

Toward Electric Field Tomography

by

Joshua Reynolds Smith

B.A., Computer Science, Philosophy
Williams College, Williamstown, MA
June 1991

B.A. Hons. Natural Sciences (Physics)
University of Cambridge, Cambridge, England
June 1993

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
in partial fulfillment of the requirements for the degree of
MASTER OF SCIENCE IN MEDIA ARTS AND SCIENCES
at the
Massachusetts Institute of Technology
August 1995

© Massachusetts Institute of Technology, 1995
All Rights Reserved

Signature of Author _____
Program in Media Arts and Sciences
22 August 1995

Certified by _____
Neil Gershenfeld
Assistant Professor of Media Arts and Sciences
Program in Media Arts and Sciences
Thesis Supervisor

Accepted by _____
Stephen A. Benton
Chair
Departmental Committee on Graduate Students
Program in Media Arts and Sciences

Toward Electric Field Tomography

by

Joshua Reynolds Smith

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning

on 22 August 1995

in partial fulfillment of the requirements for the degree of

Master of Science in Media Arts and Sciences

Abstract

This thesis addresses the problem of recovering three-dimensional location and shape information from measurements made with Electric Field Sensors. A single sensor functions as a proximity detector; two can be used as a mouse; three allow tracking of the hand in three dimensions, and each additional sensor allows us to extract additional shape information.

The ultimate goal of this work is to understand the entire hierarchy, from a single sensor up to an array of sensors. In this thesis, we take the first steps toward imaging, moving from simple proximity detection to imaging a single point in three dimensions. At every step along the hierarchy, we are interested in two questions: i) given a fixed number of sensors, how should they be arranged in space to enable us to extract the most information, and ii) given an arrangement of sensors, how should we infer “what’s out there” from the data they return? We describe a probabilistic framework that can be used to answer these questions for any sensor configuration, from a single sensor to an array.

The thesis describes the physics of Electric Field Sensing, uses this discussion to find an approximate analytical solution to the forward problem of determining the sensor values from a conductivity distribution, shows how to use this analytical model in conjunction with probability theory to design optimal electrode layouts, and presents several methods of inverting the signals to recover information about a conductivity distribution from sensor values. As an example application, we present a non-contact three-dimensional mouse.

Thesis Supervisor: Neil Gershenfeld

Title: Assistant Professor of Media Arts and Sciences

Toward Electric Field Tomography

by
Joshua Reynolds Smith

The following people served as readers for this thesis:

Reader: _____
Stephen A. Benton
Professor of Media Arts and Sciences
Program in Media Arts and Sciences

Reader: _____
Seth Lloyd
Assistant Professor
Department of Mechanical Engineering

Acknowledgement

I thank my advisor, Neil Gershenfeld, for creating the Physics and Media Group and giving me a place in it; Joe Paradiso and Tom Zimmerman, for their work on the Fish and sensor array hardware, and for sharing some of their vast knowledge of electronics; Barrett, for writing the graphics code used in the 3-d mouse demo; the other members of the group, for generosity with their expertise; my readers, Steve Benton and Seth Lloyd, for their constructive comments; Linda Peterson and Santina Tonelli, for running the Media Arts and Sciences program so smoothly; David Waxman and Eric Metois for their friendship; my family, including my Grandmother, Ruth P. Smith, for their support; and Maggie Orth, for her patience.

This work was supported in part by Hewlett-Packard, Festo, and the News in the Future Consortium.

Contents

1	Introduction	7
1.1	Electric Field Sensing	7
1.1.1	Background: The Theremin and Capacitive Sensing	7
1.1.2	Motivation for Electric Field Sensing	7
1.1.3	Recent Work on Electric Field Sensing	8
1.1.4	Alternatives to Electric Field Sensing	8
1.2	Toward Electric Field Tomography	9
1.2.1	Motivation for Electric Field Tomography	9
1.2.2	Alternatives to Electric Field Tomography	9
1.2.3	Background: Electrical Impedance Tomography	10
1.2.4	Field Sensing Hierarchy	12
1.2.5	Forward, Inverse, and Experimental Design Problems	12
1.2.6	Bayesian Framework	13
2	Physics of Electric Field Sensing	14
2.1	Hardware	14
2.1.1	Noise	15
2.1.2	Contrast to Noise and Lengthscale	16
2.2	Derivation of Circuit Model from Maxwell Equations	18
2.2.1	Maxwell's equations	18
2.2.2	Quasistatic limit	20
2.2.3	Circuit Theory	20
2.2.4	Laplace's equation in an inhomogeneous medium	21
2.2.5	Electrostatics	22
2.2.6	Capacitance	22
2.3	Lumped Circuit Model and Sensing Modes	23
2.3.1	Transmit Mode	25
2.3.2	Transmitter Loading Mode	25
2.3.3	Receiver Loading Mode	25
2.3.4	Shunt Mode	25
3	Forward Problem	28
3.1	General Case	28
3.2	Approximate effect of a small grounded object	28
3.3	Modeling the field: the dipole approximation	30
3.4	Modeling the sensor response	30
3.5	Iso-Signal Shells	32

3.6	Ground plane	32
3.6.1	Generalizations of hand model	33
4	Constructing the ambiguity class	34
4.1	Uncertainty and optimal sensor geometry	36
4.2	Entropy	36
4.3	Example: Two dimensional mouse	37
4.3.1	Recovering orientation	42
5	Example Inverse Problem: 3D Mouse	43
6	Conclusion	47
A	Backprojection Algorithm	49

Chapter 1

Introduction

1.1 Electric Field Sensing

1.1.1 Background: The Theremin and Capacitive Sensing

Electric Field sensing has existed in some form since the Theremin was invented circa 1917.[Gal91, Oth92, Nic93] Before that time, only aquatic animals had used electric fields to sense their environments.[Bas94] The Theremin combined analog sound synthesis and an early form of Electric Field sensing in a single clever circuit. It may be surprising how little research attention has been focused since then on the use of electric fields for measuring the shape and motion of the human body. With the exception of the work of Mathews[Mat90] and Vranish[V⁺92, V⁺93], little effort has been made to improve upon the early, “capacitive” form of Electric Field sensing until recently.

However, it is worth remembering that before the advent of digital electronics, there existed very few devices that could be controlled by tiny currents. The Theremin represented not only the first application of Electric Field sensing, but also one of the first examples of electronically synthesized music of any kind. Furthermore, the computation needed to perform even the simplest interpretation of the signals produced by Electric Field Sensors was not available in an inexpensive and physically small form until quite recently. In Theremin’s day, the word “computer” referred to women who tabulated ballistics tables. Thus there was little demand for methods of transducing the position of the human body into tiny currents.

In its time, the Theremin was clearly the “killer app,” or rather the “only app,” for capacitive sensing, since its analog synthesizer was one of the few existing devices that could be controlled by tiny currents, and since the signals required no inversion or gesture interpretation.

1.1.2 Motivation for Electric Field Sensing

Today, the supply of devices that can be controlled with tiny currents, in particular digital devices, is growing explosively. In one sense, controlling these devices is trivial: they can be programmed to map arbitrary messages into arbitrary actions. But the problem of translating human intention into those digital control messages is nontrivial; some would argue that it has become the main factor limiting their usefulness.

For example, the utility of 3-d graphics programs is limited by the fact that the most

sophisticated input device most people have access to is the 2-d computer mouse, which implements a straightforward mapping between motions of the mouse in a plane and motions of an on-screen pointer. In principle, arbitrary 3-d manipulations can be performed using the right combination of 2-d mouse moves (for example, clicking on-screen buttons to change rotation modes), but this is by no means a natural way to interact with a 3d environment.

The demand for more subtle and sophisticated means of physically interacting with digital devices and environments, plus the abundance of signal processing power, make this an auspicious time to develop Electric Field Sensing further. We will now survey recent efforts by members of the MIT Media Lab's Physics and Media Group to develop Electric Field Sensing further; this thesis is a continuation of this line of research.

1.1.3 Recent Work on Electric Field Sensing

In 1991, Gershenfeld[Ger91, Ger93] used a version of Electric Field sensing to measure the position of a cello bow with respect to the cello body. A transmitter on the cello body induces a signal in a resistive strip affixed to the bow. The received signal depends on the position of the bow relative to the body. Paradiso later applied these ideas in an accelerator alignment problem.[Par94, PM94]

The Physics and Media Group of the MIT Media Lab has developed this EF sensing technology further into non-contact sensors that detect the bulk conductivity of the human body[ZSP⁺95]. When part of a human body, which is typically coupled to ground, enters the field set up by a transmitter and receiver, some of the displacement current is diverted from the receiver and shunted to ground. This form of Electric Field sensing differs from previous, "capacitive" forms, because it is a three-terminal measurement, the sensed body part being the third terminal. The class of sensors usually called "capacitive," including Theremin's, Vranish's, Matthews', the cello and alignment sensors, and present industrial "capacitive sensors" are two terminal devices, in which the object being sensed functions as the second terminal.

This new form of Electric Field sensing, and variants, have been used in a variety of ways: with simple gesture recognition software to flip through the pages of an electronic newspaper[SPZG95], in a musical sculpture known as the Gesture Cube[Wax95, SPZG95], in a "musical chair" made for the magicians Penn and Teller[PGng], to fly through information spaces[ARS95, AZP⁺95], and in numerous musical installations by David Waxman[Wax95].

1.1.4 Alternatives to Electric Field Sensing

So far we have argued that the growing supply of digital devices presents a need for sophisticated sensing technologies, and that EF sensing is such a technology. In other words, EF sensing is sufficient to solve a problem of great practical interest. But to what extent is it necessary? Are there other sensing technologies that could solve the problem just as well?

There are, of course, other sensing mechanisms. Here we will briefly contrast some of the alternatives with EF sensing. Many sensors, including infra-red and pyroelectric, require a line of sight to the object being sensed. Electric Field Sensors can "see" through low conductivity materials, such as wood, paper, and cloth—they do not require a line of sight. Since EF sensors measure a bulk property of the body, impedance, surface properties of the object have no effect. By contrast, a change of clothing or lighting can change the response of video cameras and infra-red, pyroelectric, or sonar sensors. This is obviously

an advantage if one is interested in measuring surface or lighting properties. But often one is not. The fact the EF sensing functions independently of lighting conditions may be one of the reasons it is so prevalent in fish.

Magnetic sensors, such as the commercially available Polhemus, have none of the problems listed above. However, they tend to be much noisier than electric field sensors, and also are prone to hysteresis problems. One recent sensing technique, Micropower Impulse Radar,[Lab95] makes bulk measurements and does not require a line of sight to the object. But it is much more expensive than EF sensing, and while it can be used either as a simple proximity detector or an imager, it is unclear whether it could be used in regimes of intermediate complexity, such as for a non-contact 3-d mouse.

1.2 Toward Electric Field Tomography

The goal of the research begun in this thesis is to develop Electric Field sensing beyond the relatively simple applications surveyed in section 1.1.3, in which small numbers of sensors were used as proximity detectors, into a genuine imaging technology, Electric Field Tomography. At this point we have only taken the first steps, from proximity sensor to 3-d mouse, but the framework we introduce is very general should be applicable to much more complicated problems. We hope to gain a general understanding of how to extract information from a set of electric field sensors, and of the relationship between sensor geometry and the information they can provide. We are ultimately interested in the following questions: (1) Given a sensor geometry, how much information can be extracted from it? (2) To estimate a given quantity characterizing a matter distribution, what is the optimal sensor configuration (both number and geometry)? (3) How do we infer the information of interest about the distribution from the sensor values?

In this thesis, we introduce a framework for addressing these questions, and demonstrate its utility for a couple of relatively simple problems. But before continuing, let us consider some motivating applications and problems.

1.2.1 Motivation for Electric Field Tomography

Fast, inexpensive, and unobtrusive imaging of the human body will permit non-contact shape, motion, and gesture capture for hands, faces, and entire bodies, which will enable people to interact in new (and hopefully better) ways with digital devices, computer models, and other people. For example, as realtime three-dimensional computer graphics becomes more common, the need for techniques to measure 3-d gesture becomes more acute. Creating 3-d models, placing lights, and animating the models are all cumbersome procedures with present techniques. Creating and lighting models is often accomplished with a mouse, keyboard, or stylus. More sophisticated methods for capturing shape, and motion capture techniques, which we will review below, are typically more intrusive.

1.2.2 Alternatives to Electric Field Tomography

The CyberWare scanner, a laser-based shape capture device, requires a line of site between the apparatus and the object being scanned. Medical imaging techniques such as CT-scans, MRI, or PET are too costly, too large, and provide more data than is needed for human

interface and simple geometry-capture applications, since they provide images of the interior of the body.

All present motion capture or “performance animation” systems, which map human movement into that of three-dimensional graphic “puppets,” are intrusive as well.[Rob94, Cha94] Magnetic motion capture devices such as the Polhemus require the user to wear tethered magnetic sensors whose response to a static field can be measured through the cable. The other popular motion-capture technology is video.

For human interface purposes, video provides both too much and too little data. The frame rate is too slow for musical and other demanding realtime control applications, so in that sense it doesn’t provide enough data. On the other hand, when a video camera does deliver a frame, it provides as many as several million numbers. For control applications, these million numbers must be laboriously boiled down into the final control message, often just a few bytes. In this sense, a video camera provides far too much data.

To combat the second problem, video-based motion-capture systems require the user to wear black clothing and bright targets that the vision software can easily track. Video-based gesture recognition systems require large amounts of computing horsepower and often place restrictions on the background. Geometry capture systems require multiple cameras and consume large amounts of computing. Having surveyed the competition, we will now revisit the ancestors of Electric Field Tomography.

1.2.3 Background: Electrical Impedance Tomography

Electrical Impedance Tomography (EIT),[BB84, Web89] sometimes known as Applied Potential Tomography, is an imaging technique that has been used in medical and geophysical applications. It is closely related to Electric Field Tomography. In a typical medical EIT experiment, current is applied to electrodes attached to the skin, the resulting voltages are measured, and then, using a variety of inversion techniques, the conductivity distribution is recovered.

More formally, a current pattern is applied to the surface $\partial\Omega$ of body Ω , and the resulting electrical potential on $\partial\Omega$ is measured. The conductivity $\sigma \in \Omega$ is then inferred. The potential u in Ω satisfies $\nabla \cdot (\sigma \nabla u) = 0$, which is Laplace’s equation with a finite conductivity field, and may also be viewed as the microscopic form of Ohm’s and Kirchoff’s laws. The current density j on $\partial\Omega$ given by $j = -\sigma \nabla_n u$ where n is the outward normal on $\partial\Omega$.

Comparison with Electric Field Tomography

EIT is similar enough to EFT that the mathematics used to show the feasibility of inversion appears to carry over directly, as we discuss in section 1.2.3. Though Electric Field Tomography and Electrical Impedance Tomography may be identical in the most general mathematical sense, there are substantial differences in practice. In EFT, we usually apply alternating voltages at the boundary and measure currents to find capacitances, rather than applying currents and measuring voltages to find conductances.

More significantly, the practical inversion problems are somewhat different. In EIT, the electrodes are in contact with the body. The goal is to image the small variations in conductivity in the interior. EIT imaging occurs in a very low-contrast regime. In Electric Field Tomography, on the other hand, the electrodes are not in contact with the person. There are essentially two, vastly different impedances: that of air, and that of the person’s

hand. This is the high-contrast regime. Though EFT is an intrinsically easier problem (producing a 1-bit image should be easier than producing an 8-bit image), there are some advantages to the low-contrast regime: the image can be treated as a perturbation on a background conductivity, and linear approximations can be used.

Finally, note that EFT could be posed in a completely different way, in which all boundary potentials are known, but the boundary geometry is not known. This may turn out to be the most straightforward way of posing the problem.

Questions of Principle Answered in EIT

Extensive work has been done in the Electrical Impedance Tomography literature on the feasibility of the inverse problem. Basically, it has been shown that the problem is soluble in principle. Kohn and Vogelius proved the uniqueness of solutions of the inverse problem for a piecewise analytic conductivity.[KV84, KV85, KV83] Sylvester and Uhlman proved invertibility for $\sigma \in C^\infty(\Omega)$.[SU87] Though most of the EIT literature is posed in the context of DC conductivity, we will see in section 2.2.4 that there is an exact analogy between the Laplace equation for inhomogeneous media that describes the resistive forward problem, and the equation describing the capacitive forward problems, so that in the most general sense, arguments from one domain automatically apply in the other.

However, because of the different contrast regimes, algorithms for constructing the inverse are not automatically applicable in both domains. Furthermore, EFT has the additional practical complication that the object being imaged (a hand, say) is typically connected to a conductive body that is not in the sensing field, but is coupled capacitively to ground. Nevertheless, we will briefly survey the three classes of inversion algorithms used in EIT.

Backprojection

The backprojection algorithm for EIT, introduced by Barber and Brown,[BB84] was adapted from the algorithm of the same name that originally was used to reconstruct CT-scans. The classic backprojection algorithm is explained in Appendix 6. In the generalized version used in EIT, measurements are backprojected along the iso-signal surfaces, in Barber and Brown's case curved equipotential lines, since they measure voltage.

Barber and Brown's forward model makes perturbative linear approximations. It is at present unclear whether it could be adapted to our high-contrast regime.

Backprojection is reported to provide fast, high-quality, approximate reconstructions. If slower, more refined algorithms are to be employed, it may make sense to give them a head start by backprojecting first.

Spectral

Spectral methods typically provide a set of basis images and a matrix that transforms a vector of sensor readings into a vector of basis image coefficients. In EIT, these techniques rely on perturbative linear approximations. Since spectral techniques are analytical, they are more likely to give insight into algorithm stability, the ill-posedness of the inverse problem, and so forth.[AS91, Z⁺91]

Iterative

Iterative numerical techniques are the only way of solving the inversion problem exactly. Based on the applied boundary conditions and a trial conductivity distribution, the free boundary conditions (sensor values) are predicted. An optimization technique such as Newton-Raphson is used to relax the conductivity distribution so as to reduce the error between the actual sensor values and those predicted by the trial conductivity distribution.[KM90, DL81]

A very interesting variational formulation of the inversion problem, which provides explicit estimate of the quality of the reconstruction, is developed by Berryman in [Ber89] and by Berryman and Kohn in [BK90]. They show that there is a close analogy between the electrical impedance tomography problem and the seismic travel-time inversion problem, which they solve using Fermat’s principle.

1.2.4 Field Sensing Hierarchy

For many human interface applications, a full image may not be desirable. Therefore, rather than simply adopting one of the brute force imaging techniques from EIT, it would be nice to understand how to extract just the required information using only the necessary number of sensors. For example: roughly speaking, from one sensor it is possible to infer at most one number characterizing a matter distribution. The number might be an estimate of position or size, for example. With two sensors, one can estimate two position coordinates, or one size and one position. This field sensing hierarchy extends upwards, until we reach a continuum of sensors. When a large number of sensors is available, it becomes possible to extract a three-dimensional image of a conductivity distribution, as the arguments cited in section 1.2.3 show. In Chapter 4, we will show how to construct and analyze “ambiguity classes,” which will make the notion of the field sensing hierarchy more precise.

This thesis is entitled “Toward Electric Field Tomography” because, although we have not yet arrived there, imaging is the endpoint toward which we are heading, and because our goal is to understand the entire hierarchy, from a single sensor up to a continuum. At this point, we have explored the hierarchy in a practical way as far as three sensors: we have made non-contact two- and three-dimensional mice. The user’s body is the only moving part of these input devices.

1.2.5 Forward, Inverse, and Experimental Design Problems

Now that we have explained the ultimate destination of this research, we will describe the sub-problems that must be solved in order to arrive there. There are two obvious parts to the full imaging problem, and to any of the problems along the EF sensing hierarchy. The goal, of course, is to infer as much as possible about the conductivity distribution from the measurements returned by the sensors. We will refer to this as the inverse problem; before we can solve it, we must first understand the direct or forward problem of calculating the sensor values given a conductivity distribution and sensor geometry.

There is also a less obvious problem, which might be called the “experimental design” problem. When a large number of sensors is available, the problem is to decide which of the large number of measurements that could be made actually should be made. When a smaller number of sensors is available, the question is, how many are needed, and where should they be placed?

1.2.6 Bayesian Framework

We will introduce a general framework that allows us not only to solve inversion problems anywhere along the EF sensing hierarchy, but also to design optimal sensor geometries. Given a sensor geometry and data values, we will construct the ambiguity class of models that could explain the data. This ambiguity class takes the form of a probability distribution over model parameters. Once this distribution has been constructed, the problem of inverting the sensors (to find, say, the location of the hand from some sensor values) reduces to the optimization problem of maximizing the probability of the model, given the data.

This ambiguity function is also useful in choosing sensor geometries. The expected change in entropy of the ambiguity class will provide a measure of the quality of sensor geometries. By maximizing this quantity, optimal sensor geometries can be found.

Before we discuss methods of solving the inverse and experimental design problems in Chapters 4 and 5, however, we will explain the physics of electric field sensing in Chapter 2, and then, in Chapter 3, we will introduce approximations that allow us to predict efficiently the sensor values resulting from a simple matter distribution.

Chapter 2

Physics of Electric Field Sensing

In this chapter we will describe the hardware we use for Electric Field Sensing, and then discuss the physics underlying the hardware.

2.1 Hardware

The Electric Field Sensors[Ger93, ZSP⁺95] developed in the Physics Group of the MIT Media Laboratory provide high-precision, low-noise measurements of the bulk conductivity of a matter distribution. The “Fish” field sensing board¹ consists of a transmitter that can be tuned from 20kHz to 100kHz, and four receive channels that use synchronous detection. The transmitter consists of an oscillator connected to an op-amp. The op-amp defines the voltage on the transmit electrode, as specified by the oscillator, by putting out as much current as required to maintain the correct voltage. The amount of power that the user is exposed to is on the same order of magnitude as that received from a pair of stereo headphones, and is several orders of magnitude below FCC regulation.

Each receive channel consists of an op-amp gain stage, a multiplier, and another op-amp used as an integrator. The received signal is multiplied by the original transmitted signal, and the resulting function is integrated over an interval of 60 ms. The effect of these two operations is to project out, in a sense defined below, all the Fourier components of the received signal except for the component that was transmitted. The multiplier and integrator are computing (in analog electronics) the inner product of the transmitted signal function s_t and the received signal function s_r , with a window function set by the integration time. The sense in which the multiplier and integrator project out all undesirable Fourier components is the following: because all distinct pairs of Fourier components are orthogonal, the contribution to the inner product $\langle s_t, s_r \rangle$ from all the undesirable (i.e. different from s_t and therefore orthogonal) components is zero. The input stage is therefore a very sharp filter that rejects all signals not of the proper frequency and phase.

It is also possible to describe the sensing circuitry in terms of amplitude modulation. The transmitter may be thought of as a carrier whose amplitude is modulated by the motions of a person’s body. The receive multiplier mixes the carrier down to DC, and then the final

¹It is called “Fish” because electric fish use similar mechanisms to sense their environments, and because we hope that the Fish, which navigates in three dimensions, might be the successor input device to the mouse, which only navigates in two.

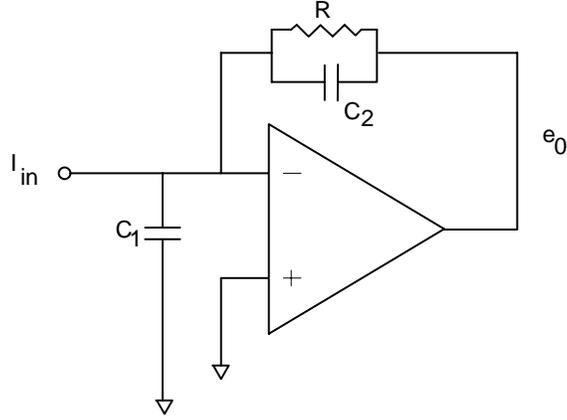


Figure 2-1: Diagram of front end gain stage for noise calculations.

lowpass filter rejects all signals other than those superimposed on the carrier.

2.1.1 Noise

There are two classes of op-amp noise to consider, *interference noise*, which is caused by external sources, such as coupling from the digital circuitry on the board, and *intrinsic noise*, which is due to the physics of the op-amp and feedback network components. We will calculate the intrinsic noise for the Fish front end, compare it to the observed noise, and conclude that interference noise dominates the intrinsic noise.

Intrinsic Noise Calculation

Intrinsic noise may be divided into voltage noise, current noise, and Johnson noise in the feedback and input resistors.[SS69] The most important component of the intrinsic noise turns out in our case to be voltage noise. For our front end gain stage, shown schematically in Figure 2-1, the output voltage noise e_o is given by $e_o = N e_n$ where e_n is this intrinsic voltage noise and the amplifier's noise gain N is given by

$$N = \frac{1}{\beta} \frac{1}{1 + \frac{1}{A\beta}}$$

Here A is the amplifier's open-loop gain, and the feedback ratio

$$\frac{1}{\beta} = 1 + \frac{R}{1 + RC_{2p}}(C_{1p})$$

with $R = 1M\Omega$ and $C1 = C2 = 1pF$. The impedance of C_1 is denoted $C_{1p} = \frac{1}{i\omega C}$. For the TL082CP op-amp used in the Fish front end, e_n is $\frac{18nV}{\sqrt{Hz}}$ [Dev95]. This yields a voltage noise of $0.12mV$ rms. The current noise is $8.31 \times 10^{-18}V$ and the Johnson noise is $9.95 \times 10^{-7}V$. The quadrature sum of all the three intrinsic noise figures is essentially the same as that of the voltage figure, $0.119mV$.

The Fish board uses the 8-bit analog to digital converter on the Motorola MC68HC11,

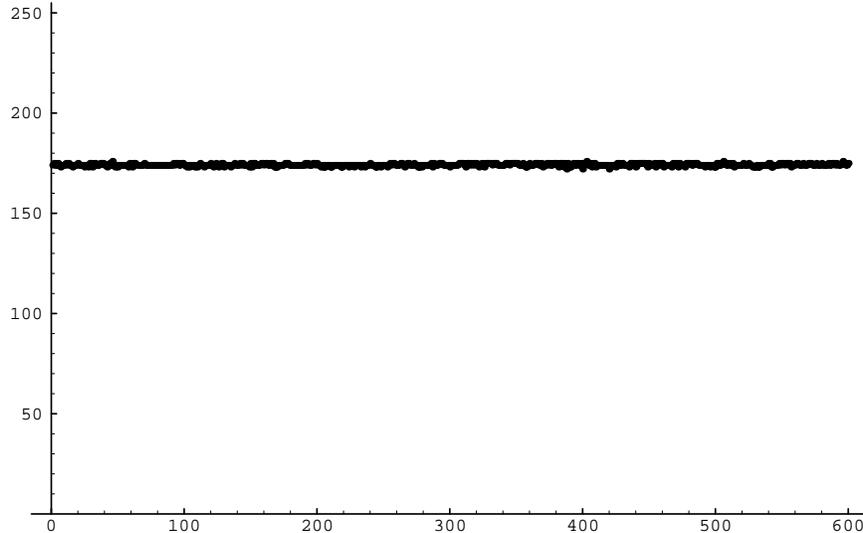


Figure 2-2: Timeseries of 600 sensor values showing quantization noise.

which has $\pm 1/2$ least significant bit accuracy[Mot91], so the quantization resolution is $5V/256 = .195V$. Since the quantization noise is 3 orders of magnitude greater than the intrinsic noise of the op-amp, we expect the readings to show the Gaussian profile typical of quantization noise.

Measured Noise

To study the noise characteristics of the final digitized signal, we chose the offset voltage so that the sensor gave a value comfortably far from clipping high or low (about 175 out of 255) when the field was unperturbed. We collected 600 samples at a sampling rate of 60 per second. The maximum likelihood estimate of the standard deviation of this data set is 0.66. Figure 2-2 shows this timeseries, Figure 2-3 shows its normalized histogram, and Figure 2-4 shows the normal distribution with the mean and standard deviation of the noise timeseries.

2.1.2 Contrast to Noise and Lengthscale

The signal received depends on the sensor geometry. Consider a parallel plate transmit-receive geometry. Since the capacitance goes as $1/d$, where d is the spacing between the plates, the unperturbed signal also falls off as $1/d$. Thus at a certain point the op-amp noise places a limit on how far we can see with the sensors, that is, on the lengthscale on which they are operable. In fact, however, the relevant quantity in determining the usefulness of the sensors is not signal to noise (the ratio of the unperturbed signal to the voltage or other noise), but *contrast* to noise, that is, the ratio of the difference between the maximum and minimum signal to the noise.

In order to maximize contrast to noise microscopically, one needs to set $\frac{d(\frac{S}{N})}{dp} = 0$, where S is signal, N is noise, and p is an adjustable parameter that affects both signal and noise.[Ros95]

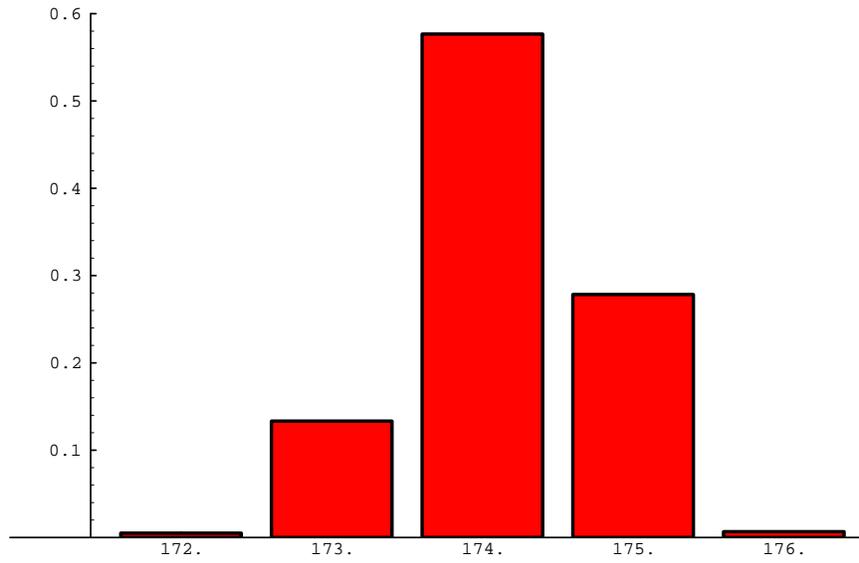


Figure 2-3: Normalized histogram of noise dataset.

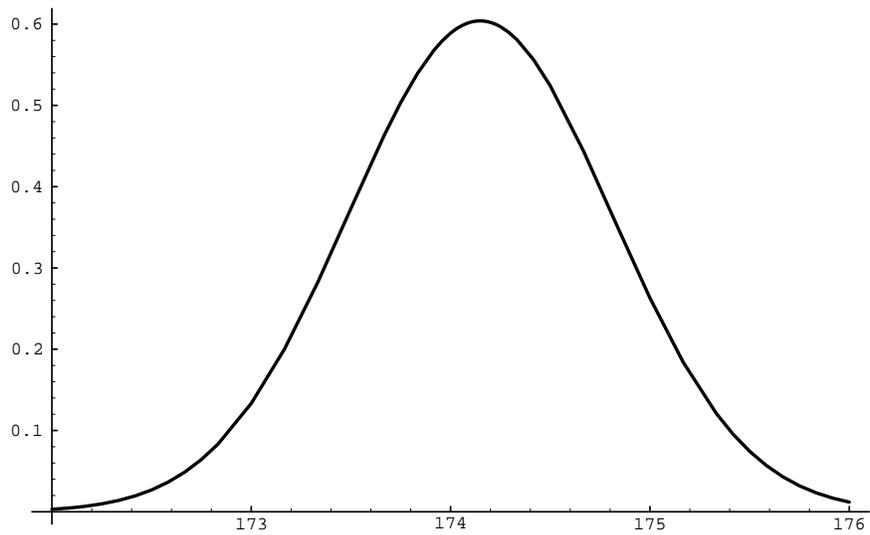


Figure 2-4: Normal distribution with mean and standard deviation of noise dataset.

2.2 Derivation of Circuit Model from Maxwell Equations

The goal of EFT imaging is to estimate a microscopic “impedance field” from macroscopic measurements provided by circuits such as the Fish. It will therefore be useful to be clear about the relationship between the circuit description and the field description of the sensing process. In the remainder of this chapter we will outline the derivation of the circuit description from the quasi-static approximation to the Maxwell equations, following the treatment in Chapter 6 of Fano, Chu, and Adler[FCA60]. In the next chapter, we will come full circle by using circuit concepts to describe fields, and introduce a hybrid circuit-field description that will allow us to predict the response of the sensors to simple conductivity distributions.

2.2.1 Maxwell’s equations

Maxwell’s equations can be written in the form

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \quad (2.1)$$

$$\nabla \times \mathbf{H} = \mathbf{J}_f + \frac{\partial \mathbf{D}}{\partial t} \quad (2.2)$$

$$\nabla \cdot \mathbf{D} = \rho_f \quad (2.3)$$

$$\nabla \cdot \mathbf{B} = 0 \quad (2.4)$$

$$\nabla \cdot \mathbf{J}_f = -\frac{\partial \rho_f}{\partial t} \quad (2.5)$$

where, for linear and isotropic media,

$$\mathbf{D} = \epsilon \mathbf{E}$$

$$\mathbf{B} = \mu \mathbf{H}$$

$$\mathbf{J}_f = \sigma \mathbf{E}$$

As explained in section 2.1, Electric Field Sensing uses low frequencies. To study the properties of low-frequency solutions of the Maxwell equations, we can introduce a time-rate parameter α and a new, scaled time $\tau = \alpha t$. Small values of α map long periods of real time t into a unit of scaled time. Thus slow or low-frequency behavior corresponds to small values of α . The low-frequency behavior is therefore described by the low order terms in an expansion of the fields in a power series in α .

We can put a rough physical interpretation on α : its value is the ratio between the time τ for an electromagnetic wave to propagate across the longest lengthscale in the problem (the characteristic time for wave behavior), and the smallest time t of interest, in our case the period of the highest frequency that our oscillator can produce. We will say this quantitatively later, but for now note that the period of our oscillator is slow compared to the wave propagation time, so α is small. The expansion of \mathbf{E} in powers of α has the form

$$\mathbf{E}(x, y, z, t) = \mathbf{E}(x, y, z, \tau, \alpha) = \mathbf{E}_0(x, y, z, \tau) + \alpha \mathbf{E}_1(x, y, z, \tau) + \alpha^2 \mathbf{E}_2(x, y, z, \tau) + \dots$$

where

$$\begin{aligned}\mathbf{E}_0(x, y, z, \tau) &= [\mathbf{E}(x, y, z, \tau, \alpha)]_{\alpha=0} \\ \mathbf{E}_1(x, y, z, \tau) &= \left[\frac{\partial \mathbf{E}(x, y, z, \tau, \alpha)}{\partial \alpha} \right]_{\alpha=0} \\ \mathbf{E}_k(x, y, z, \tau) &= \frac{1}{k!} \left[\frac{\partial^k \mathbf{E}(x, y, z, \tau, \alpha)}{\partial \alpha^k} \right]_{\alpha=0}\end{aligned}$$

When the frequency is low enough that all but the zeroth and first order terms can be neglected, the solution is called quasi-static.[FCA60]

Using the new, scaled time τ , time derivatives will be multiplied by α , for example:

$$\frac{\partial \mathbf{B}}{\partial t} = \frac{\partial \mathbf{B}}{\partial \tau} \frac{\partial \tau}{\partial t} = \alpha \frac{\partial \mathbf{B}}{\partial \tau}$$

The three Maxwell equations involving time derivatives become

$$\nabla \times \mathbf{E} = -\alpha \frac{\partial \mathbf{B}}{\partial \tau} \quad (2.6)$$

$$\nabla \times \mathbf{H} = \mathbf{J}_f + \alpha \frac{\partial \mathbf{D}}{\partial \tau} \quad (2.7)$$

$$\nabla \cdot \mathbf{J}_f = -\alpha \frac{\partial \rho_f}{\partial \tau} \quad (2.8)$$

Substituting the expanded \mathbf{E} and \mathbf{B} fields back into the scaled Maxwell equation 2.6 and grouping terms, 2.6 becomes

$$\nabla \times \mathbf{E}_0 + \alpha(\nabla \times \mathbf{E}_1 + \frac{\partial \mathbf{B}_0}{\partial \tau}) + \alpha^2(\nabla \times \mathbf{E}_2 + \frac{\partial \mathbf{B}_1}{\partial \tau}) + \dots = 0$$

Each term in the sum must equal zero individually for the equation to hold for all values of α . This defines a series of equations whose solution is the series of fields that make up our expansion. Because the \mathbf{B} term in 2.6 is multiplied by α , and the \mathbf{E} term is not, k th order \mathbf{E} terms are related in the infinite series of equations to $k - 1$ th order \mathbf{B} terms. The expansion of equation 2.7 will yield a series of equations coupling k th order \mathbf{B} fields to $k - 1$ th order \mathbf{E} fields. Since all fields are coupled only to lower order fields, any number of terms can be evaluated, by starting from the zeroth order solution, using that to find the first order, and so on. The zeroth order \mathbf{E} field equations are

$$\nabla \times \mathbf{E}_0 = 0 \quad (2.9)$$

$$\nabla \times \mathbf{H}_0 = \mathbf{J}_{f0} \quad (2.10)$$

$$\nabla \cdot \mathbf{J}_{f0} = 0 \quad (2.11)$$

The Maxwell equations that do not involve time derivatives become:

$$\nabla \cdot \epsilon \mathbf{E}_0 = \rho_{f0} \quad (2.12)$$

$$\nabla \cdot \mu \mathbf{H}_0 = 0 \quad (2.13)$$

Next we will write out the first order fields. Since all values of α correspond to physically realizable fields, any field can be viewed as the original, “unscaled” field. Therefore no loss of generality results from setting $\alpha = 1$, and writing t instead of τ :

$$\nabla \times \mathbf{E}_1 = \mu \frac{\partial \mathbf{H}_0}{\partial t} \quad (2.14)$$

$$\nabla \times \mathbf{H}_1 = \epsilon \frac{\partial \mathbf{E}_0}{\partial t} + \mathbf{J}_{f1} \quad (2.15)$$

$$\nabla \cdot \epsilon \mathbf{E}_1 = \rho_{f1} \quad (2.16)$$

$$\nabla \cdot \mu \mathbf{H}_1 = 0 \quad (2.17)$$

$$\nabla \cdot \mathbf{J}_{f1} = -\frac{\partial \rho_{f0}}{\partial t} \quad (2.18)$$

Because the curl of any vector field V equals zero if and only if V can be written as the gradient of a scalar potential, equation 2.9 implies that $\mathbf{E}_0 = \nabla \phi_0$. In a region with no sources or sinks, any vector field satisfies $\nabla \cdot \mathbf{V} = 0$, so if there are no free charges, $\nabla \cdot \nabla \phi_0 = \nabla^2 \phi_0 = 0$; that is, ϕ_0 satisfies Laplace’s equation. If free charges are present, then ϕ_0 satisfies Poisson’s equation, by a similar argument.

2.2.2 Quasistatic limit

In terms of our expansion above, the quasi-static condition holds when $\alpha = \frac{\tau}{t} = \frac{L}{ct} \ll 1$, because higher powers of α are negligible when $\alpha \ll 1$. Again, τ is the time for an electromagnetic wave to propagate across the longest lengthscale in the problem, and t is the period of the transmit oscillator. If L is 10 meters and the transmit frequency is 100kHz, so that $t = 1.0 \times 10^{-5}$, then $\alpha = 3.3 \times 10^{-3} \ll 1$, so we are comfortably in the quasistatic regime.

When α is vanishingly small, so that only the zeroth order terms are required, we are in the regime of DC circuits. For small but finite rates of change, the first order terms must also be taken into account. This is the regime of AC circuitry. In the next section we will see in more detail how the concepts and laws of circuit theory emerge naturally as the quasistatic limit of the Maxwell equations.

2.2.3 Circuit Theory

There are three basic types of solutions to the zeroth and first order Maxwell equations, which correspond to the three basic types of circuit components: capacitive, inductive, and resistive. For Electric Field Sensing and Tomography, only the capacitive solutions are relevant;² for Electrical Impedance Tomography only the resistive solutions matter (for this reason the name Electrical Resistivity Tomography would be more accurate). We will see, however, that the equations specifying the “capacitive” and “resistive” fields are identical in form, which might be guessed from the fact that resistance and capacitance can be viewed as special cases of the generalized circuit concept of impedance.

The three types of quasi-static fields can be classified according to their zeroth-order terms. The first two types arise when there is no conduction current. In these first two

²We will see later that the situation is slightly more complicated than this.

cases the right side of equation 2.10 is zero, and there is no coupling between the electric and magnetic fields, so the two can be treated separately. The first type of quasistatic solution, electrical, has no magnetic component, and will be associated with capacitance, as we will explain below. A magnetic solution with no electrical component will be associated with inductance. The solution associated with resistance arises when conduction currents are present. If $\mathbf{J}_{f0} = \sigma \mathbf{E}_0$, then equation 2.10 becomes $\nabla \times \mathbf{H}_0 = \sigma \mathbf{E}_0$. Thus in resistive solutions the zeroth order electric field is coupled to the zeroth order magnetic field through a finite conductivity.

To see why the electrical solution is associated with capacitance, first recall the circuit definition of capacitance:

$$I = C \frac{dV}{dt} \quad (2.19)$$

A capacitance couples a current to the time derivative of a voltage. Now consider the “capacitive” field. Because of equation 2.15, a zeroth order electric field induces a first order magnetic field proportional to the time derivative of the electric field. Associated with the zeroth-order electric field is a zeroth-order charge; by equation 2.18, the time derivative of this charge induces a first order current. Since the zeroth order electric field may be represented by a scalar potential, this first order current is coupled to the time derivative of the zeroth order potential. As we saw in equation 2.19, this type of coupling is referred to as capacitive in circuit theory. Similar arguments demonstrate the correspondance between the other types of fields and circuit components.

2.2.4 Laplace’s equation in an inhomogeneous medium

For an inhomogeneous medium with embedded conductors at defined potentials, for example, the medium shown in Figure 2-5, the zeroth order electric field satisfies equation 2.12, which becomes

$$-\nabla \cdot \epsilon \mathbf{E}_0 = \nabla \cdot (\epsilon \nabla \phi_0) = \epsilon \nabla^2 \phi_0 + \nabla \epsilon \cdot \nabla \phi_0 = 0 \quad (2.20)$$

in the region between the conductors. For a homogeneous medium, this reduces to the standard Laplace equation, $\nabla^2 \phi_0 = 0$. The third, resistive field solution leads to an equation analogous to 2.20, but with permittivity replaced by conductivity:

$$-\nabla \cdot \sigma \mathbf{E}_0 = \nabla \cdot (\sigma \nabla \phi_0) = \sigma \nabla^2 \phi_0 + \nabla \sigma \cdot \nabla \phi_0 = 0 \quad (2.21)$$

Because these two equations have the same form, solutions pertaining to one physical situation are solutions to the other, as long as the boundary conditions are analogous. In fact, to determine the capacitance due to a body with a complex shape, Haus[HM89] (p. 274) recommends immersing the body in a tank filled with electrolytes and making conduction measurements, which he feels are easier experimentally.

Equation 2.21 describes the general forward problem in Electrical Impedance Tomography. Technically, equation 2.20 describes the forward problem in Electric Field Tomography. However, the bodies we image in EFT tend not to be isolated permittivities, as would be the case if we only wanted to use EFT to, for example, image the surface of a wooden desk and the books on top. For this sort of problem, all the machinery of Electrical Impedance Tomography applies in the most straightforward manner. In a typical EFT problem, however, the body being imaged is a perfect conductor coupled capacitively (through the shoes) to a defined potential, ground, which changes the problem somewhat. Practical approaches

will probably model the part of the body that is outside the sensing field as a lumped circuit component. For example, the parts of the body outside the field can probably be modeled as a capacitor to ground.

2.2.5 Electrostatics

Our expansion showed that static (zeroth order) electric fields satisfy Laplace's equation. The behavior of the static fields is crucial to Electric Field Sensing, because, as we shall see in section 2.2.6, though EF sensing requires first order fields to operate, no new information is contained in the first order fields; it is all present in the zeroth order. Therefore, we will now briefly discuss methods for solving electrostatic field problems.

There are two basic ways of viewing such problems: the boundary value perspective, and the superposition integral perspective. In the boundary value perspective, we find the potential and hence the field by solving Laplace's equation subject to specified boundary conditions. There are a variety of analytical methods useful for solving Laplace's equation in particular cases, but in general, one must resort to numerical methods. In describing the boundary value perspective, I have essentially just said "solve the zeroth order Maxwell equations for the electric field."

There is another way to view electrostatics problems, however. In the superposition integral approach, we assume that a charge distribution is known. We find the total field by superposing the fields induced by each charge separately. This approach works well for finding the field around a (non-conductive) molecule, since the charge density is essentially fixed. It is less convenient for dealing with conductors, since the charges are not in fixed locations. It may not be obvious how to relate the superposition integral perspective to our quasistatic solution of Maxwell's equations, and thus to the boundary value perspective. The relationship is that the \mathbf{E} field due to a static point charge is the Green's function of Laplace's equation. The reason we can superpose the fields is that Maxwell's equations are linear.

2.2.6 Capacitance

In the previous two sections, we discussed the principles of how to find the static electric field due to a set of isolated conductors at known potentials, such as those shown in Figure 2-5, embedded in a medium (isotropic, but not necessarily homogeneous) of permittivity ϵ . We now will show how to use quasistatic field solutions to calculate macroscopic circuit quantities such as capacitance and received current.

The static charge on a conductor i is due to the \mathbf{E}_0 field:

$$Q_i = - \int_{S_i} \epsilon \mathbf{n} \cdot \nabla \phi_0 da$$

where S_i is the surface of i , \mathbf{n} is the outward normal to S_i , and ϵ may be a function of position, since the medium need not be homogeneous.

Using the standard definition, the capacitance of conductor i due to a conductor j is the ratio between the charge on Q_i and the voltage between j and a reference. Of course if we know the capacitance and voltages for a pair of electrodes, we can find the charge induced on one by the other. Because of the linearity of all the equations involved, the total charge on i induced by all the other conductors is the sum of the separately induced

charges[FCA60] (note that the capacitances are not linear functions of position):

$$Q_i = \sum_j C_{ij} V_j \quad (2.22)$$

The off-diagonal terms of this capacitance matrix C_{ij} represent the ratio between Q_i and V_j when all the other V s are zero. The diagonal “self-capacitance” terms C_{ii} represent the charge on i when it is held at V_i and all the other electrodes are at zero. In terms of our sensing hardware, the diagonal terms represent the intrinsic capacitance of the cable and the electrode. The matrix is symmetrical.

We will now see that from the capacitances, we can calculate the currents received at the electrodes. This is because equation 2.18 relates the first order current to the zeroth order charge. By charge continuity (expressed microscopically in equation 2.18), the current I_i entering receiver i is given by the time derivative of the charge on i : $I_i = \frac{dQ_i}{dt}$.

$$I_i = \frac{d}{dt} \sum_j C_{ij} V_j = \sum_j C_{ij} \frac{dV_j}{dt} \quad (2.23)$$

The currents that we measure in Electric Field Sensing are first order phenomena. However, we only use the currents to measure capacitance, the zeroth order property that is geometry dependent and therefore encodes the geometrical information that we ultimately want to extract. This tells us something about the physical limits on the time resolution of EF sensing: the “frame rate” must be much shorter than the characteristic time for first order phenomena, that is, the oscillator period.

2.3 Lumped Circuit Model and Sensing Modes

Now we will apply the framework developed in the previous section to Electric Field Sensing. We will present a lumped circuit model of a single transmit-receive pair with a single target object, whose proximity we are interested in. The various “modes” in which the Fish circuitry can be used have clear interpretations as current paths through the circuit diagram, shown in Figure 2-5. For each sensing mode, we will give a brief overview from a user’s point of view, and then explain the physics of the mode in terms of this diagram.

Figure 2-5 shows the model. There are four “terminals”: the transmitter, the receiver, the target object (shown as a hand), and ground. The $\frac{4 \times 3}{2} = 6$ distinct inter-conductor capacitances are shown. The small resistor and capacitor and C_i represent the body’s internal capacitance and resistance. Capacitor C_5 is the target object’s coupling to ground. If a person is being sensed, C_5 is usually dominated by the capacitance through the shoes. The ground terminal may either be a ground plane in close proximity to the transmitter and receiver, or the ambient room ground.

The sensors can be used in a variety of ways, explained below, each of which modifies these capacitances differently. We measure capacitance by measuring the current arriving at the receiver, as explained in section 2.1 and 2.2.6.

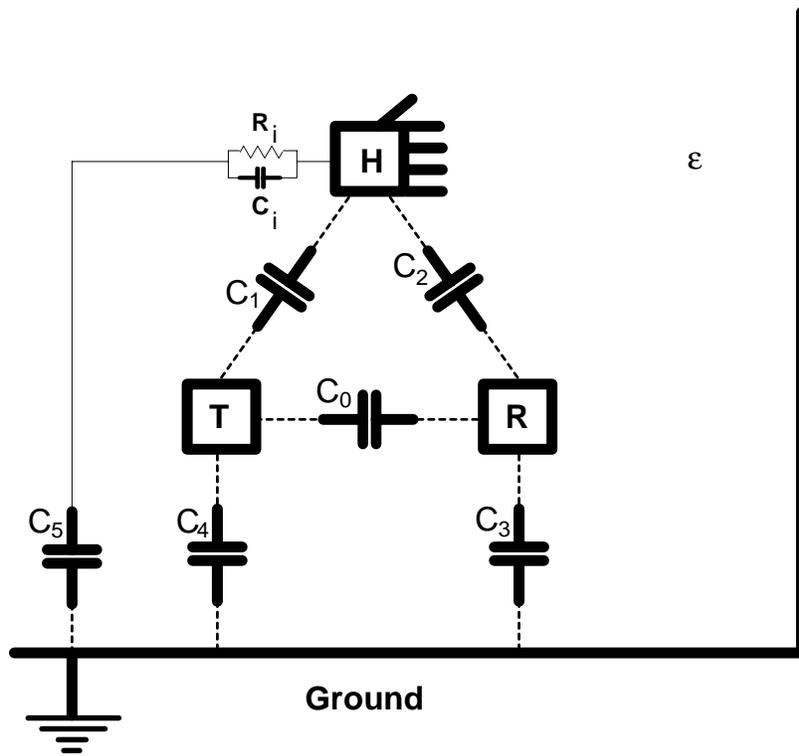


Figure 2-5: Lumped circuit model of Electric Field Sensing.

2.3.1 Transmit Mode

In transmit mode[GZA95], the transmit electrode is put in contact with the user's body, which then becomes a transmitter, either because of direct electrical connection, or AC coupling through the clothes, which is shown as current path C_1 in the circuit diagram.

When the hand moves, the spacing to the receiver changes, which changes the value of C_2 . When the spacing from the hand to the receiver is large, the received signal goes roughly as $1/r^2$, because the hand acts like a point object and the field falls off as $1/r^2$. By Gauss's law then, the induced charge on the receiver also goes as $1/r^2$. Since the potentials on the electrodes are defined by the Fish circuit, we know the capacitance to be $C = Q/V$, and the received current $I_R = 2\pi fCV$, as explained in section 2.2.6. When the hand is very close to the receiver, C_2 (typically) has the geometry of a parallel plate capacitor, and the signal goes as $1/r$.

2.3.2 Transmitter Loading Mode

The Smart Fish[PG95] has circuitry to measure the current being lost at the transmit electrode. This is known as loading mode. It has also been investigated by Vranish et al.[V+92, V+93].

If the only object in our sensing apparatus that can move is the body that we wish to sense, then any change in the amount of current leaving the transmitter indicates a change in the value of C_1 .

2.3.3 Receiver Loading Mode

Receiver loading mode is a pseudo mode, in that it relies on non-ideal characteristics of the op-amp. When a grounded object approaches the receiver, current may be shunted from the receiver, even if the receiver completely shields the transmitter from the ground.

I discovered this mode while trying to make a two-dimensional electric field mouse with touch click. The mode is useful for making very large pressure sensors that are *not* sensitive to matter until it touches the sensing surface. Additional proximity sensitive electrodes, for example those used to track the hand in the clickable two-d mouse, can be mounted on top of the touch surface; the touch sensor and the proximity sensors do not interfere with one another.

It is important for Fish users to know about this mode because it can often be an annoyance. If a receiver is placed against a wall, there is a current path from the receiver to the wall; this path appears as capacitor C_3 in the diagram. Because wall electrodes are usually placed flush against the wall, they behave like parallel plate capacitors, with a very small separation r between the plates. The reason this mode can be such a problem is that the capacitance k/r (where k is a constant) is very large. In particular, tiny changes in r , which may arise from air currents, lead to large changes in k/r , and therefore large offsets in signal. The mode can therefore introduce mysterious drifts in the sensor values, or be used to make very sensitive measurements.

2.3.4 Shunt Mode

In this thesis, we will be concerned primarily with shunt mode. Shunt mode is better suited to imaging, because it provides more informative measurements than other modes, as we

will explain below.

In shunt mode, neither the transmitter nor the receiver are in contact with the user’s body. When the user’s body is out of the field, current flows from transmitter to receiver through the effective capacitance C_0 .

When part of the user’s body, such as a hand, enters the field, it functions as a third terminal, and the capacitance matrix changes, often drastically. In particular, the value of C_0 shifts. Since the voltage between the transmitter and receiver is held constant, the change in C leads to a change in the current arriving at the receiver. From the amount of current that fails to arrive at the receiver, one can infer something—what, exactly, is the subject of the thesis—about the “amount of arm” in the vicinity of the sensor. [The rest of the body functions as a ground for the part that is in the field. It does so because, at these frequencies, the body is nearly a perfect conductor, and is a relatively large charge reservoir. The rest of the body is also capacitively coupled to an even larger charge reservoir, the earth or room ground; this current path is labeled C_5 . Both effects (the body as a ground, and the body as a wire to ground) contribute to C_5 . [Zim95]]

The sense in which shunt mode is more informative is the following: an n electrode array can make as many as $\frac{n(n-1)}{2}$ distinct measurements, since it uses pairs of electrodes. This number, rather than n^2 , arises because electrodes cannot function simultaneously as receiver and transmitter (hence the $n(n-1)$). The factor of $\frac{1}{2}$ arises because the capacitance matrix is symmetrical: we learn nothing new when we interchange a transmitter and receiver. By contrast, an n -electrode array of sensors operating in transmit or loading mode can only make n distinct measurements. To see the difference more clearly, note that each shunt mode measurement has an associated length and orientation, that of the vector between the transmitter and receiver. A pair of electrodes far apart “sees” much further than a pair of electrodes close together (because the lengthscale of the $\frac{1}{r^3}$ field dependence is set by the spacing between the electrodes), and a pair oriented horizontally responds differently than a vertically oriented pair. So while both shunt and loading mode measurements have associated positions, a shunt mode measurement has an extra geometrical parameter, this sensor orientation vector.

Relationship to other modes The modes cannot actually be separated as cleanly as I have suggested. When the hand is close to both the transmitter and receiver, both C_1 and C_2 can become large. Their sum,³ in fact, can become larger than C_0 , and the current arriving at R can be greater than the current that arrived before the field was perturbed.

We usually refer to this as “coupling mode” or “crossover mode,” meaning the pseudo mode in which the electrodes are configured for shunting, but the dominant effect is the transmit effect. As the hand approaches a transmit-receive pair in shunting configuration, the signal decreases as more displacement current is shunted from C_0 to C_1 . But as the hand approaches the transmitter and receiver more closely, the signal starts increasing as the “transmit mode” current path through C_2 grows.

Variations in C_5 can affect the sensors dramatically. The C_5 of a person wearing 10 cm thick platform shoes would be one tenth that of a person wearing shoes with 1 cm thick soles. In fact, the value of C_5 can vary even more, for example when the person is on an actual platform such as a stage.[Par] This can be a serious problem.

³Since the arm is a conductor, C_1 and C_2 are connected in parallel, and therefore add.

For mice or imaging systems to work well, they must infer the global offset C_5 , in addition to the C_1 for each receiver. We usually try to suppress C_2 , though for some applications it might be desirable to try to infer its value also.

Component Values

At the frequencies we are concerned with, the real impedance of free space is essentially infinite (capacitors block direct current), and the real impedance of the body is almost zero. Barber [BB84] gives resistivity figures on the order of $10\Omega m$ (Ohm-meters), plus or minus an order of magnitude: cerebrospinal fluid has a resistivity of $.65\Omega m$, wet bovine bone has $166\Omega m$, blood has $1.5\Omega m$, and a human arm has $2.4\Omega m$ longitudinally and $6.75\Omega m$ transverse.

Tom Zimmerman measured the capacitance between the right hand and the left foot, and found a value of 9.1 pF [Zim95]. A simple parallel plate model of feet in shoes with 1cm thick soles gives a capacitance of 35 nF, using $C = \epsilon_0 A/d$, and taking $A = 2 \text{ feet} \times 20\text{cm} \times 10\text{cm}$ and $d = 1\text{cm}$. For 10 cm thick platform shoes, the value of $C = 3.5\text{nF}$. (We have neglected the dielectric constant of the soles, and all inductive effects.)

Chapter 3

Forward Problem

3.1 General Case

As we saw in section 2.2.4, the sensor values can be determined in the most general case by solving the Laplace equation with an inhomogeneous permittivity ϵ , equation 2.20:

$$-\nabla \cdot \epsilon \mathbf{E}_0 = \nabla \cdot (\epsilon \nabla \phi_0) = \epsilon \nabla^2 \phi_0 + \nabla \epsilon \cdot \nabla \phi_0 = 0$$

However, for very simple imaging problems (such as imaging a single point, that is, making a two- or three-dimensional mouse), this model is too general. We cannot afford to solve Laplace's equation each time we move the mouse. In this chapter we will use some of the physics presented in Chapter 2 to motivate a an approximate forward model of the response of the sensor to a single point-like grounded object.

3.2 Approximate effect of a small grounded object

We want to model the effect of a small, perfect conductor h at a point \mathbf{x} in space, connected to ground through a capacitance C_5 and a wire whose effect on the field is negligible. What we want is a simple model of how the presence of the conductor affects the capacitance between the transmitter and the receiver. As a “zeroth order” approximation, we will assume that the object's effect on the geometry of the field is negligible.¹ We will be more explicit below about what is meant by the geometry of the field.

Figure 3-1 shows both the unperturbed electric field impinging on the receiver, and the perturbed field. It will be helpful to define a “two-terminal” component model of a field line. Such a component is shown in the Figure, stretching from the perturbation to the receiver. The component is defined by a small tubular surface with the property that the tangent vector of its axis is everywhere parallel to the field. Since the tube is small, its sides are also parallel to the field, and its endcaps are perpendicular to the field. Because the sides are parallel to the field, there is no flux through the sides. The only flux into or out of the tube is through the endcaps. Since the endcaps are small and perpendicular to the field, the flux through an endcap is equal to the field strength at a point on the endcap

¹It should be possible to do an expansion in the size of the perturbation, analogous to the expansion in the time-rate parameter from Chapter 2, to make the “zeroth-order” character of the approximation explicit.

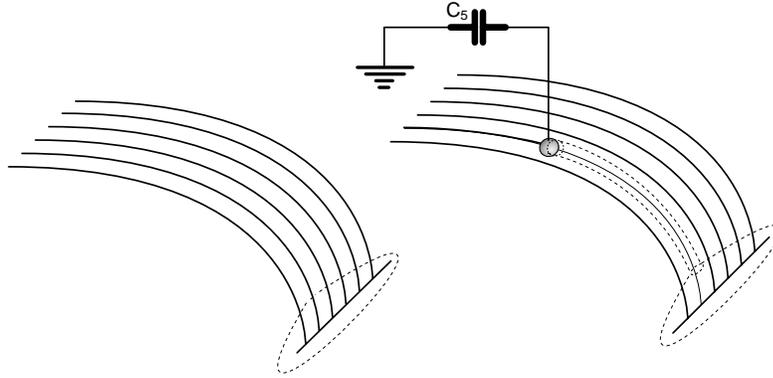


Figure 3-1: The unperturbed electric field impinging on the receiver, left, and the perturbed field, right. A two-terminal component model of a field line is shown between the perturbation and the receiver.

disk multiplied by the disk's area.

The divergence theorem says that the total flux through the tube's surface is equal to the integral of the enclosed source and sink terms ρ_f . The flux through a tube containing no sources or sinks (charge) is zero. Thus the flux leaving the bottom endcap of a sourceless tube equals the flux entering the top.

Now we can consider three contiguous tubes, the first starting at the transmit electrode (but not overlapping it) and ending just above the perturbation, the second enclosing the perturbation, and the third beginning just below and extending to the receiver. The net flux through the first tube is zero, since it contains no sources. The flux entering the small section of tube enclosing the perturbation is simply the field strength at the endcap disk multiplied by the disk's area, as explained earlier.

If the perturbation is a grounded, perfect conductor, it absorbs all the incident field lines: charge is induced on its surface in proportion to the incident flux. No flux leaves the bottom endcap, so that the net flux, rather than being zero as it would be in a charge-free region, is proportional to the induced charge.

Since none of the flux that left the first tube and entered the second arrives at the third, we can calculate an approximate change in the flux into the receiver, if we have made the "zeroth order" assumption that the object does not affect the field geometry (meaning the shapes of the tubes). The decrease in flux at the receiver is the amount that left the top tube, that is, the field strength at the perturbation, multiplied by the tube area.

If the object is connected to ground through an impedance, instead of directly, a voltage divider is a natural model to consider. The first tube would be modeled as one impedance, the third tube would be an impedance in series with it, and the connection to ground would be the load on the divider. This model is consistent with the case considered above, in which the impedance to ground is zero, and all the "current" is shunted to ground. When the impedance to ground is non-zero, some fraction of the flux entering the second tube

would also leave, and arrive at the receiver. We will be investigating the utility of this model further.

Now we have introduced a model for how an object changes the received signal that depends on the field strength at the object's location. Next we will need a model of the field itself.

3.3 Modeling the field: the dipole approximation

We will approximate the field due to a pair of small, identical, rectangular electrodes of dimension $b \times c$ and displaced from one another by a along the x -axis as a dipole with the same spacing. The dipole moment of a charge distribution is

$$\mathbf{p} = \int \mathbf{x}'\rho(\mathbf{x}')d^3x' \quad (3.1)$$

If the distribution on the electrode surfaces had a uniform value of ν ,² then 3.1 yields

$$p_x = \nu abc$$

and

$$p_y = 0$$

The expression for p_x makes sense: νbc yields the total charge on one electrode, so we could write $p_x = Qa$. Thus the pair of rectangular electrodes displaced from one another by a and charged to $+Q$ and $-Q$ has the same dipole moment as a pair of point charges $+Q$ and $-Q$ displaced by the same amount.

To justify the dipole approximation more rigorously, we would have to solve for the charge distribution on the electrode surfaces, and then perform a multipole expansion of the charge distribution. The dimensions at which the higher order terms became significant would be the limits of the approximation's validity.

3.4 Modeling the sensor response

Using the point absorber model together with the dipole approximation of the field geometry, we can model the sensor response data measured using a small grounded object as a hand phantom. Figure 3-2 shows a plot of the function $C - |E(\mathbf{x})|$, where C is a constant and $E(\mathbf{x})$ is a dipole field, given by the gradient of the dipole potential $\frac{\mathbf{p} \cdot \hat{\mathbf{r}}}{r^2}$. The dipole moment \mathbf{p} is a constant representing charge multiplied by the vector from the center of the transmitter to the center of the receiver.

Figure 3-3 shows sensor measurements along the z -axis, perpendicular to the transmit-receive axis, and the dipole response model. Since x and y are zero, the dipole model simplifies to $\mathbf{E}(0, 0, z) = \frac{1}{z^3}$. Scale and offset parameters for the distance (abscissa) and sensor value (ordinate) have been fit to the data. The function plotted is shown at the top of the graph.

²The surface charge distribution is not in fact uniform, though the surface is an equipotential. These two statements are by no means equivalent.

Comparison of dipole model with experimental data

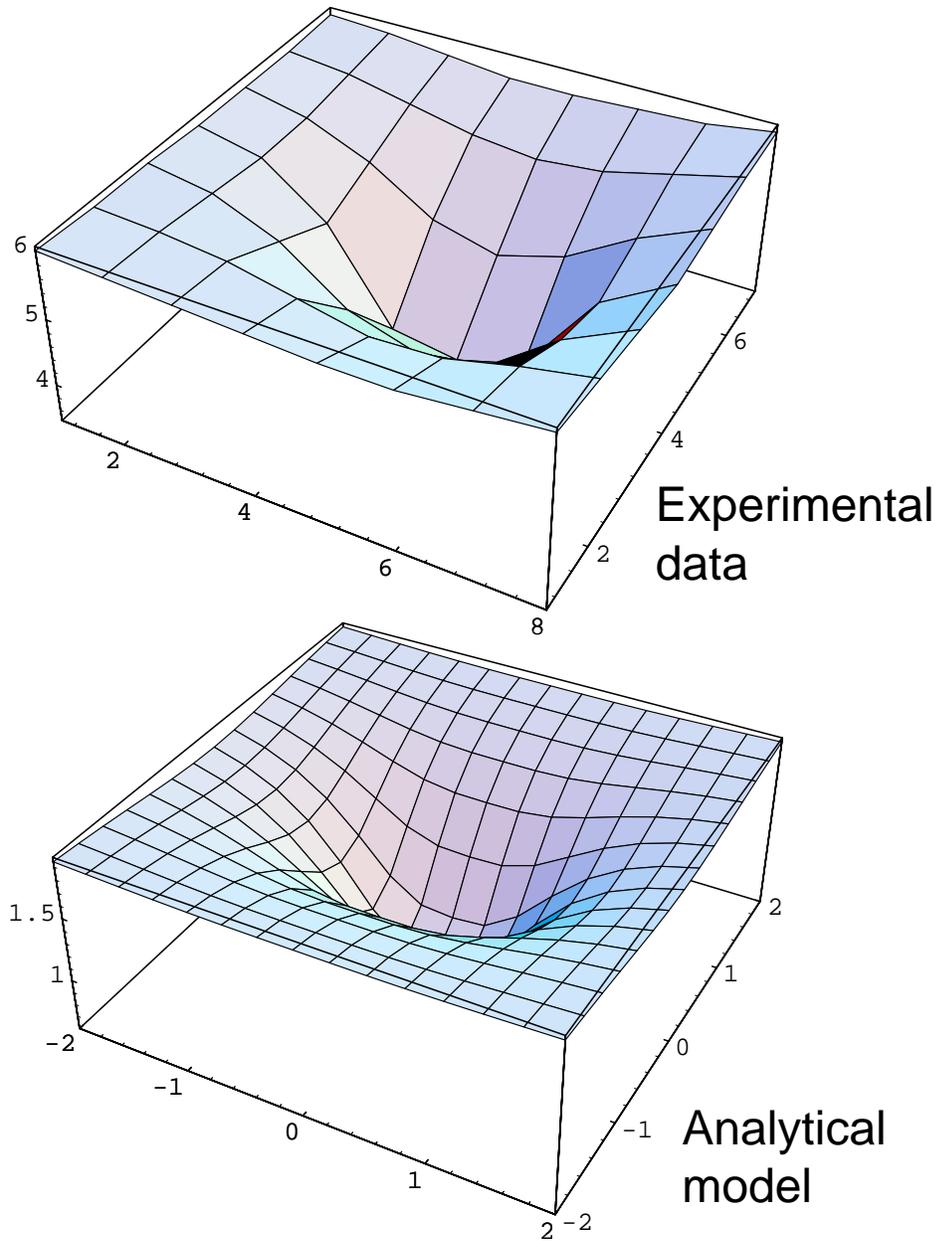


Figure 3-2: Comparison of functional form of dipole model with experimental data. Measurements were made in a plane parallel to the dipole axis, using a grounded metal cube as a hand phantom. The theoretical plot is for a plane parallel to dipole axis, at a distance of $.9a$ above the dipole axis, where a is the dipole spacing.

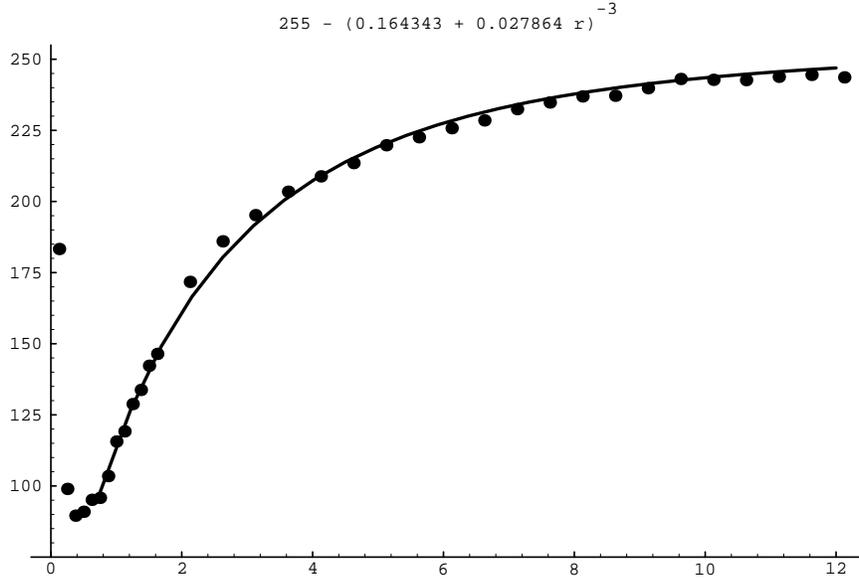


Figure 3-3: Comparison of functional form of dipole model with experimental data. Measurements were made along a line originating at the dipole origin, and extending outward, perpendicular to the dipole axis. At very short distances, transmit mode starts to dominate, and the signal rises again.

3.5 Iso-Signal Shells

Given this field model, we can plot surfaces of constant sensor readings. These plots are very helpful in shaping one’s intuition about the behavior of the sensors. The surfaces are ellipsoidal shells. The central axis of the ellipsoid is the dipole axis. Figure 3-4 shows two nested iso-signal shells, for two different sensor readings. The outer shell has been cut away to reveal the inner one. As we will see in Chapter 4, these shells, with a Gaussian thickness to represent uncertainty due to sensor noise, are the ambiguity classes for single sensor measurements.

3.6 Ground plane

The dipole approximation is also a good model of the response of a sensor in the presence of a groundplane. Note that in the presence of a groundplane, the functional form of the response of a sensor is no longer simply the field strength, since many field lines starting at the transmitter end on the groundplane, instead of on a receiver, and are therefore irrelevant to the signal measured. In the non-groundplane case, almost all the field lines leaving the transmitter arrive at the receiver, so the sensor response is proportional to the absolute field strength.

The electric field between a transmit-receive pair in the presence of a groundplane is not dipolar. The field due to a single transmitter might be modeled as a dipole, but it would be one whose axis was perpendicular to the groundplane (by the method of images). Nevertheless, the response of the sensor can be modeled almost as well by the field strength of a dipole whose axis runs between the transmitter and receiver, parallel to the groundplane.

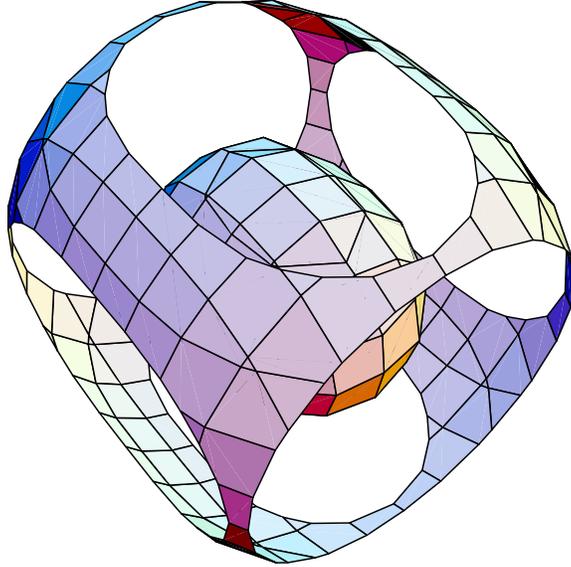


Figure 3-4: Two nested iso-signal shells, for two sensor values. The outer shell has been cut away to reveal the inner one. The dipole generating the field lies along the central axis of the ellipsoid.

Evidently the density of field lines that leave the transmitter and arrive on the receiver, rather than on groundplane, has this dipolar form, which is not surprising since this set of field lines has at least the very least the same bilateral symmetry as those produced by a pair of isolated electrodes.

We now have an analytical model for the behavior of a single transmit-receive pair in the presence of a groundplane, and this has useful consequences. Since receivers are virtual grounds, the field configuration due to, say, one transmitter, should not be changed by replacing a patch of groundplane with a receiver. In other words, the field due to a transmitter surrounded by groundplane is identical to that produced by a transmitter surrounded by a patchwork of receivers and ground. This in turn means that, in the presence of a groundplane, receivers do not affect one another—they are entirely independent dipoles. For small objects at least, we have an analytical solution to the forward problem of predicting the sensor values given a hand location, and this will be very helpful in solving the inverse problem.

3.6.1 Generalizations of hand model

It is an open research question whether the sort of approximations we have been using can be generalized and applied to more complex conductivity geometries, for example, to the case of two hands, and for the general imaging problem. If there is a way to do so, we will be able to make faster algorithms. In the worst case, we will have to solve Laplace's equation each time we need to evaluate the forward model. This is a computationally slow but well-understood procedure.

Chapter 4

Constructing the ambiguity class

In this chapter we will introduce a general probabilistic framework that will allow us solve inversion problems anywhere along the EF sensing hierarchy, and also provide a means of designing optimal sensor geometries. This approach, applied to full imaging problems, is described by Jaynes in [Jay83], Hurwitz in [HHL79], Gull and Daniell in [GD78], and Skilling and Gull in [SG83].

The essence of the approach is to view imaging (or proximity detecting, or hand finding) as an inference problem. We define a model whose parameters we wish to know, and a probability distribution over those parameters. As more data becomes available (for example, as we consider additional sensors), the volume of the feasible set of model parameters, which we will call the ambiguity class, decreases and the probability distribution becomes increasingly peaked around the “true” values of the parameters.

The inverse curvature of a peak in a particular direction gives the uncertainty of the estimate of the parameter value (or linear combination of parameter values) corresponding to that direction. The amount of information provided by a measurement can be quantified by the change in entropy of the distribution that resulted from the measurement. Ill-posed (underdetermined) problems can be made well-posed by specifying additional constraints on the feasible set. (This is the Bayesian view of regularization.) These constraints can be encoded in the prior probability distribution that defines the initial feasible set. The problem of designing sensor arrays may be posed in terms of maximizing the expected information provided by a measurement.

Since, as we saw in section 2.1.1, the sensors are subject to additive Gaussian noise, the probability of the data given some setting of model parameters is given by

$$p(D|m) = \sqrt{\pi}\sigma e^{-\frac{(D-f(m))^2}{\sigma^2}}$$

where σ is the standard deviation,¹ d is a data value, $f(m)$ is the data value predicted by our analytical forward model given a model configuration (hand position) m . This distribution is normalized: if we integrate over all values of d , we get 1. By Bayes’ theorem,

$$p(m|D) = \sqrt{\pi}\sigma \frac{e^{-\frac{(D-f(m))^2}{\sigma^2}} p(m)}{p(D)}$$

¹It will not represent conductivity in this chapter.

For the case of a two- or three-dimensional mouse, we can chose a prior $p(m)$ that renders the inversion well-posed, by, for example, restricting the possible hand positions to positive coordinate values. A useful prior for for one of the model parameters is $p(m) = \frac{c}{1+e^{-\beta x}}$, defined in some finite range of x , where c is a normalizing constant and β is a sharpness parameter. This function is a way to approximate a step function with a closed form expression. A possible advantage of using this function over a hard step function is that numerical optimization techniques are able to follow it back into the high probability region, since it is smoothly varying. The prior for our entire model is the product of the priors for x , y , and z .

$$p(m|D) \propto e^{-\frac{(D-f(m))^2}{\sigma^2}} \frac{c}{1+e^{-\beta x}}$$

Apart from the prior, which we might have chosen to be a constant over some region, the functional form of $p(m|D)$ is identical to that of $p(D|m)$. The remarkable fact that the $p(m|D)$ distribution and the $p(D|m)$ distribution, which have completely different meanings, happen always to have the same functional form is the content of Bayes' rule. However, the similarity in functional form is in some sense superficial. Consider the normalization of $p(m|D)$. Rather than performing the trivial Gaussian integral over d (trivial and Gaussian because when we integrated $p(D|m)$, m and therefore $f(m)$ was fixed), we must integrate over all values of m , which means integrating our forward model composed with a Gaussian. The difficulty of performing this integration depends on the form of f . This normalization constant, which Bayesians grandly call the evidence, is not important for finding the best setting \hat{m} of the model paramters, since a scaling of the dependent variable (probability) has no effect on the location of maxima. However, it does become important when making any sort of comparison between different functions f , or calculating entropies. The difficulty of performing this integration would be the Achilles' heel of Bayesian methods, were it not for the fact that for small σ , $p(m|D)$ can be well approximated by a Guassian, which is trivial to integrate.

Information collected by multiple sensors can easily be fused: simply multiply the $p(m|D)$ terms due to each sensor to get the joint probability of a model given all the data. Thus if we use D to denote the set of N measurements d_i , where i indexes the sensor,

$$p(m|D) \propto \prod_i^N e^{-\frac{(D_i-f_i(m))^2}{\sigma^2}} \prod_j \frac{c}{1+e^{-\beta x_j}}$$

Notice that, since log is a monotonically increasing function, if we maximize $\log p(m|D)$, we will get the same m as if we had maximized $p(m|D)$. It will be desirable in practice to work with log probabilities rather than probabilities, for several reasons: we can save computation time since exponentials disappear and multiplication and division become addition and subtraction. Furthermore, when we multiply many probabilities together, the numbers become very small, so that numerical precision can become a problem. Using log probabilities alleviates this problem, and reduces computation time, since exponentials appear so often, in each of our Gaussian probability distributions.

$$\log p(m|D) = \sum_i^N \frac{(D_i - f_i(m))^2}{\sigma^2} - \sum_j \log(1 + e^{-\beta x_j}) + c$$

This has the familiar interpretation of the sum of squared errors between the data and the data predicted by the model, with an additional error term derived from the prior.

4.1 Uncertainty and optimal sensor geometry

Once the basic degeneracies have been broken, either by collecting sufficient data or imposing constraints via a prior, so that there is a single maximum in the log probability, the uncertainty about the best setting of model parameters may be represented by the inverse Hessian matrix A^{-1} evaluated at the maximum. To see why, we will consider the Hessian and its properties. The Hessian A gives the curvature, which is a measure of confidence or certainty. In A 's eigenvector basis, in which it is diagonal, the diagonal elements (the eigenvalues) A_{ii} represent the curvature along each of the eigenvector directions (known as the principal directions). The curvatures along the principal directions are called the principal curvatures. The product of the curvatures, the Gauss curvature, which serves as a summary of the certainty at a point, is given by the determinant of A . The average curvature is given by $\frac{1}{2} \text{trace } A = \frac{A_{11} + A_{22}}{2}$. Finally, the curvature in a particular direction $v = (\cos\theta, \sin\theta)$ is given by Euler's formula:[Mor93]

$$\kappa = v^T A v = \kappa_1 \cos^2\theta + \kappa_2 \sin^2\theta$$

The inverse of A in this basis is the matrix with diagonal elements $1/A_{ii}$. Thus the inverse Hessian specifies “radii of curvature” of the probability distribution, which can be used as a measure of uncertainty. The determinant and trace of the Hessian are independent of coordinates, so we may use these as local measures of the “Gauss uncertainty” and mean uncertainty.

4.2 Entropy

The most general global measure of uncertainty is the entropy. The change in entropy of the $p(m|d)$ distribution resulting from the collection of new data measures the change in uncertainty about the values of the model parameters, including uncertainty due to multiple maxima, given a set of measurements. The change in total entropy ΔH of the ambiguity class m resulting from a measurement d_{n+1} is

$$\Delta H(m|D_{n+1}) = H(m|D_{n+1}) - H(m|D_n)$$

where

$$H(m|D_n) = \int p(m|D_n) \log p(m|D_n) dm$$

The expected change in entropy when we collect a new piece of data, that is, the change in entropy averaged over possible data values, gives a basis for comparing sensor geometries. The expected value of $H(m|D)$ is

$$I = \int p(D) H(m|D) dD$$

I is thus a measure of the quality of a sensor geometry. By analogy with coding theory, the best measurement procedure (for single measurements) reduces the entropy as much as

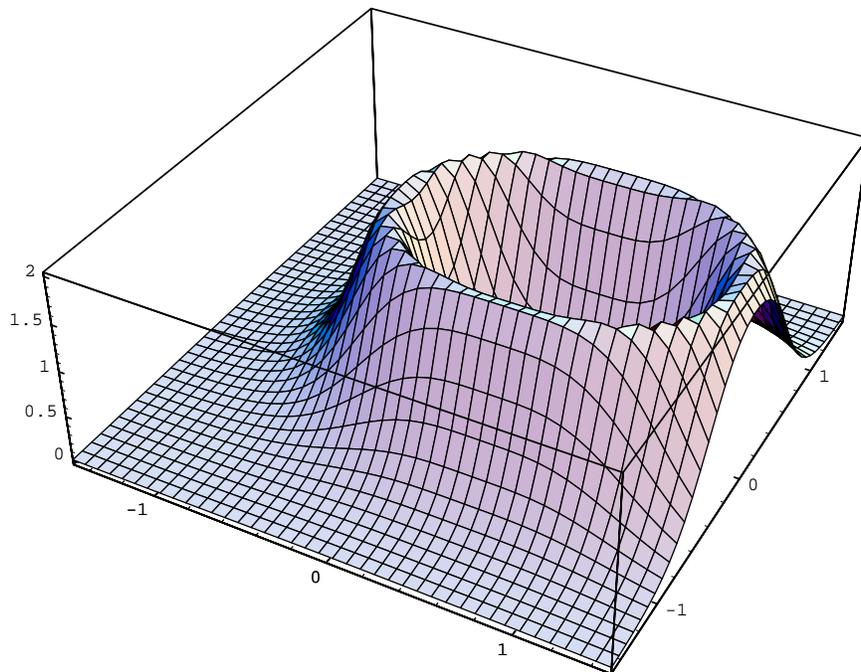


Figure 4-1: Forward probabilities for sensor 1.

possible. One could therefore search for optimal sensor geometries by maximizing I . [Mac91, Lut85, Lin56]

4.3 Example: Two dimensional mouse

Here we use this technique to construct the ambiguity class and find the most likely model parameters given two sensor readings. We want to infer the position of the hand in two dimensions from two sensor readings. So the model consists simply of two numbers, representing the position of the object purported to explain the sensor readings. The sensor axes are oriented perpendicular to one another, and the transmit electrode is shared.

Figures 4-1 and 4-2 below show the forward probability distributions $p(D|m)$ for the two sensors, oriented perpendicular to one another. To make the figure easier to view, the noise has been exaggerated dramatically. If we were to use the actual noise levels for the sensors, as measured in section 2.1.1, the features of the surfaces would be so minute that the plot routine would miss them.

Figure 4-3 shows these two distributions in the same space. Their product, the joint forward probability, $p(D1, D2|m) = p(D1|m)p(D2|m)$, is shown in Figure 4-4. The inverse probability distribution, $p(m|D1, D2)$ is the same picture, multiplied by a prior and divided by a normalizing constant.

The surfaces in Figures 4-1 through 4-3 are not normalized with respect to m (i.e. they show $p(D|m)$ not $p(m|D)$), because the heights of the two marginal distributions are not in fact the same. Their actual heights would make Figure 4-3 less clear. The important feature of that picture is the point where the straight sections of the ovals intersect perpendicularly. In Figure 4-4, which shows the normalized joint distribution (the product of the first two

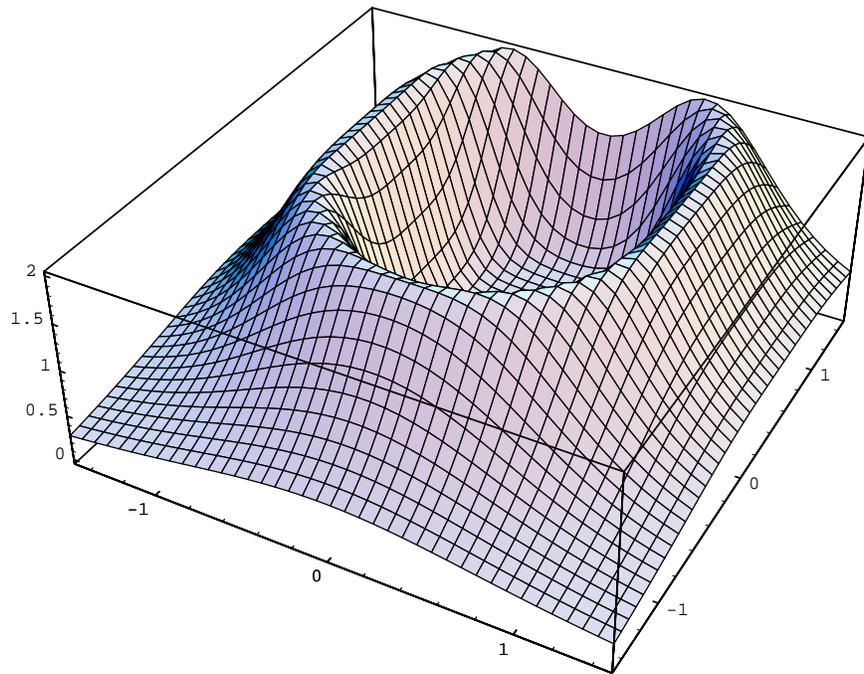


Figure 4-2: Forward probabilities for sensor 2.

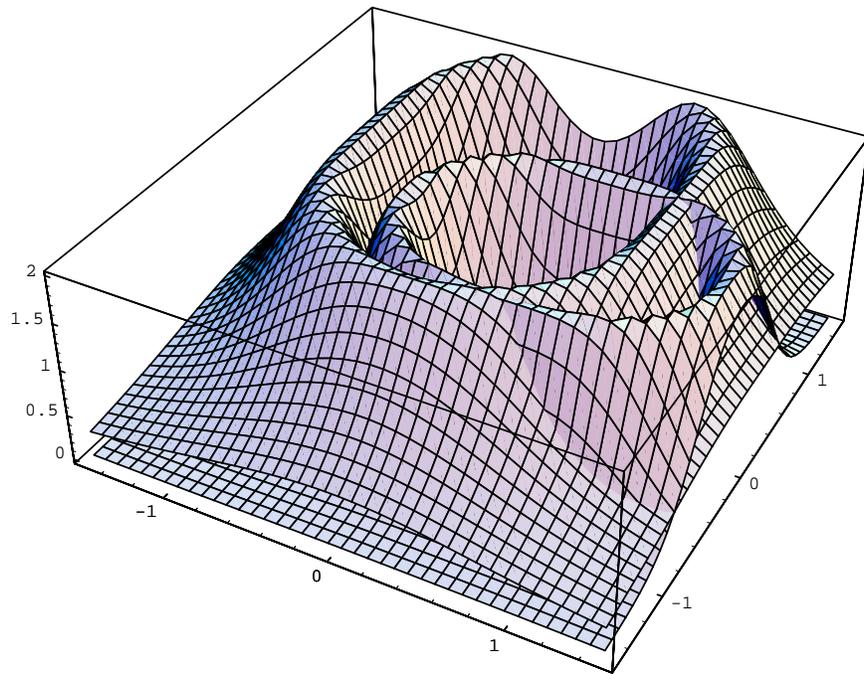


Figure 4-3: Forward probabilities for sensor 1 and sensor 2.

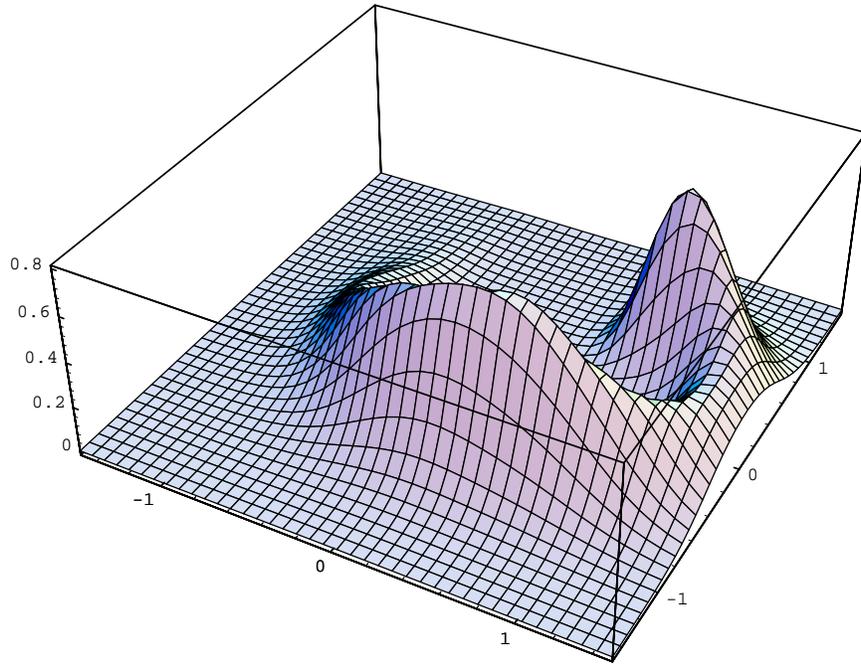


Figure 4-4: The ambiguity class: joint forward probabilities.

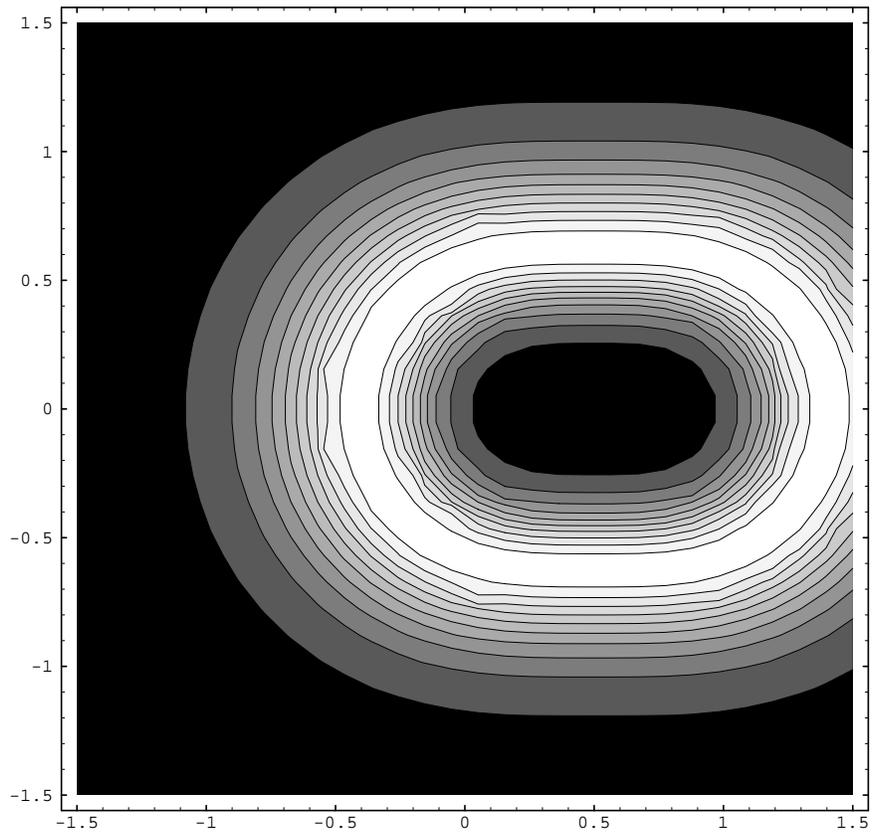


Figure 4-5: Forward probabilities for sensor 1.

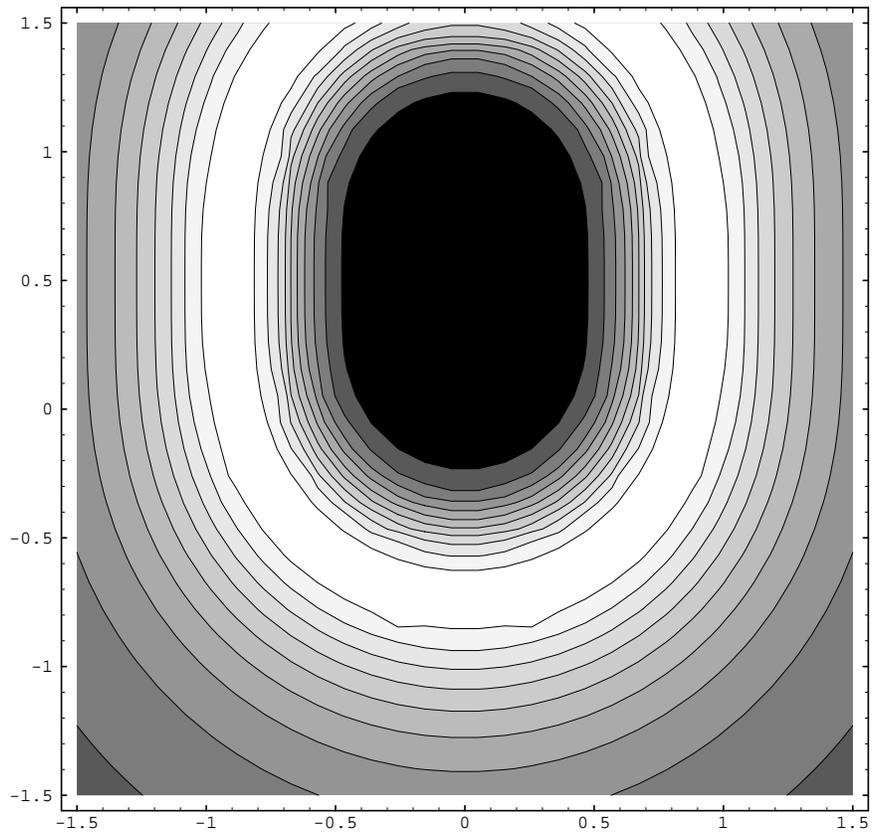


Figure 4-6: Forward probabilities for sensor 2.

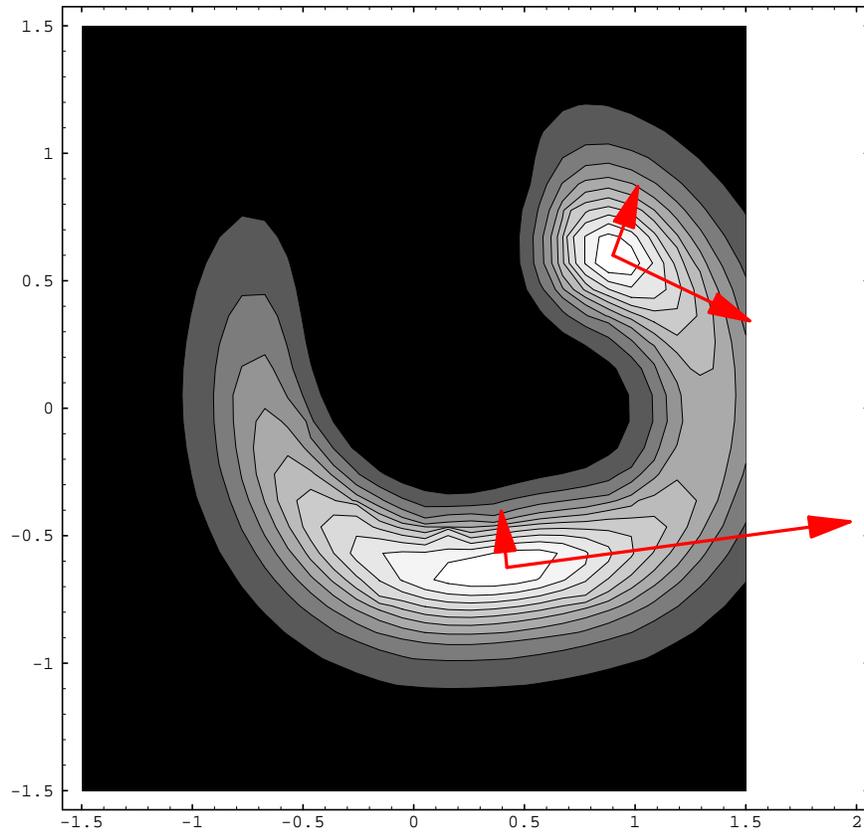


Figure 4-7: The ambiguity class: joint forward probabilities, with error bars, the principal components of the inverse Hessian evaluated at each peak. The larger error bar on the less sharply defined peak has been scaled down by $1/3$ to fit it on the page.

distributions, normalized), this intersection point appears as the sharper peak. Figure 4-7 shows a contour plot of the joint distribution, with the “principal uncertainties,” or error bars, evaluated at the two maxima, superimposed on the maxima. The three smallest arrows have been scaled up by a factor of 10 to make them more visible. The larger arrow on the less sharply defined peak has only been scaled up by $3\frac{1}{3}$, so that it fits on the page.

4.3.1 Recovering orientation

Since in our calculations we have been assuming that the hand is point-like, the uncertainties seen here are due to the field, not the hand. But if the hand were not point-like, this would presumably introduce additional uncertainties, flattening the distribution further, at least in some directions. If we “calibrated out” the intrinsic ambiguities due to the sensor layout, for example by multiplying the uncertainties at each point by the principal curvatures for the single-point ambiguity class, or by subtracting the same, then any ambiguities we did detect would presumably represent a spread in the distribution being measured. This may provide a way to estimate the size and orientation of the object: simply use the “uncertainty ellipsoid” defined by the principal curvatures.

Chapter 5

Example Inverse Problem: 3D Mouse

In this section we will present another application of the discussion in Chapter 4: a three-dimensional mouse. We will choose a sensor geometry and construct its ambiguity class for an example hand position. It is possible to check the suitability and quality of a sensor layout and prior by examining the ambiguity class: if there are multiple maxima, the inversion is ill-posed, and if the peak is not sharp (if the maximum has high radii of curvature, that is, a high value of the determinant of the inverse Hessian matrix) the value is very uncertain.

Figure 5-1 shows the layout we selected. In section 4.1 we discussed criteria for optimal sensor design. Evaluating the entropy integrals, and averaging over all possible data values, represents a substantial practical challenge. Efficient means of doing so may require sophisticated Monte-Carlo techniques, except in special cases.

Therefore, we will simply satisfy ourselves that this layout does not lead to ill-posed inversion problems by examining its ambiguity class. Figure 5-2 shows the (log) ambiguity class for our sensor layout with the hand at $.5, .5, .5$, measured in units of the sensor dipole spacing. Each image shows a slice through the three-dimensional joint probability distribution, parallel to the $X - Y$ plane.

Unlike in Chapter 4, where we plotted the probability directly because it led to a clearer figure, here we have plotted the log probability, which is the more useful quantity. In this case, the same figure made with probability instead of log probability, is not very interesting: every frame is black except one, which has a small white area in the region of the maximum. The log figure is easier to read.

Actually to invert the signals (as opposed to showing that they can be inverted), we can maximize the log probability, which corresponds to minimizing the prior term plus the sum of squared error between the measured value and that predicted by the current estimate of the hand position. We have done so for this example case, and found that the maximum was the at the expected location of $.5, .5, .5$.

For the practical mouse we have made, we have taken a cruder approach. The inversion process is computationally intensive, so we have generated a set of input-output data using the forward model, and fit to this dataset surfaces mapping sensor values to hand positions. Once the surfaces are fit, evaluating them is more computationally efficient than maximizing a function each time. However, function fitting can be very difficult. In fact, we never found a full, satisfactory map from all 3 sensor values to all 3 axes. We ended up fitting a map

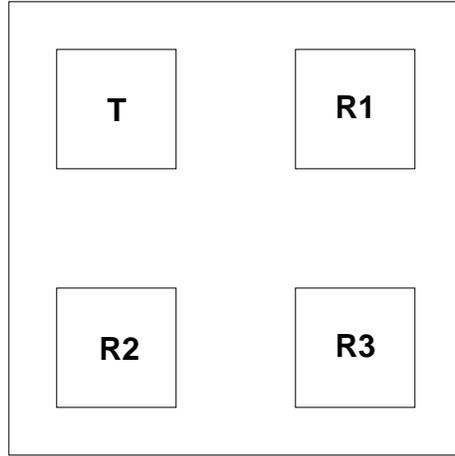


Figure 5-1: Sensor geometry for three-d mouse.

from 2 of the sensor values to x and y , and then estimating the z separately. This allows horizontal motions to be tracked well, but the approach is fundamentally flawed, and it leads to other problems: as the hand moves left and right, there is a large change in z , because the left-right change also changes the value of the third (diagonal) sensor.

The Bayesian approach, if it can be realized in real time, is free of these flaws, and gives the actual position and error, to the extent that the forward model is correct. We plan to try an alternate implementation of the mouse in which we solve the optimization problem in real time. In addition to the fitting difficulties we encountered, the fitting approach will not be applicable to harder inverse problems, in which we are trying to infer more parameters. Furthermore, possibility raised in section 4.3.1 that curvatures in the ambiguity class may reveal size and orientation information is tantalizing. Finally, even if some other inversion technique is employed, the Bayesian formalism we have introduced can be used offline to design optimal sensor configurations, by minimizing the expected change in the entropy of the ambiguity class resulting from a typical measurement.

Figure 5-3 shows a screen-shot of the mouse.¹ The user's hand motion is mapped onto the motion of the hand icon. The hand can pick up the small cube shown, move it around the space, and set it back on the floor. Because we cannot yet extract hand size, we have used a "sticky hand, sticky floor" protocol for grasping and releasing the cube. The small cube starts on the floor. When the hand first touches the cube, the hand closes and the cube "sticks" to the hand and moves with it until the hand returns to the floor, at which point the hand opens and the cube sticks to the floor, where it remains until the hand returns.

¹The 3d graphics for the mouse was created by Barrett Comisky.

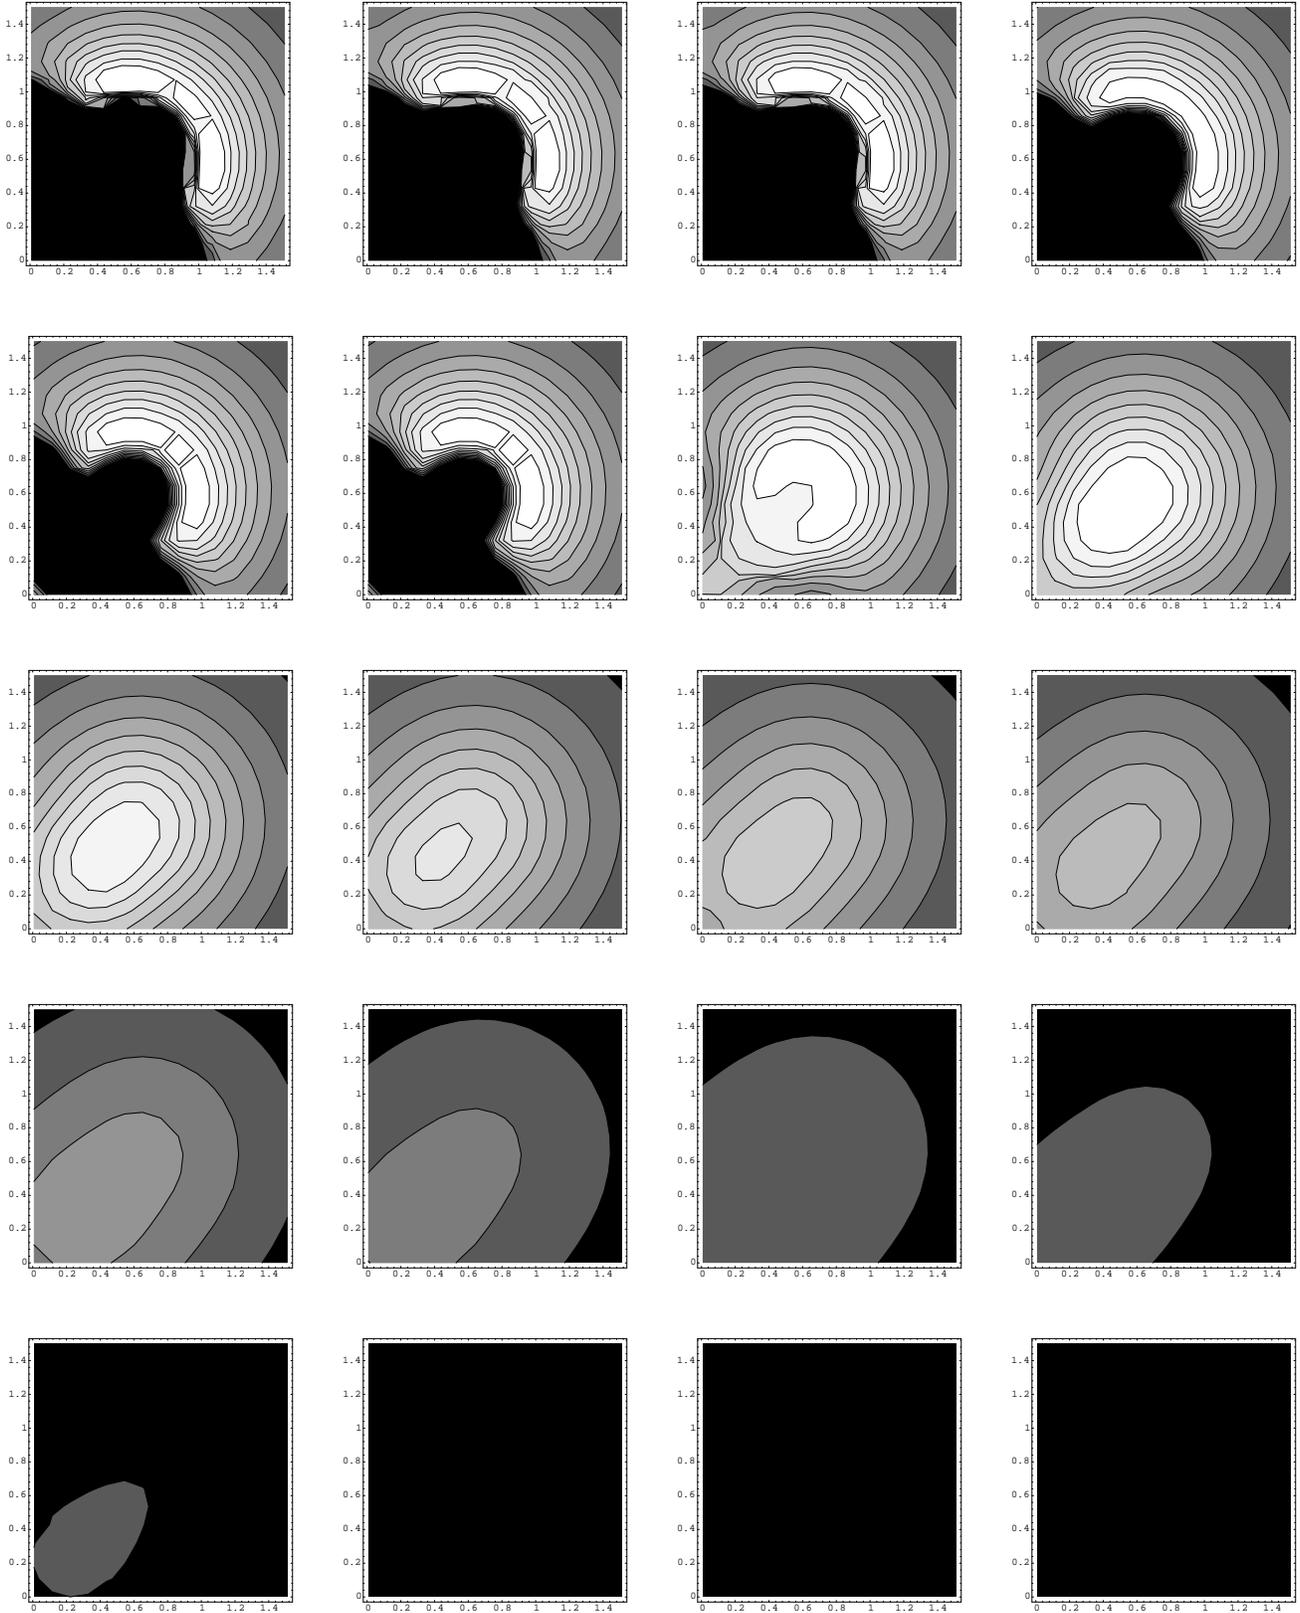


Figure 5-2: Ambiguity class for sensor geometry pictured in the previous figure, with hand at location $.5, .5, .5$.

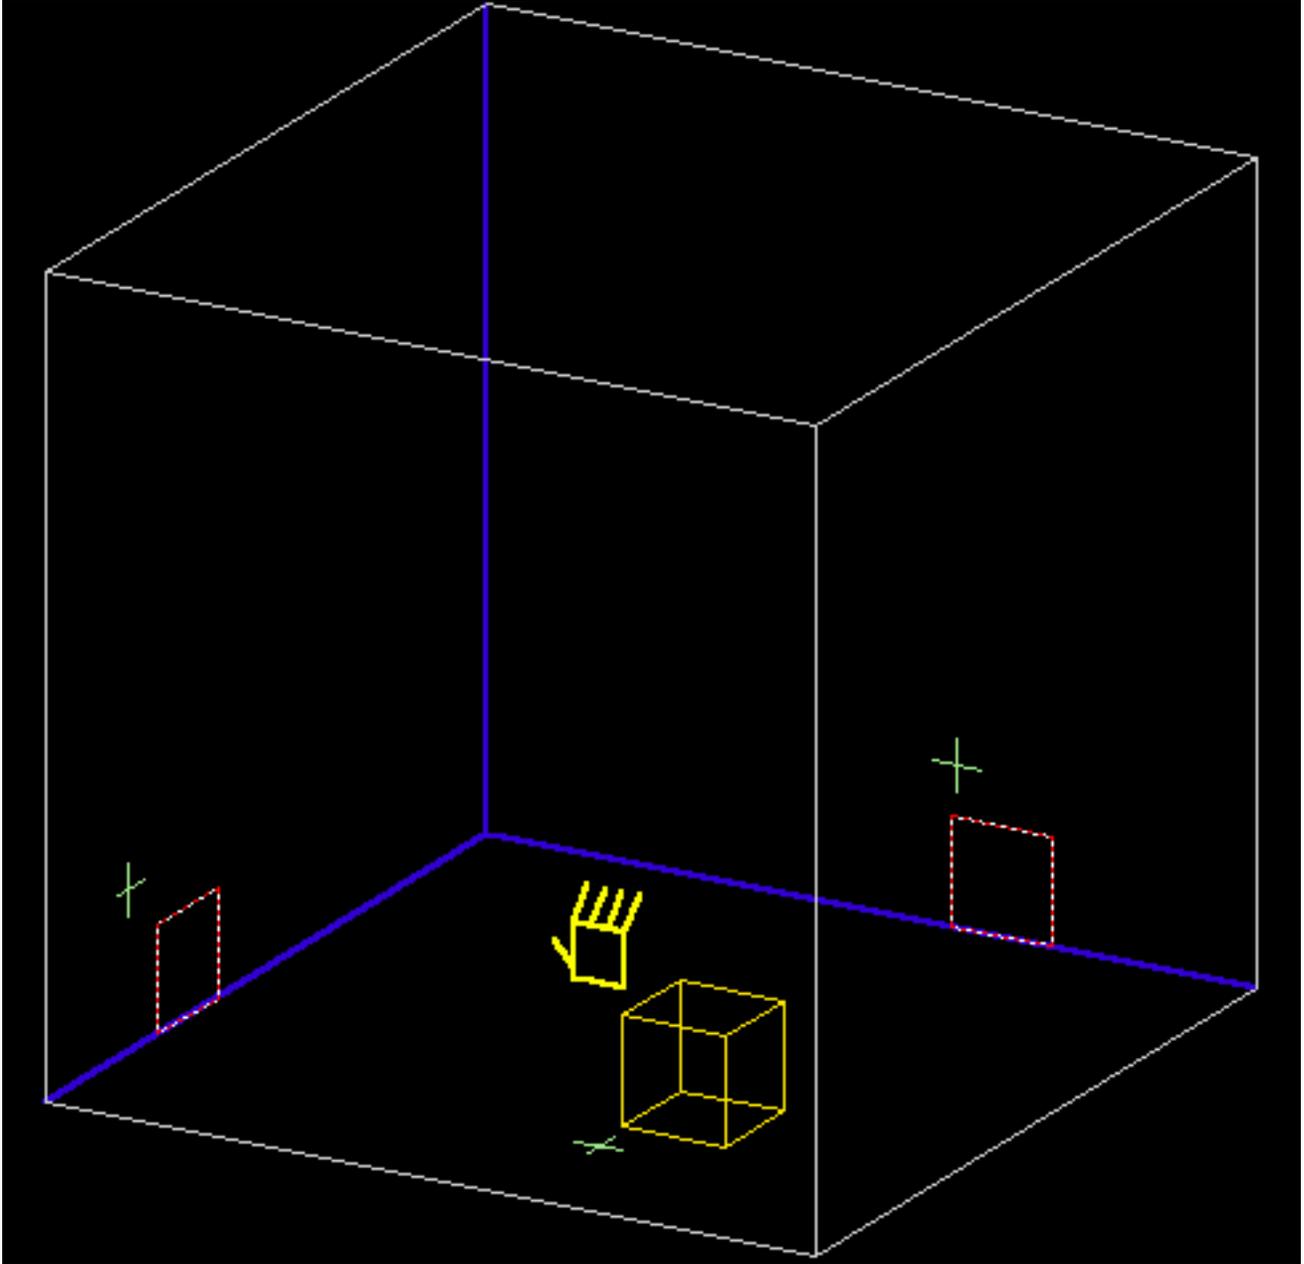


Figure 5-3: Three-d mouse.

Chapter 6

Conclusion

We have taken the first steps toward Electric Field Tomography, moving from simple proximity detector to 3-d mouse. We have introduced a general mathematical framework with which it should be possible to solve the inversion and optimal sensor design problems for any number of sensors.

In Chapter 2, we explained the physics of Electric Field Sensing. Starting from Maxwell's equations, we found that Laplace's equation with an inhomogeneous permittivity is the general solution to the forward problem of predicting the sensor values given a permittivity distribution and boundary conditions. But with an eye toward the simplified problem of imaging a single point, in Chapter 3 we introduced an approximate analytical forward model. This model has two parts: one describes how the object interacts with the field, and the other describes the field itself.

Next, in Chapter 4 we introduced a framework that allows us to discuss both inversion and optimal sensor placement. Given a forward model, in this framework we define a probability distribution over the model parameters. Additional data sharpens the distribution around the "true" parameter values. The most likely set of model parameters is found by maximizing the probability, or log probability. The uncertainty of the final estimate of model parameters (the error bars) is given by the curvature of the log probability. An optimal sensor layout can be chosen by minimizing the expected change in the entropy of the ambiguity class. Finally, we showed how to use the Bayesian framework to image a single point in two and three dimensions.

One open question is whether the forward model introduced in Chapter 3 will continue to be useful with more complicated distributions, for example with multiple small objects. As long as the approximation on which it is based remains valid, it should work with multiple objects. The only question is whether multiple objects start to distort the field so much that the zeroth order approximation is no longer good enough. Answering this question will require additional experimental and theoretical work. Another issue that must be addressed is that of incorporating a variable impedance to ground into our model. A suggestion for how to do this was presented at the end of section 3.2, but it must be tried experimentally.

If the forward model can be generalized successfully, then there is truly a path from the 3-d mouse to imaging. If it cannot be, then we will have to take a brute force approach, and solve Laplace's equation numerically in order to image more complex geometries. Even if our present forward model cannot be generalized, however, all the inference and inversion techniques introduced here are applicable, since they treat the forward model as a black

box. And even if the 3-d mouse is not a path to imaging, developing it for its own sake, perhaps by using the principle curvatures of the log probability to extract orientation and size information, will be interesting and useful. If the long-term project is successful, we will be able unobtrusively to extract 3-d shape information using electric field measurements, which could profoundly affect the way people interact with machines.

Appendix A

Backprojection Algorithm

The standard backprojection algorithm is a computationally efficient means of recovering an image of the cross section of an object from a set of projections. A projection is a line integral of some parameter of the object, for example X-Ray opacity (as in CT-scans). We will restrict the discussion in this Appendix to the case of parallel projections, that is, integrals along straight lines that are parallel to one another. Figure A-1 shows an object $f(x, y)$, and (schematically) a single parallel projection $P_\theta(t)$.

The mathematical basis of the backprojection algorithm is the Fourier Slice Theorem, which equates the Fourier transform of $P_\theta(t)$ with a sample (or slice) along the projection angle θ , of $F(u, v)$, the two-dimensional Fourier Transform of $f(x, y)$. The proof of the Fourier Slice Theorem, following Kak and Slaney [KS88], is straightforward:

Consider the coordinate system (t, s) , a rotated version of (x, y) :

$$\begin{bmatrix} t \\ s \end{bmatrix} = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \quad (\text{A.1})$$

In this coordinate system, a projection along lines of constant t is written

$$P_\theta = \int_{-\infty}^{\infty} f(t, s) ds$$

Its Fourier Transform is given by

$$S_\theta(\omega) = \int_{-\infty}^{\infty} P_\theta(t) e^{-2\pi i \omega t} dt = \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} f(t, s) ds \right] e^{-2\pi i \omega t} dt = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(t, s) e^{-2\pi i \omega t} ds dt$$

We can transform the expression for $S_\theta(\omega)$ back into the (x, y) frame using A.1:

$$S_\theta(\omega) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) e^{-2\pi i \omega (x \cos \theta + y \sin \theta)} dx dy$$

This expression represents the two-dimensional Fourier Transform at a spatial frequency of $(u = \omega \cos \theta, v = \omega \sin \theta)$, that is, along a line through the origin at angle θ :

$$S_\theta = F(\omega, \theta) = F(\omega \cos \theta, \omega \sin \theta)$$

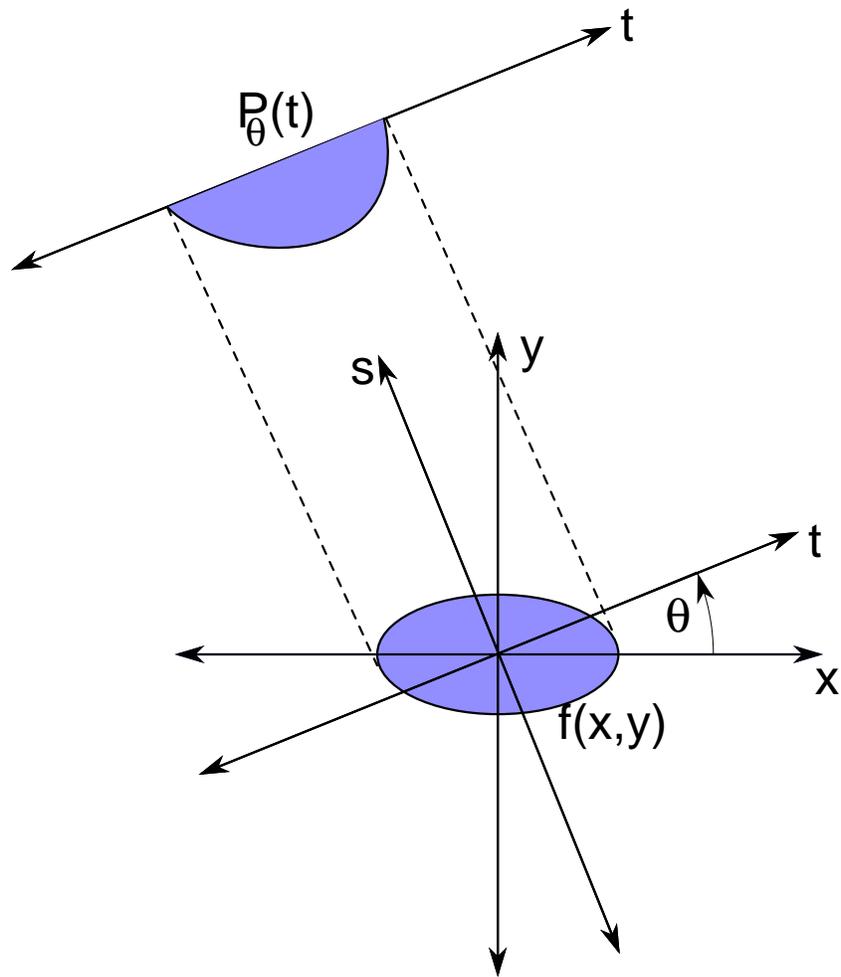


Figure A-1: The projection $P_{\theta}(t)$ of an object $f(x, y)$.

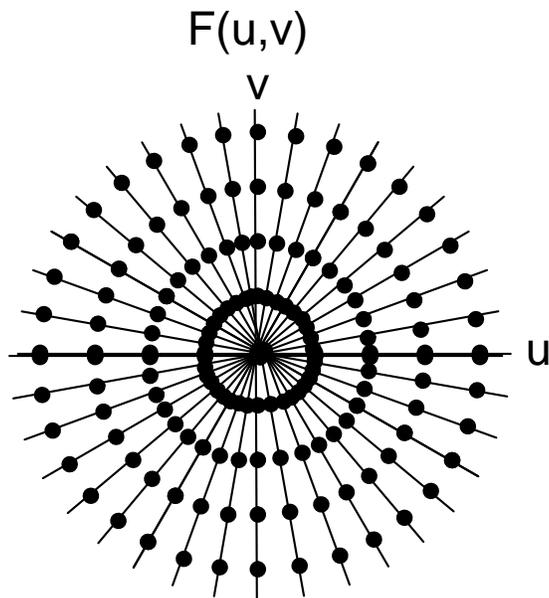


Figure A-2: The non-uniform sampling of k -space corresponding to a set of projections 10 degrees apart.

Now that we have explained the Fourier Slice Theorem, we can use it to analyse the process of collecting projections. This will lead us to the backprojection algorithm. The following explanation of backprojection is a combination of Kak[KS88], Rosen[Ros95], and my own explanation.

The Fourier Slice Theorem tells us that a series of projections in increments of 10 degrees, for example, provides a set of samples of the Fourier transform plane along a set of lines 10 degrees apart, as shown in Figure A-2. To reconstruct the image $f(x, y)$, we might naively sum all the samples of the Fourier Transform plane and then take the inverse transform. But this would introduce a systematic error. The low spatial frequencies in the center of the Fourier Transform plane have been sampled more densely than the higher spatial frequencies at the outskirts of the plane, because we sampled along lines through the origin. In particular, the origin (the spatial DC component) received a contribution from each sample line. To correct this, in the we could scale each frequency domain sample by $|k_r|$, where $|k_r|$ is the distance of the sample from the origin in k space (the Fourier Transform plane). This is a high-pass filtering operation. Notice that this removes the overemphasized DC component entirely, which is no loss, since it contained no spatial information (just a global “brightness” that would typically be adjusted to make the picture look best anyway).

The final reconstruction could be found by adding together the two-dimensional inverse Fourier Transform of each weighted slice. But to carry out the reconstruction, we do not actually have to perform the forward and inverse Fourier transforms. We can implement the filtering (weighting), by convolving each projection with the Fourier transform h of our one-dimensional “ramp” weighting/filtering function $|k_r|$. The (t, s) space counterpart of adding together the weighted k -space samples is adding the functions $h(t) * P_\theta(t)$. Since this function has no explicit or implicit s dependence, $h(t) * P_\theta(t)$ has the same value for

all s . This means that the value of pixel $p(s, t)$ is given by

$$p(s, t) = \sum_{\theta} h(t) * P_{\theta}(t)$$

This means that we are “smearing” the values $h(t) * P_{\theta}(t)$ along all values of s , or backprojecting, and summing the backprojections for each θ value. This is the filtered backprojection algorithm.

Note that if we forget the k -space picture completely, backprojection still makes some intuitive sense: we smear the measured values back along the set of pixels that could have contributed to the measurement; additional projections sharpen the image, because the pixels where the object actually is receive contributions from many projections. Nevertheless, we are always left with an unwanted spatial “DC” term, in all the pixels that received contributions from, say, just one projection, so we perform a highpass filtering operation to eliminate this DC component.

The backprojection algorithm can be generalized to handle non-parallel, non-straightline projections. The backprojection algorithm used in Electrical Impedance Tomography backprojects along curved equipotentials. This is discussed further in Chapter 1.

The set of pixels $p(s, t)$ along which we “smear” the measured value $h(t) * P_{\theta}(t)$ bears some resemblance to the ambiguity class for the projection measurement, to use the terminology of Chapter 4. However, the ambiguity class in fact is more general. It is the probability distribution over all possible states of the model. So in this case, it is the probability distribution over each possible brightness setting of the pixel values in the set $p(s, t)$. Backprojection may be viewed as an approximation of this approach in which the probability distribution over the possible brightness values of each pixel is summarized by the mean value. Using the Bayesian approach with the Gaussian approximation amounts to keeping the first two moments of the ambiguity class, instead of just the mean. This allows us to put error bars on each pixel brightness.

Bibliography

- [ARS95] D. Allport, E. Rennison, and L. Strausfeld. Issues of gestural navigation in abstract information spaces. In *CHI 95 Human Factors in Computing Systems*, pages 206–207, Denver, Co, 1995. ACM Press.
- [AS91] A. Allers and F. Santosa. Stability and resolution analysis of a linearized problem in electrical impedance tomography. *Inverse Problems*, 7:515–533, 1991.
- [AZP⁺95] D. Allport, T.G. Zimmerman, J.A. Paradiso, J.R. Smith, and N. Gershenfeld. Electric field sensing and the flying fish. Article submitted to ACM special issue on Multimedia and Multisensory Virtual Worlds, December 1995.
- [Bas94] J. Bastian. Electrosensory organisms. *Physics Today*, pages pp.30–37, February 1994.
- [BB84] D.C. Barber and B.H. Brown. Applied potential tomography. *Journal of Physics E*, 17:723–733, 1984.
- [Ber89] J.G. Berryman. Fermat’s principle and nonlinear traveltime tomography. *Physical Review Letters*, 62:2953–2956, 1989.
- [BK90] J.G. Berryman and R. V. Kohn. Variational constraints for electrical-impedance tomography. *Physical Review Letters*, 65:325–328, 1990.
- [Cha94] D. Chase. The human factor: Figuring it out with performance animation. *Millimeter*, February:pp.34–44, 1994.
- [Dev95] Analog Devices. Analog circuit data book, 1995.
- [DL81] K.A. Dynes and R.J. Lytle. Analysis of electrical conductivity imaging. *Geophysics*, 46:1025–36, 1981.
- [FCA60] R.M. Fano, L.J. Chu, and R.B. Adler. *Electromagnetic Fields, Energy, and Forces*. John Wiley & Sons, New York, 1960.
- [Gal91] B.M. Galeyev. L.S. Termen: Faustus of the twentieth century. *Leonardo*, 24(5):pp.573–579, 1991.
- [GD78] S.F. Gull and G.J. Daniell. Image reconstruction from incomplete and noisy data. *Nature*, 272:686–690, 1978.
- [Ger91] N. Gershenfeld. Sensors for real-time cello analysis and interpretation. In *Proceedings of the ICMC*, Montral, Canada, 1991.

- [Ger93] N. Gershenfeld. Method and apparatus for electromagnetic non-contact position measurement with respect to one or more axes. *U.S. Patent No. 5,247,261*, Sept 21, 1993.
- [GZA95] N. Gershenfeld, T. Zimmerman, and D. Allport. Non-contact system for sensing and signalling by externally induced intra-body currents. *U.S. Patent Application*, May 8, 1995.
- [HHL79] G.T. Herman, H. Hurwitz, and A. Lent. A Bayesian analysis of image reconstruction. In G.T. Herman, editor, *Image Reconstruction From Projections: Implementations and Applications*, pages pp.85–103. Springer-Verlag, 1979.
- [HM89] H.A. Haus and J.R. Melcher. *Electromagnetic Fields and Energy*. Prentice Hall, Englewood Cliffs, NJ, 1989.
- [Jay83] E.T. Jaynes. Prior information and ambiguity in inverse problems. *SIAM AMS Proceedings*, 14:151–166, 1983.
- [KM90] R.V. Kohn and A. McKenney. Numerical implementation of a variational method for electrical impedance tomography. *Inverse Problems*, 6:389–414, 1990.
- [KS88] A.C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*. IEEE Press, New York, 1988.
- [KV83] R.V. Kohn and M. Vogelius. Identification of an unknown conductivity by means of measurements at the boundary. *SIAM AMS Proceedings*, 14:113–123, 1983. A less formal and more accessible introduction to the approach developed in the other papers by the authors.
- [KV84] R.V. Kohn and M. Vogelius. Determining the conductivity by boundary measurements. *Communications on Pure and Applied Mathematics*, 38:289–298, 1984.
- [KV85] R.V. Kohn and M. Vogelius. Determining the conductivity by boundary measurements II, interior results. *Communications on Pure and Applied Mathematics*, 38:643–667, 1985.
- [Lab95] Lawrence Livermore National Lab. Micropower impulse radar specifications, 1995.
- [Lin56] D.V. Lindley. On a measure of the information provided by an experiment. *Annals of Mathematical Statistics*, 27:pp.986–1005, 1956.
- [Lut85] S.P. Luttrell. The use of transinformation in the design of data sampling schemes for inverse problems. *Inverse Problems*, 1:pp.199–218, 1985.
- [Mac91] D.J.C. MacKay. *Bayesian Methods for Adaptive Models*. PhD thesis, CalTech, 1991.
- [Mat90] M. Matthews. Three dimensional baton and gesture sensor. *U.S. Patent No. 4,980,519*, December 25, 1990.

- [Mor93] F. Morgan. *Riemannian Geometry: A Beginner's Guide*. Jones and Bartlett, Boston, 1993.
- [Mot91] Motorola. MC68HC11 reference manual, 1991.
- [Nic93] M. Nicholl. Good vibrations. *Invention and Technology*, Spring:pp.26–31, 1993.
- [Oth92] A.N. Other. Leon Theremin. *Keyboard*, February:pp.47–54, 1992.
- [Par] J.A. Paradiso. Personal Communication.
- [Par94] J.A. Paradiso. A simple technique for measuring axial displacement in stretched-wire alignment systems. Technical Report GEM-TN-94-60, Draper Labs GEM Collaboration, May 1994.
- [PG95] Physics and Media Group. *Smart Fish Technical Manual*, 1995.
- [PGng] J.A. Paradiso and N. Gershenfeld. Electric Field Sensing. *Computer Music Journal*, Forthcoming.
- [PM94] J.A. Paradiso and D.R. Marlow. Electronics for precision alignment of the gem muon system. In *Proceedings of the International Conference on Electronics for Future Colliders*, number GEM-TN-94-636. LeCroy Corporation, May 11-12 1994.
- [Rob94] B. Robertson. Caught in the act: Motion captured from live beings helps animators create animations faster and makes live performances of 3d characters possible. *Computer Graphics World*, September:pp.23–28, 1994.
- [Ros95] B. Rosen. Course notes for Magnetic Resonance: Analytic, biochemical, and imaging techniques. MIT course 22.561J, 1995.
- [SG83] J. Skilling and S.F. Gull. The entropy of an image. *SIAM AMS Proceedings*, 14:167–188, 1983.
- [SPZG95] J.R. Smith, J.A. Paradiso, T.G. Zimmerman, and N. Gershenfeld. Activating space with electric field sensing. *SIGGRAPH Technical Sketch*, 1995.
- [SS69] L. Smith and D.H. Sheingold. Noise and operational amplifier circuits. *Analog Dialogue*, 3(1):pp.19–31, 1969.
- [SU87] J. Sylvester and G. Uhlman. A global uniqueness theorem for an inverse boundary value problem. *Annals of Mathematics*, 125:153–169, 1987.
- [V⁺92] Vranish et al. Driven shielding capacitive proximity sensor. *U.S. Patent No. 5,166,679*, November 24, 1992.
- [V⁺93] Vranish et al. Phase discriminating capacitive array sensor system. *U.S. Patent No. 5,214,388*, May 25, 1993.
- [Wax95] D. Waxman. Digital theremins: Interactive musical experiences for amateurs using electric field sensing. Master's thesis, MIT Media Lab, 1995.

- [Web89] J.G. Webster, editor. *Electrical Impedance Tomography*. Adam Hilger, New York, 1989.
- [Z⁺91] M. Zadehkoochak et al. Spectral expansion analysis in electrical impedance tomography. *Journal of Physics D: Applied Physics*, 24:1911–1916, 1991.
- [Zim95] T.G. Zimmerman. Personal Area Networks (PAN): Near-field intra-body communication. Master’s thesis, MIT Media Lab, 1995.
- [ZSP⁺95] T.G. Zimmerman, J.R. Smith, J.A. Paradiso, D. Allport, and N. Gershenfeld. Applying electric field sensing to human-computer interfaces. In *CHI 95 Human Factors in Computing Systems*, pages 280–287, Denver, Co, 1995. ACM Press.