

September 1992

LIDS-TH-2140

Research Supported by:

*National Science Foundation
Grant NSF 9015281-MIP*

*Army Research Office
Grant ARO DAAL03-92-G-0115*

*Office of Naval Research
Grant N00014-91-J-1004*

**Estimation of Dynamically Evolving
Ellipsoids with Applications to
Cardiac Imaging**

Seema Jaggi

September 1992

LIDS-TH-2140

ESTIMATION OF DYNAMICALLY EVOLVING ELLIPSOIDS
WITH APPLICATIONS TO CARDIAC IMAGING

by

Seema Jaggi

This report is based on the unaltered thesis of Seema Jaggi, submitted to the Department of Electrical Engineering and Computer Science in partial fulfillment of the requirements for the degree of Master of Science at the Massachusetts Institute of Technology in September 1992. This research was conducted at the Laboratory for Information and Decision Systems with support provided in part by the National Science Foundation under grant 9015281-MIP, Army Research Office under grant DAAL03-92-G-0115 and the Office of Naval Research under grant N00014-91-J-1004.

Massachusetts Institute of Technology
Laboratory for Information and Decision Systems
Cambridge, MA 02139

**Estimation of Dynamically Evolving Ellipsoids
with Applications to Cardiac Imaging**

by

Seema Jaggi

Submitted to the Department of Electrical Engineering and
Computer Science

in partial fulfillment of the requirements for the degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1992

© Massachusetts Institute of Technology 1992. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
September 11, 1992

Certified by
Alan S. Willsky
Professor
Thesis Supervisor

Certified by
W. Clem Karl
Research Scientist
Thesis Supervisor

Accepted by
Campbell L. Searle
Chairman, Departmental Committee on Graduate Students

Estimation of Dynamically Evolving Ellipsoids with Applications to Cardiac Imaging

by

Seema Jaggi

Submitted to the Department of Electrical Engineering and Computer Science
on September 11, 1992, in partial fulfillment of the
requirements for the degree of
Master of Science

Abstract

This thesis deals with a method to obtain the ejection fraction of the left ventricle of the heart from a gated set of planar myocardial perfusion images. Ejection fraction is defined as the ratio of the fully contracted left-ventricular volume to the fully expanded left-ventricular volume and is known as an effective gauge of cardiac function. This method is proposed as a safer and more cost effective alternative to currently used radionuclide ventriculographic based techniques.

To formulate this estimate of ejection fraction, we employ geometric reconstruction and recursive estimation techniques. The left ventricle is modelled as a three-dimensional ellipsoid. Projections of this left ventricular ellipsoid are two-dimensional ellipsoids, which we use to model the left-ventricular outline in the observed myocardial perfusion images. The ellipsoid that approximates the left ventricle is reconstructed using Rauch-Tung Striebel smoothing which combines the observed temporal set of projection images with an evolution model to produce the best estimate of the ellipsoid at any point given all the data. Ejection fraction is calculated from the reconstructed ellipsoids. We study the error introduced in the reconstruction by imperfect knowledge of the ellipsoid dynamics and projection geometry through simulations. In addition, we use simulations to investigate model identification as a method to identify which of two hypothesized models best approximates the true ellipsoid dynamics. Finally, we test these techniques on real myocardial perfusion data.

Thesis Supervisor: Alan S. Willsky

Title: Professor

Thesis Supervisor: W. Clem Karl

Title: Research Scientist

Acknowledgments

First and foremost, I would like to thank my thesis supervisors, Alan Willsky and Clem Karl. Alan's experience, insight, and suggestions have served as an invaluable guide for the work in this thesis. There is no doubt that his high standards have shaped my approach to research. I would like to thank Clem for his guidance, suggestions and contagious energy. Above all, I would like to thank both my supervisors for their enthusiasm and encouragement.

This work would not have been possible without Charles Boucher, M.D., and David Chesler, Sc.D., who provided the myocardial perfusion data used in this thesis.

I would also like to thank my associates in the laboratory. Peyman Milanfar has always been willing to offer suggestions and share his experiences. I would like to give special thanks to Rajesh Pankaj, Stefano Casadei, Rick Barry, and Sheila Hegarty, who have all made my work more pleasant. Finally, I would like to acknowledge my friends, Rachel Learned and Mike Daniel, who have been wonderfully supportive.

Contents

1	Introduction	11
1.1	Overview	11
1.2	Contributions	13
1.3	Organization	14
2	Medical Background	15
2.1	Objective	15
2.2	Left Ventricular Ejection Fraction	16
2.3	Radionuclide Ventriculography	17
2.3.1	Procedure	18
2.3.2	First Pass Technique	18
2.3.3	Multiple Gated Blood Pool Technique	19
2.4	Proposed Approach	20
2.4.1	Myocardial Perfusion Imaging	20
2.4.2	Benefits	22
3	Geometric Reconstruction Background	23
3.1	Objective	23
3.2	Problem Formulation	24
3.2.1	Matrix Representation	24
3.2.2	Vector Representation	30
3.3	Ellipsoid Reconstruction	34
3.3.1	Static Reconstruction	34

3.3.2	Dynamic Reconstruction	34
4	Simulated Ellipsoid Reconstruction	36
4.1	Objective	36
4.2	Simulation Setup	37
4.2.1	Generation of Data	37
4.2.2	Reconstruction : Practical Considerations	39
4.3	Two-Dimensional Ellipsoid Reconstruction	45
4.3.1	Generation	46
4.3.2	Estimation	48
4.4	Three-Dimensional Ellipsoid Reconstruction	66
4.4.1	Generation	66
4.4.2	Estimation	66
4.5	2D and 3D Ejection Fractions	69
4.6	Conclusions	73
4-A	2D and 3D Ejection Fractions	75
5	Model Identification	77
5.1	Objective	77
5.2	Background	78
5.3	Simulation Setup	80
5.4	Experiment #1 : Matched Contraction	81
5.5	Experiment #2 : Contraction Mismatch	89
5.6	Experiment #3 : Contraction and Rotation Mismatch	94
5.7	Model Identification of 3D Dynamics	95
5.8	Conclusions	99
6	Real Data	100
6.1	Objective	100
6.2	Data Description	102
6.3	2D Processing	103

6.3.1	2D Model Identification	105
6.3.2	Smoothing Filter Reconstructions	107
6.4	3D Processing	116
6.4.1	3D Model Identification	117
6.4.2	Smoothing Filter Reconstructions	119
6-A	Ellipse Extraction Program	124
7	Future Work and Conclusions	137
7.1	Conclusions	137
7.2	Future Work	138

List of Figures

2-1	Sample ECG	17
2-2	Gating Process	20
2-3	A sample myocardial perfusion image	21
3-1	Example of an Ellipse and Rotated Ellipse	26
3-2	Projection of a Three-Dimensional Ellipsoid	29
4-1	Shrinking and Expanding Ellipse	48
4-2	Shrinking and Expanding Ellipse with Rotation	49
4-3	First Signal-to-Noise Ratio	51
4-4	Magnitude Error with Perfect Model	52
4-5	Area Error with Perfect Model	53
4-6	Several Signal-to-Noise Ratios	54
4-7	Example of Measurements	54
4-8	Average Magnitude Error with Perfect Model	55
4-9	Ejection Fraction Error with Perfect Model	57
4-10	Average Magnitude Error with Assumed Model	58
4-11	Ejection Fraction Error with Assumed Model	58
4-12	Average Magnitude Error with Random Walk Model	60
4-13	Ejection Fraction Error with Random Walk Model	60
4-14	Static Ejection Fraction Error.	61
4-15	Magnitude Error with Perfect Model and Rotating Ellipse	63
4-16	Ejection Fraction Error with Perfect Model and Rotating Ellipse	63
4-17	Magnitude Error with Assumed Model and Rotating Ellipse	64

4-18	Ejection Fraction Error with Assumed Model and Rotating Ellipse . .	64
4-19	Magnitude Error with Random Walk Model and Rotating Ellipse . .	65
4-20	Ejection Fraction Error with Random Walk Model and Rotating Ellipse	65
4-21	Several Signal-to-Noise Ratios for Three Dimensional Reconstruction	68
4-22	Magnitude Error with Perfect Model (3D)	69
4-23	Ejection Fraction Error with Perfect Model (3D)	70
4-24	Magnitude Error with Assumed Model (3D)	70
4-25	Ejection Fraction Error with Assumed Model (3D)	71
4-26	Magnitude Error with Random Walk Model (3D)	71
4-27	Ejection Fraction Error with Random Walk Model (3D)	72
4-28	Static Ejection Fraction Error (3D)	72
5-1	Signal to noise ratio for low noise	83
5-2	Model ID Performance	83
5-3	Signal to noise ratio for several noise	84
5-4	Model Identification Performance	85
5-5	Signal to noise ratio for several noise	86
5-6	Model Identification Performance	86
5-7	Signal to noise ratio for several noise	87
5-8	Model Identification Performance	88
5-9	Model Identification Performance	89
5-10	Summary Model Identification Performance	90
5-11	Model Identification Performance	92
5-12	Model Identification Performance	93
5-13	Model Identification Performance	96
5-14	Model Identification Experiment #1, 3D	97
5-15	Model Identification Experiment #2, 3D	98
5-16	Model Identification Experiment #3, 3D	99
6-1	Anterior areas for all patients	104
6-2	Effect of r/q on Smoothing Filter Output	108

6-3	Smoothed two dimensional reconstructions for ANT view for all patients	110
6-4	Smoothed three dimensional reconstructions for each patient	121

List of Tables

4.1	Characteristics of starting 2D ellipsoid	47
4.2	Process Noise Values	52
4.3	Measurement Noise Values	56
5.1	Measurement Noise Values	84
5.2	Values for true contraction rates and corresponding EFs	91
6.1	MUGA Ejection Fractions for Six Individuals	103
6.2	Results of 2D Model Identification	107
6.3	Values of Assumed Model EFs	109
6.4	Ejection Fractions for 2D reconstruction	111
6.5	Variation of Reconstructed Ellipse EF with Assumed EF for Patient #6112	
6.6	Combined View Approximations to 3D EF	114
6.7	Correlation Coefficients for 2D	116
6.8	Results of 3D Model Identification	118
6.9	Ejection Fractions for 3D	122
6.10	Variation of Reconstructed Ellipsoids EF with Assumed EF for Patient #6	122
6.11	Correlation Coefficients	123

Chapter 1

Introduction

1.1 Overview

This thesis deals with a method to obtain the ejection fraction of the left ventricle of the heart from a set of planar myocardial perfusion images obtained over time. Ejection fraction of the left ventricle has long been known as an effective gauge of cardiac function [14]. Myocardial perfusion imaging is a radionuclide technique that may be used to produce a sequential set of images of the heart in motion. The approach described in this thesis employs geometric reconstruction and recursive estimation techniques to track left-ventricular volume throughout the cardiac cycle, thus allowing the generation of an ejection fraction estimate.

Ejection fraction is defined by the ratio of the fully contracted left-ventricular volume to the fully expanded left-ventricular volume. Thus, it is a measure of the pumping capability of the heart and has great prognostic value to cardiologists. A reduced ejection fraction is indicative of potential cardiac malfunction. Thus, inexpensive and noninvasive techniques of determining ejection fraction are essential to the prompt diagnosis of heart disease.

Planar myocardial perfusion images are obtained by injecting the patient with a radionuclide-marked substance. The heart is imaged while the radioactive tracer is absorbed by the heart muscle or myocardium. This procedure uses a gamma camera which produces images by counting the photons emitted from the radioactive tracer.

It is possible to use the patient's electrocardiogram (ECG) to gate the counts recorded by the gamma camera into separate bins that correspond to different points in the cardiac cycle. Typically, ECG gating is used to obtain a sequential set of images at 16 equally spaced points in the cardiac cycle. These gated myocardial perfusion images are not snapshots of the heart in motion, but rather, averages of a particular cardiac phase over several cardiac cycles. Therefore, they are poor visual quality images.

We propose to use the myocardial perfusion images to estimate ejection fraction, a purpose for which they are not traditionally used. Myocardial perfusion images while of poor visual quality, contain a large amount of *physiological* information reflecting the tie between the chosen radionuclide and the biochemistry of the region under study. As a result, they are traditionally used to locate infarcts, areas in the heart muscle that are being deprived of nutrients because of an occlusion. In the myocardial perfusion images, such infarcts appear as dark regions. In addition to this traditional role, it is possible to distinguish the outline of the left-ventricular cavity from myocardial perfusion images. It is this projection-like outline of the left-ventricular cavity that we propose to use to estimate ejection fraction.

To estimate ejection fraction from the temporal set of myocardial perfusion images, we employ geometric reconstruction and recursive estimation techniques. One commonly used approximation to the true shape of the left ventricle is a three-dimensional ellipsoid. Projections of this left ventricular ellipsoid model are two-dimensional ellipsoids, which we use to model the left-ventricular outline in the observed myocardial perfusion images. Much work in geometric reconstruction [7, 11, 12] has focused on reconstructing objects such as ellipsoids from noisy lower-dimensional projections. Combining these geometric reconstruction techniques with recursive estimation procedures, such as Rauch-Tung Striebel smoothing, we have the means to formulate an estimation procedure that combines the observed temporal set of projection images with an evolution model to produce the best estimate of the ellipsoid at any point given all the data. From the reconstructed three-dimensional ellipsoids which approximate the true shape of the left ventricle, we may then approximate left ventricular volume and calculate ejection fraction.

1.2 Contributions

The work presented in this thesis contributes in the following areas. First, by defining a method to estimate left-ventricular ejection fraction from a gated set of myocardial perfusion images, the current process of diagnosing heart disease becomes more cost effective and safer. The current “gold standard” technique used to estimate ejection fraction is based on images produced through radionuclide ventriculography. These images are also produced by injecting a patient with a radionuclide-marked substance and imaging with an ECG gated gamma camera. While myocardial perfusion images are produced as the radioactive tracer is absorbed by the heart muscle itself, radionuclide ventriculographic images are produced as the tracer passes through the left ventricular cavity. Typically, a physician uses myocardial perfusion images as a preliminary screening tool. If the myocardial perfusion images show potential infarcts, further tests such as radionuclide ventriculography are performed to determine ejection fraction. By estimating ejection fraction from myocardial perfusion images, it is possible to use this preliminary step to screen for potential infarcts and estimate ejection fraction using only one procedure. Thus, the patient must undergo only one exposure to radioactive substances.

In addition, our proposed method of estimating ejection fraction improves on current techniques by integrating temporal information. In current techniques based on images produced by radionuclide ventriculography, volume is estimated only at the fully contracted and fully expanded phases. That is, only the data in these two frames are used. In our proposed approach based on gated myocardial perfusion images, we track the volume of the left ventricle through the entire cardiac cycle. Thus, we use all the available data. By incorporating this temporal set of data, we are able to more robustly estimate the fully expanded and fully contracted left ventricular volumes and thus the ejection fraction itself.

In addition, this thesis builds on previous work done in geometric reconstruction [7, 11, 12]. In Chapter 4, we reconstruct dynamically evolving ellipsoids from noisy lower-dimensional projections. While previous research concentrated on reconstruct-

ing dynamically evolving ellipsoids with known dynamics, the work in this thesis concentrates on imperfect knowledge of dynamics. We investigate several dynamic models that approximate the true dynamics of the ellipsoid and the effect of each of these dynamic models on the error in the reconstructed ellipse. In Chapter 5, we discuss a method to determine an approximation to the true dynamics of an ellipsoid using a model identification scheme. In particular, we describe a method to choose which of two hypothesized dynamic models more closely matches the true dynamics of a dynamically evolving ellipsoid. Finally, in Chapter 6, we apply these techniques to real myocardial perfusion images.

1.3 Organization

The organization of this thesis is as follows. In Chapter 2, we discuss the medical motivation for the work presented in this thesis. We discuss the differences in the types of radionuclide-based cardiac imaging and explore the motivation for a myocardial-perfusion-based estimate of ejection fraction. In Chapter 3, we discuss previous work in geometric reconstruction of ellipsoids from noisy lower-dimensional projections. Chapter 4 ties together the ideas presented in Chapters 2 and 3, by investigating the reconstruction of computer simulated, dynamically evolving ellipsoids. This chapter also discusses, in some detail, the effect of imperfect knowledge of ellipsoid dynamics on the reconstruction of dynamically evolving ellipsoids. Chapter 5 investigates a model identification scheme used to choose which of two hypothesized dynamic models best approximates the true dynamics of an evolving ellipsoid. Finally, Chapter 6 combines the model identification and smoothing filter-based reconstruction to estimate the ejection fraction of several patients from real myocardial perfusion images. Chapter 7 offers suggestions for future work and summarizes the work contained in this thesis.

Chapter 2

Medical Background

2.1 Objective

The objective of this chapter is to discuss the medical motivation for the work presented in this thesis. Ejection fraction of the left ventricle of the heart has long been known as an effective gauge of cardiac function [14]. Many techniques exist for estimating ejection fraction including angiography [4], echocardiography [3], magnetic resonant imaging (MRI) [1], and radionuclide ventriculography [4]. Of these techniques, the least expensive and most prevalent is radionuclide ventriculography. It has been shown that images produced by this technique may be used to give an accurate assessment of ejection fraction [14]. Myocardial perfusion is another relatively safe and inexpensive technique used to image the heart and diagnose cardiovascular disease. Myocardial perfusion images have not traditionally been used in the estimation of ejection fraction. We propose to use the images produced by myocardial perfusion as an alternative to radionuclide ventriculography to estimate ejection fraction.

This chapter will be organized in the following way. We begin by explaining and defining left ventricular ejection fraction. Next, we include a discussion of currently used radionuclide ventriculographic techniques for measuring ejection fraction. Finally, we provide background information on myocardial perfusion imaging and describe the benefits of estimating ejection fraction from these images as opposed to

radionuclide ventriculography.

2.2 Left Ventricular Ejection Fraction

Before explicitly defining left ventricular ejection fraction, a measure of how efficiently the heart is pumping blood to the rest of the body, let us describe the pumping action of the heart, the cardiac cycle, and the electrocardiogram. The heart consists of four chambers: the right atrium, the right ventricle, the left atrium, and the left ventricle. These four chambers act as a pump which, in conjunction with the lungs, receives deoxygenated blood and ejects oxygenated blood. Specifically, the right atrium and right ventricle are responsible for pumping deoxygenated blood to the lungs. It is the left side of the heart which then receives and pumps oxygenated blood to the rest of the body. The left atrium is simply a receiving chamber. The chamber responsible for pumping blood to the rest of the body is the left ventricle. Therefore, we may evaluate the pumping capability of the heart by evaluating the pumping capability of the left ventricle.

This pumping action is repeated in a periodic manner known as the cardiac cycle. The evolution of the left ventricle, the pumping chamber, through the cardiac cycle is as follows. During the systolic phase of the cardiac cycle, the left ventricle contracts and ejects blood. At end systole (ES) the left ventricle is fully contracted. In the diastolic phase, the left ventricle expands to receive blood. End diastole (ED) is the fully expanded phase of the left ventricle.

Phases of the cardiac cycle may be monitored using an electrocardiogram (ECG). The chambers of the heart expand and contract by a process of depolarization and repolarization of the heart muscle itself. For further details on the process of depolarization/repolarization, see [4]. At the surface of the skin, it is possible to measure voltage changes caused by the depolarization/repolarization process. A plot of this voltage versus time is known as an electrocardiogram. Figure 2-1 shows a sample ECG. The sections of the ECG are known to correspond to portions of the cardiac cycle. For example, the QRS complex on the ECG precedes each cardiac contraction




Figure 2-1: Sample of ECG from [4]. Note that sections of the wave are marked.

(the systolic phase). Thus, we can discern the location of the heart in its cycle by observing the ECG.

Having established the necessary background, we may now define ejection fraction of the left ventricle. Ejection fraction (EF) is the percent of total left ventricular volume ejected by each contraction. The formal definition of ejection fraction is given by :

$$EF = \frac{ED \text{ volume} - ES \text{ volume}}{ED \text{ volume}} \quad (2.1)$$

where ED is the fully expanded cardiac phase and ES is the fully contracted cardiac phase. A normal value for ejection fraction is 0.50 or above. Reduced ejection fraction is indicative of reduced cardiac performance.

2.3 Radionuclide Ventriculography

Radionuclide ventriculography is the most prevalent imaging technique used to estimate ejection fraction. This technique is used to produce a sequential set of images as a bolus or blood pool of radionuclide tagged blood passes through the cardiac chambers. The images corresponding to ES and ED are used to provide estimates of left ventricular volume at those phases. From these volume estimates, one may

formulate an estimate of ejection fraction.

There are two types of radionuclide ventriculography. They are the first pass technique and the multiple gated blood pool technique (MUGA). Both of these techniques have been shown to accurately assess ejection fraction [4]. Both radionuclide ventriculographic techniques give a so-called “gold standard” estimate of ejection fraction. We will give a brief description of the images produced using each technique.

2.3.1 Procedure

The radionuclide ventriculographic images are produced by, first, injecting the patient with a radiopharmaceutical such as Technetium-99m (Tc^{99m}). While a bolus or blood pool of the radiopharmaceutical passes through the chambers of the heart, a gamma or scintillation camera produces images by counting the photons emitted by the radioactive substance. Gamma cameras work by converting radiation (gamma rays) to light (scintillation) using a sodium iodide crystal.

Note that the images produced by radionuclide ventriculography are not images of the heart. Rather, the images show the pool of blood contained in the chambers of the heart.

2.3.2 First Pass Technique

First pass images are produced as the radionuclide tagged blood traverses through the cardiac chambers for the first time [4]. The gamma ray camera is used to acquire a set of images of the heart from one view. The images are acquired rapidly, usually one or more per second. The entire data acquisition lasts about 30 seconds. The gamma camera is triggered by the patients ECG so that the images correspond to specific points in the cardiac cycle.

Ejection fraction as defined in (2.1) is based on the left ventricular volume at ES and ED phases. Using the frames corresponding to ES and ED from the sequential set of first pass images, it is possible to calculate left ventricular volume at ES and ED. The left ventricular volume is assumed to be proportional to the intensity of the

area corresponding to the left ventricle on the image.

The benefit of the first pass technique is that there is low background radiation because the radiopharmaceutical has not been taken up by other organs. The delineation of the edges of the heart in a first pass image is easily achieved.

2.3.3 Multiple Gated Blood Pool Technique

Multiple gated blood pool (MUGA) images are produced by similarly viewing a bolus as it is pumped through the cardiac chambers. The MUGA images are produced not on the first pass of the bolus through the heart, but rather during several subsequent cycles. Thus, the radiopharmaceutical in the cardiac chambers during MUGA imaging is less concentrated than during first pass imaging. In fact, the concentration of the radiopharmaceutical is too low to produce enough gamma-scintillations to form a meaningful image. For this reason, a process known as gating is used. The patient's ECG is used to gate the counts recorded by the gamma camera over time into separate bins that correspond to different points in the cardiac cycle (see Figure 2-2). Thus, the MUGA images produced are the sum of images at a particular cardiac phase over several cardiac cycles.

There are some additional differences between the MUGA technique and the first pass technique. While images produced using the first pass technique are typically only produced from one view, MUGA images are produced from multiple views. Also, because the radiopharmaceutical has dissipated through the body in MUGA imaging, there is more background radiation and the edges of the heart are not so clearly delineated as in first pass images.

Again, left ventricular volume is assumed to be proportional to the intensity of the area corresponding to the left ventricle on the MUGA image. Using the estimate of left ventricular volume as based on the MUGA images corresponding to ES and ED, it is possible to calculate ejection fraction as in (2.1). Samples of a MUGA images may be found in [4, 10].




Figure 2-2: Reprinted from [6]. The gating process for producing MUGA images.

2.4 Proposed Approach

While techniques based on radionuclide ventriculographic images are known to give an accurate estimate of ejection fraction, an alternative approach based on gated myocardial perfusion images may prove to be safer, more cost effective, and more robust. Thus, we propose to use a temporal set of myocardial perfusion images to estimate ejection fraction. In this section, we present background information on myocardial perfusion imaging and the benefits of using myocardial perfusion images to estimate ejection fraction.

2.4.1 Myocardial Perfusion Imaging

Myocardial perfusion is another relatively safe and inexpensive technique used to image the heart and diagnose cardiovascular disease. As we present background information on myocardial perfusion imaging, keep in mind that we propose to use the images produced by myocardial perfusion as an alternative to radionuclide ventriculography to estimate ejection fraction.

Planar myocardial perfusion images are also obtained by injecting the patient with a radionuclide marked substance and imaging with an ECG gated gamma camera.




Figure 2-3: A sample myocardial perfusion image

While radionuclide ventriculography images are produced as the tracer passes through the left-ventricular cavity, myocardial perfusion images are produced when the tracer has been absorbed by the heart muscle or myocardium. The myocardial perfusion images produced are not snapshots of the heart in motion, but rather, the sum of a particular cardiac phase over several cardiac cycles. Therefore, they are poor visual quality images. A sample of such a myocardial perfusion image is shown in Figure 2-3.

Although myocardial perfusion images are of poor visual quality, they contain a large amount of *physiological* information reflecting the tie between the chosen radionuclide and the biochemistry of the region under study. As a result, they are traditionally used to locate infarcts, areas in the heart muscle that are being deprived of nutrients because of an occlusion. In the myocardial perfusion images, such infarcts appear as dark regions.

It is also possible to distinguish the projection-like outline of the left-ventricular cavity in myocardial perfusion images. It is this outline that we propose to use to estimate left-ventricular volume. Thus, while the images produced by myocardial perfusion are not traditionally used to estimate left-ventricular volume, we note that it may be possible to use them in this capacity.

One additional note of interest is that recent developments have made myocardial perfusion imaging even safer and less expensive. Myocardial perfusion is com-

monly used; however, there have been certain drawbacks associated with this procedure. In the past, the radiopharmaceutical of choice for myocardial perfusion imaging was thallium-201 (Tl^{201}), which has a long half life and is difficult to produce. Recently, doctors have found a way to produce myocardial perfusion images using Tc^{99m} -labelled agents [13, 5]. Because Tc^{99m} has a shorter half life and is more easily produced, myocardial perfusion imaging is now even more attractive.

2.4.2 Benefits

By defining a method to estimate left-ventricular ejection fraction from a gated set of myocardial perfusion images, the current process of diagnosing heart disease becomes safer and more cost effective. Typically, a physician uses myocardial perfusion images as a preliminary screening tool. If the myocardial perfusion images show potential occlusions, further tests such as radionuclide ventriculography are performed to determine ejection fraction. By estimating ejection fraction from myocardial perfusion images, it is possible to screen for potential occlusions and estimate ejection fraction using only one procedure. Thus, we eliminate a second procedure and limit the patient's exposure to radioactive substances.

In addition, our proposed method provides a more robust estimate of ejection fraction by integrating temporal information. In current techniques based on images produced by radionuclide ventriculography, volume is estimated only at the fully contracted and fully expanded phases. That is, only the data in these two frames are used. In our proposed approach based on gated myocardial perfusion images, we track the volume of the left ventricle through the entire cardiac cycle. Thus, we use all the available data. By incorporating this temporal set of data, we are able to more robustly estimate the fully expanded and fully contracted left ventricular volumes and thus the ejection fraction itself.

Chapter 3

Geometric Reconstruction

Background

3.1 Objective

Before moving on with the analysis of the myocardial perfusion data presented in Chapter 2, we need some additional background. As we have mentioned previously, it is our intention to model the left ventricular cavity as a dynamically evolving three-dimensional ellipsoid. Thus, our myocardial perfusion images may be viewed as noisy two-dimensional projections of this three-dimensional ellipsoid. For this reason, in this chapter we discuss the reconstruction of ellipsoids from lower dimensional projections. Much work has been done on reconstructing objects from lower dimensional projections [12, 11, 7]. In particular [7], discusses the reconstruction of an n -dimensional ellipsoid from noisy lower dimensional projections. This chapter summarizes the problem formulation and reconstruction of an n -dimensional dynamically evolving ellipsoid.

The organization of this chapter is as follows. We begin by discussing the problem formulation. We discuss mathematical representations for the ellipsoid itself, dynamics of the ellipsoid, and projections of the ellipsoid. Next, we discuss methods for reconstructing both a single ellipsoid from one measurement and a dynamically evolving ellipsoid from multiple measurements. For simplicity, the examples used

to illustrate these points are two or three-dimensional ellipsoids, but generally the information presented here is applicable to higher dimensional ellipsoids as well.

3.2 Problem Formulation

In this section, we discuss two representations for an ellipsoid, ellipsoid dynamics, and ellipsoid projections. The first ellipsoid representation is in terms of a symmetric, positive definite matrix. The dynamics and projections of an ellipsoid may be written as linear functions of the matrix that represents the ellipsoid. Alternatively, we may express an ellipsoid in terms of a vector. Again, we may express ellipsoid dynamics and projections as linear functions of the vector that represents the ellipsoid.

Although both representations for an ellipsoid, its dynamics, and its projections are equivalent, we will find the vector representation of the ellipsoid yields a simplified problem formulation. The reconstruction of a dynamically evolving ellipsoid may be reduced to tracking a dynamically evolving state vector from noisy measurements by using the vector representation of the ellipsoid. We may then call on a rich class of techniques from stochastic estimation such as Kalman filtering and Rauch-Tung Striebel smoothing [9, 15]. Therefore, the vector representation for the ellipsoid is the one we will use in Section 3.3.

3.2.1 Matrix Representation

Ellipsoid Representation

It is possible to represent the points included in an n -dimensional, origin centered, non-degenerate ellipsoid, ε , in the following way :

$$\{x|x^T E^{-1} x \leq 1, x \in R^n\} \tag{3.1}$$

where the positive definite matrix E that represents an ellipsoid is easily determined. Consider a two-dimensional ellipsoid (an ellipse) as an example. A two-dimensional ellipsoid centered at the origin is uniquely specified by two semi-axis lengths, a and

b , and an angle of rotation ϕ . This ellipsoid is represented by the matrix

$$E_2 = R^T E_{2u} R \quad (3.2)$$

where

$$R = \begin{pmatrix} \cos \phi & \sin \phi \\ -\sin \phi & \cos \phi \end{pmatrix} \quad (3.3)$$

$$E_{2u} = \begin{pmatrix} a^2 & 0 \\ 0 & b^2 \end{pmatrix} \quad (3.4)$$

Note that if the angle of rotation $\phi = 0$, $E = E_{2u}$. Hence, E_{2u} is the matrix that represents the ellipse that has its axes lined up with the coordinate system (i.e. an unrotated ellipse).

For example, an unrotated ellipse with major and minor semi-axis lengths of 8 and 6 respectively is shown in Figure 3-1. This figure also shows the same ellipse rotated by an angle of $\phi = \pi/4$. The corresponding matrix representations are given by

$$E_{2u} = \begin{bmatrix} 64 & 0 \\ 0 & 36 \end{bmatrix} \quad (3.5)$$

$$E_2 = \begin{bmatrix} 50 & 14 \\ 14 & 50 \end{bmatrix} \quad (3.6)$$

Similarly, a three-dimensional ellipsoid centered at the origin may be specified by the three semi-axis lengths and three angles of rotation. The semi-axis lengths, a , b , and c correspond to the semi-axis lengths of the unrotated ellipsoid along the x , y , and z -axes, respectively. The angles of rotation, ϕ_x , ϕ_y , ϕ_z correspond to rotation about the x , y , and z -axes, respectively. The corresponding matrix representation is




Figure 3-1: Example of an Ellipse and Rotated Ellipse

given by :

$$\begin{aligned}
 E_3 &= R^T E_{3_u} R \\
 E_{3_u} &= \begin{pmatrix} a^2 & 0 & 0 \\ 0 & b^2 & 0 \\ 0 & 0 & c^2 \end{pmatrix} \\
 R_3 &= \begin{pmatrix} \cos \phi_z & \sin \phi_z & 0 \\ -\sin \phi_z & \cos \phi_z & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \cos \phi_y & 0 & \sin \phi_y \\ 0 & 1 & 0 \\ -\sin \phi_y & 0 & \cos \phi_y \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \phi_x & \sin \phi_x \\ 0 & -\sin \phi_y & \cos \phi_y \end{pmatrix}
 \end{aligned}$$

For example, a three-dimensional ellipsoid with axis lengths 8, 6, and 8 with a rotation

about the x -axis of $\phi_x = -\pi/4$ is represented by the matrix

$$E_3 = \begin{bmatrix} 64 & 0 & 0 \\ 0 & 50 & 14 \\ 0 & 14 & 50 \end{bmatrix}$$

Note that we have made the simplifying assumption that the ellipsoid center is known. The above representations assume the ellipsoids are centered at the origin. The same representations may be extended to represent ellipsoids with centers other than the origin by translating the the coordinate system origin to the ellipsoid center.

Dynamics

Recall that we are interested in tracking dynamically evolving ellipsoids. For this reason, we will briefly discuss several types of ellipsoid evolution that we may encounter. It is possible to capture a broad range of ellipsoid dynamics in the following form :

$$E(k+1) = A(k)^T E(k) A(k) \quad (3.7)$$

where we can represent changes such as magnification, rotation, and eccentricity change in $A(k)$. For the two-dimensional case, the matrix that captures these dynamics is given by (see [12]) :

$$A(k) = \begin{pmatrix} t_k & 0 \\ 0 & t_k \end{pmatrix} \begin{pmatrix} \cos \theta_k & \sin \theta_k \\ -\sin \theta_k & \cos \theta_k \end{pmatrix} \begin{pmatrix} \alpha_k & 0 \\ 0 & 1/\alpha_k \end{pmatrix} \quad (3.8)$$

In (3.8), the parameters t_k , θ_k , and α_k represent magnification, rotation, and eccentricity change respectively. This class of dynamics does not include all possible ellipsoid dynamics; however, for the estimation of ejection fraction from myocardial perfusion images this set of dynamics will be sufficient.

Note, also, that the extension to higher dimensions is straightforward. Magnification is still a scalar times the identity matrix. Rotation is expressed as the product of several matrices where each matrix represents rotation about one axis of the coordi-

nate system. Eccentricity is also expressed as the product of several matrices, where each matrix represents the eccentricity change between two axes at a time.

Projections

The projection of an ellipsoid is an ellipsoid of lower dimension. In equation form, a single projection, Y_i , of an ellipsoid, E , is represented by

$$Y_i = C_i^T E C_i \quad (3.9)$$

where the matrix C_i captures the geometry of the projection. The rows of the matrix C_i span the space of the projection.

Projections may be thought of as the shadow cast by the ellipsoid. Figure 3-2 provides an illustration of the projection of a three-dimensional ellipsoid. The matrix that represents the projection of a three-dimensional ellipsoid onto the xy -plane is given by

$$C_i = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix} \quad (3.10)$$

The one dimensional projection of a two-dimensional ellipsoid is a line. The matrix that represents the projection of a two-dimensional ellipsoid onto the x -axis is given by

$$C_i = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad (3.11)$$

Problem Statement

We may summarize the dynamics and projections of an ellipsoid represented as a matrix in the following manner

$$E(k+1) = A(k)^T E(k) A(k) + W(k) \quad (3.12)$$




Figure 3-2: Projection of a Three-Dimensional Ellipsoid

$$Y(k) = C^T E(k)C + V(k) \quad (3.13)$$

where $W(k)$ and $V(k)$ represent process and measurement noise, respectively. Noise terms are included to represent the uncertainty in the dynamics and projections. The matrices $E(k+1)$ and $Y(k)$ represent ellipsoids and, therefore, must be positive definite. To simplify the ellipsoid reconstruction problem, we assume that the independent elements of $W(k)$ and $V(k)$ have a Gaussian distribution. Clearly, this is not the case. If the independent elements of $W(k)$ and $V(k)$ had Gaussian distribution, $E(k+1)$ and $Y(k)$ would not always be positive definite matrices. However, we choose to employ this commonly-used assumption because of the simplification it provides in the ellipsoid reconstruction problem. Our objective is then to determine $E(k)$ for all k given noisy projections (or measurements) at all k .

This problem formulation proves to be difficult to manipulate. A more convenient problem statement is obtained using a vector representation of the ellipsoid. This alternative problem formulation is examined in Section 3.2.2.

3.2.2 Vector Representation

Ellipsoid Representation : The Mapping Ξ

For simplicity, we define an equivalent vector representation for an ellipsoid. In [7], several such vector representations are discussed. One vector representation with certain very useful properties is the $\frac{n(n+1)}{2}$ vector which consists of the independent elements of the matrix E as defined in Section 3.2.1. As we will see in the sections to follow, this vector representation of an ellipsoid greatly simplifies the problem of reconstructing ellipsoids from lower dimensional projections.

The mapping Ξ defines a transformation from a symmetric $n \times n$ matrix, E , to an $\frac{n(n+1)}{2}$ dimensional vector, ϵ . The vector, ϵ , consists of multiples of the independent elements of the symmetric matrix, E . Thus, a one-to-one relationship exists between E and ϵ . We need the following definitions to specify the transformation Ξ . We define a standard orthonormal basis for the space of symmetric $n \times n$ matrices as $\{M_\ell^{(n)} | 1 \leq \ell \leq n(n+1)/2\}$ where

$$M_\ell^{(n)} = \begin{cases} e_i e_i^T & \text{if } \ell = \frac{(i-1)(2n+2-i)}{2} + 1 \quad i = 1 \dots n \\ \frac{e_i e_j^T + e_j e_i^T}{\sqrt{2}} & \text{if } \ell = \frac{i(2n+1-i)}{2} - n + j \quad 1 \leq i < j \leq n \end{cases} \quad (3.14)$$

where e_i is the i -th unit vector consisting of all zeros except a one in the i -th position. In addition, we define an inner product as

$$\langle A, B \rangle_n \equiv \text{trace}(A^T B) \quad (3.15)$$

The j -th element of the vector, ϵ , is the inner product of the matrix, E , and the j -th basis matrix. That is,

$$(\epsilon)_j = \langle E, M_j^{(n)} \rangle_n \quad (3.16)$$

Equation (3.16) defines a mapping Ξ from E to ϵ . We denote this relationship as

$$\epsilon = \Xi(E) \quad (3.17)$$

As an example, consider the two-dimensional case. An orthonormal basis for the set of symmetric 2×2 matrices consists of the following elements :

$$\begin{aligned} M_1^{(2)} &= \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \\ M_2^{(2)} &= \begin{bmatrix} 0 & 1/\sqrt{2} \\ 1/\sqrt{2} & 0 \end{bmatrix} \\ M_3^{(2)} &= \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \end{aligned}$$

Using this standard basis and the mapping Ξ , the transformation from a 2×2 matrix to a vector in R^3 takes the following form :

$$\begin{bmatrix} X_{11} & X_{12} \\ X_{12} & X_{22} \end{bmatrix} \mapsto \begin{bmatrix} X_{11} \\ \sqrt{2}X_{12} \\ X_{22} \end{bmatrix}$$

Continuing our example from above, the vector representation of the matrix in (3.6) is

$$\epsilon_2 = \begin{bmatrix} 50 \\ 19.799 \\ 50 \end{bmatrix}$$

Dynamics and Projections : The Mapping Γ

Consider an equation of the form :

$$Y = S^T X S \tag{3.18}$$

where X , Y , and S are matrices with dimensions $n \times n$, $m \times m$, and $m \times n$ respectively. If the matrices X and (thus) Y are symmetric matrices, it is possible to write the

same relationship in a vector representation of the form :

$$y = \tilde{S}x \quad (3.19)$$

where x and y are $\frac{n(n+1)}{2}$ and $\frac{m(m+1)}{2}$ dimensional vectors respectively and \tilde{S} is a $\frac{m(m+1)}{2} \times \frac{n(n+1)}{2}$ matrix. If the mapping Ξ defined in (3.16) describes the transformation from X to x and from Y to y , then a mapping Γ defining the transformation from S to \tilde{S} is induced. The following equation shows this induced relationship between the elements of S and \tilde{S} .

$$(\tilde{S})_{ij} = \langle M_i^{(m)}, S^T M_j^{(n)} S \rangle_m \quad (3.20)$$

This mapping is denoted

$$\tilde{S} = \Gamma(S) \quad (3.21)$$

We may apply the mapping Γ to matrix A that describes the ellipsoid dynamics in (3.7). Using the mappings Γ and Ξ , we rewrite (3.7) as :

$$\epsilon(k+1) = \tilde{A}(k)\epsilon(k) \quad (3.22)$$

where $\epsilon(k) = \Xi(E(k))$ and $\tilde{A}(k) = \Gamma(A(k))$. As an example, consider the two-dimensional case. Using the mapping Γ , we may express the transformation from A to \tilde{A} as

$$A \quad \mapsto \quad \tilde{A}$$

$$\begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} \mapsto \begin{pmatrix} A_{11}^2 & \sqrt{2}A_{11}A_{21} & A_{21}^2 \\ \sqrt{2}A_{11}A_{12} & A_{12}A_{21} + A_{11}A_{22} & \sqrt{2}A_{21}A_{22} \\ A_{12}^2 & \sqrt{2}A_{12}A_{22} & A_{22}^2 \end{pmatrix}$$

We may also use the mapping Γ to rewrite (3.9) in the following form

$$y_i(k) = \tilde{C}_i \epsilon(k) \quad (3.23)$$

where $\epsilon(k) = \Xi(E)$, $y_i(k) = \Xi(Y_i)$, and $\tilde{C}_i = \Gamma(C_i)$. As an example, consider the matrix C_i that defines the projection of a two-dimensional ellipsoid onto the x -axis in (3.11). The corresponding matrix \tilde{C}_i is given by :

$$\tilde{C}_i = \begin{pmatrix} 1 & 0 & 0 \end{pmatrix} \quad (3.24)$$

Note that (3.23) represents a single projection. We may have multiple projections at a given instant in time. Stacking up all such projections we obtain :

$$y(k) = \tilde{C}\epsilon(k) \quad (3.25)$$

where $y = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \end{bmatrix}$ and $\tilde{C} = \begin{bmatrix} \tilde{C}_1 \\ \tilde{C}_2 \\ \vdots \end{bmatrix}$.

Problem Statement

We may summarize the dynamics and projections of an ellipsoid represented as a vector in the following manner

$$\epsilon(k+1) = \tilde{A}(k)\epsilon(k) + w(k) \quad (3.26)$$

$$y(k) = \tilde{C}\epsilon(k) + v(k) \quad (3.27)$$

where $w(k)$ and $v(k)$ represent process and measurement noise, respectively. As in Section 3.2.1, we assume that the noises $w(k)$ and $v(k)$ are zero mean, white, Gaussian noise processes with variances Q and R respectively.

We note that the representation of ellipsoid dynamics and projections in (3.26)-(3.27) is equivalent to the representation given by (3.12)-(3.13). The representation given by (3.26)-(3.27) is a more convenient problem formulation. Now, the problem of reconstructing the three-dimensional ellipsoid represented by $\epsilon(k)$ is stated in a familiar framework. The problem is reduced to tracking a dynamically evolving state vector from noisy measurements. We may now call on a rich class of techniques from

stochastic estimation.

3.3 Ellipsoid Reconstruction

3.3.1 Static Reconstruction

It is possible to reconstruct a dynamically evolving ellipsoid at each time step given a sufficient number of projections at that time step. This estimate of the ellipsoid does not incorporate past or future measurements in the estimate of the ellipsoid at any time instant. We term this a static reconstruction of the ellipsoid because it does not incorporate dynamics in the estimate. Our static estimate is taken to be the standard least squares estimate which is equivalent to a maximum likelihood estimate under the assumption of additive Gaussian noise.

Given a system of the form of (3.27) where \tilde{C} is of full column rank, the linear static estimate which minimizes the square error is given by :

$$\hat{\epsilon}(k) = (\tilde{C}^T \tilde{C})^{-1} \tilde{C}^T y(k) \quad (3.28)$$

3.3.2 Dynamic Reconstruction

Alternatively, we may form an ellipsoid estimate that incorporates knowledge of the dynamics of the system. Our system is described by (3.26)-(3.27). We may formulate an estimate at time k that is optimal given all measurements up to and including measurements at time k . Such an estimate is known as the Kalman filter estimate [9, 15]. Another approach we may use is to formulate an estimate at time k that is optimal given all measurements in the interval $[0, T]$. This type of an estimate may be formed using a non-causal smoothing algorithm such as Rauch-Tung Striebel [9, 15]. We concentrate on the Rauch-Tung Striebel smoothing filter because it eliminates the lag intrinsic to the Kalman filter.

The Rauch-Tung Striebel smoothing filter consists of a Kalman filter and a subsequent update step. The Kalman filter processes the data running forward in time.

The estimate obtained from the forward running filter at time k is the best estimate given all measurements up to and including time k . This estimate is denoted $\hat{\epsilon}(k|k)$. The update step processes the data backwards in time. This is the smoothed estimate. We denote the smoothed estimate as $\hat{\epsilon}(k|T)$ where T is the time interval over which we have data.

The forward running filter must be initialized with an initial estimate, ϵ_0 , and an initial error covariance P_0 . The estimate from the forward running Kalman filter is given by :

$$\hat{\epsilon}(k|k) = \hat{\epsilon}(k|k-1) + K(k)[y(k) - \tilde{C}\hat{\epsilon}(k|k-1)] \quad (3.29)$$

$$K(k) = P(k|k-1)\tilde{C}^T[\tilde{C}P(k|k-1)\tilde{C}^T - R]^{-1} \quad (3.30)$$

$$P(k|k) = P(k|k-1) - K(k)\tilde{C}P(k|k-1) \quad (3.31)$$

$$\hat{\epsilon}(k+1|k) = \tilde{A}\hat{\epsilon}(k|k) \quad (3.32)$$

$$P(k+1|k) = \tilde{A}P(k|k)\tilde{A}^T + Q \quad (3.33)$$

The update step is initialized with the final error covariance $P(T|T)$ of the forward running filter. The smoothed estimate is then given by :

$$\hat{\epsilon}(k|T) = \tilde{A}^{-1}(k)\hat{\epsilon}(k+1|T) - \tilde{A}^{-1}(k)Q(k)P^{-1}(k+1|k)[\hat{\epsilon}(k+1|T) - \hat{\epsilon}(k+1|k)] \quad (3.34)$$

The combination of the forward Kalman filter and backward update step is known as the Rauch-Tung Striebel smoothing filter.

Chapter 4

Simulated Ellipsoid Reconstruction

4.1 Objective

Thus far, we have discussed two seemingly unrelated topics. In Chapter 2, we outlined medical background related to the ejection fraction of the left ventricle of the heart. In Chapter 3, we explored geometric reconstruction of ellipsoids. We have developed methods to mathematically represent ellipsoids, their projections, and their dynamics. In addition, we have introduced reconstruction techniques for a single ellipsoid and for a dynamically evolving set of ellipsoids.

Our ultimate objective is to employ the techniques developed in Chapter 3 to analyze the myocardial perfusion data presented in Chapter 2 and form an estimate of ejection fraction. To do this, we will view the left ventricular cavity as a dynamically evolving three-dimensional ellipsoid. Thus, our myocardial perfusion images will be viewed as two-dimensional projections of this dynamically evolving ellipsoid. The estimate will be based on a Rauch-Tung Striebel smoothing algorithm that combines all the data with a model to produce the best estimate at any point given all the data.

Before tackling the problem of estimating ejection fraction from real myocardial perfusion images, we test our approach by using Rauch-Tung Striebel smoothing

to reconstruct computer simulated ellipsoids. First, we use the smoothing filter to reconstruct a computer simulated, dynamically evolving two-dimensional ellipsoid from noisy measurements of the ellipse itself. This is meant to be a preliminary exercise that gives us insight into the reconstruction of a three-dimensional ellipsoid. Second, we use smoothing to reconstruct a computer simulated, dynamically evolving three-dimensional ellipsoid from noisy two-dimensional projections.

In this chapter, we describe our simulation setup and results. We begin by discussing the simulated input to the smoothing filters, namely two-dimensional ellipsoids that may be thought of as noisy projections of a computer simulated, dynamically evolving three-dimensional ellipsoid whose shape and dynamics mimic those of the ellipsoid that approximates the left ventricle of the human heart. Next, we present practical issues to be considered in implementing our smoothing algorithm. These are the limitations we will encounter in reconstructing the ellipsoid that approximates the left ventricle from myocardial perfusion data. For example, we will not know the true dynamics or the projection geometry of the ellipsoid that approximates the left ventricle. These same limitations should be reflected in our simulated ellipsoid reconstructions. In this way, our simulated results will give an indication of how well we will do in estimating ejection fraction from the myocardial perfusion images. We then present results from the Rauch-Tung Striebel smoothing filter reconstruction of a computer simulated, dynamically evolving two-dimensional ellipsoid. Finally, we present results from the Rauch-Tung Striebel smoothing filter reconstruction of a simulated, dynamically evolving three-dimensional ellipsoid.

4.2 Simulation Setup

4.2.1 Generation of Data

The true shape of the heart is irregular and not easily characterized. In addition, the shape of the heart varies from person to person and throughout the cardiac cycle. We would like to formulate an approximation to the true shape of the human heart

(specifically the left ventricle) for use in generating simulated data. One such commonly used simple approximation to the shape of the left ventricular cavity is a time varying, three-dimensional ellipsoid. This type of ellipsoid heart model is commonly used in medical applications [4, 8]. The justification for this type of ellipsoid heart model is as follows. If the left ventricle really were a three-dimensional ellipsoid, then its projections on a plane would be two-dimensional ellipsoids (ellipses). Our myocardial perfusion data show (see Chapter 2) that the projection of the left ventricle is roughly an ellipse. Therefore, the ellipsoidal model seems adequate. Thus, our approach will be to reconstruct dynamically evolving three-dimensional ellipsoids based on two-dimensional projections. From these reconstructed three-dimensional ellipsoids, we will determine an approximation to the ejection fraction of our simulated data.

The true dynamics of the left ventricular ellipsoid model are again irregular and not easily characterized. From the end diastolic phase to the end systolic phase, the ellipsoid that approximates the left ventricle contracts and rotates. Similarly, from the end systolic to the end diastolic phase the left ventricular ellipsoid model expands and rotates. There may be additional dynamics such as eccentricity change. The rates of these dynamic changes vary from person to person and from step to step. One rough approximation to the dynamics of the left ventricular ellipsoid model is simple contraction/expansion and rotation where the rates of change are constant over time. These dynamics capture the coarse motion of the heart. In this chapter, we generate and reconstruct an ellipsoid with these coarse dynamics (fixed rate contraction/expansion and rotation). This exercise will give us insight for Chapter 6 where we apply these reconstruction techniques to real data.

To summarize, our simulated data will be noisy projections of a dynamically evolving three-dimensional ellipsoid. The eccentricity of the simulated ellipsoid is chosen to approximate the eccentricity of the ellipsoid that is a best fit to the left ventricle. In addition, the ejection fraction of our simulated heart will be 50. This value for ejection fraction is used because it is the expected ejection fraction of an average human heart. The simulated ellipsoid will have a period of 16 steps. In

the first eight steps, the ellipsoid will contract at a constant rate. Similarly, the ellipsoid will expand at a constant rate for the last eight steps. In addition, we include simulation results for an two-dimensional ellipsoid whose dynamics include rotation.

In terms of the notation developed in Chapter 3, the generation of simulated data is summarized as

$$e(k+1) = \tilde{A}(k)e(k) \quad (4.1)$$

where $e(k)$ is the vector representation for the ellipsoid, $\tilde{A}(k)$ captures the dynamics of the ellipsoid. Recall, $\tilde{A}(k) = \Gamma A(k)$ (see 3.2.2) where the matrices $A(k)$ which represent the shrinking/expanding dynamics we described above are given by :

$$A(k) = \begin{cases} t(k)I & \text{for } k = 1 \dots 8 \\ \frac{1}{t(k)}I & \text{for } k = 9 \dots 16 \end{cases}$$

As we mentioned above, we will also include reconstruction of a two-dimensional ellipsoid whose dynamics also include shrinking/expanding and rotation, $\phi(k)$. These dynamics are represented by

$$A(k) = \begin{cases} t(k)I \begin{pmatrix} \cos \phi(k) & \sin \phi(k) \\ -\sin \phi(k) & \cos \phi(k) \end{pmatrix} & \text{for } k = 1 \dots 8 \\ \frac{1}{t(k)}I \begin{pmatrix} \cos \phi(k) & -\sin \phi(k) \\ \sin \phi(k) & \cos \phi(k) \end{pmatrix} & \text{for } k = 9 \dots 16 \end{cases}$$

where the fixed rate contraction/expansion and rotation described above implies $t(k)$ and $\phi(k)$ are constants.

4.2.2 Reconstruction : Practical Considerations

We have outlined a method to reconstruct a dynamically evolving ellipsoid from lower dimensional projections in Chapter 3. We will apply the Rauch-Tung Striebel smoothing filter to reconstruct our simulated ellipsoids and then use these reconstructed ellipsoids to formulate an estimate of ejection fraction.

Let us address a few of the difficulties that we expect to encounter in using the smoothing filter to reconstruct the ellipsoid that approximates the left ventricle from the myocardial perfusion data. This foresight will enable us to include these limitations in our reconstruction of simulated data. Thus, our results from simulations will be easily extended to the real data as well. Valuable insight about the strengths and weaknesses of the smoothing filter approach of reconstructing ellipsoids is gained through this discussion.

Data Extraction

The smoothing filter we have sketched out in (3.29)-(3.34) expects as its input ellipses that are noisy projections of a dynamically evolving three-dimensional ellipsoid. To use the smoothing filter approach to reconstruct the ellipsoid modelling the left ventricle from the myocardial perfusion images, we would have to include some type of a preprocessing step that extracts ellipses from the raw images. Automation of the process of extracting ellipses from the raw data is in and of itself a difficult problem. We will use an interactive method to extract ellipses from the myocardial perfusion images. Thus, for our simulations, we may assume that the observations of the smoothing filter are noisy ellipses.

Dynamic Models

The smoothing filter assumes knowledge of the true dynamics of the system. When reconstructing the ellipsoid that approximates the left ventricle, we will not know the exact dynamics of evolution. The true dynamics of a left ventricular ellipsoid model would vary from person to person. In Chapter 3, the types of ellipsoid dynamics we discussed included magnification, rotation, and eccentricity change. The dynamics of the left ventricular ellipsoid model for a particular person may not fit into these categories. For these reasons, it becomes necessary to study different dynamic models and the error introduced by model mismatch. We will use simulations to address these issues.

In the discussion to follow, we would like to keep in mind that our goal is to deter-

mine the simplest model that still gives an improved estimate of ejection fraction. We will judge improvement in the following way. Recall that in Chapter 2, we outlined a method that physicians attempted to use to estimate ejection fraction using myocardial perfusion images. This method used only the end systolic and end diastolic frames for volume estimates. The results obtained did not correlate well with the “gold standard” ejection fraction estimates. The analogous estimate for simulated data is determined by using static estimates for volume at the fully expanded and fully contracted phases. By static estimates, we mean reconstruction of a single ellipsoid based on the set of projections at one time step. We will compare our smoothed estimates, with each type of dynamic model, to these static estimates to determine which models give us improvement on current methods.

In reconstructing our simulated ellipsoid, we will investigate model mismatch (i.e. a mismatch between the dynamics used to generate the data and the dynamics used in the smoothing filter) by using prior knowledge to formulate models that capture the coarse motion of the heart. This approach is equivalent to replacing the true ellipsoid dynamics, $\tilde{A}(k)$ in (3.29)-(3.34) with an approximation of the true dynamics, $\tilde{M}(k)$.

That is, recall the system which is used to generate the data is given by :

$$\epsilon(k+1) = \tilde{A}(k)\epsilon(k)$$

The smoothing filter assumes that the system is described by

$$\epsilon_s(k+1) = \tilde{M}(k)\epsilon_s(k) \tag{4.2}$$

where $\tilde{M}(k) = \tilde{A}(k) + \Delta(k)$. Thus, $\Delta(k)$ is the error in modelling the true dynamics. Equation (4.2) may be rewritten as

$$\epsilon_s(k+1) = \tilde{A}(k)\epsilon_s(k) + w(k) \tag{4.3}$$

the term $w(k) = \Delta(k)\epsilon_s(k)$ may be thought of as process noise. We assume that the process noise is Gaussian white noise with variance, $Q = qI$. Clearly, the process noise

is neither Gaussian nor white. This is just a simplified model we use for convenience. We use the standard filter equations to implement the Rauch-Tung Striebel algorithm.

Now the question arises as to the optimal value for q . One could imagine experimentally finding the variance of $w(k)$ and using this value for the variance of the process noise, q , in the smoothing filter equations. Another approach might be to evaluate filter performance as we vary q . We hope that this type of an investigation will yield a wide range of values of q where the error in the estimate is small.

The smoothed estimate is a weighted average of the measured and predicted states. In the smoothing filter, the variance of the process noise may be interpreted as a measure of our trust in the dynamic model and thus the predicted state. If q is very high, the estimate will be based on the measurement alone. If q is low, the estimate will be based on the predicted state alone. We will investigate how sensitive the error in the filtered estimate is to changes in the variance of the process noise.

For each of our sets of experiments, we consider three classes of dynamic models for the smoothing filter that is used, which we term : perfect model, assumed model, and random walk model. Note that the model classes we have chosen are meant to represent a trade off between complexity and accuracy. The perfect model captures complicated ellipsoid dynamics. There is no model mismatch in a smoothing filter based on the perfect model. Thus, a smoothing filter based on the perfect model of dynamics is most accurate. The assumed model is less accurate because it captures only the simplified, nominal dynamics of the ellipsoid. The assumed model represents a small model mismatch. Finally, the random walk model is the simplest of the three dynamic models and consequently the least accurate.

We begin by considering the perfect model of dynamics. This is not really a model class, but rather a best case smoothing filter where the dynamics used to generate the data match those used in the smoothing filter. One may think of the output of a smoothing filter based on perfect model of dynamics as a best case performance. The best results we can hope for are by using a smoothing filter based on the perfect model. The results from the other two model classes may be compared to the perfect model results to evaluate their performance. For the perfect model, $\tilde{M}_p(k) = \tilde{A}(k)$

(i.e. there is no model mismatch). Through our simulations, we will also investigate the effect of including process noise variance in the smoothing filter equations using the perfect model of dynamics. A high value of q reflects low confidence in the model. Therefore, we expect the error to increase as q increases.

Next, we consider a smoothing filter based on the assumed model. This model incorporates limited prior knowledge of the true dynamics of the ellipsoid. Suppose we know that in general the ellipsoid is shrinking and rotating, but we do not know the exact rates of change. We might use an assumed nominal model that gave an approximation to these dynamics. For example, if the true dynamics (in terms of matrix notation) are given by

$$A(k) = \begin{pmatrix} t(k) & 0 \\ 0 & t(k) \end{pmatrix} \begin{pmatrix} \cos \phi(k) & \sin \phi(k) \\ -\sin \phi(k) & \cos \phi(k) \end{pmatrix}$$

then we could use an assumed model of the form

$$M_a(k) = \begin{pmatrix} t_a & 0 \\ 0 & t_a \end{pmatrix}$$

where t_a is a time invariant approximation to $t(k)$, the true time varying rate of magnification. Note our approximation of the rotation angle is $\phi_a(k) = 0$. In terms of the vector notation that we developed in Chapter 3, $\tilde{M}_a(k) = \Gamma(M_a(k))$. When using the assumed model, neither the measurement nor the predicted state is exact. For this reason, we would expect that a range of q exists that minimizes the error in the estimate. We will investigate the possibility of “tuning” the filter with the value of q .

Finally, consider the simplest dynamic model, the random walk. This model implies we have no knowledge of the true ellipse dynamics, so we pick the simplest possible evolution. In other words, we assume the ellipse evolves as a random walk, or on average is constant. For this approach, we use $\tilde{M}_r(k) = I$. The filter used is given by (3.29)-(3.34). Again, we hope to compensate for the error introduced by the random walk assumption through the process noise. We will again investigate

the possibility of “tuning” the filter with the value of q . There is a greater model mismatch in a smoothing filter based on the random walk model than in one using the assumed model for dynamics. Therefore, we would expect that the range of q that minimizes the error in the estimate for the random walk model is higher than the corresponding range for the assumed model.

We will be able to construct dynamic models that are similar to those we have just discussed when we reconstruct the ellipsoid that approximates the left ventricle from myocardial perfusion data. The closest we can come to the perfect model case is by adding a preprocessing step to “learn” the dynamics of each individual patient from the myocardial perfusion data itself. This might be accomplished using some type of a system identification approach. Similarly, we may be able to use the assumed model approach in a model identification context which will be discussed in Chapter 5. A bank of Kalman filters would be used to determine the optimal matrix $\tilde{M}(k)$ for the patient.

Projection Angles

A third problem we might encounter in the real data is lack of knowledge of the projection angles. The smoothing filter formulation described by (3.29)-(3.34) requires knowledge of the matrix, \tilde{C} , which captures the geometry of the ellipsoidal projections. Recall that the imaging procedure used to obtain the myocardial perfusion images is highly inexact. The cardiac views obtained are roughly anterior, lateral, and left anterior oblique. Optimally estimating the three-dimensional ejection fraction requires exact knowledge of the projection angles (and hence, the matrix \tilde{C}). In Section 4.4, we will evaluate filter performance for three-dimensional ellipsoid reconstruction when these angles are known. Further discussion of the problems that arise from lack of knowledge of the \tilde{C} matrices is deferred to Chapter 6.

Measurement Noise

Finally, the smoothing filter we propose to use requires that the measurement noise has a normal distribution with zero mean and variance R . As we explained above,

the observations for reconstructing the ellipsoid that approximates the left ventricle are ellipses that have been extracted from the myocardial perfusion images. Therefore, measurement noise is the error introduced in fitting ellipses to the myocardial perfusion images. This measurement noise is due to many factors which include the noise in the image itself and human or algorithm error in fitting ellipses. It is difficult to characterize the true measurement noise introduced by fitting ellipses to the myocardial perfusion data. One commonly used approach is to assume the measurement noise is Gaussian.

In our ellipsoid reconstruction simulations described in the next two sections, the noise added to observations will have a zero mean normal distribution with variance $R = rI$. This corresponds to adding Gaussian noise to the elements of the vector that represents the ellipsoid. There will be no mismatch in the value for noise variance used to produce the measurement and that used in the smoothing reconstruction algorithm.

The problem of reconstructing the left ventricular ellipsoid model from myocardial perfusion images differs from the simulations. First, the noise in the elliptical observations from the myocardial perfusion images will almost certainly not be Gaussian. Second, the value of R which best approximates the variance of the distribution of the measurement noise in the extracted ellipses will be unknown. While the measurement noise in the simulated observations differs from the noise in the extracted ellipses, we hope the simulation results give some indication as to how well the filter performs on the myocardial perfusion data and how sensitive the filter is to noise.

4.3 Two-Dimensional Ellipsoid Reconstruction

We study the reconstruction of a computer simulated, dynamically evolving two-dimensional ellipsoid from noisy measurements of the ellipsoid itself as a preliminary exercise. This exercise yields insight into the three-dimensional ellipsoid reconstruction problem because two-dimensional ellipsoids are more easily visualized. One major difference between the two- and three-dimensional ellipsoid reconstructions is that the

input to the smoothing filter for the two-dimensional case does not consist of noisy projections of the two-dimensional ellipsoids, but rather, noisy observations of the ellipsoids themselves.

Generally in reconstructing a three-dimensional ellipsoid from two-dimensional projections, the projection geometry may be unknown or imprecisely known. For this reason, we investigate using the two-dimensional ejection fraction of the projections of a three-dimensional ellipsoid as an approximation to the true ejection fraction of the three-dimensional ellipsoid. In addition, for some very specific cases, the two-dimensional ejection fraction of the projections of a three-dimensional ellipsoid is exactly equal to the ejection fraction of the three-dimensional ellipsoid. For further discussion see Section 4.5.

4.3.1 Generation

In the generation of a computer simulated two-dimensional ellipsoid, note the following points. The characteristics of the starting ellipse as well as the rate of ellipse contraction will be chosen to reflect the characteristics of the heart. We will study two types of ellipse dynamics that are meant to cover the range of motion we might see in the heart. One simplification to note is that while true heart dynamics have irregular rates of change, our ellipses will have constant rates of change. The two types of ellipse dynamics we will investigate are shrink/expand and shrink/expand/rotate. The heart definitely has some rotation associated with its dynamics, but it is difficult to identify the exact degree of rotation. The two classes of dynamics we have chosen are meant to represent the extremes of rotation. The shrinking/expanding ellipse shows no rotation which is less than we would expect to see in the dynamics of the heart. The shrinking/expanding/rotating ellipse shows a total rotation of approximately $\pi/2$ which is more rotation than one would expect. Thus, we hope to draw conclusions on the performance of our reconstructions when the true rotation lies somewhere between these two extremes.

Shrinking and Expanding Ellipse

The first two-dimensional reconstruction we consider is the shrinking and expanding ellipse. In light of our knowledge about the shape and motion of the heart, we choose the following characteristics for our starting ellipse.

Characteristic	Symbol	Value
long axis	a	8
short axis	b	7.2
eccentricity	$\frac{b}{a}$.9
rotation	θ	0

Table 4.1: Characteristics of starting 2D ellipsoid

This starting ellipse may be represented by the 2×2 matrix in the form of (3.2)

$$E = \begin{pmatrix} 64 & 0 \\ 0 & 51.84 \end{pmatrix}$$

Our ellipse shrinks for eight steps at a rate $t = .957$ and expands for eight steps at rate $\frac{1}{t}$. That is, our dynamic matrices in (3.7) are given by

$$A(k) = \begin{cases} tI & \text{for } k = 1 \dots 8 \\ \frac{1}{t}I & \text{for } k = 9 \dots 15 \end{cases}$$

This set of ellipses is shown in Figure 4-1. The rate of contraction $t = .957$ is meant to give a true ejection fraction of 50, which is the ejection fraction of a average human heart.

Shrinking and Expanding Ellipse with Rotation

The second type of ellipse dynamics we will attempt to track is a shrinking and expanding ellipse with rotation. Our starting ellipse is the same as given above. The dynamics now also include a rotation of $\pi/15$ counterclockwise at each step in the contraction phase and a rotation of $\pi/15$ clockwise at each step in the expansion




Figure 4-1: Shrinking and Expanding Ellipse

phase. Thus, our new dynamic matrices are given by

$$A(k) = \begin{cases} \begin{pmatrix} 0.9361 & 0.1990 \\ -0.1990 & 0.9361 \end{pmatrix} & \text{for } k = 1 \dots 8 \\ \begin{pmatrix} 1.0222 & -0.2173 \\ 0.2173 & 1.0221 \end{pmatrix} & \text{for } k = 9 \dots 15 \end{cases}$$

This set of ellipses is shown in Figure 4-2.

4.3.2 Estimation

Now we consider the reconstruction of these dynamically evolving two dimensional ellipsoid. There are many directions our investigation may follow. We would like to evaluate the performance of the smoothing algorithm outlined in (3.29)-(3.34) as we vary measurement noise, dynamics of the true underlying ellipse, the models that approximate the true dynamics, and “process noise”. We will evaluate the smoothing




Figure 4-2: Shrinking and Expanding Ellipse with Rotation

filter performance by holding all but one of these factors constant. Thus, we hope to map out the performance of our reconstruction with respect to all variables.

But first, a few comments on definitions and initializations. The reconstructions will be based on noisy measurements of the ellipses themselves for reasons that we discussed in Section 4.2.2. That is, the measurements are given by

$$y(k) = \epsilon(k) + v(k)$$

where $\epsilon(k)$ is the vector representation of the ellipse and $v(k)$ is zero-mean, Gaussian white noise with variance $R = rI$. In addition, we define signal-to-noise ratio as

$$\text{SNR}(k) = \sqrt{\frac{\|\epsilon(k)\|^2}{n(n+1)r/2}} \quad (4.4)$$

where n is the dimension of the ellipsoid. Note, the ellipse we are trying to track shrinks, but the noise remains at a constant level. Thus, the signal-to-noise ratio will

decrease as the ellipse contracts. One last issue we must deal with is initialization of the forward Kalman filter. As is the case with the real data, we would like to set the initial error covariance, P_0 to be infinity to indicate that our confidence in the initial estimate is low. We set the initial estimate $\hat{\epsilon}_0$ to be the linear least square estimate. If our initial error covariance is set high enough, the initial guess, $\hat{\epsilon}_0$ is more or less ignored and the estimate at time $k = 1$ is equal to the first measurement. We have chosen to implement these initial conditions in the following manner

$$P_0 = 5R = pI$$

$$\hat{\epsilon}_0 = (\tilde{C}^T \tilde{C})^{-1} \tilde{C}^T y(1)$$

We choose the initial error covariance P_0 to be five times the measurement noise which is not infinity, but still high enough to indicate very low confidence in the initial estimate.

Shrinking and Expanding Ellipse Results

Let us begin with the simplest possible problem. We will fix the strength of measurement noise to be

$$r_1 = 10^{-3} m \tag{4.5}$$

$$\text{where } m = \frac{\|\epsilon(1)\|^2}{3} \tag{4.6}$$

The true ellipse dynamics will be shrinking and expanding (no rotation) as defined above. In addition, we will start with a smoothing filter using the perfect model for the dynamics so that $\tilde{M}_p(k) = \tilde{A}(k)$. So the only remaining variable is the “process noise”. For these conditions, the signal-to-noise ratios through the ellipse cycle are as shown in Figure 4-3. For this level of noise, the measurements are very close to the true ellipses.

In light of the facts that our measurement noise is low, our true ellipse dynamics are as simple as possible, and our model is highly accurate, we expect the estimate




Figure 4-3: Signal to Noise Ratio for Shrinking and Expanding Ellipse and Low Level Noise (r_1)

error to be very low. Also, recall that we are using “process noise” $Q = qI$ to indicate our level of confidence in the dynamic model. For the perfect model, we expect the estimate error will increase as q increases because there is no noise in the generation of our ellipses and there is no model mismatch in the smoothing filter. Our simulations support these expectations.

We define percent magnitude error as

$$100 \frac{\|\epsilon(k) - \hat{\epsilon}(k|T)\|}{\|\epsilon(k)\|}$$

where $\epsilon(k)$ is the vector that represents the true ellipse at time k and $\hat{\epsilon}(k|T)$ is the vector that represents the reconstructed ellipsoid at time k given all measurements. Figure 4-4 illustrates the percent magnitude error at each time step. Each curve corresponds to a different value of q and is the average of 100 realizations of the smoothing algorithm. In this figure, the values of q used are shown in Table 4.2 where m is defined in (4.6). Our estimate error ranges from less than 1% for low




Figure 4-4: Percent magnitude error using perfect model of dynamics. Each curve corresponds to a different value of process noise, q .

q_1	q_2	q_3	q_4	q_5	q_6	q_7	q_8
$10^{-8}m$	$10^{-7}m$	$10^{-6}m$	$10^{-5}m$	$10^{-4}m$	$10^{-3}m$	$10^{-2}m$	$10^{-1}m$

Table 4.2: Process Noise Values

values of q to approximately 5.5% for higher values of q . The error increases as the value of q increases.

Ultimately, we are interested in ejection fraction which is a function of the ellipse area. Therefore another quantity that is of interest to us is the percent area error which is defined as

$$100 \frac{\text{area}(\epsilon(k)) - \text{area}(\hat{\epsilon}(k|k))}{\text{area}(\epsilon(k))}$$

Figure 4-5 shows the variation of percent area error with q . Again, this error is low ranging from .5% to less than 3%. Percent area error also increases as q increases. Note that percent area error is smaller than percent magnitude error. This is due to

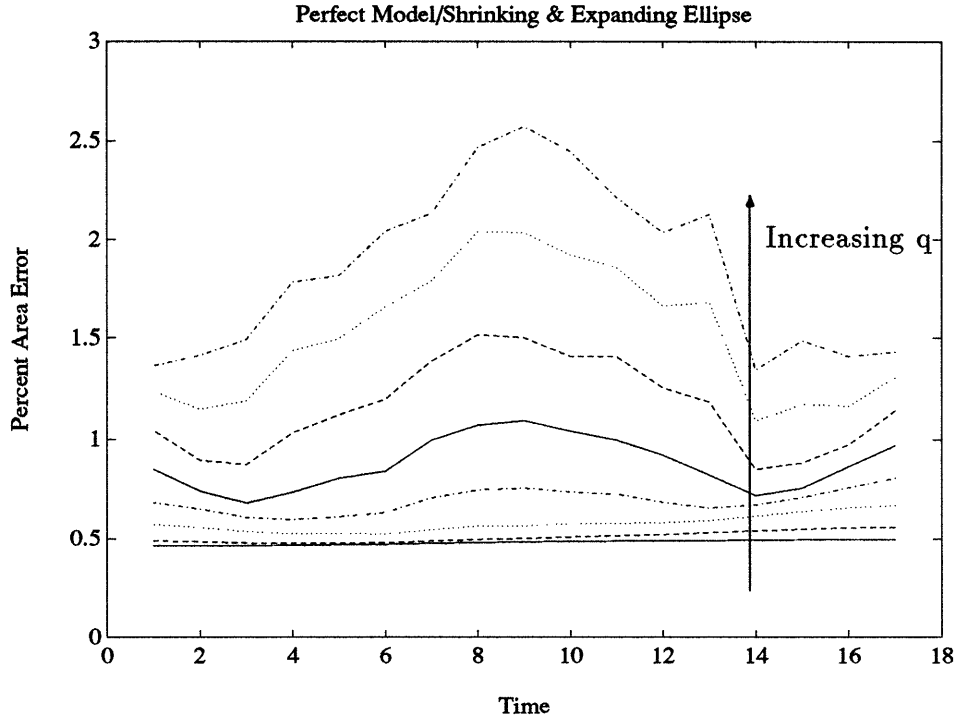


Figure 4-5: Percent area error using perfect model of dynamics. Each curve corresponds to a different value of process noise, q .

the fact that an error in orientation affects magnitude error but has no effect on area error. Note also that both types of error are largest at the same point in the ellipse cycle when the actual ellipse size is smallest. This is a result of the fact that signal to noise ratio is smallest at this point in the ellipse cycle.

Now we will add another level of complexity. We would like to see the effect of varying the strength of the measurement noise, $R = rI$. The values of r used are shown in Table 4.3. Also, Figure 4-7 shows several ellipses in the ellipse cycle and one set of corresponding measurements with measurement noise r_4 . At this level of noise, the measurements poorly reflect the true shape of the ellipse. The signal-to-noise ratios for these values of measurement noise are plotted in Figure 4-6. The highest noise level is chosen so that none of the measurements generated is a degenerate ellipse (i.e. all measurements correspond to matrices that are positive definite). In Figure 4-8, we have plotted the average percent magnitude error as we vary q and r . Each individual curve is the percent magnitude error over time averaged over 100 sample




Figure 4-6: Signal to Noise Ratios for several noise levels




Figure 4-7: Ellipses #1,3,6, and 8 from set of sixteen shrinking and expanding ellipses. Also shown are the corresponding measurements. True ellipses are shown in solid lines and the measurements are shown in dotted lines




Figure 4-8: Average of percent magnitude error over time and over 100 runs. Uses perfect model of dynamics. Each curve corresponds to a different value of measurement noise, r

r_1	$10^{-3}m$
r_2	$10^{-2.47}m$
r_3	$10^{-1.94}m$
r_4	$10^{-1.40}m$

Table 4.3: Measurement Noise Values

paths versus q with r fixed. As we would expect, the average percent magnitude error increases as the assumed strength of measurement noise is increased. For all values of r considered, the error will increase as q increases.

Of course, the real quantity of interest is the percent ejection fraction error. Figure 4-9 shows the percent ejection fraction error as a function of q . Again, each curve corresponds to a value of r defined above. The two-dimensional version of the definition for ejection fraction (EF) is

$$\text{EF} = \frac{\text{largest area} - \text{smallest area}}{\text{largest area}}$$

Our ellipse cycle as defined in Section 4.3.1 has two states where the area corresponds to a maximum. Therefore, we will take our largest area to be the average of the areas of the two maximum states. Our definition for the smoothed ejection fraction estimate is then

$$\widehat{\text{EF}} = \frac{\text{area}(\hat{e}(16)) + \text{area}(\hat{e}(1)) - 2\text{area}(\hat{e}(9))}{\text{area}(\hat{e}(16)) + \text{area}(\hat{e}(1))} \quad (4.7)$$

Once again, we note that because this is a function of area, which is not affected by errors in orientation, the percent ejection fraction error is smaller than the percent average magnitude error. The percent ejection fraction error increases as the strength of measurement noise increases. At the lowest measurement noise, the lowest value for percent ejection fraction error is less than .5%. At the highest measurement noise, the lowest value for percent ejection fraction error goes up to approximately 1.5%. Also, for all values of r , the percent ejection fraction error increases as q increases. At the lowest measurement noise, the percent ejection fraction error increases from ap-




Figure 4-9: Percent ejection fraction error using perfect model of dynamics. Each curve corresponds to a different value of measurement noise, r

proximately .5% to 3% as q increases. At the highest measurement noise, the percent ejection fraction error increases from approximately 1.5% to 12% as q increases.

Now let us turn our attention to the next class of dynamic models. A smoothing filter based on the assumed model (as described in Section 4.2.2) uses an approximation to the true dynamic matrices. We will consider an assumed model such that

$$M_a(k) = \begin{cases} t_a I & \text{for } k = 1 \dots 8 \\ \frac{1}{t_a} I & \text{for } k = 9 \dots 15 \end{cases} \quad (4.8)$$

where $t_a = .947$ which gives an ejection fraction of roughly 60. Figure 4-10 shows the average percent magnitude error versus q for the same four values of r as above. Figure 4-11 shows the percent ejection fraction error versus q . Of course, these figures show that both the average percent magnitude error and percent ejection fraction error are higher than the corresponding quantities for the perfect model reconstruction. This comes as no surprise as the assumed model is not matched to the true dynamics




Figure 4-10: Average percent magnitude error using assumed model. Each curve corresponds to a different value of measurement noise, r .




Figure 4-11: Average ejection fraction error using assumed model of dynamics. Each curve corresponds to a different value of measurement noise, r .

of the system. Again, each curve corresponds to a different value of r . As we would expect, the errors increase for increasing values of r . The shape of the curves seem to indicate again that we might “tune” the filter with q . That is to say, our best performance for ejection fraction estimates for each value of r comes for $1 < q < 10^2$. In this range of values of q , the average percent magnitude error increases from 2.5% to 8% as measurement noise increases from r_1 to r_4 . Also, the percent ejection fraction error increases from 3% to 8% as measurement noise increases. The important point to note is that, for both the average percent magnitude error and the percent ejection fraction error, the error is minimized by range of values of q that is not equal to zero.

The final model class to consider is the random walk model. The smoothing filter based on the random walk model of dynamics assumes that our approximation to the true dynamic matrices is the identity :

$$\tilde{M}_r(k) = I \text{ for } k = 1 \dots 15$$

Figures 4-12 and 4-13 show the percent average magnitude error and average ejection fraction error respectively.

To summarize, the average percent magnitude error is minimized for $1 < q < 100$ for all levels of measurement noise. The average percent magnitude error in this range of values of q increases from 2.5% to 14% as measurement noise increases. The percent ejection fraction error is minimized for a slightly different range $10 < q < 1000$. In this range of values of q , the percent ejection fraction error increases from 3% to 12% as measurement noise increases. As we had hoped, there is a range of q that minimizes both types of error over a broad range of measurement noise intensities. These errors are the highest of the three model classes because the random walk model introduces the most model mismatch. However, we still see that the filter may still be tuned to minimize error by changing the value of q .

We stated earlier that one of the questions we hoped to answer was what is the simplest model class that improves upon the current methods of ejection fraction estimation. In estimating ejection fraction from myocardial perfusion images, physi-




Figure 4-12: Average percent magnitude error using random walk model of dynamics. Each curve corresponds to a different value of measurement noise, r

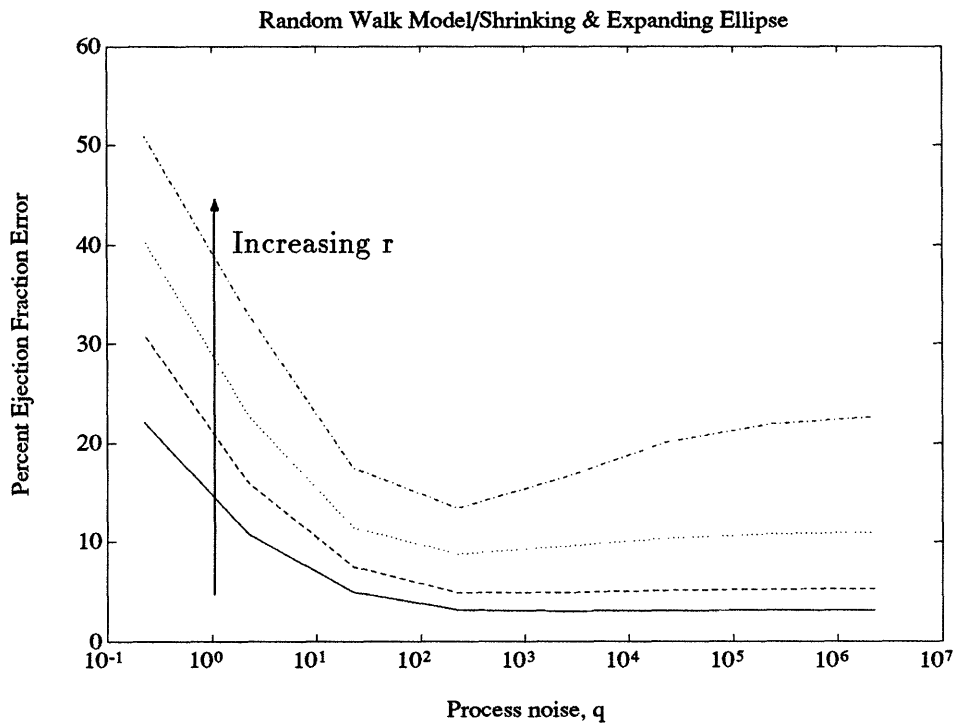


Figure 4-13: Average ejection fraction error using random walk model of dynamics. Each curve corresponds to a different value of measurement noise, r




Figure 4-14: Static Ejection Fraction Error.

cians have attempted to estimate ejection fraction using only the end systolic and end diastolic frames. The analogous procedure for simulated data is to use only the measurements at the fully contracted and fully expanded phases. This is effectively a static estimate where the dynamic model is ignored. We compare the ejection fraction obtained from the smoothing filter reconstructed ellipse to the static ejection fraction. The static ejection fraction will be based on a linear least square estimate for area only at the fully contracted and fully expanded phases. Thus, we should compare the results for each of the model classes to this static estimate. Again for the same values of r as above, Figure 4-14 shows the percent ejection fraction error.

Comparing this to earlier results, we conclude that even our simplest model (the random walk) may be “tuned” with q to give a significant error reduction at high measurement noise levels. The error may be reduced by more than 10%. At low noise levels, there is not much improvement over static estimation. As we saw in Chapter 2 our images are of such poor quality that it is reasonable to assume that we will be dealing with high levels of noise. We can conclude that it will be possible to

significantly improve on current methods of ejection fraction estimation using even the random walk model.

Shrinking and Expanding Ellipse with Rotation Results

We treat the reconstruction of the shrinking and expanding ellipse with rotation in a similar manner as the shrinking and expanding ellipse without rotation. We use a smoothing filter that is based on each of the dynamic model classes we discussed in Section 4.2.2. The same values for measurement noise intensity as in Table 4.3 are used; the same values for process noise as in Table 4.2 are used. As we hoped, these results are similar to the results for the no rotation case.

To start with, consider a smoothing filter reconstruction based on the perfect model of dynamics. In this case, $M_p(k) = A(k)$. Figures 4-15 and 4-16 show the average percent magnitude error and the percent ejection fraction error respectively. Next, consider a smoothing filter based on the assumed model of dynamics. The assumed model uses the dynamic matrix \tilde{M}_a of (4.8). Now we have not only a mismatch in the contraction rates, but also a mismatch in rotation rates. The model assumes no rotation, but the true ellipse dynamics does have rotation. Figures 4-17 and 4-18 again illustrate the variation of average percent magnitude error and percent ejection fraction error. These errors are slightly higher than the corresponding errors for the no rotation case. However, important qualitative features remain the same. The errors may be reduced by tuning the filter with q for all levels of r . Finally, the corresponding results for a smoothing filter based on the random walk model are shown in Figures 4-19 and 4-20. These results are very similar to the no rotation case. Again, when these errors are compared to the static estimation errors, we find that it is possible to significantly improve the accuracy of the ejection fraction estimate by using the Rauch-Tung Striebel algorithm. Our results show that the smoothing filter approach is valid for the two extremes of rotational dynamics that we expect to encounter.




Figure 4-15: Average percent magnitude error using perfect model for dynamics. Each curve corresponds to a different value of measurement noise, r




Figure 4-16: Ejection fraction error using perfect model for dynamics. Each curve corresponds to a different value of measurement noise, r




Figure 4-17: Average percent magnitude error using assumed model for dynamics. Each curve corresponds to a different value of measurement noise, r




Figure 4-18: Ejection fraction error using assumed model for dynamics. Each curve corresponds to a different value of measurement noise, r




Figure 4-19: Average percent magnitude error using random walk model for dynamics. Each curve corresponds to a different value of measurement noise, r




Figure 4-20: Ejection Fraction error using random walk model for dynamics. Each curve corresponds to a different value of measurement noise, r

4.4 Three-Dimensional Ellipsoid Reconstruction

4.4.1 Generation

For the three-dimensional ellipsoid reconstruction problem, we will consider only the case of a shrinking and expanding ellipsoid with no rotation. This is not exactly the type of dynamics we expect to see in the heart. However, the insight gained by studying this simplified problem should guide our approach with the real data. The starting ellipsoid has semi-axis lengths of 8, 8, and 7.2. The ellipsoid is oriented with the coordinate axes; there is no rotation. This starting ellipse may be represented by the 3×3 matrix :

$$E = \begin{pmatrix} 64 & 0 & 0 \\ 0 & 51.84 & 0 \\ 0 & 0 & 64 \end{pmatrix}$$

Our ellipse shrinks for eight steps at a rate $t = .9715$ and expands for eight steps at rate $\frac{1}{t}$. That is, our dynamic matrices are given by

$$A(k) = \begin{cases} tI & \text{for } k = 1 \dots 8 \\ \frac{1}{t}I & \text{for } k = 9 \dots 15 \end{cases}$$

This rate of contraction again gives an ejection fraction of 50.

4.4.2 Estimation

Reconstruction for the three-dimensional ellipsoid problem is accomplished using the smoothing algorithm outlined in (3.29)-(3.34). We are assuming that the projection angles are known. That is, that we know the matrix \tilde{C} . The measurements are given by

$$y(k) = \tilde{C}\epsilon(k) + v(k) \tag{4.9}$$

$$C = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (4.10)$$

where $\tilde{C} = \Gamma(C)$ and $v(k)$ is zero-mean, Gaussian white noise with variance $R = rI$. The projections are onto the xy , xz , and yz -planes.

The three-dimensional reconstruction may be treated in the same way the two-dimensional reconstruction was treated. We will use a smoothing filter based on the same three model classes: perfect model, assumed model, and random walk. The measurement noise intensities are the same as for the two-dimensional case. The signal-to-noise ratios are somewhat different than in the two-dimensional case because the magnitude of the state is different. Figure 4-21 shows the signal-to-noise ratios through the ellipse cycle for the values of noise considered. First using a smoothing filter based on the perfect model, we set the filter model $\tilde{M}_a(k) = \tilde{A}(k)$. Figures 4-22 -4-23 show the average percent magnitude and percent ejection fraction error. The results show that, even for the highest levels of measurement noise, it is possible to obtain average percent magnitude error less than 6% and percent ejection fraction error less than 1.5% by choosing q to be low. For a smoothing filter based on the assumed model, the filter model is given by :

$$M_a(k) = \begin{cases} t_a I & \text{for } k = 1 \dots 8 \\ \frac{1}{t_a} I & \text{for } k = 9 \dots 15 \end{cases}$$

where $t_a = .9615$ gives an ejection fraction of roughly 60. The corresponding average percent magnitude errors and percent ejection fraction errors are shown in Figures 4-24-4-25. These results show that average percent magnitude errors are less than 10% even at high measurement noise levels for $q < 10$. Also, percent ejection fraction errors is less than 6% for all values of noise if $1 < q < 100$. Finally, for a smoothing filter




Figure 4-21: Several Signal-to-Noise Ratios for Three Dimensional Reconstruction

based on the random walk model, $\tilde{M}_r(k) = I$. These results are shown in Figures 4-26 - 4-27. Even at high levels of measurement noise, average percent magnitude error will be less than 11% for $1 < q < 100$. Percent ejection fraction is less than 15% for all noise levels if $q > 10$. As in the two-dimensional case, the error increases as we go from the perfect model to the assumed model to the random walk model. For a smoothing filter based on the perfect model, both types of error increase as q increases. The smoothing filter using an assumed model results show that we can tune the filter with q to minimize both the average percent magnitude error and the percent ejection fraction error. The smoothing filter based on a random walk model results are not quite as good as we had hoped. It is possible to minimize the average percent magnitude error by adjusting q . There is not a similar result for the percent ejection fraction error. The percent magnitude error represented by the curves in Figure 4-27 starts out high for $q = .1$, decreases steadily for $.1 < q < 100$, and then remains constant at higher values of q . These results indicate that we will not see significant improvement over static estimates. Indeed, when we compare these results




Figure 4-22: Average percent magnitude error using perfect model of dynamics. Each curve corresponds to a different value of measurement noise, r .

to the static estimate results in Figure 4-28, we see that we have not reduced percent ejection fraction error. Our conclusion is that the simplest model that improves on static estimation techniques is the assumed model class of dynamics. We will be able to extend these results to the real data in Chapter 6.

4.5 2D and 3D Ejection Fractions

As we mentioned earlier, for some very specific cases, the ejection fraction of the two-dimensional projections of a three-dimensional ellipsoid (ellipses) is equivalent to the ejection fraction of the original ellipsoid. Before showing this explicitly, we discuss our approach more qualitatively. Our ultimate objective is to determine the ejection fraction of a dynamically evolving three-dimensional ellipsoid. Recall that ejection fraction is a measure of contraction along all axes. Picture a contracting three-dimensional ellipsoid. The projection of this ellipsoid on a plane is a contracting




Figure 4-23: Ejection fraction error using perfect model of dynamics. Each curve corresponds to a different value of measurement noise, r .




Figure 4-24: Average percent magnitude error using assumed model of dynamics. Each curve corresponds to a different value of measurement noise, r .




Figure 4-25: Ejection fraction error using assumed model of dynamics. Each curve corresponds to a different value of measurement noise, r .

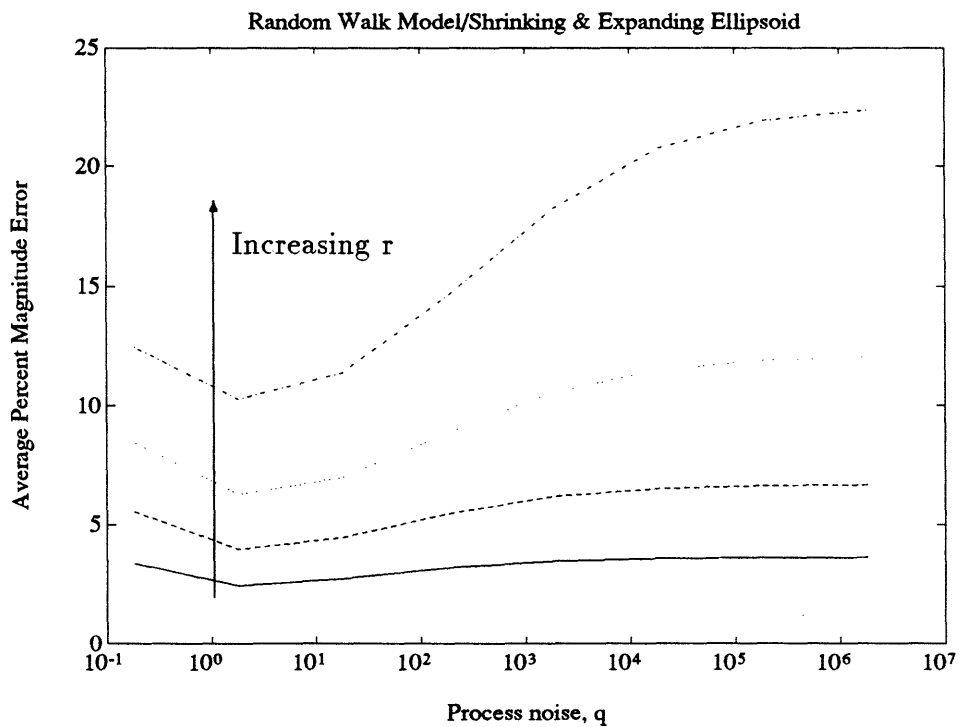


Figure 4-26: Average percent magnitude error using random walk model of dynamics. Each curve corresponds to a different value of measurement noise, r .

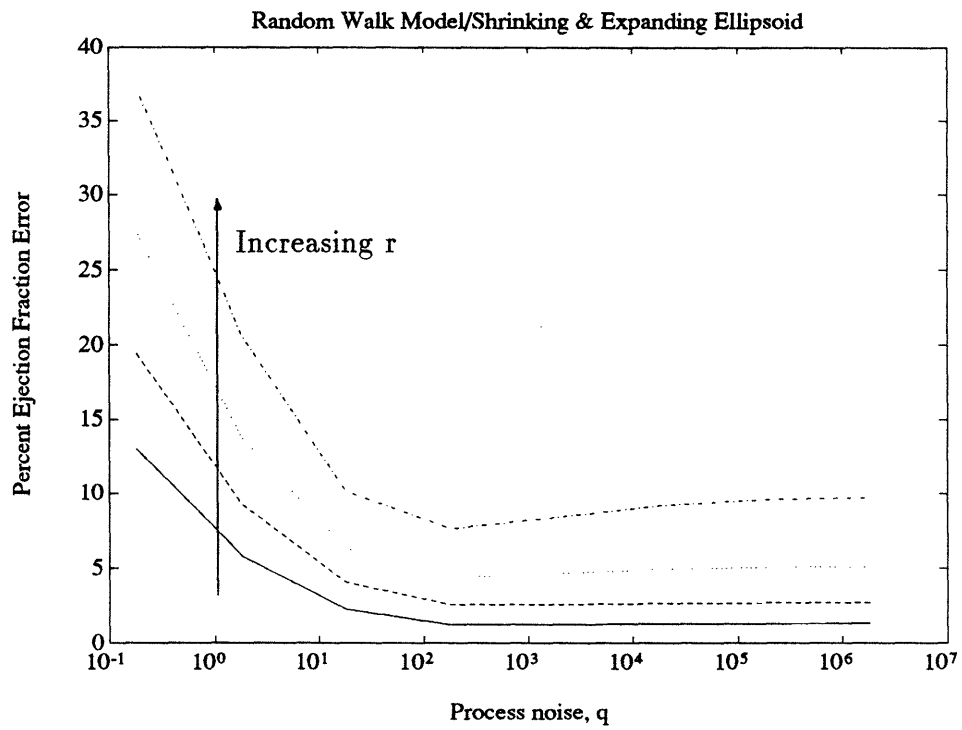


Figure 4-27: Ejection fraction error using random walk model of dynamics. Each curve corresponds to a different value of measurement noise, r .




Figure 4-28: Static Ejection Fraction Error (3D).

two-dimensional ellipsoid (an ellipse). From this projection, any contraction along an axis perpendicular to the plane of projection is imperceptible. Contraction along an axis that is perpendicular to the plane of the projection is information that is lost in looking at the projection. If, however, the three-dimensional ellipsoid does not contract along the axis that is perpendicular to the plane of the projection, no information is lost. It seems intuitive that there exists a simple relationship between the ejection fraction of a three-dimensional ellipsoid that contracts only along two axes and the ejection fraction of the two-dimensional ellipsoid that is the projection of the three-dimensional ellipsoid on to a plane that is perpendicular to the axis of no contraction. As it turns out, these two quantities are equivalent. Thus, in some cases, determining the ejection fraction of the three-dimensional ellipsoid is equivalent to the simpler problem of determining the ejection fraction of the two-dimensional ellipsoidal projection. Appendix 4-A shows this result more explicitly.

This, of course, is a specialized case. Generally, we will not encounter such a scenario. We may, however, encounter projections that are roughly perpendicular to an axis along which the ellipsoid does not contract very much. We investigate the possibility of using the two-dimensional ejection fraction as an approximation to the three-dimensional ejection fraction of the ellipsoid that approximates the left ventricle in Chapter 6.

4.6 Conclusions

In this chapter, we have taken the first steps toward using the techniques explored in Chapter 3 to estimate the ejection fraction of the heart. We have used the Rauch-Tung Striebel smoothing algorithm to reconstruct computer simulated, dynamically evolving, two and three-dimensional ellipsoids. The problem was formulated in such a way that our results might be easily applicable to real data. First, we explored a few model classes that approximated the true ellipsoid dynamics. Our objective was to determine which was the simplest model class that still gave improvement over currently used static estimation techniques. For the two-dimensional case, we

found that we were able to significantly improve the estimate using the simplest possible model, the random walk. In the three-dimensional case, we found it was necessary to use a more complicated dynamic model class to gain improvement over the static estimate results. We discussed methods of adjusting the filter to account for mismatches between the true and modelled dynamics. Second, we addressed techniques of estimating ejection fraction in spite of lack of knowledge of projection angles. Our results give us a good deal of insight into the generalized dynamic ellipsoid reconstruction problem that will be applicable to the real data problem at hand.

4-A 2D and 3D Ejection Fractions

We now show the equivalence of the two-dimensional and three-dimensional ejection fraction for this special case more explicitly. Consider an ellipsoid whose axes are aligned with the coordinate system. For this ellipsoid, the semi-axis lengths, a , b , and c are the extent of the ellipsoid along the x , y , and z axes respectively. Suppose this three-dimensional ellipsoid contracts only along two axes, say the x and y axes. Recall that the volume of a three-dimensional ellipsoid is given by

$$\text{Vol} = \frac{4}{3}\pi abc$$

where a , b , and c are the semi-axis lengths. Thus, the three-dimensional ejection fraction is given by

$$\text{3D-EF} = \frac{a_1 b_1 c_1 - a_2 b_2 c_2}{a_1 b_1 c_1} \quad (4.11)$$

where a_1 , b_1 , and c_1 correspond to the semi-axis lengths of the fully expanded ellipsoid and a_2 , b_2 , and c_2 correspond to the semi axis lengths of the fully contracted ellipsoid. The ellipsoid we have specified does not contract along the z -axis; therefore, $c_1 = c_2$. We may rewrite (4.11) as

$$\text{3D-EF} = \frac{a_1 b_1 - a_2 b_2}{a_1 b_1} \quad (4.12)$$

The projection of this ellipsoid on the xy -plane (i.e. the plane that is perpendicular to the axis of no contraction) is an ellipse with semi-axes a and b that are aligned with the x and y -axes respectively. Recall that the area of a two-dimensional ellipsoid is given by :

$$\text{Area} = \pi ab$$

where a and b are the semi-major and minor axis lengths. The two-dimensional version of ejection fraction for this ellipsoid is given by

$$\text{2D-EF} = \frac{a_1 b_1 - a_2 b_2}{a_1 b_1} \quad (4.13)$$

where a_1 and b_1 are the semi-axis lengths of the fully expanded ellipse and a_2 and b_2 are the semi-axis lengths of the fully contracted ellipse. Comparing (4.12) and (4.13), we see that the ejection fraction for the three-dimensional ellipsoid equals the ejection fraction for the two-dimensional ellipsoid for this special case.

Chapter 5

Model Identification

5.1 Objective

Recall that Chapter 4 included a discussion of the error created by model mismatch. That is, we investigated the errors introduced when the true ellipsoid dynamics differ from those used by the smoothing filter to reconstruct the ellipsoid. We discussed smoothing filters based on three types of dynamic models : the perfect model, the assumed model, and the random walk model. As we showed in Chapter 4, it is possible to improve on static estimates of ejection fraction by using a smoothing filter based on the assumed model. However, we did not discuss how to choose the assumed model dynamics.

In this chapter, we discuss a method of choosing the assumed model dynamics based on a model identification or parameter estimation scheme. The model identification scheme described is used to choose which of several hypothesized dynamic models best approximates the dynamics used to create the ellipsoid. Model identification employs a bank of Kalman filters based on each of the hypothesized models to determine which model is most likely given all measurements. In this chapter, we use model identification to choose between two hypothesized models. Thus, we use a Kalman filter bank that consists of two filters. Note that the objective of this coarse model identification procedure is to determine which of the hypothesized models is the best approximation to the true ellipsoid dynamics rather than to determine the

best overall estimate to the true ellipsoid dynamics. Thus, our model identification gives a coarse approximation to the true ellipsoid dynamics and provides a means to determine an assumed model to be used in the smoothing filter based reconstructions described in Chapter 4 which in turn are used to estimate ejection fraction.

In addition, model identification may be thought of as a hypothesis test used to determine whether ejection fraction is normal or below normal. Each hypothesized model used in the model identification scheme has an associated hypothesized ejection fraction. The true ejection fraction tends to lie within a range around the hypothesized ejection fraction that is associated with the model chosen by the model identification scheme. Thus, model identification may be used to determine a range of the probable values for the true ejection fraction.

This chapter includes a summary of model identification background. In addition, we include the results of three experiments. Our input will be noisy measurements of simulated dynamically evolving two-dimensional ellipsoids. We start by investigating the performance of the model identification scheme when one of the hypothesized models exactly matches the true ellipsoid dynamics. We also investigate the performance of model identification as the difference in the ejection fractions of the two hypothesized models becomes smaller. Next, we present simulation results when the true and hypothesized dynamics differ in contraction rate only. Finally, we present simulation results where the true dynamics include contraction and rotation, but the hypothesized dynamics include only contraction.

5.2 Background

As we mentioned previously, the model identification approach that we describe in this section is used to determine which of several models most closely approximates the true dynamics of the system given all the measurements. We concentrate on the case where model identifier chooses between two hypothesized models, but a similar approach may be used to include several hypothesized models. For further details see [15]. Recall that the system of equations that describes the evolution and

measurement of a vector, $\epsilon(k)$, that represents an ellipsoid is given by :

$$\epsilon(k+1) = \tilde{A}(k)\epsilon(k) \quad (5.1)$$

$$y(k) = \tilde{C}\epsilon(k) + v(k) \quad (5.2)$$

where $v(k)$ is a zero mean, Gaussian white noise process with covariance R .

The matrices $\tilde{M}_1(k)$ and $\tilde{M}_2(k)$ are two hypothesized models of the true dynamics of the system $\tilde{A}(k)$. We choose $\tilde{M}_i(k)$ as the better approximation to the true system dynamics if model i is the more likely model given all the measurements thus far. We define

$$p_i(k) = \Pr(\text{model } i \text{ is correct} | Y_k) \quad (5.3)$$

where Y_k is the set of measurements up to and including the measurement at time k .

Applying Bayes' rule, we rewrite (5.3) in the following recursive form

$$p_i(k) = \frac{p(y(k)|Y_{k-1}, \text{model } i \text{ is correct})p_i(k-1)}{\sum_{j=1}^2 p(y(k)|Y_{k-1}, \text{model } j \text{ is correct})p_j(k-1)} \quad (5.4)$$

Thus, to determine $p_i(k)$ we must calculate $p(y(k)|Y_{k-1}, \text{model } i \text{ is correct})$. This quantity is obtained from the Kalman filter based on model i . Recall the Kalman filtering equations,

$$\hat{\epsilon}_i(k|k) = \hat{\epsilon}_i(k|k-1) + K_i(k)\nu_i(k) \quad (5.5)$$

$$\nu_i(k) = y(k) - \tilde{C}\hat{\epsilon}_i(k|k-1) \quad (5.6)$$

$$K_i(k) = P_i(k|k-1)\tilde{C}^T[\tilde{C}P_i(k|k-1)\tilde{C}^T - R]^{-1} \quad (5.7)$$

$$V_i(k) = \tilde{C}P_i(k|k-1)\tilde{C}^T - R \quad (5.8)$$

$$P_i(k|k) = P_i(k|k-1) - K_i(k)\tilde{C}P_i(k|k-1) \quad (5.9)$$

$$\hat{\epsilon}_i(k+1|k) = \tilde{M}_i\hat{\epsilon}_i(k|k) \quad (5.10)$$

$$P_i(k+1|k) = \tilde{M}_iP_i(k|k)\tilde{M}_i^T + Q \quad (5.11)$$

The subscript i is used to indicate that these Kalman filter equations correspond to

the filter based on model i . We know

$$p(y(k)|Y_k, \text{model } i \text{ is correct}) = N(\nu_i(k); 0, V_i(k)) \quad (5.12)$$

as shown in [15]. Thus, by substituting (5.12) in (5.4), we determine the probability that model i is correct given all measurements up to and including time k .

For each time k , model $\tilde{M}_i(k)$ is the more likely approximation to the true dynamics of the system $\tilde{A}(k)$ if $p_i(k) > p_j(k)$ for $j \neq i$. Of course, the final model chosen by the model identifier is the hypothesized model that is most likely given all the data. That is, the final decision of the model identifier is $\tilde{M}_i(k)$ if $p_i(T) > p_j(T)$ where T is the time interval over which we have data. Thus, the model identification technique described in this section consists of two Kalman filters and a recursive probability calculation.

5.3 Simulation Setup

We generate a computer simulated, dynamically evolving ellipsoid and its measurements as described by (5.1) and (5.2). For these simulations, we use a two-dimensional ellipsoid whose characteristics are chosen to reflect those of projections of the ellipsoid that approximates the left ventricle. The characteristics of the starting ellipse are as given in Table 4.1. This simulated ellipse shrinks for eight steps at a rate t_a and expands for eight steps at rate $\frac{1}{t_a}$. The matrices that describe this type of dynamics as in (3.7) are represented as $\tilde{A}(k) = \Gamma(A(k))$ where $\Gamma(\cdot)$ is defined in Section 3.2.2 and

$$A(k) = \begin{cases} t_a I & \text{for } k = 1 \dots 8 \\ \frac{1}{t_a} I & \text{for } k = 9 \dots 16 \end{cases} \quad (5.13)$$

These dynamics yield the actual ejection fraction which we denote EF_a . Measurements, as described by (5.2), are noisy observations of the ellipsoid itself (i.e. $\tilde{C} = I$). The variance of the measurement noise, $v(k)$, is given by $R = rI$.

In addition, we generate two hypothesized dynamic models which are described

by the matrices $\tilde{M}_1(k) = \Gamma(M_1(k))$ and $\tilde{M}_2(k) = \Gamma(M_2(k))$ where

$$M_i(k) = \begin{cases} t_i I & \text{for } k = 1 \dots 8 \\ \frac{1}{t_i} I & \text{for } k = 9 \dots 16 \end{cases} \quad (5.14)$$

The ejection fractions corresponding to the hypothesized models will be denoted as EF_1 and EF_2 .

Before continuing, a few comments on initializations and definitions are in order. For reasons discussed in Section 4.3.2, the Kalman filter is initialized with

$$P_0 = 5R = pI$$

$$\hat{\epsilon}_0 = (\tilde{C}^T \tilde{C})^{-1} \tilde{C}^T y(1)$$

In addition, the model identification scheme must be initialized with the prior probability of each model. We assume that both hypothesized models are equally likely and set $p_1(0) = p_2(0) = 1/2$. Finally, note that signal-to-noise ratio is defined as in (4.4).

To summarize, our simulation procedure is as follows. First, we generate a dynamically evolving ellipsoid and its measurements as described in Equations (5.1)-(5.2). The measurements described by (5.2) are the input to a model identifier which consists of two Kalman filters, KF_1 and KF_2 , based on hypothesized dynamic models $\tilde{M}_1(k)$ and $\tilde{M}_2(k)$, respectively. Based on the output of the Kalman filters, $\nu_1(k)$, $\nu_2(k)$, $V_1(k)$, and $V_2(k)$, we calculate $p_1(k)$ and $p_2(k)$ as defined in (5.4). Finally, we choose model i as the better approximation to the true dynamics of the system given all the data if $p_i(T) > p_j(T)$ where T is the interval over which we have data.

5.4 Experiment #1 : Matched Contraction

The objective of this first model identification experiment is to investigate the performance of the model identification technique when one of the hypothesized dynamic models is exactly equal the dynamics used to create the set of ellipses. In general,

we will not encounter such a scenario in the real world. Indeed, we need to know the true dynamics to construct a hypothesized model that exactly matches true dynamics. This exercise is meant to give us insight into the performance of the model identification scheme.

In this section, we conduct a series of model identification simulations where the hypothesized model $\tilde{M}_1(k)$ equals the true ellipsoid dynamics $\tilde{A}(k)$. We investigate the effect of noise on the performance of the model identification scheme. In addition, we evaluate the performance of model identification as the two hypothesized models become more similar.

We begin with the simplest problem. The true rate of ellipsoid contraction is $t_a = .936$ which yields an actual ejection fraction $EF_a = .65$. We will fix the strength of measurement noise to be

$$r_1 = 10^{-3}m \quad (5.15)$$

$$\text{where } m = \frac{\|\epsilon(1)\|^2}{3} \quad (5.16)$$

For this rate of contraction and noise level, the signal-to-noise ratio is as shown in Figure 5-1. The first hypothesized model has a contraction rate equal to the actual contraction rate (i.e. $t_1 = t_a = .936$) and therefore, $EF_1 = .65$. The second hypothesized model has a contraction rate $t_2 = .986$ which yields an ejection fraction, $EF_2 = .2$. Note that the separation between ejection fractions of the two models is large, 0.45. For 100 realizations of this simulation, we obtain the results shown in Figure 5-2. This figure shows percent of our realizations that correctly choose $\tilde{M}_1(k)$ as the best approximation to the true ellipse dynamics $\tilde{A}(k)$ at each time step (i.e. the number of realizations where $p_1(k) > p_2(k)$). For this low noise and high separation between hypothesized ejection fractions, we get 100% classification after the third time step.

Next, we investigate the performance of the model identification as the strength of measurement noise is increased. The same values as above are used for true and hypothesized contraction rates. We investigate the performance of model identification




Figure 5-1: Signal to Noise Ratios for low noise, r_1




Figure 5-2: Percent correctly identified for low noise r_1 and high degree of ejection fraction separation




Figure 5-3: Signal to Noise Ratios for several noises with true EF = .65

for the measurement noise strengths given in Table 5.1. Note that m is as defined by

r_1	$10^{-3}m$
r_2	$10^{-2.53}m$
r_3	$10^{-2.06}m$
r_4	$10^{-1.60}m$

Table 5.1: Measurement Noise Values

(5.16). The signal-to-noise ratios for these values of noise are shown in Figure 5-3. Again, we conduct 100 realizations of the simulation. Figure 5-4 shows the percent of realizations that correctly identify $\tilde{M}_1(k)$ as the model which more closely approximates the true ellipse dynamics for each time step as measurement noise increases. As we expect, an increase in measurement noise results a lower percentage of correctly identified realizations. However, because the separation of the hypothesized ejection fractions is large, 100% of our realizations correctly identified $\tilde{M}_1(k)$ at time $k = 17$.




Figure 5-4: Percent correctly identified for several noises and model separation = .45

Next, consider the case where the true rate of ellipsoid contraction is $t_a = 0.951$ which yields an actual ejection fraction $EF_a = 0.55$. Again, the first hypothesized model has a contraction rate which equals the contraction rate of the true ellipsoid dynamics (i.e. $t_1 = t_a = 0.951$). The second hypothesized model has a contraction rate $t_2 = 0.978$ which yields an ejection fraction, $EF_2 = 0.3$. Now the separation between the two hypothesized ejection fraction is 0.25. We consider the same set of measurement noises specified in Table 5.1. The signal-to-noise ratios at each time step for the different values of measurement noise are shown in Figure 5-5. Note that the signal-to-noise ratios differ from those shown in Figure 5-3 because the ellipsoid dynamics we are considering now have a different rate of contraction. In Figure 5-6, we show the percent of the realizations that correctly choose $\tilde{M}_1(k)$ at each time step and for increasing measurement noise. As we expect, we see a degradation in the performance of the model identification scheme because of the decrease in the separation between the two hypothesized ejection fractions. By time $k = 17$, we have correctly chosen $\tilde{M}_1(k)$ in over 98% of our realizations for all levels of noise.




Figure 5-5: Signal to Noise Ratios for several noises with true EF = .55




Figure 5-6: Percent correctly identified for several noises and model separation = .25




Figure 5-7: Signal to Noise Ratios for several noises with true EF = .5

Similarly, we conduct 100 realizations of the same simulation with $t_a = t_1 = 0.957$ which yields $EF_a = 0.5$ and $t_2 = 0.968$ which yields $EF_2 = 0.40$. Thus, the separation between the two hypothesized ejection fractions is .10. For the values of noise strength specified in Table 5.1 and $t_a = 0.957$, the signal-to-noise ratios at each time step are shown in Figure 5-7. In Figure 5-8, we show the percent of the realizations that correctly choose $\tilde{M}_1(k)$ at each time step and for increasing measurement noise. For low levels of noise, 100% of the realizations correctly identify model #1 as the more likely of the two hypothesized models at time $k = 17$. For higher noise levels, only about 80% of the realizations correctly conclude that model #1 is more accurate at time $k = 17$.

Finally, we conduct another 100 realizations of this simulation with $t_a = t_1 = 0.957$ yielding $EF_a = 0.5$ and $t_2 = 0.963$ which yields $EF_2 = .45$. Now, we have reduced the separation between the two hypothesized ejection fractions to just 0.05. Again, we use the measurement noise strengths specified in Table 5.1; the corresponding signal-to-noise curves are the same as given in Figure 5-7. Figure 5-9 shows the percent of




Figure 5-8: Percent correctly identified for several noises and model separation = .10

realizations that correctly identify $\tilde{M}_1(k)$ at each time step for increasing measurement noise. At the lowest level of measurement noise, 100% of the realizations correctly conclude that that model #1 is the best approximation to the true dynamics; at the highest level of measurement noise, about 65% of the realizations correctly conclude that model #1 is the best approximation to the true dynamics.

We may summarize the results of this experiment by plotting the percent of realizations that are correctly identified at time $k = 17$ as a function of the separation between hypothesized ejection fractions. Figure 5-10 shows such a plot where each curve represents the different values of measurement noise strength considered. The results from this experiment show the model identifier performs well when one of the hypothesized dynamic models matches the true ellipsoid dynamics. If the separation between the ejection fractions of the hypothesized models is high (i.e. > 0.45), 100% of our realizations are correctly identified by time $k = 17$. As the separation between hypothesized ejection fractions decreases to 0.05, the model identifier still correctly classifies more than 65% of the realizations by time $k = 17$. These results illustrate




Figure 5-9: Percent correctly identified for several noises and model separation = 0.05

the trade off between the separation of hypothesized ejection fractions and noise. At low levels of noise, the model identifier performs well even if the separation between hypothesized ejection fractions is small. At high levels of noise, the model identifier performs better as the separation between hypothesized ejection fractions increases.

In addition, we repeated this experiment with the true ejection fraction equal to EF_2 , the smaller of the two hypothesized ejection fractions. The results were similar to the case described above where the true ejection fraction was equal to EF_1 .

5.5 Experiment #2 : Contraction Mismatch

We will almost certainly be faced with identifying the ellipsoid dynamics which differ from the two hypothesized models. In this experiment, we evaluate the performance of model identification when neither of the two hypothesized dynamic models has a contraction rate that equals the contraction rate of the true ellipsoid dynamics. To do this, we establish two hypothesized dynamic models with ejection fractions, EF_1 and




Figure 5-10: Percent correctly identified as a function of separation of the hypothesized ejection fractions for each of the noises listed in Table 5.1

EF_2 . Then, we vary the true ellipsoid dynamics such that the true ejection fraction EF_a varies between EF_1 and EF_2 . In addition, we will investigate the effect of noise on the performance of the model identification scheme.

The hypothesized models are fixed for this experiment. The first hypothesized model has a contraction rate of $t_1 = 0.936$ which yields an ejection fraction of $EF_1 = 0.65$. This hypothesized model corresponds to an above normal ejection fraction. The second hypothesized model has a contraction rate of $t_1 = 0.986$ which yields an ejection fraction of $EF_2 = 0.20$ which corresponds to a below normal ejection fraction. These values for ejection fraction are used so that on the basis of which hypothesized model is chosen by the model identifier, we may classify the true ejection fraction as above or below normal.

In addition, we generate the true dynamically evolving ellipsoids in the following way. We generate fourteen different dynamically evolving ellipsoid sets with contraction rates t_a^i chosen so that the true ejection fraction EF_a^i vary between EF_1 and EF_2 .

Table 5.2 lists the contraction rates used and their corresponding ejection fractions. Measurements of each of these fourteen dynamically evolving ellipsoid sets are processed by the model identifier as described in Section 5.2 which determines which of the hypothesized dynamic models is the better approximation to the true ellipsoid dynamics.

t_a^i	EF_a^i
0.9838	.23
0.9814	.26
0.9788	.29
0.9762	.32
0.9734	.35
0.9706	.38
0.9676	.41
0.9644	.44
0.9611	.47
0.9576	.50
0.9539	.53
0.9500	.56
0.9458	.59
0.9413	.62

Table 5.2: Values for true contraction rates and corresponding EFs

Once again, we start with the low noise case. Let the variance of the measurement noise be $R = r_1 I$ where r_1 is defined as in (5.15). This simulation was repeated 100 times and the results are presented in Figure 5-11. The plot shows the percent of realizations that picked model #1 at time $k = 17$ for each value of true ejection fraction. Of course for this two model case, the percent of realizations that choose model #2 is just one hundred minus the percent that choose model #1. As we had hoped, if the true ejection fraction is high (i.e. $EF_a > 0.5$), the model judged to be most accurate is model #1. If the true ejection fraction is low (i.e. $EF_a < 0.4$), model #2 is judged to be the more accurate estimate of the true dynamics. Note that there is a transition region between 0.4 and 0.5. If the true ejection fraction lies in this transition region it is difficult to determine what the decision of the model identifier

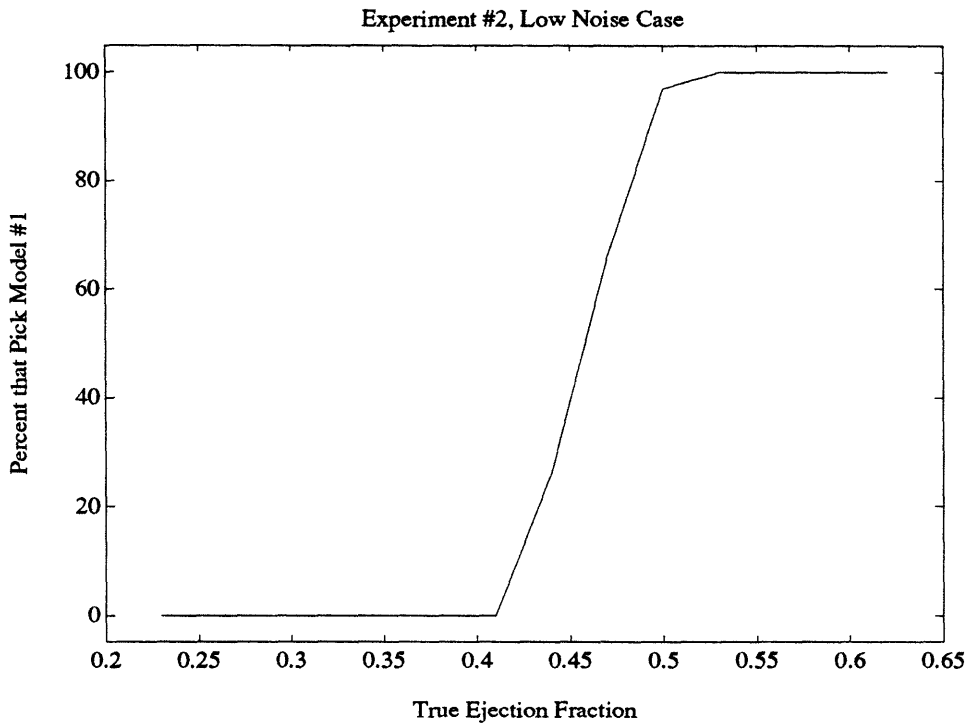


Figure 5-11: True ejection fraction vs. percent of realizations that choose model #1 at $k = 17$ for low noise

will be.

Thus, it is possible to determine a range in which the true ejection fraction lies based on the outcome of the model identifier. Suppose the true ejection fraction of the ellipsoid is unknown and the model identification chooses model #1 as the more likely of the two models. We would guess that the true ejection fraction lies between 0.5 and 1. Thus, it is possible to determine if the ejection fraction lies in a normal or below normal range based on the outcome of the model identifier.

In addition, we may consider the case where the strength of the measurement noise is increased. In Figure 5-12, we show the ejection fraction of the true ellipsoid dynamics versus the percent of realizations that picked model #1 at time $k = 17$ for each of the measurement noise strengths in Table 5.1. Each curve corresponds to a different value of measurement noise strength. As the measurement noise increases, the transition region where the outcome of the model identification scheme is uncertain becomes wider. Thus, even at high measurement noise, the model identifier




Figure 5-12: True ejection fraction vs. percent of realizations that choose model #1 at $k = 17$ for several noises. The curves with the wider transition region correspond to higher noise.

chooses the hypothesized model with the ejection fraction that more closely matches the true ejection fraction in a majority of the realizations that we considered.

5.6 Experiment #3 : Contraction and Rotation Mismatch

It is reasonable to expect that we will need to identify the dynamics of an ellipsoid which rotates as well as contracts/expands. The objective of this final experiment is to evaluate the performance of the model identification scheme when the true ellipsoid dynamics include contraction/expansion and rotation and the hypothesized models include only contraction/expansion at a rate different from the true contraction/expansion. To study this case, we use the same procedure as in Experiment #2 where the ejection fractions of the hypothesized models are fixed and the true contraction rate varies such that the true ejection fraction varies between the two hypothesized ejection fractions. Again, we study the effect of noise on model identification when the mismatch in the models and truth includes contraction and rotation.

Once again, the hypothesized models are fixed. The first hypothesized model has a contraction rate of $t_1 = 0.936$ which yields an ejection fraction of $EF_1 = 0.65$. The second hypothesized model has a contraction rate of $t_1 = 0.986$ which yields an ejection fraction of $EF_2 = 0.20$.

The matrix that represents the true ellipsoid dynamics differs from that given in (5.13). The true ellipsoid dynamics are now given by $\tilde{A}(k) = \Gamma(A(k))$ where

$$A(k) = \begin{cases} t_a I \begin{pmatrix} \cos \phi & \sin \phi \\ -\sin \phi & \cos \phi \end{pmatrix} & \text{for } k = 1 \dots 8 \\ \frac{1}{t_a} I \begin{pmatrix} \cos \phi & -\sin \phi \\ \sin \phi & \cos \phi \end{pmatrix} & \text{for } k = 9 \dots 16 \end{cases} \quad (5.17)$$

where t_a varies as shown in Table 5.2 and $\phi = \pi/54$, which implies a total rotation of approximately $\pi/6$.

We consider the same four levels of measurement noise strength as in Section 5.5. Let the variance of the measurement noise be $R = r_i I$ for $i = 1 \dots 4$ where r_i is defined as in Table 5.1.

Again, we perform 100 realizations of this simulation and present the results in Figure 5-13 which shows the percent of realizations that picked model #1 at time $k = 17$ for each value of ejection fraction of the true ellipsoid dynamics. The results are similar to those presented in Figure 5-12. Each curve corresponds to a different value for measurement noise. For low measurement noise, the transition region, where the decision of the model identifier is uncertain, is narrow. As the measurement noise is increased, the transition region becomes wider.

These results indicate that, even for high levels of measurement noise and when the true dynamics include rotation, the model identifier will choose the hypothesized model whose ejection fraction is closest to the true ejection fraction. In addition, these results are encouraging in that they do not significantly differ from the results of Experiment #2 where the true dynamics did not include rotation.

5.7 Model Identification of 3D Dynamics

We repeat these model identification experiments to determine which of two hypothesized dynamic models best approximates the true dynamics of a three-dimensional ellipsoid. For the three-dimensional dynamic model identifier, the input is a set of noisy two-dimensional projections onto orthogonal planes. The matrix that captures the geometry of the projections is given by (4.10).

We begin by investigating the performance of the three-dimensional model identifier when one of the hypothesized models exactly matches the true dynamics of the ellipsoid. Figure 5-14 shows the number realizations out of 100 that correctly identify model #1 by time $k = 17$ as a function of separation between the hypothesized ejection fractions. Each curve corresponds to different level of measurement noise, r_i , where the measurement noise strengths are given in Table 5.1. The results indicate that we are able to accurately identify the dynamics of the three-dimensional ellip-




Figure 5-13: True ejection fraction vs. percent of realizations that choose model #1 at $k = 17$ for several noises. The curves with the wider transition region correspond to higher noise.




Figure 5-14: Percent of realizations that choose model #1 as a function of the separation of the hypothesized ejection fractions for several values of noise

soid for 100% of the realizations if the measurement noise is low and for 75% of the realizations if the measurement noise is high.

Next, we investigate the performance of the three-dimensional model identifier when neither of the two hypothesized dynamic models has a contraction rate that equals the the contraction rate of the true ellipsoid dynamics. Just as we did in Section 5.5, we fix the two hypothesized ejection fractions at EF_1 and EF_2 and vary the true ejection fraction EF_a between EF_1 and EF_2 . In Figure 5-15, we show the percent of realizations that picked model #1 at time $k = 17$ as the measurement noise increases. The true ejection fraction of the ellipsoid again varies as shown in Table 5.2. Again, these results indicate that even for high levels of measurement noise, the model identifier chooses the hypothesized model with the ejection fraction that more closely matches the true ejection fraction in a majority of the realizations that we considered.

Finally, we evaluate the performance of the model identification scheme for a three-

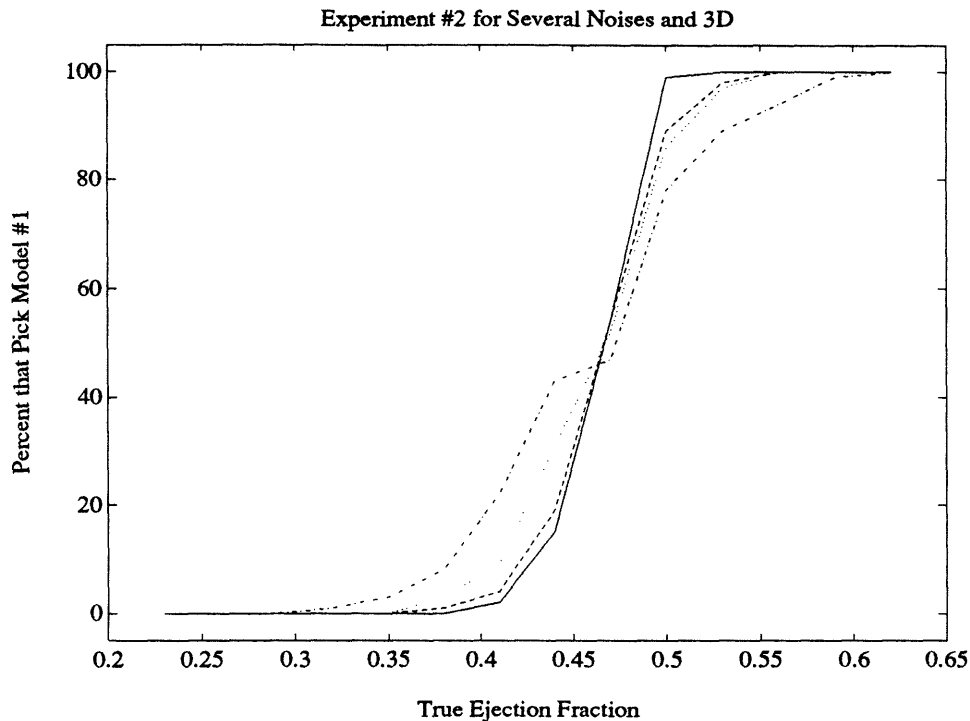


Figure 5-15: Percent of realizations that choose model #1 as a function of true ejection fraction value for several levels of measurement noise

dimensional ellipsoid when the true ellipsoid dynamics include contraction/expansion and rotation and the hypothesized models include only contraction/expansion at a rate different from the true contraction/expansion rate. The true ellipsoid dynamics include contraction/expansion at the rates listed in Table 5.2 and a rotation about the x -axis of $\pi/54$ at each time step for a total rotation of approximately $\pi/6$. Again, we perform 100 realizations of this simulation and present the results in Figure 5-16 which shows the percent of realizations that picked model #1 at time $k = 17$ for each value of true ejection fraction and for several measurement noises. These results illustrate that the model identification scheme accurately chooses the hypothesized model with the ejection fraction that more closely matches the true ejection fraction in a majority of realizations for three-dimensional ellipsoids. These results are encouraging because they show that the effect of rotation on the model identifier is minimal.




Figure 5-16: Percent of realizations that choose model #1 as a function of true ejection fraction value for several levels of measurement noise

5.8 Conclusions

In this chapter, we have evaluated the performance of model identification. We have studied three cases : the case where one of the hypothesized models equals the true dynamics; the case where the neither hypothesized dynamic model matches the true contraction rate of the ellipsoid dynamics; and the case where the neither hypothesized model matches the true contraction or rotation rate of the ellipsoid dynamics. We have studied the performance of the model identifier in each of these scenarios as the measurement noise is increased. As our simulated results indicate, the model identifier performs well for these scenarios.

Chapter 6

Real Data

6.1 Objective

Thus far, our investigation has focused on simulated data. In Chapter 4, we focused on reconstructing a computer simulated, dynamically evolving ellipsoid based on the Rauch-Tung Striebel smoothing algorithm. We reconstructed a two-dimensional ellipsoid from noisy observations of the ellipse itself and a three-dimensional ellipsoid from noisy two-dimensional projections. In Chapter 5, we investigated the performance of a model identification scheme to determine which of two hypothesized dynamic models best matched the dynamics of a computer generated, dynamically evolving ellipsoid. Again, we applied model identification to both two- and three-dimensional ellipsoids. Of course, the true test of the work done in this thesis is to apply similar approaches to myocardial perfusion data actually taken from patients with potential heart disease.

In this chapter, we test the model identification scheme and the smoothing-filter-based ellipsoid reconstruction on ellipses extracted from real myocardial perfusion data and estimate ejection fraction of the left ventricle. Myocardial perfusion data is typically taken as a set of three views for each patient. Ellipses extracted from these three views maybe thought of as noisy two-dimensional projections of the three-dimensional ellipsoid that approximates the left ventricle. Just as we have done in previous chapters, we may process these extracted ellipses with a model identifier or a smoothing filter.

In this chapter, we process the real data in two ways. First, we may process the three views together as a three-dimensional model identification or reconstruction problem. That is, we may identify the dynamics of and reconstruct the three-dimensional ellipsoid that approximates the left ventricle based on the set of noisy two-dimensional projections in the three views. The outcome of the model identifier is used to determine a range of the probable values of the true ejection fraction and provide an assumed model for the smoothing-filter-based reconstruction. Based on the volume of three-dimensional ellipsoids which are the output of the smoothing filter, we calculate an estimate of the ejection fraction of the left ventricle. We investigate the correlation between the ejection fraction calculated this way and the “gold standard” MUGA based ejection fraction values for each patient.

Alternatively, we may calculate approximations to the three-dimensional ejection fraction by processing each view individually as a two-dimensional problem. That is, we may identify the dynamics of the two-dimensional ellipsoidal projections and use a smoothing filter to reconstruct the two-dimensional ellipsoidal projections. For each view, the outcome of the model identifier is used to determine a range of the probable values of the true two-dimensional ejection fraction and provide an assumed model for the smoothing-filter-based reconstruction. Based on the reconstructed two-dimensional projections which are the output of the smoothing filter, we may calculate a two-dimensional ejection fraction for each view. This type of processing of individual projection data gives a measure of the contraction of the three-dimensional ellipsoid in the plane of the projection. We investigate approximations of the three-dimensional ejection fraction based on the two-dimensional ejection fractions for each view. Again, we investigate the correlation between these approximations to the three-dimensional ejection fraction and the MUGA based ejection fraction values for each patient.

This chapter is organized as follows. We begin with a description of the real myocardial perfusion data that is available to us. Next, we employ model identification and smoothing filter based reconstruction to estimate the two-dimensional ejection fraction of each view for each patient. Finally, we use model identification and smoothing filter based reconstruction to combine the data of the three views and

directly estimate the three-dimensional ejection fraction.

6.2 Data Description

We are fortunate to have access to the raw myocardial perfusion images and the “gold standard” MUGA based ejection fraction estimates for 28 patients. The data set for each individual consists of three views (anterior, lateral, left anterior oblique) of sixteen frames each. Anterior (ANT) refers to a frontal view; lateral (LAT) refers to a side view from under the left arm; left anterior oblique (LAO) refers to a frontal view skewed down and to the left. In general, the ANT view shows the smallest cross-sectional area. The LAT view shows the largest cross-sectional area. The LAO view shows a cross-sectional area that lies between the other two views. The set of sixteen frames for each view is produced using ECG gating as described in Chapter 2. Thus, each frame is not a snapshot of the heart in motion, but rather the sum of images corresponding to a particular cardiac phase over several cardiac cycles. These raw images are several years old and were taken using the radionuclide Thallium. As we mentioned in Chapter 2, recent developments make it feasible to use Technetium as the imaging agent for myocardial perfusion. The work described in this chapter is applicable to myocardial perfusion based on either of these imaging agents.

As we have previously mentioned, both the model identification scheme and the smoothing filter based reconstruction expect as input ellipses that have been extracted from the raw myocardial perfusion images. For consistency, we employ a semi-automated method to extract ellipses from each frame. The Matlab program used to extract these ellipses is included as Appendix 6-A. This method has proven to be robust for the following reasons. First, there is minimal human intervention. Second, the ellipse extraction is not sensitive to occlusions which appear as dark regions in the myocardium. This point is important because most patients undergoing this procedure have some history of cardiovascular disease.

For this chapter, we concentrate on testing our model identification and smoothing filter based ejection fraction estimation techniques on six individuals. In Figure 6-1,

Patient	#1	#2	#3	#4	#5	#6
MUGA EF	0.26	0.56	0.20	0.25	0.28	0.40

Table 6.1: MUGA Ejection Fractions for Six Individuals

we have plotted the areas of the 16 extracted ellipses from the ANT view for each of our six patients. Similar plots may be generated for the LAT and LAO views. In comparing the six patients, we conclude that the maximum of the cycle occurs at the first and sixteenth time step for each individual, but the minimum point varies from person to person. In addition, we know the “gold standard” MUGA based ejection fraction numbers for each of the six individuals. These values are given in Table 6.1. These ejection fractions range from 0.20 to 0.56. Thus, our sample set consists of individuals whose ejection fractions range from below normal to normal.

6.3 2D Processing

In this section, we use model identification and smoothing filter-based reconstruction to process each view individually as a two-dimensional problem. That is, we apply these techniques to extracted ellipses from one view at a time. First, we use the model identification scheme to determine which of two hypothesized models best approximates the dynamics of the two-dimensional ellipsoid projection. The outcome of the model identifier is of interest, in and of itself, because it indicates whether the patient’s ejection fraction is in a normal or below normal range. Second, we employ the Rauch-Tung Striebel smoothing filter to reconstruct the two-dimensional ellipsoidal projections. We investigate a smoothing filter based on the model chosen by the model identifier as well as a smoothing filter based on a random walk model of dynamics. We also calculate the two-dimensional ejection fraction based on these reconstructed ellipses.

In addition, we investigate approximations to the three-dimensional ejection fraction based on the two-dimensional ejection fractions of the three views. One rough

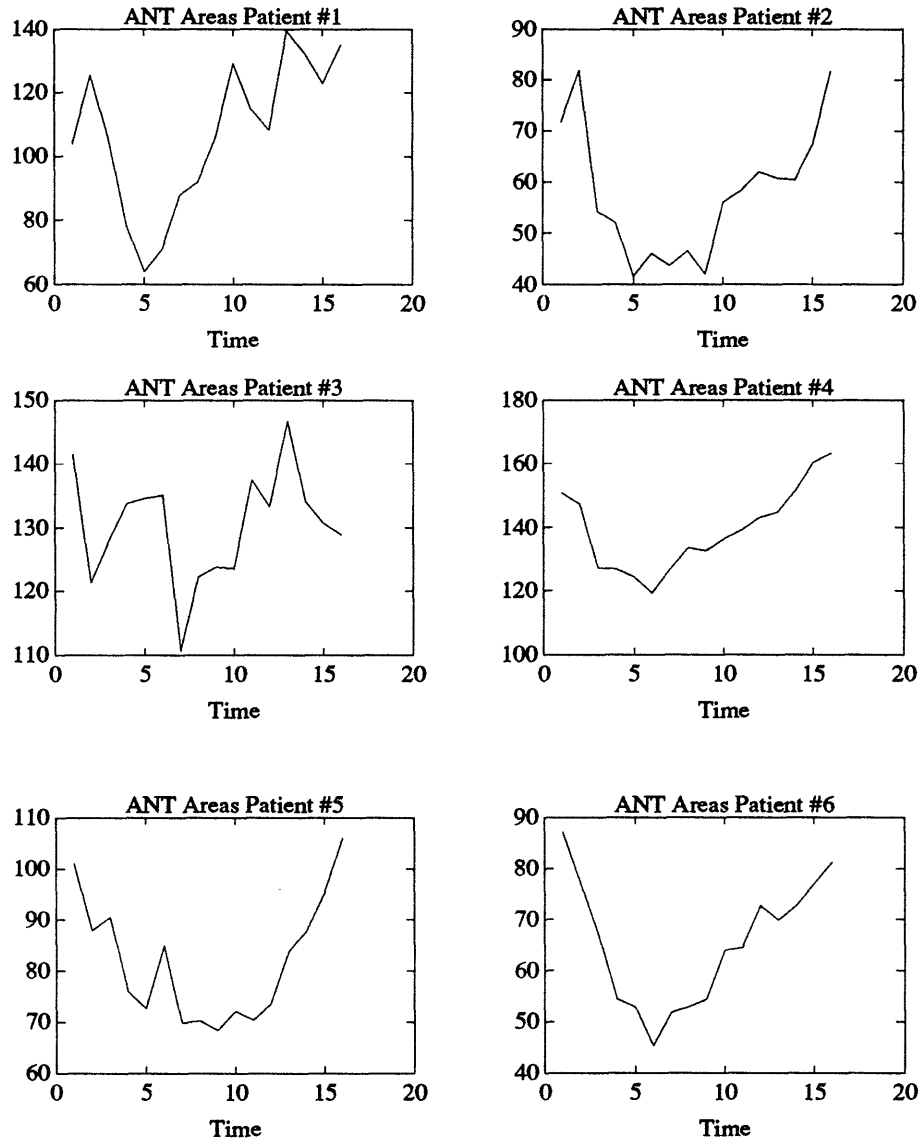


Figure 6-1: Anterior areas for all patients

approximation to the three-dimensional ejection fraction is the two-dimensional ejection for each view. A second, more sophisticated, approximation to the the three-dimensional ejection fraction is a function that combines the two-dimensional ejection fraction for each view. In this section, we investigate the correlation between these approximations to the three-dimensional ejection fraction and the MUGA based ejection fraction values for each patient.

6.3.1 2D Model Identification

We use the model identifier to process the data for each view of each individual to determine which of the hypothesized models we should use as an assumed model in the smoothing filter. Specifically, the model identification scheme is implemented in the following way. For each of the three views, the input to the model identifier consists of the set of 16 ellipses extracted from raw data. Because we are reconstructing the ellipses based on noisy observations of the ellipses themselves, the matrix that captures the geometry of the measurements, \tilde{C} , is the identity. The two hypothesized models have ejection fractions $EF_1 = 0.6$ and $EF_2 = 0.2$ respectively. These two hypothesized models correspond to above and below normal values for ejection fraction. As we mentioned in Section 6.2, the point corresponding to the minimum in the cycle of the extracted ellipses varies from person to person. For this reason, the hypothesized models used by the model identifier have a variable minimum point, T_m , that is chosen to correspond to the apparent cycle minimum for each individual. Thus, the matrices that represent the hypothesized dynamics are given by $\tilde{M}_i(k) = \Gamma(M_i(k))$ where

$$M_i(k) = \begin{cases} c_i I & \text{for } k = 1 \dots T_m \\ e_i I & \text{for } k = T_m + 1 \dots 15 \end{cases} \quad (6.1)$$

where $c_i = (1 - EF_i)^{\frac{1}{(T_m-1)^2}}$, $e_i = (1 - EF_i)^{\frac{-1}{(16-T_m)^2}}$, and T_m is the apparent cycle minimum specified by the user. It is also possible to implement an automated method of determining the cycle minimum, although we did not do so here.

In addition, we set the value for the variance of the process noise, q , to zero for

both Kalman filters, indicating a high degree of confidence in the models. This value of q accentuates the difference between the two models. The last parameter to be set in the model identifier is the variance of the measurement noise, r . In dealing with real data, the true value of the variance of the measurement noise is unknown. We use $20 < r < 120$ because for this range the decision of the model identifier is accurate. A smaller value of r indicates that the measurements are less noisy. Specifically, for the six cases considered in this chapter, a value of $r < 20$ indicates that the variance of the measurement noise is smaller than the true variance of the measurement noise in the observations. For this reason, the outcome of the model identifier becomes less reliable for values of $r < 20$.

The decisions of the model identifier for each patient and each view when $r = 100$ are shown in Table 6.2. We get the same results as in Table 6.2 for $20 < r < 100$ for Patients #1-5. For Patients #1-5, it is clear which dynamic model to use as the assumed model in the smoothing filter because the outcome of the model identifier is the same for each view.

Patient #6 is a special case because the MUGA ejection fraction value for this patient is equidistant from the two hypothesized ejection fractions. That is, the MUGA ejection fraction lies in the transition region we discussed in Chapter 5. This may be one reason why the decision of the model identifier differs for the three views. The outcome of the model identifier for this patient also varies as we change the value of r . Because it is not clear which hypothesized model we should use as an assumed model in the smoothing filter, we use a second model identifier where the two hypothesized ejection fractions are $EF_1 = 0.50$ and $EF_2 = 0.20$. This model identifier is also implemented with $r = 100$ and $q = 0$. The second model identifier chooses model #1 for ANT and LAO views, but still chooses model #2 for LAT view for Patient #6. Since the decision of the model identifier is still inconsistent between the three views, we implement a third model identifier where the two hypothesized ejection fractions are $EF_1 = 0.40$ and $EF_2 = 0.20$. For all three views, this model identifier chooses model #1. Thus for Patient #6, we use an assumed model that has an ejection fraction of 0.40.

Patient	#1	#2	#3	#4	#5	#6
MUGA	0.26	0.56	0.20	0.25	0.28	0.40
ANT	Model #2	Model #1	Model #2	Model #2	Model #2	Model #1
LAT	Model #2	Model #1	Model #2	Model #2	Model #2	Model #2
LAO	Model #2	Model #1	Model #2	Model #2	Model #2	Model #1

Table 6.2: Results of 2D Model Identifications for $r = 100$

Note that this result indicates another potential use for the model identification scheme. That is, we may be able to implement several successive model identifiers to directly estimate ejection fraction. We investigate this type of scheme further with three-dimensional processing.

6.3.2 Smoothing Filter Reconstructions

In this subsection, we use the Rauch-Tung Striebel smoothing filter to reconstruct the two-dimensional ellipsoidal projections of the three-dimensional ellipsoid that approximates the left ventricle. That is, for each view, we reconstruct a two-dimensional ellipsoid based on noisy observations of the ellipses themselves. We implement one smoothing filter based on an assumed model that was chosen by a method such as the model identification scheme. In addition, we implement a second smoothing filter based on the random walk model of dynamics.

Implementation

We implement two Rauch-Tung Striebel smoothing filters to process the data for each view of each individual. The first smoothing filter is based on an assumed dynamic model. The matrix that represents the assumed dynamic model is as given in (6.1) where the value for EF_i for each patient is given in Table 6.3. Note that the same assumed model is used for each of the three views for any given individual. The second smoothing filter is based on the random walk model of dynamics. The dynamic model used by the smoothing filter is defined by the matrix $\tilde{M}(k) = I$. Thus, the ejection fraction associated with the random walk model of dynamics is

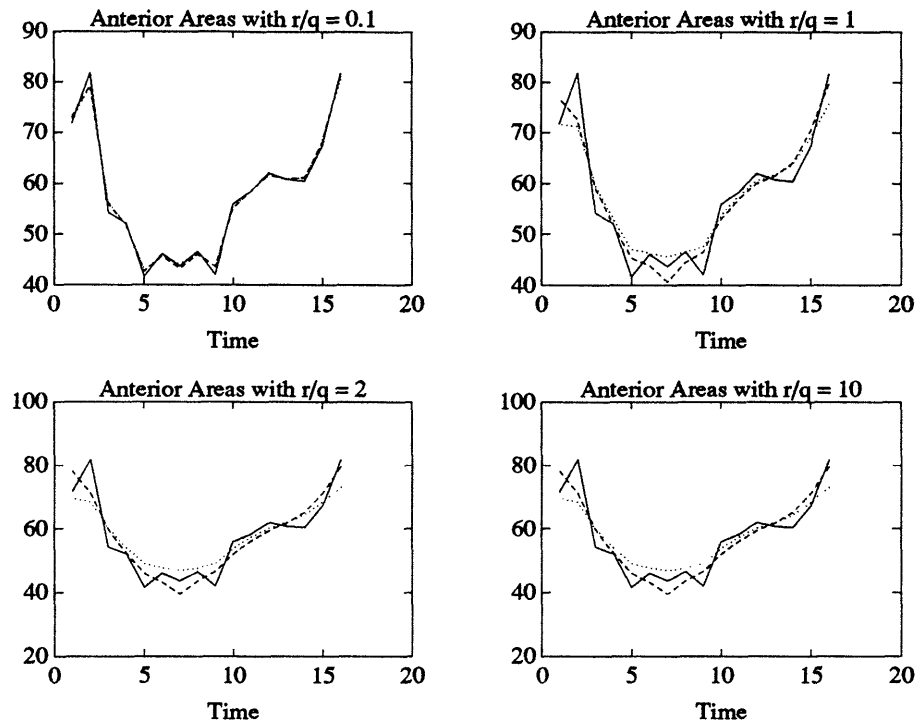


Figure 6-2: Effect of varying r/q on smoothing filter output for ANT view for Patient #2

zero. To implement either of these smoothing filter, we need to set values for the process noise variance, $Q = qI$, and the measurement noise variance, $R = rI$. A low value of q indicates a high degree of confidence in the model and results in the smoothing filter output following the model. A low value of r indicates a high degree of confidence in the measurements and results in the smoothing filter output following the measurements. Thus, we may set the degree of smoothing by adjusting the value of r/q . We studied the output of several smoothing filter implemented with different values for r/q . Figure 6-2 shows the effect of changing this ratio value on the smoothing filter output for the ANT view of Patient #2. This figure shows the areas of the extracted (unsmoothed) ellipses, areas of the assumed model smoothed ellipses, and areas of the random walk model smoothed ellipses. After examining plots such as Figure 6-2, we chose the ratio $r/q = 2$ because for this value the filter provides some smoothing without completely ignoring the measurements.

The output of the smoothing filter are reconstructed ellipses. In Figure 6-3, we

Patient	#1	#2	#3	#4	#5	#6
EF _i	0.2	0.6	0.2	0.2	0.2	0.4

Table 6.3: Values of Assumed Model EFs

show the areas of the unsmoothed ellipses, the areas of the reconstructed ellipses from the assumed model smoothing filter, and the areas of the reconstructed ellipses from the random walk smoothing filter for the ANT view for each of our six patients.

2D Ejection Fraction Calculation

For each of the three views, we now calculate the ejection fraction of the two-dimensional ellipsoids which are projections of the three-dimensional ellipsoid that approximates the left ventricle. For each projection, the two-dimensional ejection fraction gives a measure of the contraction of the three-dimensional ellipsoid in the plane of the projection. The general definition of ejection fraction is as given in (2.1). We note that maximum area of the cycle occurs at time $k = 1$ and $k = 16$ and the minimum area of the cycle occurs at the user specified time $k = T_m$. We take the average of the two maximum points as the maximum value for area. Thus, the formula for ejection fraction of the each two-dimensional projections is given by

$$EF = \frac{\text{area}(1) + \text{area}(16) + 2\text{area}(T_m)}{\text{area}(1) + \text{area}(16)} \quad (6.2)$$

Using this definition, we may calculate the ejection fraction for each individual and each view.

In Table 6.4, we have listed the ejection fractions for each view from static reconstruction, smoothing filter reconstruction based on the assumed model and a smoothing filter based on the random walk model. The ejection fractions of each of the two-dimensional ellipsoidal projections tends to be lower than the MUGA estimates. This is as it should be because the two-dimensional ejection fraction ignores any contraction along the axis perpendicular to the plane of the projection. Note that the static estimates of ejection fraction are based on the unsmoothed areas. Static

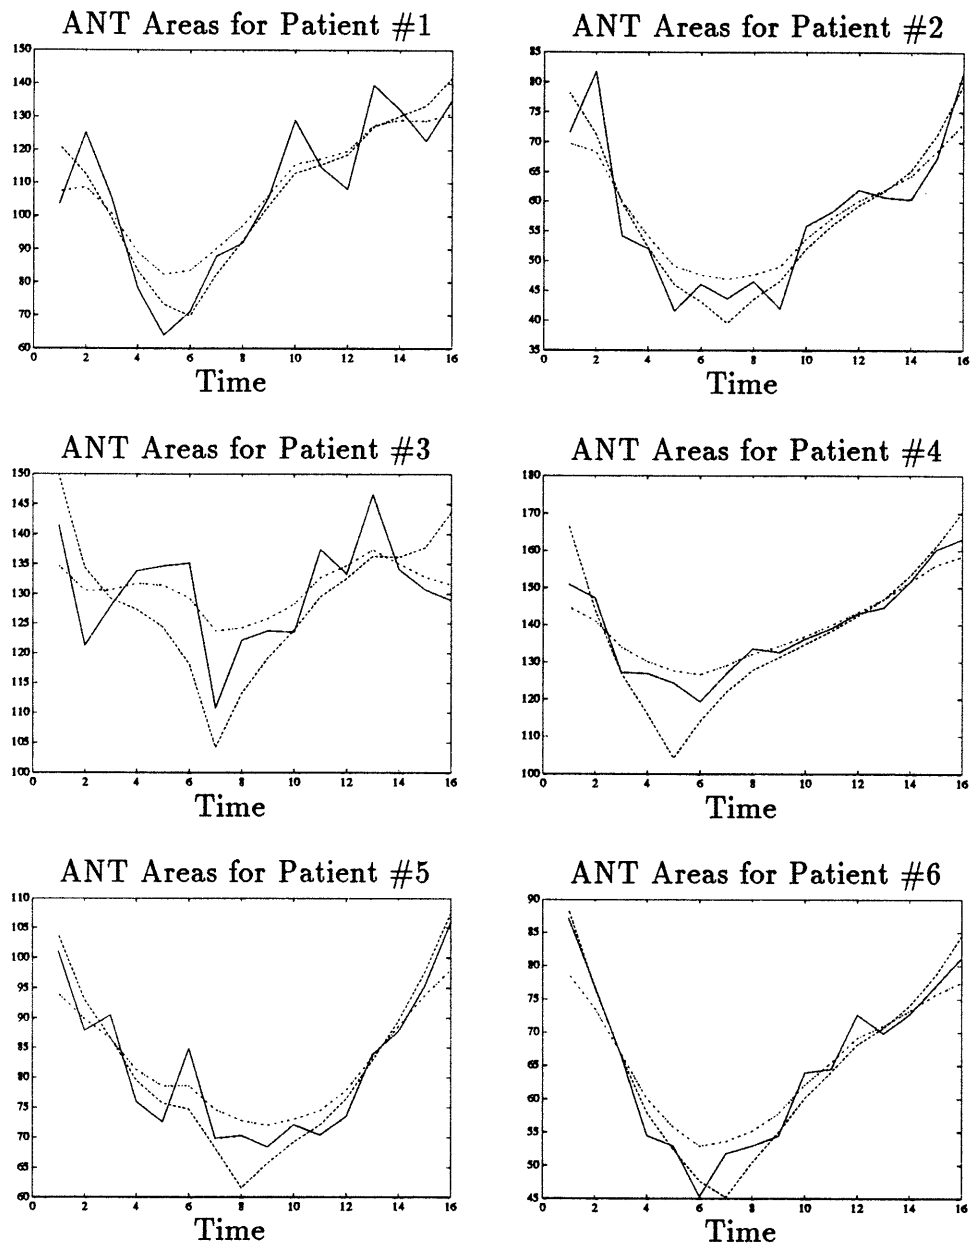


Figure 6-3: Smoothed two dimensional reconstructions for ANT view for all patients

Patient	#1	#2	#3	#4	#5	#6
MUGA	0.26	0.56	0.2	0.25	0.28	0.40
ANT static	0.4045	0.4306	0.1808	0.2079	0.3208	0.3843
LAT static	0.1045	0.4083	0.2353	0.1842	0.1198	0.2375
LAO static	0.1877	0.4603	0.1620	0.1531	0.1961	0.5194
ANT assumed model	0.3462	0.4984	0.1308	0.2222	0.2890	0.4113
LAT assumed model	0.1535	0.4925	0.1832	0.1727	0.1498	0.3247
LAO assumed model	0.2006	0.5064	0.1647	0.1496	0.2052	0.4756
ANT rw model	0.2987	0.3428	0.0703	0.1586	0.2417	0.3141
LAT rw model	0.0904	0.3322	0.1275	0.1041	0.0938	0.2127
LAO rw model	0.1421	0.3530	0.1078	0.0810	0.1526	0.3887

Table 6.4: Ejection Fractions for 2D reconstruction

estimates of ejection tend to show a large variation from view to view because these estimates are the most likely to be effected by noise. For each view, the ejection fraction estimate calculated from reconstructed ellipses using the assumed model of dynamics is consistently higher than the ejection fraction estimate calculated from reconstructed ellipses using the random walk to dynamics. This result is also expected. Recall that the output of the smoothing filter is the best estimate given the dynamic model and all the data. Thus, the reconstructed ellipses which are the output of a smoothing filter based on the assumed model tend to have an ejection fraction that is skewed towards the ejection fraction associated with the assumed model. Similarly, the reconstructed ellipses which are the output of a smoothing filter based on the random walk model tend to have an ejection fraction that is skewed towards the ejection fraction associated with the random walk model (i.e. zero).

We have touched upon another inherent flaw in smoothing filter reconstruction based on the assumed model. That is, the reconstructed ellipses which are the output of a smoothing filter based on the assumed model tend to have an ejection fraction that is skewed towards the ejection fraction associated with the assumed model. To illustrate this point, we process the ellipses extracted for Patient #6 with smoothing filters based on several different assumed models. Table 6.5 shows the ejection fractions associated with each of the assumed models used and the ejection fraction

Assumed EF	0.20	0.30	0.40	0.50	0.60
ANT EF	0.3587	0.3838	0.4113	0.4419	0.4767
LAT EF	0.2642	0.2930	0.3247	0.3599	0.3999
LAO EF	0.4286	0.4510	0.4756	0.5030	0.5343

Table 6.5: Variation of Reconstructed Ellipse EF with Assumed EF for Patient #6

calculated for each view using the different assumed models for Patient #6. Thus, the ejection fraction of the reconstructed ellipses is skewed by the assumed model used in the smoothing filter. We note that for each view the estimated ejection fraction varies by approximately 0.10 while the hypothesized ejection fraction varies by 0.40. Therefore, it is important to determine which assumed model best approximates the true ellipsoid dynamics.

Approximating the 3D Ejection Fraction

Of course, the ultimate objective is to calculate an estimate for ejection fraction of the three-dimensional ellipsoid that approximates the left ventricle. We consider two types of approximations to the ejection fraction of the three-dimensional ellipsoid based on the ejection fractions of the two-dimensional ellipsoidal projections.

One type of approximation to the ejection fraction of the three-dimensional ellipsoid is the ejection fraction of the two-dimensional ellipsoidal projections. Thus, each of the two-dimensional ejection fractions in Table 6.4 are approximations to the three-dimensional ejection fraction. As we showed in Appendix 4-A, the ejection fraction of a three-dimensional ellipsoid is equal to the ejection fraction of its two-dimensional projections if there is no contraction of the three-dimensional ellipsoid along an axis perpendicular to the plane of the projection. If the heart contracts in a manner similar to this idealized case, the ejection fraction of the two-dimensional projections would approximate the ejection fraction of the three-dimensional ellipsoidal model of the left ventricle.

A second, more sophisticated type of approximation to the ejection fraction of the three-dimensional ellipsoid that approximates the left ventricle is given by the

formula

$$EF_{3d} = 1 - \sqrt{(1 - EF_{ANT})(1 - EF_{LAT})(1 - EF_{LAO})} \quad (6.3)$$

This approximation is based on the assumption that the three views are projections onto orthogonal planes. That is, we assume that the three views ANT, LAT, and LAO correspond to projections onto the xy -, yz -, and xz -planes respectively.

The ejection fraction of the two-dimensional projections on the xy -plane may be written as

$$EF_{xy} = 1 - \frac{x_c y_c}{x_e y_e} \quad (6.4)$$

where x_c and y_c correspond to the axis lengths of the fully contracted ellipse and x_e and y_e correspond to the axis lengths of the fully expanded ellipse. We rewrite (6.4) as

$$EF_{xy} = 1 - \alpha_x \alpha_y$$

where α_x and α_y are the ratio of fully contracted axis lengths to fully expanded axis lengths along the x and y axes respectively. Similarly, we express the ejection fraction of the two-dimensional ellipsoidal projections on the xz and yz -planes as

$$EF_{xz} = 1 - \alpha_x \alpha_z \quad (6.5)$$

$$EF_{yz} = 1 - \alpha_y \alpha_z \quad (6.6)$$

The ANT, LAT, and LAO views correspond to projections onto the xy -, yz -, and xz -planes respectively. That is,

$$EF_{ANT} = EF_{xy} \quad (6.7)$$

$$EF_{LAT} = EF_{yz} \quad (6.8)$$

$$EF_{LAO} = EF_{xz} \quad (6.9)$$

We write the ejection fraction of the three-dimensional ellipsoid as

$$EF_{3d} = 1 - \alpha_x \alpha_y \alpha_z \quad (6.10)$$

Patient	#1	#2	#3	#4	#5	#6
MUGA	0.26	0.56	0.2	0.25	0.28	0.40
Combined View (am)	0.3348	0.6455	0.2299	0.2603	0.3068	0.5434
Combined View (rw)	0.2602	0.4671	0.1493	0.1677	0.2369	0.4255

Table 6.6: Combined View Approximations to 3D EF

The formula of (6.3) obviously follows.

Employing(6.3), we may calculate a combined-view approximation to the three-dimensional ellipsoid ejection fraction. These values are given in Table 6.6. These values are approximations to the three-dimensional ejection fraction; therefore, these values are higher than the two-dimensional ejection fractions for each view given in Table 6.4. Again, as expected, the approximation to the three-dimensional ejection fraction based on reconstructed ellipses output by a smoothing filter that uses the assumed model is higher than the corresponding approximation based on reconstructed ellipses output by a smoothing filter that uses the random walk model of dynamics. The set of approximations based on ellipses reconstructed using an assumed model smoothing filter tend to overestimate the true ejection fraction.

Performance Evaluation

Let us discuss a method to measure the accuracy of each of the approximations to the three-dimensional ejection fraction that were described in this chapter. As we have previously mentioned, some of the approximations we have considered underestimate the true ejection fraction and some overestimate the true ejection fraction. We would like to determine which of the approximations described in this chapter give the best overall estimate of the three-dimensional ejection fraction.

One can imagine plotting the MUGA ejection fraction values for each patient versus our calculated approximation of three-dimensional ejection fraction. We can also determine the slope and offset of the line that is best fit to these six points. If the best-fit line has a slope of one and an offset of zero, then our calculated approximation closely matches the MUGA ejection fraction values.

In addition, we may calculate the correlation coefficient of the MUGA values and our calculated approximation. The correlation coefficient gives a measure of the spread of the points about the best-fit line. A correlation coefficient close to one indicates a high degree of correlation. Let the random variables x and y denote the estimated ejection fraction and the MUGA ejection fraction, respectively. Thus, the sample pairs, (x_i, y_i) , of the random variables, x and y , represent the estimated and MUGA ejection fraction values for Patient $\#i$. Recall that the correlation coefficient of two random variables, $\rho(x, y)$, is defined as

$$\rho(x, y) = \frac{\text{cov}(x, y)}{\sigma_x \sigma_y} \quad (6.11)$$

where $\text{cov}(x, y)$ refers to the covariance of x and y and σ_x and σ_y are the standard deviations of x and y , respectively. From the sample pairs (x_i, y_i) , we may calculate statistical approximations to the covariance and standard deviation of the random variables x and y .

$$\text{cov}(x, y) \approx \frac{1}{m-1} \sum_{i=1}^m (x_i - \bar{x})(y_i - \bar{y}) \quad (6.12)$$

$$\sigma_x \approx \sqrt{\frac{\sum_{i=1}^m (x_i - \bar{x})^2}{m-1}} \quad (6.13)$$

$$\text{where } \bar{x} = \sum_{i=1}^m x_i / m \quad (6.14)$$

$$m = \text{sample size} \quad (6.15)$$

For the results presented in this chapter $m = 6$. In Table 6.7, we list the slope and offset of the best-fit line as well as the correlation coefficients for each of the approximations of the three-dimensional ejection fraction. The static estimate of ejection fraction did not correlate well with the MUGA ejection fraction values. The approximations to the three-dimensional ejection fraction based on smoothing filter reconstruction using an assumed model had correlation coefficients above 0.9. Recall that the assumed models used for each individual were well matched to the true dynamics of the ellipsoid. Thus, these high correlation coefficients reflect the additional

	slope	offset	ρ
ANT static	0.9219	0.0286	0.7301
LAT static	0.9644	0.1177	0.7981
LAO static	0.6970	0.1300	0.8627
ANT assumed model	0.9123	0.0365	0.9057
LAT assumed model	0.9302	0.0961	0.9599
LAO assumed model	0.7643	0.1082	0.9328
Combined View (am)	0.7694	0.0274	0.9726
ANT rw model	0.9681	0.0949	0.7631
LAT rw model	1.3127	0.1148	0.9456
LAO rw model	0.8769	0.1459	0.8711
Combined View (rw)	0.9416	0.0572	0.9398

Table 6.7: Correlation Coefficients for 2D

(correct) information incorporated by the smoothing filter based on the assumed model. The approximations to the three-dimensional ejection fraction based on the smoothing filter reconstruction using the random walk model of dynamics had lower correlation coefficients for the ANT and LAO views.

6.4 3D Processing

In this section, we apply model identification and smoothing filter reconstruction to combine the set of three, two-dimensional projections as observations of the three-dimensional ellipsoid that approximates the left ventricle. Based on the reconstructed three-dimensional ellipsoids, we calculate the ejection fraction of the ellipsoid that approximates the left ventricle. While for the simulated model identification and ellipsoid reconstruction of previous chapters the projection geometry was known, for model identification and ellipsoid reconstruction based on real data the exact projection geometry is unknown. In this section, we assume a projection geometry that is an approximation to the true (unknown) projection geometry.

6.4.1 3D Model Identification

We use the model identifier to combine the three projected views and determine which of two hypothesized models best matches the dynamics of the three-dimensional ellipsoid that approximates the left ventricle. The model identification scheme is implemented in the following way. The input to the model identifier is the set of ellipses extracted from each of the three views. The ANT and LAT views are assumed to be projections onto orthogonal planes. The LAO view is assumed to be projections onto a plane that is tilted by 45° from the plane that is orthogonal to both the ANT and LAT planes. Thus, the corresponding matrix that captures the geometry is given by :

$$\tilde{C} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0.5 & 0 & -0.7071 & 0 & 0 & 0.5 \\ 0.5 & 0.5 & 0 & 0 & -0.5 & -0.5 \\ 0.25 & 0.5 & 0.3536 & 0.5 & 0.5 & 0.25 \end{pmatrix} \quad (6.16)$$

The two hypothesized models have ejection fractions $EF_1 = 0.6$ and $EF_2 = 0.2$ respectively. Once again the user must provide a point that corresponds to the minimum of the cycle. The matrices that define the dynamics are again given by (6.1). The strength of the process noise, q , is set to be zero. We test this model identification scheme for a wide range of values for the variance of the measurement noise. The decision of the model identifier for $20 < r < 200$ are shown in Table 6.8. For this three-dimensional model identification the model identifier chose the correct model for Patients #1-5. Again, the decision for Patient #6 is ambiguous because the MUGA ejection fraction for this patient lies in the transition region between the two hypothesized ejection fractions.

Patient	#1	#2	#3	#4	#5	#6
MUGA	0.26	0.56	0.20	0.25	0.28	0.40
Mod ID	Model #2	Model #1	Model #2	Model #2	Model #2	Model #1

Table 6.8: Results of 3D Model Identifications for $20 < r < 2000$

Model Identification Used to Determine EF

As we have mentioned, it may also be possible to determine the ejection fraction using successive model identifiers. We use ellipses extracted from the raw data for Patient #6 to illustrate this point. The model identifiers used were implemented with $q = 0$, $r = 100$ and \tilde{C} as given in (6.16). We begin by implementing a model identifier with hypothesized ejection fractions $EF_1 = 0.6$ and $EF_2 = 0.2$. For Patient #6, the model identifier chooses model #1. Drawing on the simulated model identification examples of Chapter 5, we would conclude that the actual ejection fraction $EF_a > 0.4$. Next, we implement a model identifier with hypothesized ejection fractions $EF_1 = 0.6$ and $EF_2 = 0.3$. Again, the decision of the model identifier is model #1. This indicates that the actual ejection fraction $EF_a > 0.45$. Now, we implement a model identifier with hypothesized ejection fractions $EF_1 = 0.6$ and $EF_2 = 0.4$. This model identifier also chooses model #1 indicating $EF_a > 0.5$. Finally, we implement a model identifier with hypothesized ejection fractions $EF_1 = 0.6$ and $EF_2 = 0.5$. This final model identifier chooses model #1 indicating $EF_a > 0.55$. We stop at this point where the separation between the hypothesized ejection fractions is 0.10 because as our simulations have shown the model identifier is not reliable as the separation between the hypothesized ejection fractions becomes smaller. This value of ejection fraction does not agree with the MUGA ejection fraction value for Patient #6. Since we have tested this approach on only one patient, it is difficult to determine the reason for this discrepancy. In future work, we will use a technique similar to that illustrated by this exercise to determine the ejection fraction for more patients.

6.4.2 Smoothing Filter Reconstructions

In this subsection, we use the Rauch-Tung Striebel smoothing filter to reconstruct the three-dimensional ellipsoid that approximates the left ventricle. That is, we combine the data of all three views to reconstruct one three-dimensional ellipsoid. Once again, we implement one smoothing filter based on an assumed model and another based on a random walk model of dynamics.

Implementation

The Rauch-Tung Striebel smoothing filter based on the assumed model is implemented in the following way. Similar to approach in the two-dimensional version, we define two assumed dynamic models $\tilde{M}_1(k)$ and $\tilde{M}_2(k)$ with ejection fractions $EF_1 = 0.20$ and $EF_2 = 0.60$. The matrices $\tilde{M}_1(k)$ and $\tilde{M}_2(k)$ are again defined as in (6.1) where T_m is the user entered cycle minimum. As for the two-dimensional reconstruction, the smoothing filter reconstruction for Patients #1,3,4, and 5 are based on $\tilde{M}_2(k)$. Similarly, the smoothing filter reconstruction for Patients #2 and 6 are based on $\tilde{M}_1(k)$. The Rauch-Tung Striebel smoothing filter based on the random walk model of dynamics is implemented in the a similar manner. The matrices that represent the random walk model of dynamics are given by $\tilde{M}(k) = I$.

In addition, both smoothing filters are implemented with the following parameters. The ratio of measurement noise variance r to process noise variance q is again used to determine the amount of smoothing. After using several different values for r/q , we chose $r/q = 10$ because this value yields a filter that provides some smoothing without completely ignoring the measurements. The matrix that captures the geometry of the projections is given by (6.16).

Note that it is possible to perform a static reconstruction (i.e. estimate at time k is based only on measurements at the same time) of the ellipsoid using the formula

$$\hat{\epsilon} = (\tilde{C}^T \tilde{C}) \tilde{C}^T y(k) \quad (6.17)$$

where $y(k)$ is the vector representing the three measurements or projections at time

k .

The output of the smoothing filters are reconstructed three-dimensional ellipsoids. In Figure 6-4, we show the volumes of the unsmoothed ellipsoids, the volumes of the reconstructed ellipsoids from the assumed model smoothing filter, and the volumes of the reconstructed ellipsoids from on the random walk smoothing filter for each of our six patients. We also suggest that the plots shown in Figure 6-4 that show dynamic estimates of volume versus time may have potential prognostic value, in and of themselves. We note that for patients with low ejection fraction and what appear to be occlusions these plots show unusual variations for estimates of volume over time.

Ejection Fraction

Again the maximum volume for the reconstructed three-dimensional ellipsoid occurs at time $k = 1$ and $k = 16$ and the minimum volume of the cycle occurs at the user specified time $k = T_m$. The definition for ejection fraction is given by

$$EF = \frac{\text{vol}(1) + \text{vol}(16) + 2\text{area}(T_m)}{\text{vol}(1) + \text{vol}(16)} \quad (6.18)$$

This is the ejection fraction for the three-dimensional ellipsoid that approximates the left ventricle.

Table 6.9 lists the ejection fractions for the static reconstruction case, the assumed-model-based smoothing filter, and the random-walk-model-based smoothing filter. The ejection fraction calculated from reconstructed ellipsoids using an assumed model are again skewed by the assumed ejection fraction. For example, the assumed ejection fraction for Patient #6 is 0.50. As a result, the ejection fraction estimate based on a smoothing filter that uses the assumed model is higher than the MUGA or random walk based ejection fraction. As we expected, ejection fractions calculated using ellipsoids reconstructed using the random walk model tend to be lower than those calculated reconstructed using the assumed model.

As we have noted, there exists an inherent flaw in the assumed model based reconstruction; the reconstructed ellipses which are the output of a smoothing filter

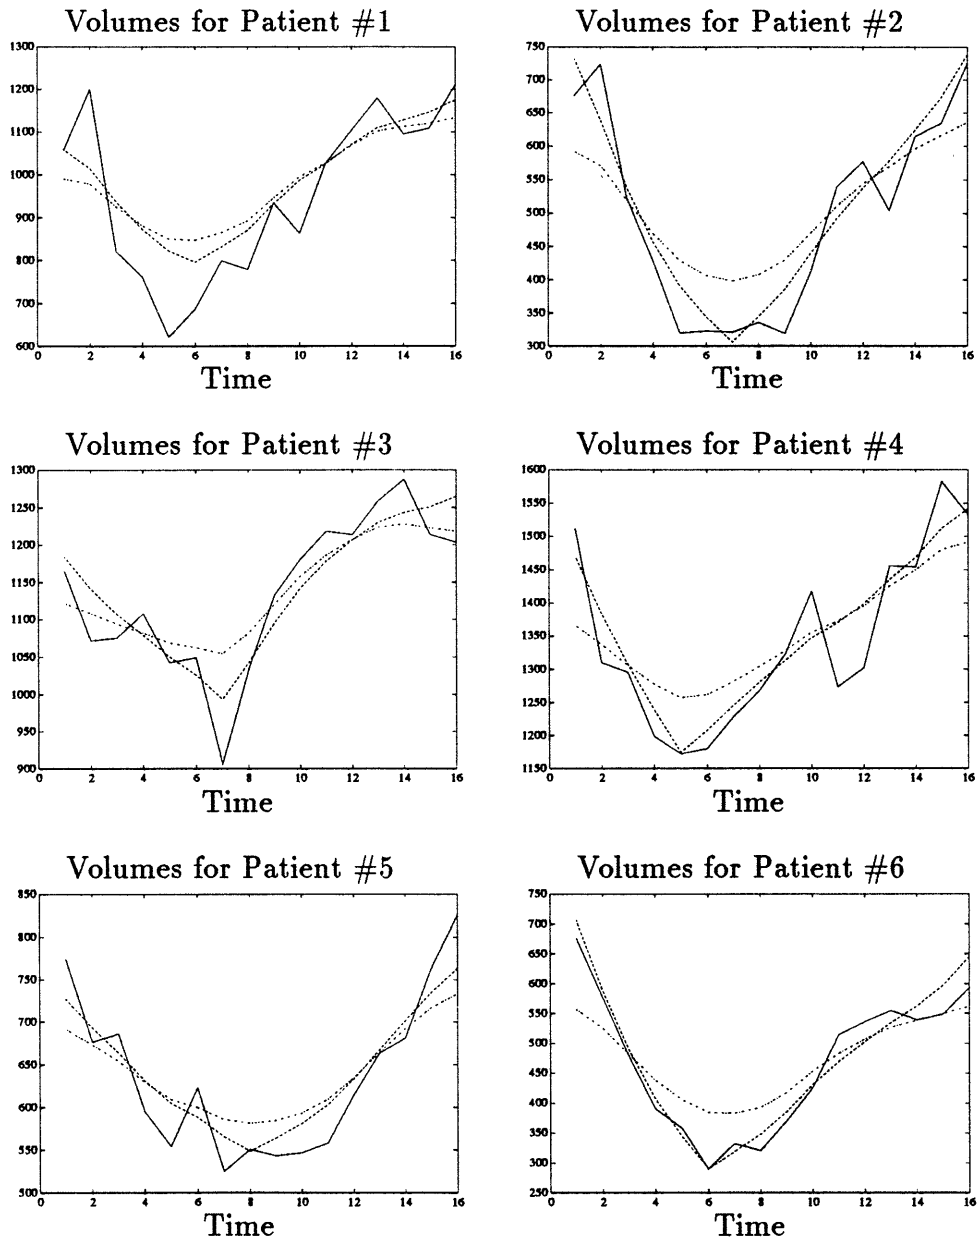


Figure 6-4: Smoothed three dimensional reconstructions for each patient

Patient	#1	#2	#3	#4	#5	#6
MUGA	0.26	0.56	0.2	0.25	0.28	0.40
Static	0.3938	0.5417	0.2347	0.2300	0.3113	0.5429
Assumed model	0.2872	0.5834	0.1888	0.2200	0.2634	0.5707
Random Walk	0.2016	0.3515	0.0988	0.1199	0.1837	0.3118

Table 6.9: Ejection Fractions for 3D

Assumed EF	0.20	0.30	0.40	0.50	0.60
Calculated EF	0.3857	0.4246	0.4664	0.5119	0.5620

Table 6.10: Variation of Reconstructed Ellipsoids EF with Assumed EF for Patient #6

based on the assumed model tend to have an ejection fraction that is skewed towards the ejection fraction associated with the assumed model. Just as we did for two-dimensional processing, we illustrate this point by processing the ellipses extracted for Patient #6 with smoothing filters based on several different assumed models. Table 6.10 shows the ejection fractions associated with each of the assumed models used and the ejection fraction calculated for each view using the different assumed models for Patient #6. Thus, the ejection fraction of the reconstructed ellipses is skewed by the assumed model used in the smoothing filter. We note that the estimated ejection fraction varies by approximately 0.18 as we vary the assumed ejection fraction by 0.40.

Performance Evaluation

Once again, one can imagine plotting the MUGA ejection fraction values versus the estimates of ejection fraction listed in Table 6.9. Just as with the two-dimensional based estimates, we determine the slope and offset of the line that is best fit to these six points. We also calculate the correlation coefficient of the MUGA ejection fraction value and the three-dimensional estimates of ejection fraction. The values for the slope and offset of the best fit line as well as the correlation coefficients are

	slope	offset	ρ
Static	0.8094	0.0209	0.8658
Assumed Model	0.6940	0.0805	0.9265
Random Walk	1.2254	0.0662	0.9368

Table 6.11: Correlation Coefficients

given in Table 6.11. The static estimates give a much lower degree of correlation than the smoothing-filter-based estimates. The correlation coefficient for the random walk ejection fraction estimates is higher than that for the assumed model ejection fraction estimates. This may be accounted for by the fact that the assumed model estimates are skewed by the assumed ejection fraction. Thus, we note once again the importance of using the correct assumed model in the smoothing filter reconstruction. In addition, we note that there is error introduced by our assumed approximation to the true (unkown) projection geometry.

6-A Ellipse Extraction Program

```
function [ae,be,thetae,Xce,xe,ye,out,out2,out3]=ellpit(A,view,PLOTIT,Xc)
% [ae,be,thetae,Xce,xe,ye,out,out2,out3]=ellpit(A,view,PLOTIT,Xc)
%
% A      : Raw Heart field of interest.
% view   : Opt. Projection. 1=ant,2=lat,3=lao. Default: User is prompted.
% PLOTIT : Opt. Verbose Plotting Toggle. 1=plot, 0=Don't. Default=1.
% Xc     : Opt. Center of heart. Default: User is prompted for center.
%
% Extracts an ellipse fitting the inner boundary of the heart data given
% in A. Performs the extraction as follows:
%
% 1) The field is smoothed by a 3x3 Gaussian window.
% 2) From given center, finds a polar version of the heart field. Theta is
%    assumed to go from pi/2 to -pi/2 (neg theta is CCW), r from 12 to -12.
% 3) Maybe average some slices into wedges (not current used).
% 4) Fits an 8-th order polynomial to each slice
% 5) For each slice finds extrema of fitted polynomial (=roots of 1st deriv).
% 6) Dump roots that are complex, too far (>12) or too close (<2) to the center.
% 7) Only roots corresponding to maxima are kept (2nd deriv < 0)
% 8) Only the single largest max in each of the pos and neg directions are kept.
% 9) The inner radius vaules for the wedge are defined to be a fraction of
%    the distance to these maximum values from the center (currently rfrac=.75).
% 10)Next those points in certain, view dependent excluded regions are
%     dropped from the list.
% 11)The final list of radius points is then median filtered to remove
%     outliers. (Median filter currently has NFTR = 4, so window size = 9.)
% 12)An ellipse is then fit to the remaining filtered (r,theta) pairs.
%
```

```

% View options:
% 1 = anterior
% 2 = lateral
% 3 = left-anterior oblique
%

% W. C. Karl 8/92

% TODO
%
% 1) Allow arbitrary user segment exclusion.

figure(1);colormap(gray);

if nargin < 2; % Get view from user
    GETVIEW=1;
else;          % User has supplied view
    GETVIEW=0;
end;

if nargin < 3;
    PLOTIT=1;
end;

if nargin < 4; % Get center from user
    GETCEN=1;
else;          % User has supplied center
    GETCEN=0;
end;

```

```

% Filters for possible use in program
gauss=[1 4 1;4 12 4;1 4 1];
gauss=gauss/sum(sum(gauss));    % Normalize Gaussian filter
vgrad3x3=[-1 -2 -1;0 0 0;1 2 1];
vgrad3x3=vgrad3x3;

% Get view from user if necessary
% View options:
% 1 = anterior
% 2 = lateral
% 3 = left-anterior oblique
%
if GETVIEW;
    view=menu('Which view is input?','Anterior','Lateral','LAO');
end;

% Smooth field
As=filter2(gauss,A);

%%%%% Find polar field %%%%%

% Find theta
nslice=1;                % Number of slices in each smoothed wedge
nwedge=40;               % Number of averaged wedges total
deltath = pi/(nslice*nwedge-1);

%%%%% View dependent parameter selection %%%%%
%
% theta:      Where to get slices
% thetaexcl: Matrix of exclusion regions. Each row is a theta region to

```



```

% exclude.
%           [thmin1 thmax1]
% thetaexcl = [thmin2 thmax2]
%           [ :      :      ]
% rfrac: Percentage of max to use for inner boundary detection

if view==1;      % Anterior view.

% theta=-.75*pi:deltath:.25*pi;
theta=-.5*pi:deltath:5*pi;
dth=pi/10;
thetaexcl=[-3*pi/4-2*dth -3*pi/4];      % Theta wedge around valve plane
rfrac=.75;

elseif view==2; % Lateral view.

% theta=-.25*pi:deltath:.75*pi;
theta=-.5*pi:deltath:5*pi;
dth=pi/10;
thetaexcl=[-pi/4-dth -pi/4+dth];      % Theta wedge around valve plane
rfrac=.75;

elseif view==3; % LAO

theta=-.5*pi:deltath:5*pi;
thetaexcl=[];      % Theta wedge to exclude
rfrac=.75;

end;

```

```

% Find r
r = -12:12;

% Actually find the polar field
if GETCEN          % Get center interactively from user
    [B,Xc]=polarit(As,r,theta);
    drawnow
else;              % User has supplied center
    [B,Xc]=polarit(As,r,theta,Xc);
end;
cx=Xc(1);
cy=Xc(2);

% Median filter polar field radially (NOT!)
%for i=1:size(B,2)
%    clc;
%    B2(:,i)=medftr(B(:,i),1);
%end;
B2=B;

% Combine slices along theta; avg every nslice slices
for i=1:nwedge;

    % Candidate slice indices for the wedge
    I = (1:nslice)+(i-1)*nslice;

    % Avg slices in the wedge
    B3(:,i) = sum([B2(:,I),zeros(size(B2,1),1)]')'/length(I);
    thetai(i) = mean(theta(I));          % Corresp theta value
end;

```

```

end;
out3=B3;

if PLOTIT
    % Plot slices used
    image(As);
    hold on;
    Ithp=[];
    Ithn=[];
    for i=1:size(thetaexcl,1); % For each exclusion region
        % Indices for pos r regions
        Ithp = [Ithp;find(thetaexcl(i,1) < thetai & thetai < thetaexcl(i,2))];

        % Indices for neg r regions
        thn = thetai+pi; % thn = thetas for negative r's
        Iflip = find(thn>pi); % Wrap around
        thn(Iflip) = thn(Iflip)-2*pi;
        Ithn = [Ithn;find(thetaexcl(i,1) < thn & thn < thetaexcl(i,2))];

    end;

    thetap=thetai;
    thetap(Ithp)=[];
    thetan=thetai;
    thetan(Ithn)=[];

    Irn=find(r(:)<0);
    Irp=find(r(:)>=0);

    xn=[r(Irn)'.*cos(thetan)];
    xp=[r(Irp)'.*cos(thetap)];

```

```

    yn=[r(Irn)']*sin(thetan)];
    yp=[r(Irp)']*sin(thetap)];

    plot(xn+cx,yn+cy,'r',xp+cx,yp+cy,'r');
    hold off;
    title('Sample vectors');
    pause(1)
end;

% Fit polynomial to each slice
polyord=8; % Polynomial order (>4 since we'll take 2 derivatives)
for i=1:nwedge
    V(i,:)=polyfit(r',B3(:,i),polyord);           % Coeff of best fit poly
    Vd(i,:)=polyord:-1:1.*V(i,1:polyord);       % Coeff of 1st deriv of poly
    Vdd(i,:)=polyord-1:-1:1.*Vd(i,1:polyord-1); % Coeff of 2nd deriv of poly
    B4(:,i) = polyval(V(i,:),r');               % Matrix of function curves
    B4d(:,i) = polyval(Vd(i,:),r');            % Matrix of 1st deriv curves
    B4dd(:,i) = polyval(Vdd(i,:),r');          % Matrix of 2nd deriv curves
end;

out2=B4dd;

%%%% Find the inner boundary %%%%

OUTLIER=12; % Value of abs(r) above which we declare an outlier.
rthetai=[];
rthetaip=[];
rthetain=[];

for i=1:nwedge

```

```

% Use inflection points
if 0
    lam=roots(Vdd(i,:));          % Find zeros of 2nd deriv
    lam(find(lam~=real(lam)))=[]; % Discard complex roots
    lam(find(abs(lam)>OUTLIER))=[];% Discard any radius points larger than OUTLIER
    [foo,I]=sort(abs(lam));      % Find closest roots to 0 (the center)
    lam=lam(I)
    rthetai=[rthetai;[lam(1:2),ones(2,1)*thetai(i)]];
end

% Use fn of min and max using 1st deriv
if 0
    lam=roots(Vd(i,:));          % Find zeros of 1st deriv
    lam(find(lam~=real(lam)))=[]; % Discard complex roots
    lam(find(abs(lam)>OUTLIER))=[];% Discard radius points larger than OUTLIER
    lam(find(abs(lam)<2))=[];     % Discard radius points smaller than 2
    [foo,I]=sort(abs(lam));      % Find closest roots to 0 (the center)
    lam=lam(I);
    cen=lam(1);                  % Identify center point and remove it
    lam(1)=[];
    lam
    nlam=min(2,length(lam));     % Number of lambda values to keep
    ri=(.6*lam(1:nlam)+.4*cen);  % Inner pts = fn of max and min
    rthetai=[rthetai;[ri,ones(nlam,1)*thetai(i)]];
end

% Use fn of user given center and poly max using 1st deriv
if 1
    lam=roots(Vd(i,:));          % Find zeros of 1st deriv

```

```

lam(find(lam~=real(lam)))=[]; % Discard complex roots
lam(find(abs(lam)>OUTLIER))=[];% Discard radius points larger than OUTLIER
lam(find(abs(lam)<2))=[]; % Discard radius points smaller than 2

% [foo,I]=sort(abs(lam)); % Find closest roots to 0 (the center)
% lam=lam(I); % Sorted roots, smallest is now first
% if length(lam)>0;
% rcen=lam(1); % r value of center
% end;

% Keep only maxima (drop mins)
lam(find(polyval(Vdd(i,:),lam)>=0))=[]; % Drop if 2nd deriv is >=0

% Split into pos and neg parts
plam=lam(find(lam>0)); % Find positive extrema values
nlam=lam(find(lam<0)); % Find negative extrema values

% Find the r values of largest max in both pos and neg dir along slice
rmaxp=[]; % Default condition: No point on the pos line
rmaxn=[]; % Default condition: No point on the neg line
if length(plam)>0; % pos direction
    [foo,Ip]=sort(polyval(V(i,:),plam));
    Ip=flipud(Ip);
    rmaxp=plam(Ip(1));
end;
if length(nlam)>0; % neg direction
    [foo,In]=sort(polyval(V(i,:),nlam));
    In=flipud(In);
    rmaxn=nlam(In(1));
end;

```

```

rmax=[rmaxn,rmaxp];

% Find R values of inner boundary points found for current slice
rip=(rfrac*rmaxp);    % Pos Inner pts = fn of max and center
rin=(rfrac*rmaxn);    % Neg Inner pts = fn of max and center
ri=[rip;rin];         % Inner radius point for plotting

% rthetaip contains the pairs of identified pos inner boundary points
% rthetain contains the pairs of identified neg inner boundary points
% rthetai=[ri thetai];
rthetaip=[rthetaip;[rip,ones(length(rip),1)*thetai(i)]];
rthetain=[rthetain;[rin,ones(length(rin),1)*thetai(i)]];

end

% Plotting stuff
if PLOTIT
    disp(['Frame: ',int2str(i)])
    rmax
    lam
    figure(3);colormap(gray) % Plot r values
    hold off
    if length(rthetain)>0
        plot(rthetain(:,2)*180/pi,rthetain(:,1),'r');
        hold on
    end;
    if length(rthetaip)>0
        plot(rthetaip(:,2)*180/pi,rthetaip(:,1),'g');
    end;
    title('Extracted Radius')

```

```

xlabel('Theta (degrees)')
ylabel('rmaxn (red) and rmaxp (green)')
hold off
figure(2);colormap(gray) % Plot slice and values found
plot(r,B4(:,i),r,B3(:,i),lam,polyval(V(i,:),lam),'rx');
hold on;plot(ri,polyval(V(i,:),ri),'bo');hold off;
hold on;plot(rmax,polyval(V(i,:),rmax),'b*');hold off;
title('Slice profile. x=inflct pts, *=max, o=inner bddry')
xlabel('Radius')
figure(1);image(As); % Plot image with boundary points
hold on;
plot(Xc(1),Xc(2),'r+');
plot(lam*cos(thetai(i))+Xc(1),lam*sin(thetai(i))+Xc(2),'rx')
plot(ri*cos(thetai(i))+Xc(1),ri*sin(thetai(i))+Xc(2),'bo')
plot(rmax*cos(thetai(i))+Xc(1),rmax*sin(thetai(i))+Xc(2),'b*')
hold off
drawnow;
disp('Spacebar to continue..')
pause
end;

end

%%% Post processing of extracted boundary points %%%

% Exclude those wedges/samples in specified regions
for i=1:size(thetaexcl,1); % For each exclusion region
    % Indices for pos r regions
    Ithp = find(thetaexcl(i,1)<rthetaip(:,2) & rthetaip(:,2)<thetaexcl(i,2));
    rthetaip(Ithp,:)=[]; % Delete the points

```



```

% Indices for neg r regions
thn = rthetain(:,2)+pi;
Iflip = find(thn>pi); % Wrap around
thn(Iflip) = thn(Iflip)-2*pi;
Ithn = find(thetaexcl(i,1) < thn & thn < thetaexcl(i,2));
rthetain(Ithn,:)=[]; % Delete the points
end;

% Median filter the remaining radius points
NFTR=4;
np = length(rthetaip(:,1)); % Number of pos radius points
nn = length(rthetain(:,1)); % Number of neg radius points
rtot = [rthetaip(:,1);-rthetain(:,1)];% Total r vector
rtotf = medftr(rtot,NFTR);

% Place filtered radius points in rthetai
% rthetai contains the r,theta boundary pairs
% rthetai=[ri thetai];
rtotf=rtotf(:);
rthetai=[[rtotf(1:np),rthetaip(:,2)];[-rtotf(np+1:np+nn),rthetain(:,2)]];

% Find the x,y coordinates of points on the boundary
for i=1:size(rthetai,1)
    xi(i,1)=rthetai(i,1)*cos(rthetai(i,2))+Xc(1);
    yi(i,1)=rthetai(i,1)*sin(rthetai(i,2))+Xc(2);
end;

[ae,be,thetae,Xce]=fitellp3(xi,yi);
if ae~=real(ae)|be~=real(be)

```

```

disp('WARNING: complex values obtained');
end;
ae=real(ae);
be=real(be);
[xe,ye]=mkellp(thetae,ae,be,Xce);

if PLOTIT
figure(1)
image(As);
hold on;
% plot(xi,yi,'rx')
plot(xe,ye);
hold off;

% Plot extracted radius values
figure(3);
plot(rthetaip(:,2)*180/pi,rthetaip(:,1),'g',rthetain(:,2)*180/pi,rthetain(:,1),
hold on;
plot(rthetaip(:,2)*180/pi,rtotf(1:np),rthetain(:,2)*180/pi,-rtotf(np+1:np+nn))
hold off;
xlabel('Theta (degrees)');
ylabel('rmaxn (red) and rmaxp (green)');

figure(1)
end;

```

Chapter 7

Future Work and Conclusions

7.1 Conclusions

In this thesis, we have presented a method to apply geometric reconstruction techniques to estimate dynamically the ejection fraction of the left ventricle of the heart from a temporal set of myocardial perfusion images. This method provides an alternative to current radionuclide ventriculography-based methods to estimate ejection fraction. By estimating ejection fraction from myocardial perfusion images, the process of diagnosing heart disease becomes safer and more cost effective. In addition, by employing recursive estimation techniques, we are able to estimate ejection fraction more robustly.

In Chapter 4, we used Rauch-Tung Striebel smoothing to reconstruct computer simulated, dynamically evolving ellipsoids. Our simulations included the reconstruction of two-dimensional ellipsoids from noisy observations of the ellipses themselves and the reconstruction of three-dimensional ellipsoids from noisy lower-dimensional projections. The simulations were formulated in such a way that our results might be easily applicable to real data. We discussed methods of adjusting the filter to account for mismatches between the true and modelled dynamics. The results from the simulations indicate that it is possible to improve on static estimates of ejection fraction by using a simplified model of the true ellipsoid dynamics. In addition, we explored techniques to estimate ejection fraction when our knowledge of projection

geometry is imperfect.

Chapter 5 explored a method of choosing which of two hypothesized dynamic models best approximates the true dynamics of a computer simulated, dynamically evolving ellipsoid. We investigated the idealized case where one of the hypothesized models exactly matches the true ellipsoid dynamics. We also considered the case where the true and hypothesized dynamics differ in contraction rate only. Finally, we considered the case where the true and hypothesized dynamics differ in rotation rate as well as contraction rate. Our results illustrate that we are able identify the hypothesized model that more closely approximates the true ellipsoid dynamics.

Finally, in Chapter 6, we applied the model identification and smoothing filter based reconstruction techniques to ellipses extracted from real myocardial perfusion data for six individuals with potential heart disease. We used model identification and smoothing filter-based reconstruction to process each of the three projected views individually. In addition, we used model identification and smoothing filter based reconstruction to process the three views together using an assumed projection geometry. We are able to show a high degree of correlation between our smoothing filter based estimates of ejection fraction and MUGA based estimates of ejection fraction.

The results presented in this thesis indicate that a dynamic estimate of ejection fraction based on myocardial perfusion images may provide an alternative to currently used techniques based on radionuclide ventriculography, thereby making the process of diagnosing heart disease safer and more cost effective. In addition, this thesis has explored in great detail the effect of dynamic model mismatch in the more general problem of reconstruction of dynamically evolving ellipsoids.

7.2 Future Work

This section deals with areas that may be explored in future research. Although our current results show that our proposed method provides an accurate estimate of ejection fraction, we may be able to provide a more robust ejection fraction estimate by incorporating some of the suggestions listed in this section.

In future work, we suggest a more detailed model identification scheme that is used to directly estimate the ejection fraction. Thus, the smoothing filter-based reconstruction would be eliminated and the process of estimating ejection fraction would be reduced to one step.

It is possible to use model identification to estimate the minimum point in the cardiac cycle which, as we noted in Chapter 6, varies from person to person. We might propose several hypothesized models with the same contraction rate, but with varying minimum points. Then, the model identification scheme could be used to identify which of the hypothesized minimum points most closely matches the true cycle minimum of the data.

In the three-dimensional reconstructions of Chapter 6, we assumed a projection geometry which incorporated our knowledge about the general orientation of the set of projection planes used to generate the three views for the myocardial perfusion images. We may obtain a more accurate estimate of ejection fraction by assuming only two of the projection angles and attempting to estimate the third projection angle using a similar model identification or parameter estimation scheme.

Future work may also investigate the possibility of a joint ellipse extraction and ellipsoid estimation scheme. This type of technique might use the estimated three-dimensional ellipsoid as feedback for the ellipse extraction routines.

We also suggest that the plots shown in Chapter 6 that show dynamic estimates of volume versus time may have potential prognostic value, in and of themselves. We have noted that for patients with low ejection fraction and what appear to be occlusions these plots show unusual variations for estimates of volume over time.

Finally, recall that in Chapter 6 we used model identification and smoothing filter based reconstruction to estimate the ejection fraction of six patients. The same model identification and smoothing filter based reconstruction techniques should be explored on a larger data set.

Bibliography

- [1] Alex Aisen and Andrew Buda. Magnetic resonance imaging of the heart. In *Digital Cardiac Imaging*. Martinus Nijhoff Publishers, 1985.
- [2] C. A. Boucher, F. J. Wackers, B. L. Zaret, and I. G. Mena. Technetium-99m sestamibi myocardial imaging at rest for the assessment of myocardial infarct and first pass ejection fraction: A multicenter trial. Unpublished.
- [3] Andrew Buda and Edward Delp. Digital two-dimensional echocardiography. In *Digital Cardiac Imaging*. Martinus Nijhoff Publishers, 1985.
- [4] Robert H. Eich. *Introduction to Cardiology*. Harper & Row, 1980.
- [5] Jaekyeong Heo, George Hermann, Abdulmassih Iskandrian, Alan Askenase, and Bernard Segal. New myocardial perfusion imaging agents: Description and applications. *American Heart Journal*, 1988.
- [6] Jack Juni and Andrew Buda. Radionuclide imaging of the heart. In *Digital Cardiac Imaging*. Martinus Nijhoff Publishers, 1985.
- [7] W. Clem Karl. *Reconstruction Objects from Projections*. PhD thesis, Massachusetts Institute of Technology, February, 1991.
- [8] Israel Mirsky, Dhanjoo Ghista, and Harold Sandler. *Cardiac Mechanics: Physiological, Clinical, and Mathematical Considerations*. John Wiley & Sons Inc., 1974.
- [9] Technical Staff of The Analytic Sciences Corporation. *Applied Optimal Estimation*. The MIT Press, 1988.

- [10] L. Opie. *The Heart*. Grune & Stratton, Inc., 1984.
- [11] J. L. Prince. *Geometric Model-Based Estimation from Projections*. PhD thesis, Massachusetts Institute of Technology, January, 1988.
- [12] D. J. Rossi. *Reconstruction From Projections Based on Detection and Estimation of Objects*. PhD thesis, Massachusetts Institute of Technology, October, 1982.
- [13] Albert Sinusas, George Beller, William Smith, Ellen Vinson, Valerie Brookeman, and Denny Watson. Quantitative planar imaging with technetium-99m methoxyisobutyl isonitrile: Comparison of uptake patterns with thallium-201. *The Journal of Nuclear Medicine*, 1989.
- [14] S. R. Underwood, S. Walton, P. J. Laming, P. H. Jarrit, P. J. Ell, R. W. Emanuel, and R. H. Swanton. Left ventricular volume and ejection fraction determined by gated blood pool emission tomography. *British Heart Journal*, 1985.
- [15] A. S. Willsky. Recursive estimation. Class notes for 6.433. MIT.