*, 178(3):397–411, 1996.
- [188] B. Rasnow. Measuring and visualizing EOD fields. In *Communication in Fishes*. Science Publishers, Enfield, NH, USA, 2006.
- [189] R. Ratnam and M. E. Nelson. Nonrenewal statistics of electrosensory afferent spike trains: Implications for the detection of weak sensory signals. *Journal of Neuroscience*, 20(17):6672–6683, 2000.
- [190] K. S. Roberts. Robot active touch exploration: constraints and strategies. In *IEEE International Conference on Robotics and Automation, 1990. Proceedings*, Cincinnati, OH, USA, 1990.
- [191] D. A. Robinson. Electrical properties of metal microelectrodes. *Proceedings of the Institute of Electrical and Electronics Engineers*, 56(6):1065, 1968.
- [192] N. Roy, W. Burgard, D. Fox, and S. Thrun. Coastal navigation – mobile robot navigation with uncertainty in dynamic environments. In *IEEE International Conference on Robotics and Automation (ICRA)*, Detroit, MI, 1999.

- [193] S. Roy. Active recognition through next view planning: a survey. *Pattern Recognition*, 37(3):429–446, 2004.
- [194] D. B. Rubin. Using the sir algorithm to simulate posterior distributions. In K. M. Bernardo, D. V. Lindley, and A. Smith, editors, *Bayesian Statistics 3*. Oxford University Press, Oxford, UK, 1988.
- [195] G. J. Saulnier, R. S. Blue, J. C. Newell, D. Isaacson, and P. M. Edic. Electrical impedance tomography. *IEEE Signal Processing Magazine*, Nov.:31–43, 2001.
- [196] N. B. Sawtell, A. Williams, and C. C. Bell. From sparks to spikes: information processing in the electrosensory systems of fish. *Current Opinion in Neurobiology*, 15(4):437–443, 2005.
- [197] H. Scheich, T. H. Bullock, and R. H. Hamstra. Coding properties of 2 classes of afferent nerve fibers - high-frequency electroreceptors in electric fish, eigenmannia. *Journal of Neurophysiology*, 36(1):39–60, 1973.
- [198] J. L. Schneiter and T. B. Sheridan. An automated tactile sensing strategy for planar object recognition and localization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(8):775–786, 1990.
- [199] M. Schut, M. Wooldridge, and S. Parsons. On partially observable MDPs and BDI models. *Foundations and Applications of Multi-Agent Systems*, 2403:243–259, 2002.
- [200] C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423, 623–656, 1948.
- [201] C. A. Shumway. Multiple electrosensory maps in the medulla of weakly electric gymnotiform fish. I. Physiological differences. *Journal of Neuroscience*, 9(12):4388–4399, 1989.
- [202] C. A. Shumway. Multiple electrosensory maps in the medulla of weakly electric gymnotiform fish. II. Anatomical differences. *Journal of Neuroscience*, 9(12):4400–4415, 1989.
- [203] R. G. Simmons and S. Koenig. Probabilistic robot navigation in partially observable environments. In *International Joint Conference on Artificial Intelligence*, pages 1080–1087, 1995.
- [204] E. Simonin, J. Diard, and P. Bessiere. Learning Bayesian models of sensorimotor interaction: from random exploration toward the discovery of new behaviors. In

- IEEE/RSJ International Conference on Intelligent Robots and Systems*, Edmonton, Alberta, Canada, 2005.
- [205] O. Simsek, A. P. Wolfe, and B. A. G. Identifying useful subgoals in reinforcement learning by local graph partitioning. Technical Report UM-CS-2004-073, University of Massachusetts, November 2004.
- [206] K. D. Skeldon, L. M. Reid, V. McNally, B. Dougan, and C. Fulton. Physics of the theremin. *American Journal of Physics*, 66(11):945–955, 1998.
- [207] J. Smith, T. White, C. Dodge, J. Paradiso, N. Gershenfeld, and D. Allport. Electric field sensing for graphical interfaces. *Computer Graphics and Applications, IEEE*, 18(3):54–60, 1998.
- [208] J. R. Smith. *Electric field imaging*. PhD thesis, MIT, 1999.
- [209] R. C. Smith and P. Cheeseman. On the representation and estimation of spatial uncertainty. *The International Journal of Robotics Research*, 5(4):56–68, 1986.
- [210] R. C. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In I. J. Cox and G. T. Wilfong, editors, *Autonomous Robot Vehicles*, pages 167–193. Springer-Verlag, 1990.
- [211] J. B. Snyder, J. W. Burdick, M. E. Nelson, and M. A. MacIver. Matching of sensory and movement volumes in an active sensing system. *Submitted*, 2007.
- [212] J. R. Solberg. Sensor-based control in weakly electric fish with applications to robotics. Technical report, Northwestern University, 2005.
- [213] J. R. Solberg, K. M. Lynch, and M. A. MacIver. Active electrolocation for underwater target localization. *Submitted to the International Journal of Robotics Research*, 2007.
- [214] J. R. Solberg, K. M. Lynch, and M. A. MacIver. Robotic electrolocation: Active underwater object localization with electric fields. In *IEEE International Conference on Robotics and Automation*, Rome, Italy, 2007.
- [215] J. H. Solomon and M. J. Hartmann. Robotic whiskers used to sense features. *Nature*, 443(7111):525–525, 2006.
- [216] J. Spletzer. Dynamic sensor planning and control for optimally tracking targets. *The International journal of robotics research*, 22(1):7–20, 2003.

- [217] M. Srinivasan and A. Ruina. Computer optimization of a minimal biped model discovers walking and running. *Nature*, 439(7072):72–75, 2006.
- [218] C. Stachniss, D. Hahnel, and W. Burgard. Exploration with active loop-closing for fastSLAM. In *IEEE International Conference on Intelligent Robots and Systems*, Sendai, Japan, 2004.
- [219] S. A. Stansfield. A robotic perceptual system utilizing passive vision and active touch. *International Journal of Robotics Research*, 7(6):138–161, 1988.
- [220] R. F. Stengel. *Optimal Control and Estimation*. Dover Publications, Inc., New York, second edition, 1994.
- [221] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. Bradford Books. MIT Press, Cambridge, MA, 1998.
- [222] H. Takeda, C. Facchinetti, and J. C. Latombe. Planning the motions of a mobile robot in a sensory uncertainty field. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(10):1002–1017, 1994.
- [223] G. Theodorou and S. Mahadevan. Approximate planning with hierarchical partially observable Markov decision process models for robot navigation. In *International Conference on Robotics and Automation*, volume 2, pages 1347–1352, Washington, DC, USA, 2002.
- [224] L. S. Theremin and O. Petrishev. The design of a musical instrument based on cathode relays. *Leonardo Music Journal*, 6:49–50, 1996.
- [225] S. Thrun. The role of exploration in learning control. In D. White and D. Sofge, editors, *Handbook for Intelligent Control: Neural, Fuzzy and Adaptive Approaches*. Van Nostrand Reinhold, Florence, Kentucky, 1992.
- [226] S. Thrun. Exploration in active learning. In M. Arbib, editor, *Handbook of Brain and Cognitive Science*. MIT Press, 1995.
- [227] S. Thrun. Probabilistic algorithms in robotics. *AI Magazine*, 21(4):93–109, 2000.
- [228] S. Thrun. Particle filters in robotics. In *Proceedings of the 17th Annual Conference on Uncertainty in AI (UAI)*, 2002.
- [229] S. Thrun. Robotic mapping: A survey. Technical Report CMU-CS-02-111, Carnegie Mellon University, February 2002.

- [230] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. Intelligent Robotics and Autonomous Agents. The MIT Press, Cambridge, Mass, 2005.
- [231] S. Thrun, D. Fox, and W. Burgard. Probabilistic mapping of an environment by a mobile robot. In *Robotics and Automation, 1998 IEEE International Conference on*, volume 2, pages 1546–1551, 1998.
- [232] S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust Monte Carlo localization for mobile robots. *Artificial Intelligence*, 128(1-2):99–141, 2001.
- [233] S. Thrun, Y. F. Liu, D. Koller, A. Y. Ng, Z. Ghahramani, and H. Durrant-Whyte. Simultaneous localization and mapping with sparse extended information filters. *International Journal of Robotics Research*, 23(7-8):693–716, 2004.
- [234] S. Thrun, M. Montemerlo, D. Koller, B. Wegbreit, J. Nieto, and E. Nebot. Fast-SLAM: An efficient solution to the simultaneous localization and mapping problem with unknown data association. *Journal of Machine Learning Research*, 2004.
- [235] P. A. Tipler. *Physics for Scientists and Engineers*, volume 2. Worth Publishers, Inc, New York, NY, 3 edition, 1991.
- [236] E. Todorov. Optimality principles in sensorimotor control. *Nature Neuroscience*, 7(9):907–915, 2004.
- [237] E. Todorov and M. I. Jordan. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5(11):1226–1235, 2002.
- [238] E. Todorov and M. I. Jordan. Optimal feedback control as a theory of motor coordination: Supplementary notes. *Nature Neuroscience: supplement*, 5(11):1226–1235, 2002.
- [239] E. Todorov and M. I. Jordan. A minimal intervention principle for coordinated movement. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in neural information processing systems*, Vancouver, British Columbia, Canada, 2003. MIT Press.
- [240] E. C. Tolman. Cognitive maps in rats and men. *Psychological Review*, 55:189–208, 1932.
- [241] N. Trawny and T. Barfoot. Optimized motion strategies for cooperative localization of mobile robots. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 2004.

- [242] R. J. van Beers, P. Baraduc, and D. M. Wolpert. Role of uncertainty in sensorimotor control. *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences*, 357(1424):1137–1145, 2002.
- [243] R. R. D. vanSteveninck and S. B. Laughlin. The rate of information transfer at graded-potential synapses. *Nature*, 379(6566), 1996.
- [244] G. Venezian. On the resistance between 2 points on a grid. *American Journal of Physics*, 62(11):1000–1004, 1994.
- [245] G. von der Emde. Capacitance discrimination in electrolocating, weakly electric pulse fish. *Naturwissenschaften*, 80(5):231–233, 1993.
- [246] G. von der Emde. Active electrolocation of objects in weakly electric fish. *Journal of Experimental Biology*, 202(10):1205–1215, 1999.
- [247] B. Webb. Neural mechanisms for prediction: do insects have forward models? *Trends in Neurosciences*, 27(5):278–282, 2004.
- [248] J. G. Webster. *Electrical Impedance Tomography*. Adam Hilger, Bristol and New York, 1990.
- [249] G. Welch and G. Bishop. An introduction to the Kalman filter. Technical Report TR 95-041, University of North Carolina, March 11, 2002 2002.
- [250] R. Wessel, C. Koch, and F. Gabbiani. Coding of time-varying electric field amplitude modulations in a wave-type electric fish. *Journal of Neurophysiology*, 75(6):2280–2293, 1996.
- [251] P. Whaite and F. P. Ferrie. Autonomous exploration: driven by uncertainty. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(3):193–205, 1997.
- [252] A. G. Witney, S. J. Goodbody, and D. M. Wolpert. Predictive motor learning of temporal delays. *Journal of Neurophysiology*, 82(5):2039–2048, 1999.
- [253] D. Wolpert. Computational approaches to motor control. *Trends in Cognitive Science*, 1(6):209–16, 1997.
- [254] D. Wolpert, R. C. Miall, and M. Kawato. Internal models in the cerebellum. *Trends in Cognitive Science*, 2(9):338–47, 1998.
- [255] D. M. Wolpert, Z. Ghahramani, and M. I. Jordan. An internal model for sensorimotor integration. *Science*, 269(5232):1880–1882, 1995.

- [256] D. M. Wolpert and R. C. Miall. Forward models for physiological motor control. *Neural Netw*, 9(8):1265–1279, 1996.
- [257] H. Zhou and S. Sakane. Sensor planning for mobile robot localization based on probabilistic inference using Bayesian network. In *International Symposium on Assembly and Task Planning*, Fukuoka, Japan, 2001.
- [258] H. Zhou and S. Sakane. Sensor planning for mobile robot localization - a hierarchical approach using Bayesian network and particle filter -. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Edmonton, Alberta, Canada, 2005.

APPENDICES

APPENDIX A

A Simple Example Using a Particle Filter

Figure A.1 demonstrates using a particle filter to construct a belief of the location of a target. In A.1(A) it is only known that the target is somewhere within the 20 m by 20 m workspace. In this simple example it is assumed that even though the sensors are noisy, each observation actually corresponds to the noise-free observation. In A.1(B) the robot makes a single observation with its range sensor. The red dashed lines show the standard deviation of the observation mapped into the target state space. The green dashed line shows the locus of target configurations that would generate the same range observation. The particles in Figure A.1(B) represent the posterior belief after the range sensor observation is updated with the prior belief from Figure A.1(A). In Figure A.1(C) the robot uses a bearing measurement to update the belief. The bearing sensor's noise is 45° , as shown by the red dashed lines. Finally, Figure A.1(D) depicts the updated belief after a GPS observation.

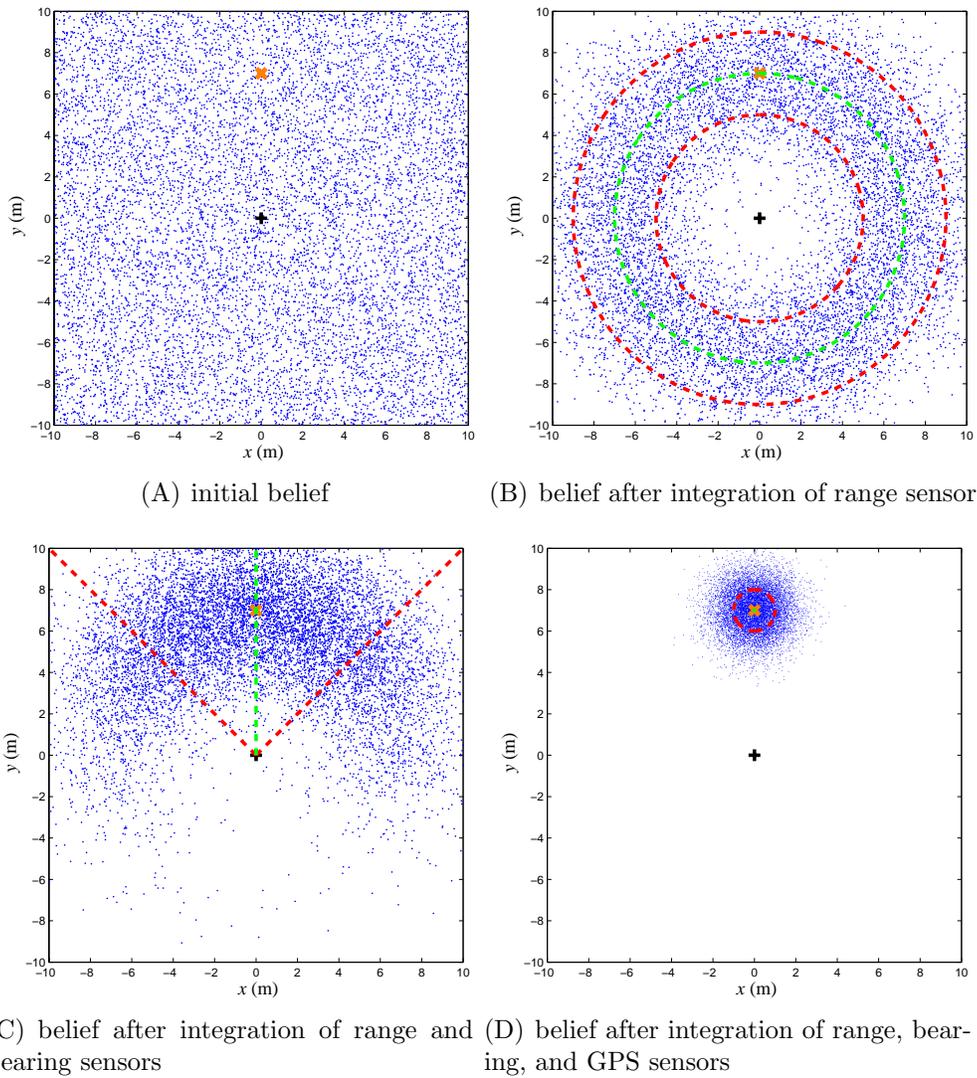


Figure A.1. A simple example using a particle filter. A robot is shown as a black cross located at $(x = 0, y = 0)$, and is trying to locate the target position $(x = 0, y = 7)$, which is shown as an orange “x.” The particles (i.e., target belief) are shown as blue dots.

APPENDIX B

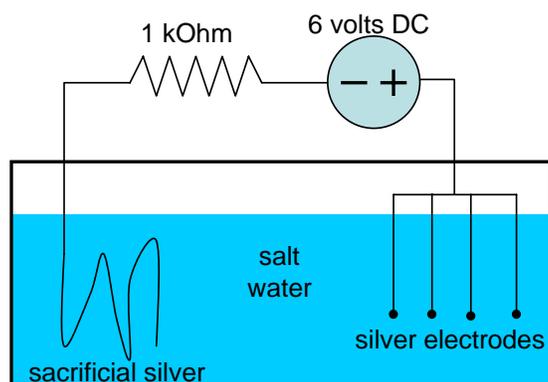
Electroplating the Electrodes

Figure B.1. Electrode electroplating setup.

To promote favorable impedance properties of the four electrodes the exposed tips of each electrode were first formed into a sphere of about 2 mm, and then each was electroplated with silver-chloride (AgCl). When heating silver wire to its melting point, intermolecular forces pull the molten silver into a sphere. Once the spheres are formed, the electrodes are placed in the electroplating setup as shown in Figure B.1. The electrodes and a bare strand of silver wire are placed in a bath of salt water (sodium-chloride). When the electrodes are held at a higher potential (6 volts used here) than the bare silver wire, the electrodes become the cathode and they become plated with the AgCl , and the bare wire becomes the sacrificial anode.

APPENDIX C

The Electric Field of the Electrosenster

The Electrosenster's electric field is generated by a pair of electrodes submerged in the fish tank. One electrode is tied to ground (0 volts). The voltage at the other electrode is a biphasic square wave (see Section 6.5). The RMS voltage of this square wave is 2 volts unless otherwise stated. In this appendix the spatial characteristics of the unperturbed electric field is examined.

The analysis of the electric field of the Electrosenster begins by making the following two assumptions:

- (1) presence of magnetic fields can be ignored
- (2) non-dispersive, isotropic media (ε and μ are time-independent scalars)

Both of these assumptions are accurate for the operating conditions of the Electrosenster.

If the assumption of time-invariant electric fields could be made, then the analysis could be simplified even further. To examine the feasibility of this assumption we begin with a temporal analysis of the square wave (The voltage drop across the electrodes of the Electrosenster is a square wave). The Fourier series of a square wave of period T is:

$$(C.1) \quad sq(t) = \sum_{n \in \{1, 3, 5, \dots, n_{max}\}} \frac{4}{\pi n} \sin\left(\frac{2\pi n t}{T}\right)$$

For the Electrosenster $T = 0.001$ sec = 1/(1 kHz). The bandwidth of the amplifier used by the Electrosenster to generate the wave is approximately 1 MHz, which means $n_{max} = 999$.

Because of assumption (1) from above, Faraday's law of induction (Equation 2.3) says the electric field is irrotational, i.e., $\nabla \times E = 0$. Thus, the electric field is the gradient of the voltage, V , i.e.,

$$(C.2) \quad E = -\nabla V$$

The total current density (typical units are A/m^2) is defined as the sum of the current density due to conduction, J_c plus the current density due to displacement, J_d :

$$(C.3) \quad J_{tot} = J_c + J_d = \sigma E - i\omega\varepsilon E$$

If it can be shown that the proportion of displacement current is insignificant relative to the amount of conduction current, then the effects of time-varying fields can be ignored.

The permittivity of water is approximately $\varepsilon = \varepsilon_r \varepsilon_0 = 7.1 \times 10^{-10} \text{ F/m}$, where ε_0 is the permittivity of free space $\varepsilon_0 = 1/(c^2 \mu_0) \approx 8.85 \times 10^{-12} \text{ F/m}$; ε_r is the relative permittivity of water (also called the dielectric constant). The conductivity of fresh and marine water are approximately $0.02 \text{ } \Omega^{-1}m^{-1}$ and $4 \text{ } \Omega^{-1}m^{-1}$, respectively. I begin with contribution of displacement current in fresh water:

$$(C.4) \quad J_{tot} = E(\sigma - i\omega\varepsilon E) = E(0.02\Omega^{-1}m^{-1} - i\omega \cdot 7.1 \times 10^{-10} \text{ Fm}^{-1})$$

where ω would be the frequency of excitation if it were a sine wave. But, since we are dealing with a pseudo-square wave (because 1 MHz bandwidth), all the relevant harmonics must be considered.

From Equations (C.1) and (C.2), we can show that the relative contribution of each odd harmonic to the total displacement current decays as $1/n$ for n equal to the odd integers. In addition, Equation (C.3) reveals that the displacement current associated with each frequency is linear with that frequency. If $\omega_0 = 2\pi f$, where $f = 1000$ Hz, then ω can be redefined as:

$$(C.5) \quad \omega = \sum_{n \in \{1, 3, 5, \dots, n_{max}=999\}} \frac{n \cdot \omega_0}{n} = 500\omega_0 = 2\pi 10^6$$

If the result of the Equation (C.5) is plugged into Equation (C.4), then it can be shown that about 18% of the total current is due to displacement current. But, when the same analysis is performed on water conductivities near sea water, less than 0.1% of the total current density is due to displacement current. Thus, displacement currents for the salt water environment can be ignored, but the displacement currents cannot be ignored for the fresh water conditions.

By invoking the assumptions from the beginning of this appendix, the electric field generated by two equal and opposite charges (i.e. dipole) in an infinite, isotropic medium can be found by Coulombs law with superposition:

$$(C.6) \quad E(\mathbf{x}) = \sum_i \frac{k \cdot q_i}{\|r_{i0}\|^2} \cdot \hat{r}_{i0}$$

where q_i is the magnitude of the charge (in Coulombs), \hat{r}_{i0} is a unit vector pointing from the charge to \mathbf{x} , $\|r_{i0}\|^2$ is the square of the magnitude of that vector; and k is the Coulomb constant ($8.988 \times 10^9 N \cdot m^2 / C^2$). But, the ElectroSensiter uses two electrodes of prescribed voltage to generate the electric field, not idealized point charges. The tips of the electrodes

are conducting spheres of radius, $r_e = 1$ mm, which would generate an identical electric field as a point charge of $q_+ = Vr_e/k = 1.11 \cdot 10^{-13}$ C, where $V = 1$ is the voltage of the electrode and k is the Coulomb constant. If q_- is equal in magnitude, but opposite in sign to q_+ , then a 2-D slice of the electric field generated by q_+ and q_- in an infinite, purely conductive medium (e.g. two point charges in the middle of the ocean) would look like Figure C.1.

Also shown in Figure C.1 are the iso-voltage contours in mV. Note that the contours are not at constant increments, and the voltage rapidly decays with the distance from the electrode. Figure C.1 is a model of an electric field in an infinite medium. But, the Electroster operates in a glass-walled fish tank of approximately 750 mm by 750 mm and 160 mm deep. This means that an electrode is usually no closer than 325 mm to any of the 4 vertical glass walls, but is always 80 mm from both the surface of the water and the glass floor. The effects of the vertical glass walls of the tank on the electric field near the middle of the tank (only the middle 80 mm by 80 mm section of the tank is used in experiments) is extremely small and can be completely ignored for our purposes. The effects of the floor and water surface as insulating boundaries are a little more noticeable. These boundaries compress the electric to stay within in the conducting medium, thus increasing the current density in the water since it cannot flow (as easily) through the air or glass bottom. In reality, these boundaries 80 mm above and below the electrodes have virtually no effect on the qualitative structure of the electric field depicted in Figure C.1. But, the magnitude of the electric field is slightly greater ($< 1\%$) than if these horizontal non-conducting boundaries were not present.

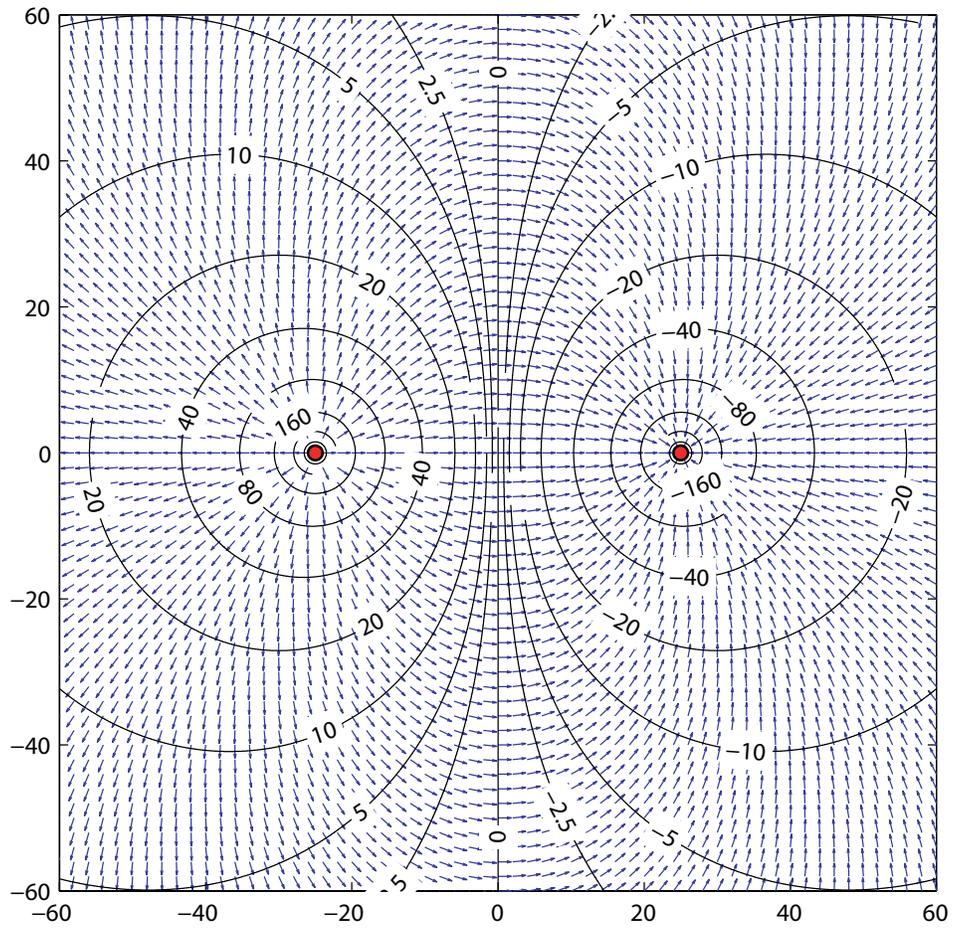


Figure C.1. Electric field of the Electrosonster. Arrows show the direction of the electric field. Contours are at constant voltages.

APPENDIX D

Some Examples of Measured EEVs of Cylinders

The purpose of this appendix is to expose the reader to EEVs of targets other than spheres. Early on in this thesis the commitment was made to model observations via Rasnow's model of simple objects (i.e., Equation (2.13)). Unfortunately, this model is invalid for cylinders. Thus, we forgo any formal analysis on the structure of the measured EEVs presented in this appendix.

We consider eleven different EEVs of cylinders (all approximately 80 mm in height) under various conditions. The 80 mm heights of the cylinders were chosen such that the bottom of the emitting and sensing electrodes were just able to pass over them without touching. See Section 8.2.1.1 on page 189 for the protocol for constructing the EEVs of cylinders using the Electrosenster. The EEVs are grouped into three sets:

- (1) The three experimental conditions for electrolocation experimental set 1 (Figure D.1).
- (2) Plastic cylinders under 4 different experimental conditions (Figure D.2).
- (3) Brass cylinders of four different diameters (Figure D.3).

In the following EEV figures, “fresh” water refers to $160 \mu\text{S}/\text{cm}$ electrical conductivity water, and “salt water” is $56,000 \mu\text{S}/\text{cm}$ water. “Plastic” refers to Delrin plastic. Red squares indicate the location of the electric field dipoles. Green diamonds indicate the

location of the voltage sensors. The color map goes from white for +10 V to gray for intermediate values to black for -10 V.

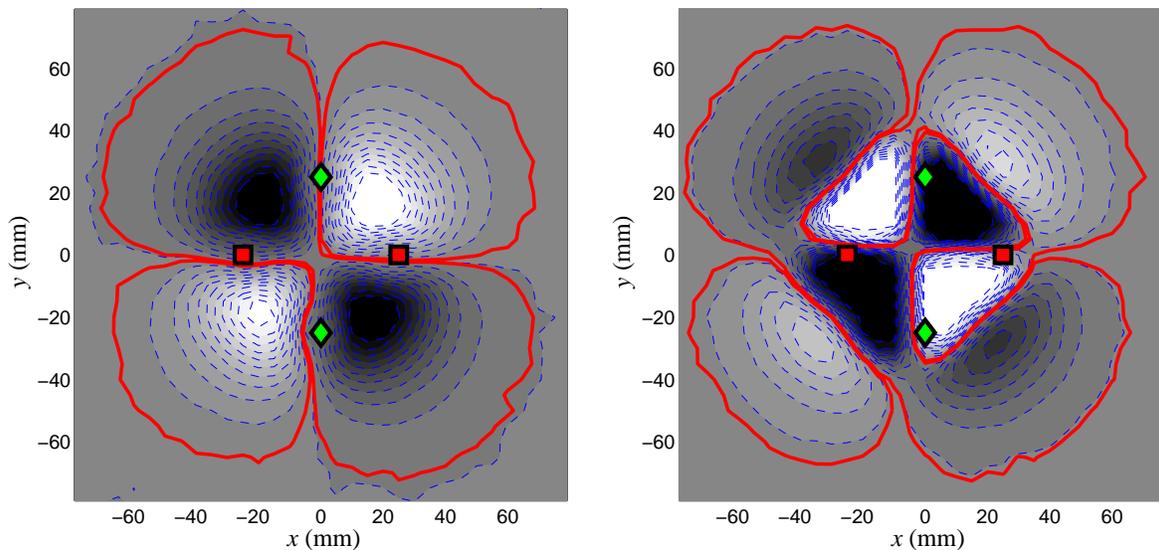
Figure D.1 depicts the EEVs for the three conditions in experimental set 1. The results of the corresponding electrolocation experiments are in Section 8.2. The three experimental conditions are:

- (1) 12.7 mm (1/2") diameter aluminum (electrical conductor) cylinder in freshwater (Figure D.1(A)).
- (2) 25.4 mm (1") diameter aluminum (electrical conductor) cylinder in freshwater (Figure D.1(B)).
- (3) 25.4 mm (1") diameter plastic (electrical insulator) cylinder in freshwater (Figure D.1(C)).

Figure D.2 depicts four more EEVs of cylinders (not used in electrolocation experiments).

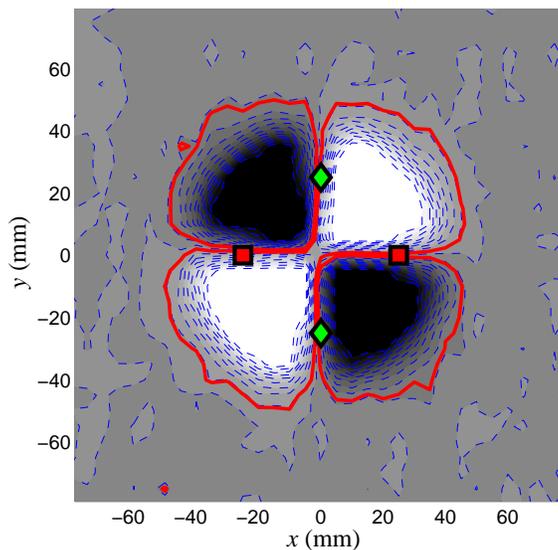
Almost all of the EEVs shown in this document prior to this appendix have been for spheres (e.g., the modeled EEVs in Chapter 2 and the measured EEVs in Chapters 7 and 8). Figures D.1 and D.2 give us a glimpse of what EEVs of cylinders under various conditions look like. They look similar to the EEVs of spheres with two major differences, which are described below.

Probably the most obvious difference is the extra set of lobes seen in Figure D.1(B). Since it was introduced in Chapter 2, we have relied on Rasnow's model of simple objects (i.e., Equation (2.13) on page 48) to model electrosensory observations. This simple model does not apply to cylinders. Thus, we cannot use the simple model to investigate these extra lobes. The numerical methods developed in Section 2.1 failed to explain the extra set of lobes due to the 2-D simplification that was enforced in the model formulation



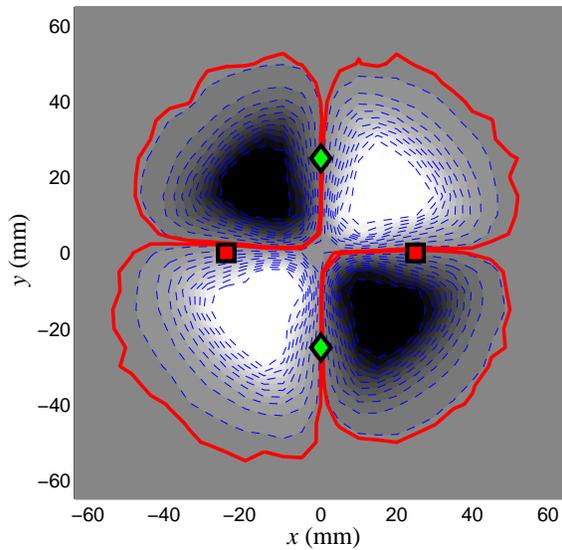
(A) 1/2" dia. AL cylinder in fresh water

(B) 1" dia. AL cylinder in fresh water

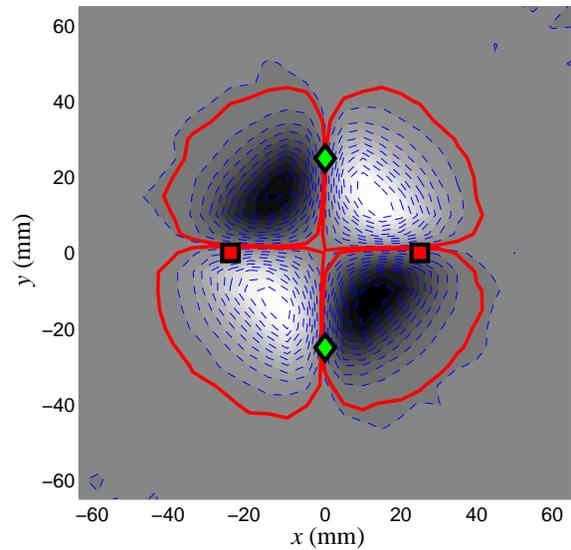


(C) 1/2" dia. plastic cylinder in fresh water

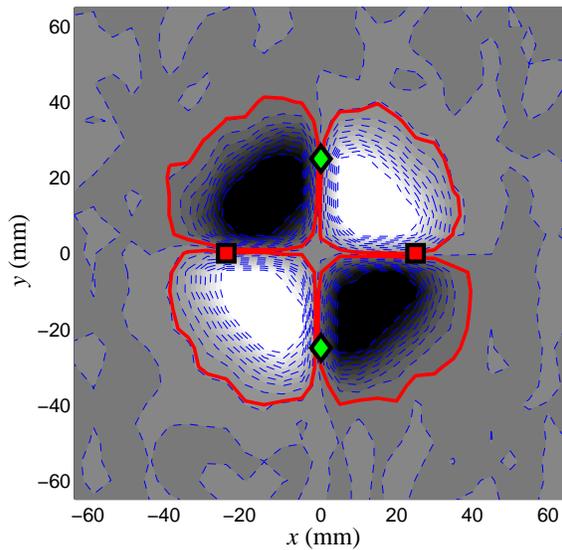
Figure D.1. The EEVs for the three conditions in experimental set 1. Blue dashed lines are at 1 V increments. The solid red contours represent an observation that results in detection of the target with 95% confidence (see Section 7.2 for details).



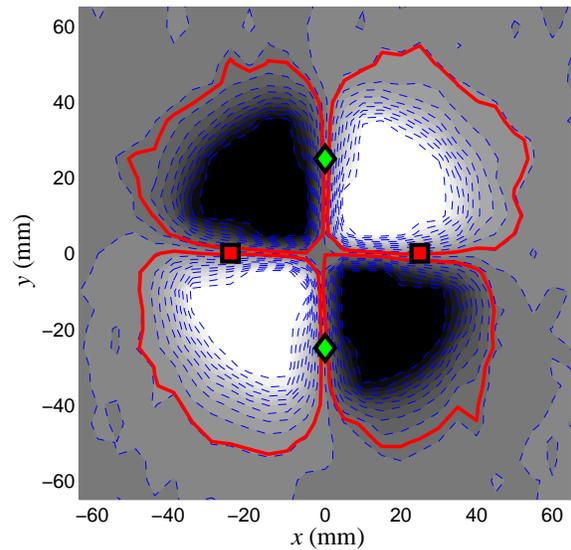
(A) 1" dia plastic cylinder in fresh water



(B) 1/2" dia plastic cylinder in fresh water



(C) 1/2" dia plastic cylinder in salt water



(D) 1" dia plastic cylinder in salt water

Figure D.2. EEVs for plastic cylinders under 4 different experimental conditions. Blue dashed lines are at 1 V increments. The solid red contours represent an observation that results in detection of the target with 95% confidence (see Section 7.2 for details).

(assumption (4) from Section 2.1). A 3-D analysis using the same methodology could be performed with access to the appropriate software package such as COMSOL Multiphysics (COMSOL, Inc., Burlington, MA). This approach is beyond the scope of this thesis.

These extra sets of lobes were only observed in conductors and not in the smaller cylinders—it is present in the one inch diameter cylinder but not the half inch cylinder. Thus, there is a significant qualitative change between the EEVs of small and large diameter conductive cylinders. In Figure D.3 EEVs of brass (electrical conductor) cylinders of intermediate diameters ($5/8''$ and $3/4''$) were constructed. Figure D.3(A) is similar to Figure D.1(A) and Figure D.3(D) is similar to Figure D.1(B). In Figure D.3(B) we see the extra set of lobes just starting to form, and in Figure D.3(C) they are well developed.

In Figures D.3(C) and D.3(D) the orange ‘x’s correspond to target locations on the boundary between lobes (these locations render an observation of $w = 0$ volts). The blue dashed circles in these two subfigures represent the edge of the cylindrical target. The boundary between the near lobes and far lobes is nearly a straight line with the endpoints corresponding to target locations when the edge of the cylinder just touches either an emitter or sensor. This does not provide an explanation, merely an observation.

The second major difference between the EEVs of cylinders and spheres is that the polarity of the lobes does not change when the conductivity of the target changes. As predicted by Rasnow’s model (Equation (2.13)) the sign of the observation depends on the relative conductivity between the target and water (specifically, it depends on the sign of the contrast factor, χ). This effect is seen in the EEVs of spheres, but it is not seen in the EEVs of cylinders (e.g. compare Figures D.1(A) and D.1(C)). Thus, the sign of the lobes is the same for metal and plastic cylindrical targets.

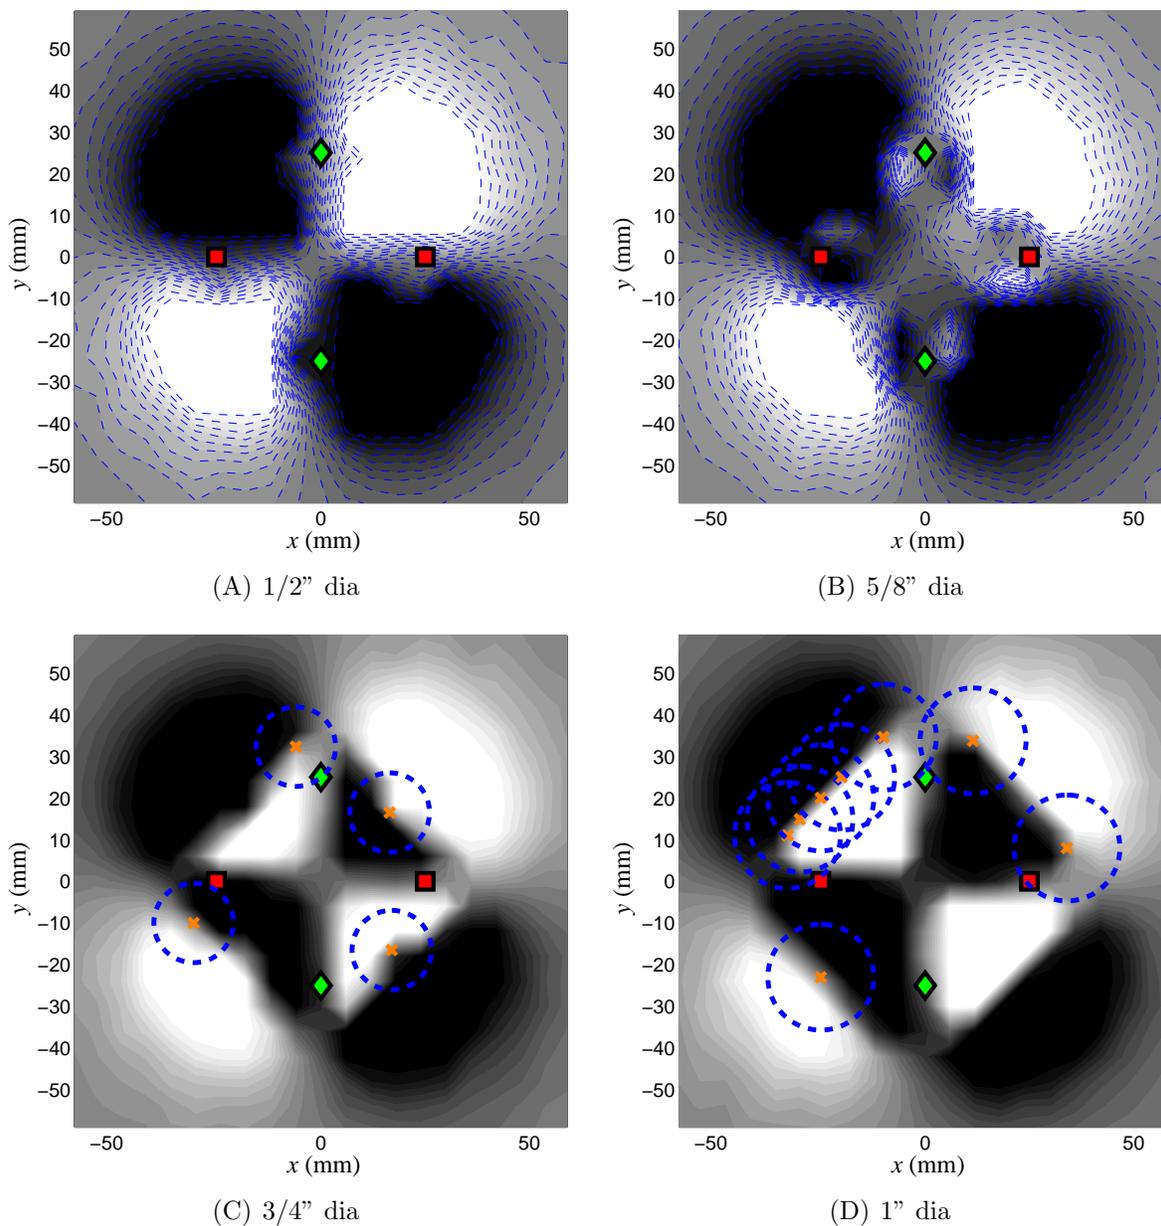


Figure D.3. EEVs for brass cylinders of four different diameters in fresh water. (A) and (B): Blue dashed lines are at 1 V increments. (C) and (D): The orange “x”s are example target locations (center of cylinder) on the boundary between the two lobes. The blue dashed circles are the edges of the corresponding cylinders projected into the 2-D plane. The dashed-line contours have been removed for clarity.

APPENDIX E

Closed-Form Solution for Voltage Perturbation

Here, we derive Equation (7.1), which is also displayed at the end of this appendix. Recall Equation (2.13) (Rasnow's model from page 48) is an analytical solution to the observed voltage perturbation, δV , as a function of the relative target position, \mathbf{x}_t , target radius, a , the unperturbed electric field at the center of the target, \mathbf{E}_f , and the contrast factor, χ . We wish to find a closed form solution of Equation (2.13), which is an explicit function of the experimental parameters.

If we assume the electric field is generated by two point charges of $+q$ and $-q$ located at positions \mathbf{x}_p and \mathbf{x}_n , respectively, then the electric field at position \mathbf{x} can be computed as

$$(E.1) \quad \mathbf{E}(\mathbf{x}) = \frac{qk}{|\mathbf{x} - \mathbf{x}_p|^3}(\mathbf{x} - \mathbf{x}_p) - \frac{qk}{|\mathbf{x} - \mathbf{x}_n|^3}(\mathbf{x} - \mathbf{x}_n)$$

where k is Coulomb's constant. It is usually inconvenient to specify the source of the electric field in units of charge, but if the emitters are assumed to be spherical conductors, then the equivalent charge for a emitter of radius E and voltage V_e is $q_{eq} = V_e E/k$.

Let us consider a simple emitter/detector layout that consists of a single voltage detector located at the origin and the two emitters separated by the distance L and at positions $\mathbf{x}_p = \{-L/2, 0, 0\}$ and $\mathbf{x}_n = \{L/2, 0, 0\}$. We wish to determine the detection

distance, R , of a spherical target of radius A along a 45° line passing through the origin (This setup is shown in Figure 7.3). We first define the dimensionless parameters $a = A/L$, $e = E/L$, $r = R/L$, and $\delta v = \delta V/V_e$. We then substitute q_{eq} into Equation (E.1) and then substitute that equation of the electric field into Equation (2.13) to arrive at a closed-form solution of the dimensionless voltage perturbation as a function of the dimensionless target radius (a), contrast factor (χ), emitter radius (e), and distance to the center of target from the detector (r) (this is Equation (7.1)):

$$\delta v(a, \chi, e, r) = -2 a^3 \chi e \frac{N_1(r) + N_2(r)}{D(r)}$$

where,

$$\begin{aligned} N_1(r) &= \left(\sqrt{2} - 4\sqrt{2}r^2 + 16r^3 \right) \sqrt{1 - 2\sqrt{2}r + 4r^2} \\ N_2(r) &= \left(\sqrt{2} - 4\sqrt{2}r^2 - 16r^3 \right) \sqrt{1 + 2\sqrt{2}r + 4r^2} \\ D(r) &= r^2 \left(1 - 2\sqrt{2}r + 4r^2 \right)^{\frac{3}{2}} \left(1 + 2\sqrt{2}r + 4r^2 \right)^{\frac{3}{2}}. \end{aligned}$$

Likewise, we can perform the same analysis for the emitter/detector configuration used in the experimental setup. The location of the emitters are the same as above, but the locations of the detectors are now $(x = 0, y = \pm L/2, z = 0)$. In the experimental setup the location of the center of target was in the plane $z = -(E + O)$ relative to the plane containing the emitters and detectors, where O is an additional offset that enabled the electrodes to pass safely over the target. In our setup $O = 2$ mm. When these new values for the positions of the detectors and target are substituted into Equations (E.1)

and (2.13) the closed-form solution of dimensionless voltage perturbation can be obtained.

The length of this equation precludes its inclusion here.

APPENDIX F

More Complex Models of the Electrosensory System of Weakly Electric Fish

In this document, only very simplified models of the weakly electric fish's sensing and motion have been considered. For example the electroreceptors of the fish were modeled as simple volt sensors. The signal processing pathways were modeled as Bayes filters. In this appendix more complex models that are grounded in the morphology, anatomy, and neurobiology of the weakly electric fish are considered. Many aspects of the weakly electric fish have been thoroughly studied, so many different models for different parts of the weakly electric fish are available (See [171] and [172] for an introduction to relevant models of the weakly electric fish.) In this section we discuss some relevant models that may help give some insight to how such an organism could perform electrolocation of prey.

The electrolocation task can be subdivided into three different phases: detection, characterization, and localization. The *detection* phase is simply determining whether or not a potential prey is in the vicinity. The *characterization* phase entails determining if the potential target is a prey. The *localization* phase will be the main focus of this appendix. This is the phase where the fish determines the position of the prey [246].

F.1. Modeling the Weakly Electric Fish

The South American black ghost knifefish (*Apteronotus albifrons*) captures prey in the dark using its active electric sense [140, 157] and also possibly its ability to detect the

bioelectric fields around prey (passive electric sense) [157] and fluid disturbances due to the prey (mechanosense) [175] (See [109] and [53] for an overview of the electric sense in weakly electric fish.). In *active electrosense*, black ghost knifefish sense their surroundings using a weak, self-generated electric field (≈ 1 mV/cm near the body). Nearby objects that differ in electrical conductivity from the surrounding water create localized voltage perturbations across the skin that are sensed by about 14,000 tuberous electroreceptor organs scattered over the body surface [61]. In *passive electrosense*, nearby objects that emit a bioelectric field stimulate about 700 ampullary electroreceptor organs [61]. In *mechanosense*, flow accelerations caused by the prey's movement are detected by about 200 canal neuromasts [175], and flow velocity caused by relative motion between the fish and its medium are detected by about 100 superficial neuromasts [61].

Do to the extreme complexities of vertebrate systems, only a few components of the weakly electric fish will be examined here, and we will mainly focus on the sensory acquisition and processing of electrosensory information from the active electrosense in the weakly electric fish. We consider models of the electroreceptor afferent¹, specifically how it transduces transdermal potentials (voltages across the skin of the fish) into *spike trains* (the communication protocol for the nervous system) to be used by the central nervous system (CNS) of the weakly electric fish. We also begin to look at how the CNS processes the electrosensory information. The CNS uses spike trains to encode information to be sent to neurons within the CNS. Nerves transmit voltages and the “spike” refers to the impulse of voltage in the nerve. Spike trains are not unlike binary voltages in computer circuitry (e.g., both represent information as streams of binary voltages).

¹In the nervous system, afferent nerves carry nerve impulses *towards* the brain. “Afferent” can also refer to a sensory signal that goes *towards* the brain.

F.2. Models of Neural Coding in Primary Electrosensory Afferents

Of the three sensing modalities used by the black ghost knifefish to localize prey in the dark, only the active electrosense modality is considered here. This is believed to be the most dominant modality of the prey capture task, and it is the most widely studied. The active electrosensory system is composed of the electric organ, which emits the electric field (referred to as the electric organ discharge, or EOD), and the *tuberosus* organs which are the electroreceptor organs for this electrosensory modality (See [8] for modeling the electric organ discharge.). Tuberosus organs are sensitive to high-frequency (typically 0.1 to 1 kHz) electric fields, and typically tuned to frequencies near the peak of the EOD spectrum.

For the nervous system to carry out the computations necessary to support electrolocation, information about the transdermal voltage (transdermal voltage is the voltage across the skin) patterns on the skin must be converted into a neural representation. The information needed for subsequent neural processing is encoded in the spike train data. The coding strategies implemented by these organs are rather sophisticated, involving various forms of input filtering and noise suppression during the encoding process.

The black ghost knifefish has two subtypes of tuberosus receptor units [197]. One type conveys information about stimulus amplitude, while the other type conveys information about stimulus timing. The ongoing quasi-sinusoidal oscillation of the fish's EOD provides a natural "clock" for the system. In wave-type fish, tuberosus afferents fire at most one spike per clock cycle. Amplitude information is conveyed by a spike probability code [23, 197]. When no target object is present, probability-coding (P-type) units fire a spike

with some baseline probability per EOD cycle, typically around 0.3. When a conducting object approaches the receptor organ, it causes an increase in the local transdermal voltage and an increase in the per-cycle firing probability of the P unit. Similarly, when a nonconducting object approaches the receptor, the local transdermal potential and the P-unit firing probability will decrease. The interspike interval distribution is irregular, so P-type units are sometimes described as “sputtering.”

The baseline EOD oscillation serves as a carrier signal and a target object induces an amplitude modulation (AM) of this carrier signal. P-type units are tuned to the carrier frequency of the fish’s own EOD [113]. The frequency content of an AM signal induced by a target is related to its speed and distance. Nearby, fast-moving targets will cause higher-frequency AMs compared to distant, slow-moving targets. *P-type units act as high-pass filters* in the AM frequency domain, giving the strongest response to fast components of the AM signals (See [20], [250], and [176] for details.).

Encoding of transdermal potentials into spike trains is now examined. In [47] a simple model of spike generation is described that gives rise to negative correlations in the interspike interval (ISI) sequence and leads to long-term spike train regularization (this model was basically derived from more complex models of [69] and [68]). Such regularizing effects have been observed in the spike trains of electrosensory afferent nerve fibers and can lead to dramatic improvements in the detectability of weak signals encoded in the spike train data [189]. The model from [47] can be described by four update rules, which are evaluated in the following order at each time step n :

$$(F.1) \quad v[n] = c i[n] + \delta[n]$$

$$(F.2) \quad \theta = \theta[n - 1] - (b/a)$$

$$(F.3) \quad s[n] = H(v[n] - \theta[n]) = \begin{cases} 1 & \text{if } v[n] \geq \theta[n] \\ 0 & \text{otherwise} \end{cases}$$

$$(F.4) \quad \theta[n] = \theta[n] + b s[n] = \begin{cases} \theta[n] + b & \text{if } s[n] = 1 \\ \theta[n] & \text{otherwise} \end{cases}$$

where H is the Heaviside function, defined as $H(x) = 0$ for $x < 0$ and $H(x) = 1$ for $x \geq 0$. The voltage v is the product of the input resistance c and the instantaneous input current i plus random noise δ , where δ is zero-mean Gaussian noise with variance σ^2 . When the voltage v rises above a threshold level θ , a spike is generated ($s = 1$), and the threshold level is elevated by an amount b . The threshold subsequently decays linearly with a slope of $-b/a$ until the next spike is generated.

Current, $i[n]$, from the EOD reaches the skin of the fish and is impeded by a resistance, c , and noise, $\delta[n]$, is superimposed on top the resultant voltage (This is Equation F.1). The voltage experienced by the fish's electroreceptor is $v[n]$. This voltage must be transduced to a spike train so the CNS can process the electrosensory information. At each "clock cycle" of the fish (i.e. at each time step n) the electroreceptor organ must choose whether

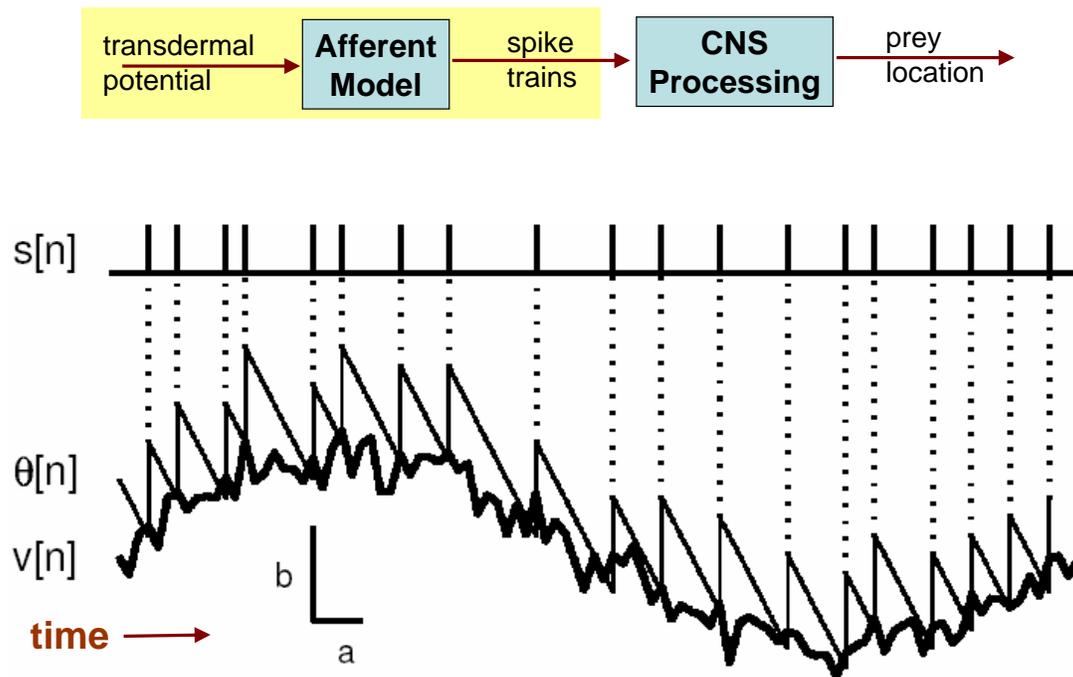


Figure F.1. Modeling the transformation from transdermal potential to spike trains in the fish. This is an example of Equations F.1 through F.4 in action. The input signal, $i[n]$, is a sinusoid with a period of 100 time steps: $i[n] = \sin(2\pi n/100)$. After implementation of Equation F.1, $v[n]$ takes the form of the thick (noisy) line in the figure. A spike occurs (i.e. $s[n] = 1$) whenever the voltage, $v[n]$ meets the threshold, $\theta[n]$. After a spike does occur the threshold is bumped back up by an amount dictated by Equation F.4 and then proceeds to decay as prescribed by Equation F.2. This figure is taken from of [47].

or not to propagate a spike. This choice is dictated by Equation F.3. If a spike does occur (i.e. $s[n] = 1$), then the threshold, $\theta[n]$, for subsequent spikes is bumped up (see Equation F.4), but the threshold will decay at each time step (see Equation F.2). Figure F.1 demonstrates the implementation of these equations for a given input function, $i[n]$.

F.3. Electrosensory Processing Models

Ultimately, weakly electric fish must be able to extract and interpret useful information from these spike trains. Once the electrosensory information is transduced into spike trains, it is sent upstream to the hindbrain electrosensory lateral line lobe (ELL) [34]. The encoding properties of these afferent signals have been the subject of much study [105, 22, 97]. Of particular interest is the research on the information content of the encoded spike trains [173, 189, 68, 70].

The ELL is the sole recipient of all the primary afferents of the electrosensory system. This means that all the spike trains that propagate from the electrosensory organs first pass through the ELL before anywhere else in the CNS. We are particularly interested in the three separate somatotopically organized² maps for the tuberous electrosensory system [201, 202]. See Figure F.2 for a schematic of the ELL.

Signals from the tuberous electroreceptors trifurcate, and each somatotopic map in the ELL gets its own copy of the spike train. Each of the maps in the ELL has a different rate of convergence of the spike trains. The rate of convergence dictates spatial resolution and the sensitivity of the map. For example, if a map has high convergence (many

²*Somatotopically organized* means that there is a topological (homeomorphic) relationship between sensory receptive fields and their corresponding activation areas in the brain. So, electroreceptors that are near each other on the skin will have sensitive neurons in the brain that are near each other (the homunculus is an example of a somatonic map).

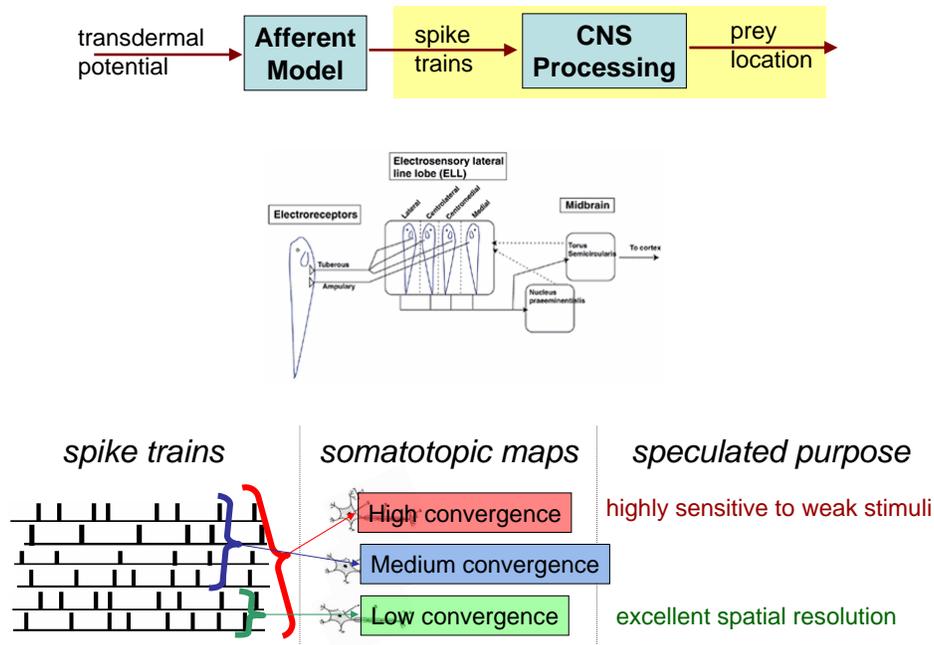


Figure F.2. The transformation of spike trains to prey location. This figure depicts the ELL with multi-resolution maps. Part of this figure is taken from [110].

spike trains going to a single nucleus of neurons), then this nucleus will have information from many electroreceptors, so it can easily determine if something has perturbed the electric field (i.e. high sensitivity); but, because of high convergence, it is poor at spatial discrimination of the target because it cannot differentiate which spike train came from which electroreceptor.

F.4. Prey Localization

Previously, Bayes filters have been used to model the signal processing in the fish, but we really do not know how the fish encodes the spike trains into the location of the prey in 3-D space. Information theory, statistics, and digital filtering techniques have

been applied to spike trains in an attempt to understand the information contained in them [41], [147], [134], [147], [68]. There has been limited progress with this due to the extreme complexities of a typical nervous system. But electrosensory systems provide a great platform for understanding principles of neural encoding [131]. Behaviorally relevant sensory stimuli are well understood and easy to control experimentally. And, of course, there is a wealth of knowledge on sensory acquisition and processing in weakly electric fish. So, many have used electrosensory models to decipher neural codes [162], [98]. See [196] for an overview of this work.

APPENDIX G

Simple Models of Sensorimotor Transformations in Fish and Robots

By “Simple” model we mean that the differential or difference equations governing the system dynamics and the observations need to be linear and time invariant (LTI). “Simple” also means there are constraints on the nature of the uncertainties. The most common assumption is Gaussian noise on a sensor and Gaussian disturbances on the system dynamics. While these constraints seem very restricting, many systems have been successfully modeled like this. And, these assumptions greatly reduce compute time.

G.1. The Kalman filter

One of the most popular state estimators is the Kalman filter [126]. It is a simple algorithm, but can deliver excellent results under the appropriate conditions. For introductions to Kalman filtering see [249], [46], [101], or [159]. The system dynamics and observation map must fit the following linear stochastic difference equations:

$$(G.1) \quad x_k = Ax_{k-1} + Bu_k + \delta_{k-1}$$

$$(G.2) \quad w_k = Hx_k + \nu_k$$

where δ_{k-1} and ν_k are normally distributed random variables with means of zero and covariance matrices of Q and R , respectively. Equations G.1 and G.2 give us our first glimpse of *feedforward model* (i.e. predictive model) vs. *feedback* (i.e. observed). Equation G.1 makes a *prediction* one timestep into the future on what the state, x_k will be, given the current state, x_{k-1} and the input signal, u_k . In fact, since δ is guaranteed to be zero-mean and normally distributed, the *expected value* of x_k , given x_{k-1} and u_k is given by Equation G.1 with δ_{k-1} set to zero (because the expected value of δ is zero; i.e. $E[\delta] = 0$). So, with an accurate model of the plant (The matrices A and B are linear models of the plant) and no disturbances (i.e. $\delta = 0$ for all time), there is no need to observe the system because all states can be known for all time *a priori*.

If Equation G.1 represents the feedforward (Specifically, this is called a *forward model*), then Equation G.2 models the feedback in the system. H transforms the states of the system into the expected observations (Again, we can say “expected” because ν is characterized as Gaussian and unbiased). So, if we wanted to estimate the state, $x_{i, est}$ from a single observation, w_i , we could simply say $x_{i, est} = H^{-1}w_i$, where $x_{i, est}$ is the expected value and the most likely state, given this single observation.

The Kalman filter is a state estimator. This means it forms a belief of the state. Previously, we’ve represented beliefs as probability distributions over all possible states. The Kalman filter does the same thing, but because we are only dealing with Gaussians, the belief can be completely represented as the mean, \hat{x} , and the covariance, P .

The Kalman filter is special case of a Bayes filter, thus a *prior* and *likelihood* must show up somewhere in the algorithm. Here, the prior belief is the belief immediately before the observation, w is made and can be represented as the set: $\{\hat{x}^-, P^-\}$. The

priors are computed as:

$$(G.3) \quad \hat{x}_k^- = A\hat{x}_{k-1} + Bu_k$$

$$(G.4) \quad \hat{P}_k^- = A\hat{P}_{k-1}A^T + Q$$

The *likelihood* is embedded in the sensor model and observation, w . The actual integration of the feedforward model (i.e. Equation G.3) and feedback (the observation, w) is by the following equation, which can be considered the posterior estimate:

$$(G.5) \quad \hat{x}_k = \hat{x}_k^- + K_k(w_k - H\hat{x}_k^-)$$

where K_k (often called the *Kalman gain*) is essentially a weighting factor that dictates how much to trust the feedforward model vs. the feedback information. K_k is a function of the R and P covariance matrices and is a square matrix that varies from the zero matrix (K_k approaches 0 as P approaches zero; this means to trust the model more) to H^{-1} (This happens when the measurement noise, P , is zero, thus, we can perfectly trust the observations). For details on combining Gaussian beliefs, see [209] and [210].

How can this be applied to the fish? Two possibilities are: 1) Make the models linear. We may be able to make linear models of the dynamics, but the sensor models seem pretty hard to force into a linear structure. 2) Use the *extended Kalman filter* (EKF) [100]. This linearizes the models at each time step.

G.2. SLAM

Simultaneous Localization And Mapping (SLAM) is an algorithm that addresses the following question: “Is it possible for a robot to start in an unknown location in an unknown environment and then to incrementally build a map of this environment while simultaneously using this map to compute its own location?” The term was probably coined in [143], but has been investigated by many others since then [42], [67], [75], [88], [234], [218], [4], [233], [145].

SLAM uses data fusion techniques from the Kalman filter, where the state variables, \mathbf{x} , represent the robot’s own coordinates (for localization) and the coordinates of landmarks to be mapped. SLAM has been said to be analogous to the “chicken and egg” paradox, since in order to move precisely you need a good map of the environment, but in order build an accurate map, you need to know your sensing locations. In the LTI case with Gaussian disturbances the algorithm addresses this dilemma by optimally combining (*à la* Kalman filter) model information (e.g., pose¹ and landmark beliefs) with incoming observations to construct beliefs of both robot pose and landmark pose *simultaneously*.

Weakly electric fish have been known to map landmarks in their environment [56]. In fact, a *cognitive map* (i.e. an internal representation of the environment, see [240]), which is constructed from sensory information has been demonstrated by many species including birds, mammals and other fish species; [32], [108], and [186]. We could use SLAM in a simulated fish to map out these landmarks [144, 229]. Also, we have always assumed that there has only been a single prey in the fish’s sensory field. SLAM provides a framework for tracking multiple prey simultaneously; each prey could be a “landmark”.

¹Pose refers to the configuration (both position and orientation) of the robot.

And, with minor modifications, we could even assume moving prey, provided we have a motion model of the prey (i.e. We need to model the prey as in Equation G.1) [30]. In fact, we should be able to keep track of landmarks, prey positions, and the fish’s own pose with a single framework.

G.3. The Utility of Uncertainty in Simple Models

Previously, we used entropy or the “spread” of particles (`parts-std`) as the uncertainty metric because it can be applied to any general probability distribution. Here, we will see a much more simple representation. Recall that in a Kalman filter the beliefs of the states can be compactly represented as $\{\hat{x}, P\}$, which represent the mean and covariances of the state beliefs, respectively. The matrix P represents how much variance there is in the estimate, \hat{x} , and thus, is a measure of *uncertainty* in the estimate. Numerically “reducing” P means reducing the uncertainty in the beliefs of x . In fact, P is the most common uncertainty metric (when assuming Gaussian disturbances) and has been used in many analyses; [25], [83], [124], and [92].

The problem with using P as an uncertainty metric is that P is an $n \times n$ matrix, where n is the number of states. In order to compare two P ’s we need a mapping from a matrix to a scalar: $\mathbb{R}^{n \times n} \rightarrow \mathbb{R}^1$. Some common methods are provided in: [92], [103], [102], [158], and see [165] for the most comprehensive review.

G.4. Sensor-Based Control with Simple Models

We can now talk about a sensor-based controller for our simple model. In summary, “simple” means linear models for the plant and observation as well as unbiased, normally distributed disturbances of those models. With these assumptions we can use the LQG

framework (**L**inear model; **Q**uadratic utility function; **G**aussian disturbances). See [220] or [238] for an introduction to LQG. The previous section gave us a utility function: we wish to minimize a scalar metric of P .

Perhaps the most simple algorithm is to just choose the action that will decrease P by the most at the next time step. The discrete-time *greedy* approach requires computing an expected value of P_{k+1} for all possible controls at timestep k . Then, choose the control, u_k , that decreases the expected value of P_{k+1} by the most:

$$(G.6) \quad u_k = \min_u E [\|P_{k+1}\|]$$

where $E[\cdot]$ is the expected value, and $\|\cdot\|$ is some generic norm on P .

APPENDIX H

More Generalized Models for Sensorimotor Transformations in Fish and Robots

In this section we will relax the LTI system and Gaussian perturbation constraints that were established in the previous appendix. Of course, this opens the door to a multitude of possible approaches. Here, we choose to frame the problem as a Markov decision process (MDP). Then we describe a framework capable of handling uncertainties in motion and sensing, and then we show how to devise controllers fit for sensor-based control. Some of the procedures were originally developed for localization in robotics (determining the robot's own pose), but many of these procedures can be easily adapted to the weakly electric fish.

H.1. (Active) Markov Localization

The objective of Markov localization is to form a belief of the pose of the robot, but the problem can be easily restated as forming a belief of the prey relative to the fish. The Markov localization algorithm estimates a posterior belief of the state, given the actions of the robot, observations of the sensors, and models of the world (see [93], [231], and [227]). Because of the Markov assumption (i.e. the past is independent of the future, given knowledge of the current state), the solution can be formulated recursively such that new information need only be integrated once [93]. At the guts of the Markov localization method is a Bayes filter used for localization.

Active Markov localization [94, 93] takes it one step further by implementing a controller that will choose appropriate actions to improve beliefs (*Active* here is the same as in the active sensing paradigm of moving to get better information). The control policy is very similar to the one presented in Equation G.6, but entropy is used as the metric for reducing the belief to a scalar. Conceptually, active Markov localization is pretty straightforward:

- (1) Make Markov assumption to allow recursive update of belief.
- (2) Apply Bayes' rule to any new information to update belief (i.e., Bayes filter).
- (3) When it is time for control, choose actions that will decrease the expected entropy after the action (this is the greedy case).

This framework is an attractive approach to modeling sensorimotor transformations because of its generality. For instance, “action” could be a movement primitive chosen from a finite motion repertoire. For the fish, a movement primitive could be heave, roll, or pitch, or even a heave-roll-pitch. The analysis stays exactly the same.

As like many generalized formulations, one of the major downfalls of Markov localization is the tractability of the solution for any useful application. Analytic solutions are impossible for all but very simple cases. This means in almost all instances, the system must be approximated via discretization. One of the most powerful approximations is known as *importance sampling* [194], and the resulting algorithms are known as particle filters [89, 96, 182], condensation algorithms [119], and Monte Carlo localization [85, 93, 232]. These methods make smart choices about which samples of the space should be examined.

H.2. Markov Decision Processes

A MDP is a discrete framework that can model how an agent (e.g. fish or robot) interacts with the world. While the effects of the agent’s actions may be uncertain, the state of agent is always known (This is in contrast to a POMDP; we will discuss these later). For complete introductions to Markov decision processes see [43], [44], [99], [125], [150], and [169]. An MDP can be described as the tuple $\{\mathcal{X}, \mathcal{U}, \mathbf{T}, \mathbf{R}\}$, where:

- \mathcal{X} is a finite set of states of the world;
- \mathcal{U} is a finite set of actions;
- $\mathbf{T} : \mathcal{X} \times \mathcal{A} \rightarrow p(x)$ is the *state-transition function*, giving for each world state and agent action, a probability distribution over world state. We have previously referred to this as the *motion model*, and was written $p(x'|x, u)$ (in lay term: “What is the likelihood of transitioning to state x' , given that I am in state x and will execute control u ?”).
- $\mathbf{R} : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}^1$ is the reward function that maps a state and action to a scalar reward (utility). We have seen this before as the objective function (sometimes called J).

The Markov assumption forces the choice of state representation to completely describe the system up to the current time. So, as a system evolves through time, the state vector is all that is needed. In terms of *information space*, this means that there is an *efficient information mapping*, κ , from all the information ever received, \mathcal{I}_{hist} , to the current state vector of the system $\kappa : \mathcal{I}_{hist} \rightarrow \mathcal{X}$.

This framework is quite general; the two major assumptions are:

- (1) **discretization of states and actions** (and we will see shortly that observations need to be discrete for POMDPs).
- (2) **the Markov assumption.** This is really just boils down to the choice of state vector. Any system could be made Markovian if the state vector is sufficient for completely describing the complete history of the system. In the worst case, one could always use the *history information space*, \mathcal{I}_{hist} , as the state space (see [142] for details on information spaces).

In general it is not obvious what the reward function, \mathbf{R} , should be, and it is always specific to the problem one is trying to solve. But, once \mathbf{R} is defined, it should be clear how one could implement a greedy controller: simply choose the action, u , that gives the expected most favorable reward, $E[r]$ (The common convention in MDP is to maximize some reward rather than minimizing some cost). In fact it is not much of a jump to implement a finite-horizon controller by maximizing the summed rewards over k steps:

$$(H.1) \quad \pi(x) = \underset{u}{max} E \left[\sum_{t=0}^{k-1} r_t \right]$$

where r_t is the reward received on step t ; and π is the conventional symbol for a *policy* of a MDP. For an MDP, a policy is defined as a mapping from state to action: $\pi : \mathcal{X} \rightarrow \mathcal{U}$.

A MDP does not say anything about observations, and if we wish to talk about observations in an MDP, then there are two options:

- (1) **There are no observations.** This is a pure feedforward process and the state transition function, \mathbf{T} , must be completely accurate in predicting state transitions; i.e. the image of \mathcal{X} from \mathbf{T} must be the actual state and not a PDF over state space.

- (2) **Perfect observations.** The MDP assumes the state is exactly known at the time of the decision, so if \mathbf{T} cannot tell the decision maker the exact state, then a perfect observer must do the job.

H.2.1. Partially Observable Markov Decision Processes

If we want to explicitly address the uncertainties associated with observations, then we will need to frame the problem as a *partially observable* Markov decision process (POMDP; often pronounced “pom-dee-pee”). A POMDP also needs \mathcal{W} and \mathbf{O} added to the $\mathcal{X}, \mathcal{U}, \mathbf{T}$, and \mathbf{R} needed to define the MDP, where:

- \mathcal{W} is a finite set of observations the agent can experience of its world; and
- $\mathbf{O} : \mathcal{U} \times \mathcal{X} \rightarrow p(w)$ is the *observation function*, which gives, for each action and resulting state, a probability distribution over all possible observations. We have previously referred to this type of relationship as the *sensor model*, which was written $p(w|x, u)$ (in lay term: “What is the likelihood of making the observation w , given the control action, u was able to transition the system to state x ?”).

For more complete descriptions of POMDPs, see [65], [62], [76], [122], [153], [168], and [199]. POMDPs have been recently used in robotics (see [121], [130], [177], [203], and [223]). POMDPs are attractive due to their generalized framework and their ability to explicitly address observation uncertainty in the planning process. But as would be expected with this type of framework, solutions can become quickly intractable. Efficient solution methods have been the subject of quite a bit of research (see [66], [63], and [153]). Solution methods are usually a variant of *dynamic programming* or some type of gradient descent search in policy space.

The structure of a POMDP contains two basic components: a state estimator and a policy. The *state estimator* uses current observations, previous actions, and the previous belief to generate current state belief. This fusion process is usually done via Bayes' rule. Let \mathcal{B} represent the set of all possible beliefs. The *policy* uses this belief to generate a control action: $\pi : \mathcal{B} \rightarrow \mathcal{U}$. So, a POMDP-based controller gathers new information on the world via observations, and attempts to influence the world via actions.

The reward function, \mathbf{R} , dictates what the decision maker (e.g. controller or policy) cares about. Previously, we have discussed minimizing uncertainty. For Kalman-filter-type problems, we have used a scalar metric on the error covariance matrix of the state estimate, P ; and for more generalized systems, we've reduced belief PDFs to scalars via Shannon entropy, H .

The POMDP framework provides essentially a continuum of options as far as deciding the level of detail we wish to model in the weakly electric fish. But, blindly setting the weakly electric fish model in the POMDP framework without giving the problem much thought could result in an unnecessarily complex model. With the weakly electric fish we have the luxury of a well understood system (relative to most biological systems). It would thus be highly advantageous to give considerable thought to 1) what we choose for our finite sets \mathcal{U} , \mathcal{X} , and \mathcal{W} ; 2) how we define our likelihood models, a.k.a. \mathbf{T} and \mathbf{O} ; and 3) what do we use for the reward function, \mathbf{R} . Also, since we are assuming temporal discretization, we will need to be smart on what we use for a Δt .

We know that weakly electric fish have preferred, stereotyped motions. Biological systems such as the weakly electric fish, have had the luxury of millions of years of evolution. So, we know that these commonly observed motions are probably some of the

better choices a fish could make. So, we could greatly reduce the complexity of a POMDP modeled fish by setting the set of possible actions, \mathcal{U} , to be the observed motion primitives from actual behavior motion capture data.

Coming up with the set of possible observations, \mathcal{W} , will be a much more challenging task. As we know, the *Apteronotus albifrons* has about 14,000 tuberous electroreceptors covering its body. Each electroreceptor transduces an analog voltage that could take a continuum of values. If a POMDP is used on the fish, special consideration for the choice of \mathcal{W} is needed.

H.3. Reinforcement Learning

Reinforcement learning (RL) is the process of an agent learning what do in order to maximize some predefined reward function, that is, it is learning the mapping from situations to actions. The good news about using reinforcement learning is that RL is an easy extension of the MDP framework [221]. So, much of background needed to understand RL is understanding MDP.

In the MDP framework, the observation function ($\mathbf{O} : \mathcal{X} \times \mathcal{U} \rightarrow p(\mathcal{W})$) and the state transition function ($\mathbf{T} : \mathcal{X} \times \mathcal{R} \rightarrow p(\mathcal{X})$) are given in the problem definition. But RL provides a framework for the agent to learn those functions [225, 226]. The agent can do this by visiting states in \mathcal{X} , trying actions, $\{u \mid u \in \mathcal{U}\}$ and gathering statistics that can estimate \mathbf{O} and \mathbf{T} .

It can then use these learned models for planning [204]. “Planning” involves deriving a feedback plan using some of the methods discussed in the POMDP section. There are generally two methods of doing this:

- (1) **policy iteration:** Searching in policy space to find the policy that maximizes the expected accumulation of rewards.
- (2) **value iteration:** Using a variant of dynamic programming [19]. A *value function* determines the expected reward to come for each state. It defines the expected accumulation of rewards from the current state to the goal or horizon time. Whereas *rewards* determine the immediate, intrinsic desirability of states, *values* indicate the long-term desirability of states after taking into account the states that are likely to follow, and the rewards available in those states (see [40] or [181] for examples).

H.3.1. Exploiting Natural Dynamics to Ease Computation

Relevant applications of RL for our purposes include powered passive dynamic walkers. *Passive dynamic walkers* (PDWs) are simple mechanical devices that perform stable bipedal locomotion down an incline without any sort of actuators [160]. Considerable attention is given to the design of the links, joints, and contact surfaces such that the dynamics of the connected falling links under gravity match the kinematics of bipedal gait. While these devices are clever in their own right, they set the stage for developing efficient walkers. If PDWs can locomote unactuated when traversing a slight incline, then it should only take minimal actuation to walk on flat ground [217]. Such a robot was investigated in [78]; small actuators picked up the slack of the change of potential energy from the robot moving down the incline.

What was particularly neat about [78], was that they used an RL algorithm to teach the robot how to walk. And, as an added bonus of using PDW, they had a much smaller

search space than if dealing with a robot that was not mechanically designed like a PDW. Under most circumstances, RL algorithms succumb to the curse of dimensionality [18] and ultimately become computationally intractable for even lower order systems. But, because the natural dynamics of the system, and because the placement of the actuators complement the task, very little effort is needed for successful gait. Because the space of feasible controls is smaller, the RL algorithm had less space to search. Also, the PDW-like robot probably requires fewer actuators than a similar robot where the passive dynamics were not taken into consideration for its design. Fewer actuators mean a lower dimensional control space and fewer dimensions the RL algorithm has to search. This should be taken into consideration when we are modeling the weakly electric fish. The natural dynamics of the fish have evolved over millions of years to complete the necessary tasks for survival (e.g. prey capture). *By understanding the natural dynamics of the fish in water, we can limit our control space by choosing efficient motion primitives.* This will require models of the fish morphology, locomotion systems, and its surrounding fluid. Or, we can cheat by extracting motion primitives from actual motion capture data of fish doing relevant tasks.

H.3.2. Exploration vs Exploitation in RL

RL algorithms iteratively improve models while they are using those models to choose appropriate actions. RL algorithms need good models to make accurate control choices, but they are also compelled to choose actions that return the best rewards. These are often competing effects and create a comment dilemma in RL called the *exploration vs exploitation compromise*. The controller can choose to *exploit* the current models and

choose the action that will maximize rewards; or, the controller can choose actions that will generate information (e.g. observations) to best improve the models (i.e. *explore* the environment).

The k armed bandit is the classical example of exploration vs exploitation. Let's say you are put in a room with k slot machines, each of which has a percentage pay-out that is unknown to you. You have a limited number of total tries that can be distributed across the k slots at your discretion. Since you do not know any of the pay-outs (each machine is different), you need to develop an estimate for each slot's payout. If you had an unlimited number of tries, you'd play each slot for a large number of tries, get an accurate payoff percentage, and then play the slot with the largest payoff until the end of time; this would maximize your total gain. But, with a limited number of tries, you must balance exploration vs exploitation, where: (1) **exploration**: try each slot as much as possible to get an accurate model of each slot's payoff; (2) **exploitation**: only use the slot with the best perceived payoff.

APPENDIX I

Feedback-Feedforward Hybrid Controllers

Controllers have the job of choosing control commands for their system. Their choice of command can be determined based incoming information from the world—i.e. *feedback*—or it could base its control choice on previously compiled knowledge of the world—i.e. *feedforward*. A simple feedback controller maps an incoming observation to a control signal: $\pi_{FB} : w \rightarrow u$. A simple feedforward controller maps a state (or belief) to a control signal: $\pi_{FF} : x \rightarrow u$.

Let’s think back to the Kalman filter (Appendix G.1); Equation G.3 used a model of the world to predict the state. At the same time, an observation is made. Equation G.5 shows the resultant state estimate as essentially a weighted sum of the two modalities of state estimation¹. A paper by Art Kuo used similar ideas to investigate “the relative roles of feedforward and feedback in the control of rhythmic movements” [138]. He used an LQG framework (i.e. Kalman-filter-type state estimation) with the dynamics of a pendulum to model rhythmic limb movements. A *CPG feedback index* (CFI; where CPG stands for central pattern generator) determines the relative weight to give feedback vs feedforward control. The optimal CFI is a function of the amount of noise in the observations and in the uncertainties in the forward model. Noisy observations call for a

¹While equation G.5 isn’t exactly a control law, my point is to demonstrate how observations and models can be used synergistically. Plus, with the LQG framework the result from the state estimator (i.e. Equation G.5) is used as an input to the controller; this is a result of the *separation principle*, so it is an indirect part of the controller

higher weighting on the feedforward component, while large model uncertainties should give more credibility to the observations.

So we've seen how a "feedback-feedforward-dial" can be used to give relative weights to different information modalities (observations vs internal models). But many systems don't have the luxury of this continuous dial; changes in relative needs for feedback or feedforward can create qualitatively different controls. I will borrow Kevin Lynch's favorite exemplifying scenario to illustrate this phenomenon: an outfielder catching a baseball that will land behind his current position. The objective of the outfielder is to be at the ball when it lands. The outfielder has two options: (1) back peddle (slower motion) while keeping his eye on the ball; (2) turn around and run to where he thinks the ball will land. Option (1) is a slower motion, but the outfield is getting new information about the ball, thus, his belief of the ball's landing spot is improving. But, the outfielder may not get to the ball before it hits the ground. Option (2) will get him there on time, but he has a very poor estimate of "there" and will most likely overshoot or undershoot. This is a simple 1-D problem, but illustrates some of the issues associated with feedforward-feedback-hybrid controllers with finite motion choices. For instance, a controller should map the state belief ($bel[x]$) and the uncertainty of that belief (i.e., $\mathbf{Uncert}(bel[x])$ from page 113) to a control signal, u , i.e. $\pi : bel[x] \times H(bel[x]) \rightarrow u$. If we could somehow plot out this mapping, then we'd see regionalized control choices with boundaries that specify when to switch control options. So, an infinitesimal change in the belief-space could qualitatively change the control choice (e.g., from back-peddle to turn-and-run).

So, what is the relevance of this baseball example to the weakly electric fish? Two similarities are:

- (1) the competing effects of getting information and achieving the objective.
- (2) the use of discrete motion options to solve a problem in continuous time.

Let's look at similarity (1). In prey strike in the weakly electric fish, the objective of the fish could be to capture the prey in the least amount of time. If the fish had no motor or sensory uncertainties, then it could just plan the time optimal path to prey *a priori* (This is analogous to the outfielder just turning around and running to where the ball will land). But, with greater uncertainty in the belief of the prey position comes a greater chance the fish will miss the prey when it goes to capture it. Thus, it must balance between getting information (i.e. choose motions that will decrease its uncertainty of the prey position) and achieving the objective (choosing motions that will get it to the prey the quickest).

Now consider similarity (2): the use of discrete motion options to solve a problem in continuous time. We know fish generally choose from a finite set of motion primitives [157]. Why might they do this? One answer is because of sensorimotor delays. There is a time delay of about 110 milliseconds between the time a sensory event happens and the time the resultant muscle actuation is realized. So, in the case of a simple feedback controller ($\pi_{FB} : w \rightarrow u$), there is a significant delay between observing w and the system responding to u . So, the fish's motor system will be reacting to events that happened some time ago. If the fish chose to control its actuators using feedback control, then it would be responding to events that happened some time ago. And, if it chose to wait for appropriate feedback before making a control choice, and if it made very small movements, then it would have to wait several hundred milliseconds before it could make another one of those very small movements. This would result in a very, very slow fish. So, to counteract this the fish constructs a "maneuver library" where each maneuver may

be well learned and relatively complex (say 500 milliseconds long). In previous sections (e.g., MDP) we needed to discretize the motions for the formulation of the problem. Now we have some justification.

APPENDIX J

Predictive Controllers in Biological Systems

We now focus on applying some of the concepts from the previous appendix to biological control systems. We initially investigate control models in the primate motor control system because of the exorbitant number of studies done on humans and monkeys in this area. We then extrapolate to the weakly electric fish system.

We are fairly certain that any nontrivial model of the weakly electric fish's control system will need both feedforward and feedback components. Control theory provides many flavors of such architectures, but we want our models to have at least some neurobiological justification. Thus, we consider some relevant theories in neuroscience and then see if can be applied to the weakly electric fish model.

While most of the literature we draw from is based on the primate motor control system, we argue the weakly electric fish is a better system to study if we wish to model sensorimotor control schemes. So, we attempt to carry over lessons from the primate central nervous system (CNS) to the weakly electric fish CNS. In particular we investigate the roles of feedback and feedforward in the *cerebellum*¹. The cerebellum is a brain structure located in the lower-bottom section of the cranial cavity and completely separate from the cerebrum². In classical neuroscience, the cerebellum was thought to be responsible

¹Both feedback and feedforward controllers are almost certainly also essential to animals without cerebellums [247]

²The *cerebrum* (also called the *telencephalon*) is the largest section of the brain and associated with conscious thought.

for the regulation and coordination of complex voluntary muscular movement as well as the maintenance of posture and balance. But, more recently, much of the research on the cerebellum has focused on its processing of sensory information and its ability to model projected sensory effects (as well as perform state estimation [180]). Weakly electric fish have a cerebellar-like structure that is disproportionately large relative to any other type of animal. It has been speculated that it aids in processing all the electrosensory information impinging its CNS. Since we know so much about the weakly electric fish's sensory, motor, and information processing systems, we think it provides a perfect platform to investigate the functions of the cerebellum.

J.1. Relevant Models of the Primate Motor Control System

The intricacies and complexities of the primate motor control system are simply mind-boggling. This is evident by looking at the number of theories that are continuously being published, refined, and outright rejected on almost a weekly basis. The primate motor control system is so incredibly multi-dimensional and non-linear that any engineer would tremble at the thought of characterizing such a juggernaut. But, no matter how complex the controller may be, it can be broken down into functional feedforward and feedback sensorimotor transformations. Almost any theory on motor control will define what roles one or both of these control schemes play [90, 115, 123, 166, 253, 254, 256, 255, 17, 163]. A feedback control scheme uses state-related information (e.g. afferent signals) in the computation of the control signal. Feedforward control relies on an intrinsic model of the plant in its estimate of the control signal. While most researchers would attest to the fact that biological motor control systems inevitably take advantage of both control

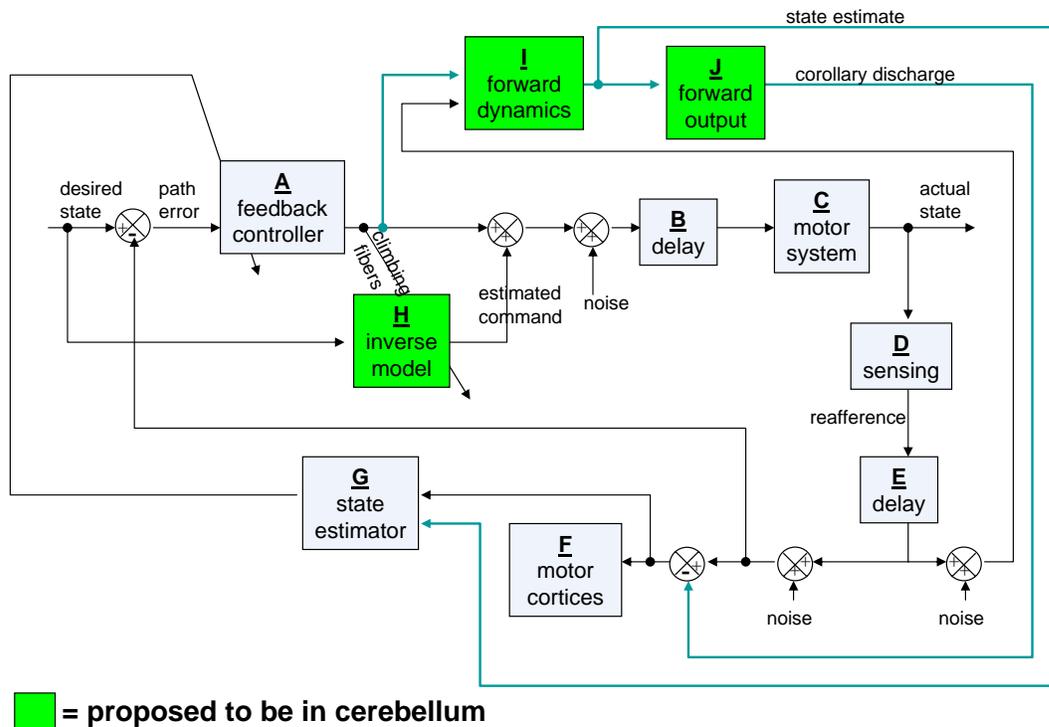


Figure J.1. A very crude block diagram of biological controller. Each block is labeled with a letter (A-J).

schemes, few provide any convincing physiological evidence on the sources of them in the brain [236].

Figure J.1 is a very crude block diagram representation of some of the components in a typical control scheme in a biological system. The green boxes are abstract information transformation modules that are thought to be in the cerebellum. The A-B-C-D-E loop is a typical model of a feedback controller where block A is the controller, and the path of blocks B, C, D, and E is the plant. As in real motor control systems, noise and delays are present. It is these elements that truly make control a difficult task. The I-J path through the cerebellum is capable of predicting both of the following: 1) the resultant state of the muscular-skeletal system that is determined from the efferent copy;

and 2) the expected sensory feedback that would result from that predicted state. The H block computes the inverse model. The result of the inverse model is integrated together with the output of the “feedback” controller that is said to be in the cerebrum (within the motor cortices).

Does the motor control system really use both feedback and feedforward control schemes? If so, does it really need both? A paper by Todorov and Jordan [237] explained in their introduction that there is a body of evidence that supports open loop (feedforward) control of planned trajectories. But Todorov et al chose a different path by proposing that the human CNS in fact uses an optimal feedback control scheme for motor coordination (specifically, they proposed an LQG modeled system). The Todorov and Jordan model of optimal control relies on sensory feedback for the state estimation needed to determine the optimal corrective movement to accomplish the task [239]. This is illustrated as the A-B-C-D-E-F-G loop in Figure J.1. For a motor control system, feedback is an extremely powerful tool. But, if used improperly, it can cause catastrophic instabilities. Let’s suppose that once the control signal is computed, block B (Figure J.1) induces a 50 msec transmission delay in the signal. The signal then recruits muscles that move the skeleton (block C). Afferent signals are encoded (block D), and block E induces another 50 msec to the signal. Upon arriving at the state estimator (block G), let’s assume a small overhead computation time of only 25 msec. So, without even considering the phase lags (e.g., mechanical) induced by block C or the additional computation that would probably be needed by blocks G and A, there is a pure delay of 125 msec. Let’s say that a particular task requires a movement path that contains frequency components up to 4 Hz (a bandwidth of 4 Hz). By the time the 4 Hz component makes the A-B-C-D-E-F-G loop,

it will have a 125 msec delay (just from blocks B, C, E, and G), which corresponds to a 180° phase lag. The system is now unstable. In reality, time delays can be longer, and phase lags (from muscular-skeletal mechanics) are non-zero and significant. This analysis would lead one to believe that any quick motion would cause the controller to go unstable and exhibit uncontrollable oscillations (using only feedback control). Control theory tells us that a controller can become more stable (larger range of stability) by taking advantage of an *inverse dynamics* compensator (Figure J.1, block H) and a forward dynamics model (block I). A feedforward transformation is able to induce a phase *lead* that could cancel an anticipated phase lag [252]. A forward dynamics transformation is able to predict the dynamics of the muscles from the state of the system and a copy of controller output. This means the CNS can predict the state variables and the errors without having to wait for the sensorimotor delay. The stability of such a system will strongly rely on the accuracy of the forward model and the ability to cancel the response from the remote system. It should be now evident that motor control systems need (at least) feedforward control.

Wolpert (and many others) have put forth many theories of the cerebellum computing inverse dynamics [253, 254, 256, 255, 163], canceling predicted sensory input [254, 38, 35, 37], and utilizing a Smith-predictor-type scheme to cancel time delays in the control loop [163]. These researchers have presented an abundance of anecdotal and circumstantial evidence for the existence of such computational units in the cerebellum. But, they have fallen short of pinpointing the exact neural avenues and representations of these modules. Wolpert and others classify these predictive models within the cerebellum

into two distinctions [254]: *inverse internal models* (e.g. block H in Figure J.1) and *forward internal models* (e.g. blocks I-J in Figure J.1) (also see [35]).

One study used the age-old question of “Why can we not tickle ourselves?” to demonstrate the cerebellum’s ability to predict sensory consequences [37, 36]. They showed (via behavior experiments and MRIs) that we can predict sensory consequences of tactile stimulation, and the neural location of this prediction is in the cerebellum. Prediction of sensory consequences of one’s own actions can greatly reduce the computational load of a controller or state estimator [84]. This is because attention only needs to be given to sensory consequences that are unexpected (Attention equates to use of computational resources). Also, if a sensory event is accounted for, then that sensory event cannot be erroneously attributed to an external stimuli. We could improve our model of the fish by utilizing such forward models.

Barto, Houk, Strick, and others have approached the cerebellum’s computational powers in a slightly different perspective [115, 17, 164]. In the Barto-Houk model [17], by including realistic mossy fiber signals, as well as realistic conduction delays in afferent and efferent pathways, the model facilitated the investigation of timing and predictive processes relevant to cerebellar involvement in the control of movement. Their model is much simpler (than say Wolpert’s) and without the need for extensive calculations. But, while their model is based on established properties of the various cerebellar neurons, they provided no evidence of their model predicting actual neural recordings. In theory, reinforcement learning algorithms could benefit greatly by utilizing such models of predicted reward.

J.2. Relevance to Weakly Electric Fish

What does all this mean for the weakly electric fish? Three things that could be used in a model of the weakly electric fish are:

1: sensory cancelation.

It had been stated that the cerebellum is capable of predicting sensory consequences. This is based off experiments on human subjects. But, an important paper by Bell and others [28, 29] on the “generation of expectations about sensory inputs and the subtraction of such expectations from actual input” in cerebellum-like structures in the weakly electric fish has suggested that these theories also apply to weakly electric fish (Also see [157] for electrosensory reafference compensation in weakly electric fish). Implementation of sensory cancelation in our weakly electric fish model would be both computationally attractive and more realistic. There are two major advantages to using sensory cancelation in a weakly electric fish model:

- (1) *attention to only novel information.* If the controller already knows what the sensory system is telling it, then the information is redundant, and it is a waste of computational resources. The controller really only needs to listen to sensory information that contradicts expected sensory consequences [114].
- (2) *electrosensory reafference compensation.* Consider the electrolocation problem. The fish processes minute voltage perturbations to compute prey location. Very minor changes in the configuration of the fish can corrupt the voltage perturbation signature of the prey by a factor of ten (i.e. create a signal-to-noise ratio of worse than 0.1) [157]. But if the fish compensates via forward model of sensory consequences, then it can greatly improve this signal-to-noise ratio.

2: inverse dynamics.

Previously, we presented evidence for constructing a finite repertoire of maneuvers for the fish. Such learned ballistic maneuvers are needed to compensate for the inherent sensorimotor delays. But, if there is an accurate inverse dynamics model, then there is no longer a need to store these maneuvers in a type of “lookup table”; we can now explicitly predict what control signals will be needed for the desired maneuver. Instead of a finite set of maneuvers, there are a continuum of possibilities.

3: predicted reward.

Reinforcement learning relies on rewards to improve models. But, for the prey strike task an external reward is received only once the fish has captured the prey. So, how can the fish use reinforcement learning on en route to capturing the prey? This is a common problem in RL. One way to solve the problem is through predicted rewards [45, 205]. The fish can receive virtual rewards for taking actions that are likely to result in rewards in the future. In order for this to work, there needs to be a predictive mechanism present. Some of the predictive mechanisms presented in this appendix are applicable.

APPENDIX K

Information in the Brain

Information processing and redundancy considerations are important for any intelligent system, artificial or biological [9]. Information theory tells us that redundancy results in inefficient use of bandwidth on communication channels, and/or inefficient encoding algorithms for occupying extra memory [33, 79]. It may seem logical that redundancy is a burden, but most biological systems, and even many engineering, systems benefit from redundancies.

Researchers have examined spike trains using information theoretic techniques and have been surprised at how low the encoding efficiency is. A major difference between biological sensory systems and engineering sensory systems is the number of sensors [170]. Engineers strive to build highly precise transducers and typically try to keep the number of sensors to a minimum [82]. Biological systems take a different approach by populating sensory arrays with many sensors. Sensory transduction and information circuits in the biological systems are often unreliable, and thus redundancy is necessary (see [16] for a discussion on the necessity of redundancy).

In this document we've assumed beliefs take the form of probability distributions of states and parameters. And, we've assumed that these beliefs are updated via a Bayes filter type process. This is the way engineers tackle the problem, but who's to say the weakly electric fish actually uses this format for belief construction and representation? There have been a few recent papers that have pointed the reader to the conclusion that

the brain does in fact use Bayesian techniques to combine sensory data with priors [107, 128, 133, 132, 167, 178]. But much of this has been inference of behavior anecdotes, so we still cannot draw any hard conclusions. How the brain represents information is a very complex issue and neuroscience has yet to draw a consensus on this.

So, what does this mean for modeling the weakly electric fish?:

- (1) We should not assume engineering techniques can model biological systems.
- (2) Biological systems play by different rules than engineering systems; things like sensory, actuation, and information redundancies may be crucial components.
- (3) “State variables” in the engineering sense may have no meaning in the brain. No one really knows how the brain represents beliefs of the world.
- (4) Bayesian techniques are really the only game in town for engineers updating priors. And, studies have claimed that Bayesian techniques are at least possible in brain. But, as with many theories in neuroscience, these results are not conclusive.

Vita

James Solberg earned his BS and MS in mechanical engineering from the University of Illinois (Urbana-Champaign) in 1997 and 1999, respectively. For his MS research he developed a novel sensor that was used as the feedback signal for an electronic valve. The control algorithm, along with this new sensor, allowed automotive air conditioners to operate more effectively. In 2000 James received a Master degree in computer systems engineering from the University of Sydney (Australia). At the University of Sydney Medical School he participated on an ambitious project of developing an active functional electrical stimulation system that would allow paraplegic patients to use their legs. His primary responsibility was developing the instrumentation and controllers. From 2001 through 2003 he worked as a research engineer at the Northwestern University Medical School, where he was in charge of the engineering operations in the Department of Physical Therapy and Human Movement Sciences. In fall of 2003 he began a PhD program at Northwestern. His current research is on the development and implementation of a novel sensing modality and complementary controllers. The sensory system is inspired by the electrosense of the weakly electric fish. These types of fish emit a weak electric field and sense perturbations to locate prey. He has developed a robotic electrolocation system capable of locating underwater targets. The system utilizes probabilistic belief-maintenance algorithms, a novel sensing technology, and various active-sensing control schemes.