September 1993 **LIDS-TH-2194**

Research Supported By:

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

National Science Foundation
Grant NSF 9015281-MIP

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

Office of Naval Research Fellowship

Tomographic Reconstruction of Polygons from Knot Locations and Chord Length Measurements

Lori-Ann Belcastro

Tomographic Reconstruction of Polygons from Knot Locations and Chord Length Measurements

by

Lori-Ann Belcastro

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Science in Computer Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1993

Author .... Lori-Ann Belcastro ....
Department of Electrical Engineering and Computer Science
September 3, 1993

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

Tomographic Reconstruction of Polygons from Knot Locations and Chord Length Measurements

by

Lori-Ann Belcastro

Submitted to the Department of Electrical Engineering and Computer Science
on September 3, 1993, in partial fulfillment of the
requirements for the degree of
Master of Science in Computer Science and Engineering

Abstract

Computerized Tomography has a wide variety of applications in diverse fields such as medical imaging, non-destructive testing, oceanography, and security inspection. Much of the past focus in tomographic imaging has been to develop algorithms that produce high resolution images from the projection data. For a variety of reasons, including limited sensor availability and noisy projection data, it is often impossible to produce a high resolution image. This thesis focuses on those situations when only sparse, noisy projection data are available. We present finite parameter estimation algorithms to reconstruct the vertices of a binary polygonal object from geometric information extracted directly from the projection data. Specifically, we develop three algorithms to reconstruction binary polygonal objects from measurements of knot locations (location of the projection of a vertex of the object in the projection data) and measurements of chords (thickness of the object). Both of these measurements can be extracted directly from the projection data. The chords correspond to the magnitude of the projection data and the knot locations correspond to the positions of abrupt change in the slope of the linear spline function that results from the projection of a binary polygonal object. The first algorithm incorporates each view sequentially, in increasing angular order. The second algorithm allows views to be incorporated in any angular order. The third, and final, algorithm is a nonlinear optimization algorithm that uses the output of either the first or second algorithm as its initial guess. Additionally the final chapter addresses situations where the extraction of the knot location data from the projections is incomplete (i.e. missing knot location measurements). At each stage of our analysis we demonstrate the performance of the algorithms with simulations of various reconstruction scenarios.

Thesis Supervisor: Alan S. Willsky
Title: Professor

Thesis Supervisor: W. Clem Karl
Title: Research Scientist

Acknowledgments

I would like to take this opportunity to express my gratitude to my thesis advisors, Alan Willsky and Clem Karl. I thank Alan Willsky for giving me the opportunity to work in his research group. His insight and experience have been invaluable to me throughout this work. I would also like to thank Clem Karl for always having an "open door" to listen to new ideas. His guidance and suggestions will always be greatly appreciated.

I also gratefully acknowledge the generous support of the Office of Navy Research fellowship which I have held for the past two years.

In addition, I would like to thank the people who have supported me both before I came to MIT and while I was here. Most importantly, I would like to thank my family and both of my parents for always being a constant source of encouragement. I also owe special thanks to Don for always being understanding and supportive in any decisions that I have made (and for all the dinners during the long summer of thesis writing) and Michelle for being a great friend and getting me out to take a break from work when I really needed it. In addition, thanks to all of the officemates that I have had while I've been at MIT for their friendliness and kindness. In particular, I have appreciated Mickey's humor, Seema's friendliness, and Peyman's suggestions and support. Thanks to all of 1W at East Campus for making MIT a great place to live. And last, but not least, Fred the cat for always making life interesting.....

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

Many algorithms have been developed to reconstruct a multi-dimensional function from its projections. In tomographic imaging, widely used methods such as filtered backprojection and Fourier methods [3,4] are used to reconstruct high resolution images in a variety of applications [4]. Although these methods produce high quality images, they require a large number of projections and a relatively large signal-to-noise ratio (SNR).

In applications such as oceanography [10] and non-destructive testing, the number of projection angles is often severely constrained. The availability of sensors to obtain the data often limits both the number of projections and the angular spacing of the views. Often data can only be collected over a limited angular range. In addition, the projection data are usually noisy. In cases of low SNR or incomplete data, a high resolution image is virtually impossible because classical reconstruction algorithms introduce artifacts during the reconstruction process. The goal in this case may be to try to estimate the boundaries of the object. Because the traditional high resolution algorithms break down when only sparse, noisy projection data are available, these techniques are no longer effective in this situation.

Additionally, in a number of computerized tomography problems, the ultimate goal is not a high resolution image, but simply to characterize the size and shape of the object being imaged. The approach in the past has been to use a classical image reconstruction algorithm to obtain the image and then extract the edges of the object with some kind of edge detector. More recently, however, a number of algorithms have been developed to extract the object directly from the projection data [5,9,12,14,16,17]. In cases of low SNR or incomplete data, the second methodology is the preferred approach because the artifacts introduced by the classical reconstruction algorithms during the reconstruction process make the application of standard signal processing techniques ineffective. Further, post-processing of a reconstructed image is also difficult because the noise in the reconstructed image is colored even when the noise in the projection data is white.

A robust algorithm to detect object size and shape may also be used as a pre-processor to provide *a priori* information to a high resolution image reconstruction algorithm. For example, the reconstruction algorithms developed by Prince [12] or Milanfar [9] assume prior geometric information about the object to be reconstructed. Thus, the emphasis in the pre-processor is not on a high resolution image, but on a robust estimate of the object's shape. This emphasis is especially true if the projection data are noisy.

Therefore, in some applications it is either not desirable or not possible to produce a high resolution image from projection data. Classical methods of tomographic image reconstruction may not be the best approach in these situations. Instead, it is desirable to implement an estimation algorithm that characterizes the shape of the object from sparse and noisy projection data.

The objective of this thesis is to investigate finite parameter, geometric-based reconstruction algorithms for binary polygonal objects from sparse and noisy projection data. The primary focus will be to formulate a robust algorithm to reconstruct the vertices of a polygon from knot location and chord length measurements of the projection data.

1.2 Contributions and Organization

The work presented in this thesis makes contributions in the following areas:

- Introduces several algorithms that extend the work of Milanfar [9], Prince [12], Rossi [14], etc. in geometric-based reconstruction techniques. Specifically, by extracting geometric information related to the vertices and chord length (thickness of the object) of the object directly from the noisy projection data, we estimate the vertices of a convex, binary polygonal object.

- Supports the work of Prince [12] and Milanfar [9] in showing the usefulness of geometric information extracted directly from the noisy Radon transform data to reconstruct a convex, binary polygonal object.
- Links the multitarget tracking problem of radar to the tomography problem of reconstructing the vertices of a polygonal object, allowing many of the techniques and algorithms developed in the radar context to be implemented in a tomography problem. In addition, this work recognizes that the geometric constraints of the tomography problem can be used to simplify the multitarget tracking algorithmic approaches.
- Develops several algorithms to reconstruct objects from sparse, noisy data. In addition, the algorithms allow limited angle or non-uniformly spaced projection data.
- Develops an algorithm that deals robustly with missed detections, or missing data measurements.

This thesis is organized as follows. Chapter 2 begins with background information and definitions utilized in the remainder of the thesis. Chapter 3 presents relevant previous work and the problem geometry considered in the remainder of this thesis. Each of the following three chapters treats a different reconstruction algorithm. Chapter 7 provides methods for dealing with data inconsistencies (missing data measurements). Finally, conclusions and future work are summarized in Chapter 8.

Chapter 2

Definitions and Problem Statement

This chapter provides a brief overview of concepts and definitions that will be used in subsequent chapters. We begin with a review of the major concepts of M-ary Maximum Likelihood (ML) hypothesis testing (Section 2.1.1), ML estimation (Section 2.1.2), and the Generalized Likelihood Ratio Test (Section 2.1.3) . Next, in Section 2.1.4 we define some of the major concepts of tomography including the Radon transform and support lines. Finally, we review definitions of knot locations and chord length and relate these values to the projection data available in tomography problems in Section 2.1.5. Having completed a review of the background definitions and concepts, Section 2.2 is a statement of the reconstruction problem considered in this thesis.

2.1 Definitions

2.1.1 M-ary Hypothesis Testing

An M-ary hypothesis testing problem assumes there are M hypotheses H_0, H_1, ... , H_{M-1} with *a priori* probabilities $P_i = Pr(H_i)$. Given the set of conditional distributions $p_{y|H_i}(\underline{Y}|H_i)$ of the observation data, $\underline{y}$, the objective is to specify a decision rule that chooses the optimal hypothesis based on measurement data and a specified cost criterion.

In general, the expected cost of choosing hypothesis k is defined:

$$E(C_k) = \sum_{j=0}^{M-1} C_{kj} Pr(H_j | \underline{y} = \underline{Y}) \tag{2.1}$$

where C_{kj} is the cost of choosing hypothesis k if hypothesis j is the true hypothesis and $Pr(H_j | \underline{y} = \underline{Y})$ is the probability that hypothesis j is the correct hypothesis given the measurements $\underline{y} = \underline{Y}$. The optimal decision rule, which minimizes the cost function defined in Equation (2.1), is given by:

$$\begin{aligned} d(\underline{Y}) &= H_k \text{ if for all } H_i \neq H_k \\ E(C_k) &\leq E(C_i) \end{aligned} \tag{2.2}$$

When all the hypotheses are considered equally likely *a priori* and all possible errors are penalized equally by letting:

$$C_{kj} = \begin{cases} 0 & \text{if } i = j \\ 1 & \text{if } i \neq j \end{cases}$$

the problem is termed a *Maximum Likelihood* (ML) rule. The result of Equation (2.2) then reduces to:

$$\begin{aligned} d(\underline{Y}) &= H_k \text{ if for all } H_i \neq H_k \\ \sum_{j \neq k} Pr(H_j | \underline{y} = \underline{Y}) &\leq \sum_{j \neq i} Pr(H_j | \underline{y} = \underline{Y}) \end{aligned} \tag{2.3}$$

or

$$\begin{aligned} d(\underline{Y}) &= H_k \text{ if for all } H_i \neq H_k \\ Pr(H_k | \underline{y} = \underline{Y}) &\geq Pr(H_i | \underline{y} = \underline{Y}) \end{aligned} \tag{2.4}$$

Using Bayes rule, we note that

$$Pr(H_k | \underline{y} = \underline{Y}) = \frac{p_{y|H_k}(\underline{Y} | H_k) Pr(H_k)}{\sum_{j=0}^{M-1} p_{y|H_j}(\underline{Y} | H_j) Pr(H_j)} \tag{2.5}$$

which can be substituted into Equation (2.4) to obtain

$$\frac{p_{y|H_k}(\underline{Y}|\ H_k)Pr(H_k)}{\sum_{j=0}^{M-1} p_{y|H_j}(\underline{Y}|\ H_j)Pr(H_j)} \geq \frac{p_{y|H_i}(\underline{Y}|\ H_i)Pr(H_i)}{\sum_{j=0}^{M-1} p_{y|H_j}(\underline{Y}|\ H_j)Pr(H_j)} \tag{2.6}$$

Using the fact that the denominator of Equation (2.6) is a normalization constant independent of H_k and that all hypotheses are equally likely *a priori*, the ML decision rule reduces to

$$\begin{aligned} d(\underline{Y}) &= H_k \text{ if for all } H_i \neq H_k \\ p_{y|H_k}(\underline{Y}|\ H_k) &\geq p_{y|H_i}(\underline{Y}|\ H_i) \end{aligned} \tag{2.7}$$

or

$$\begin{aligned} d(\underline{Y}) &= H_k \text{ if for all } H_i \neq H_k \\ \mathcal{L}(\underline{Y}) &= \frac{p_{y|H_k}(\underline{Y}|\ H_k)}{p_{y|H_i}(\underline{Y}|\ H_i)} \geq 1 \end{aligned} \tag{2.8}$$

where $\mathcal{L}(\underline{Y})$ is termed the *likelihood ratio.* This decision rule is used to obtain the optimal ML estimate of the hypothesis given the measurements $\underline{y} = \underline{Y}$.

2.1.2 ML Estimation

In a ML estimation problem[1], the value to be estimated, $\underline{x}$, is assumed to be an unknown quantity (i.e. non-random). Also, the ML estimation problem assumes that the probabilistic model relating the value to be estimated, $\underline{x}$, and the observations, $\underline{y}$, is known. Under these assumptions, the ML estimate is defined as the value of $\underline{x}$ that makes the observations $\underline{y} = \underline{Y}$ most likely

$$\hat{\underline{x}}_{ML}(\underline{Y}) = \arg\max_{\underline{X}} \ p_{y|x}(\underline{Y}|\ \underline{X}) \tag{2.9}$$

In addition to problems where $\underline{x}$ is truly a non-random but unknown vector (Willsky [19] uses the example of chemical flow rates), this assumption is often made in estimation problems where $\underline{x}$ is random but its probabilistic distribution $p_x(\underline{X})$ is unknown.

---

[1] Complete derivations of these results can be found in [19].

For the specific case of a linear Gaussian problem, the observation equation is defined

$$\underline{y} = A\underline{x} + \underline{v} \tag{2.10}$$

where $\underline{y}$ is a vector of observations, A is the system matrix relating the observations to the desired vector $\underline{x}$, and $\underline{v} \sim N(\underline{0}, R)$ is zero-mean Gaussian noise. From Equation (2.10), with the additional assumptions that $\underline{x}$ is unknown and $\underline{v}$ is jointly Gaussian, it is apparent that $\underline{y}$ is also jointly Gaussian with

$$\underline{m}_{y|x} = A\underline{X} \tag{2.11}$$

$$\Lambda_{y|x} = R \tag{2.12}$$

In this case, Equation (2.9) becomes

$$\hat{\underline{x}}_{ML}(\underline{Y}) = \arg\max_{\underline{X}} N(\underline{Y}; A\underline{X}, R) \sim \arg\min_{\underline{X}} [(\underline{Y} - A\underline{X})^T R^{-1} (\underline{Y} - A\underline{X})] \tag{2.13}$$

where $N(\underline{Y}; A\underline{X}, R)$ denotes the normal distribution evaluated at $\underline{y} = \underline{Y}$ with mean $A\underline{X}$ and covariance R. This reduces to

$$\hat{\underline{x}}_{ML}(\underline{Y}) = (A^T R^{-1} A)^{-1} A^T R^{-1} \underline{Y} \tag{2.14}$$

with the corresponding error covariance

$$\Lambda_{ML} = (A^T R^{-1} A)^{-1} \tag{2.15}$$

2.1.3 Generalized Likelihood Ratio Test

In this section we extend the work of the Section 2.1.1 to include an ML hypothesis test based on a *Generalized Likelihood Ratio*[2]. In addition to the assumptions made previously (Section 2.1.1), this formulation assumes that each of the proposed hypotheses is also a function of a non-random vector $\underline{\theta}_k$ that may be different for each hypothesis H_k being considered. In the Generalized Likelihood Ratio Test, ML estimates of $\underline{\theta}_k$ are calculated (using the equations of the previous section) for each hypothesis by assuming that it is the

---

[2] Complete derivations of these results can be found in [18].

true hypothesis and these estimates are used in a likelihood ratio test as if they were the correct values. Thus, the likelihood ratio becomes

$$\mathcal{L}_g(\underline{Y}) = \frac{p_{y|H_k,\hat{\theta}_k}(\underline{Y}|\ H_k, \underline{\hat{\theta}_k})}{p_{y|H_i,\hat{\theta}_i}(\underline{Y}|\ H_i, \underline{\hat{\theta}_i})} \tag{2.16}$$

and the decision rule becomes

$$\begin{aligned} d(\underline{Y}) &= H_k \text{ if for all } H_i \neq H_k \\ \mathcal{L}_g(\underline{Y}) &= \frac{p_{y|H_k,\hat{\theta}_k}(\underline{Y}|\ H_k, \underline{\hat{\theta}_k})}{p_{y|H_i,\hat{\theta}_i}(\underline{Y}|\ H_i, \underline{\hat{\theta}_i})} \geq 1 \end{aligned} \tag{2.17}$$

given measurements $\underline{y} = \underline{Y}$.

2.1.4 Radon Transform and Support Functions

Tomographic imaging is concerned with the reconstruction of a two-dimensional (2D) image from a set of 1D projections. The Radon Transform relates a 2D function $f(x, y)$ to its 1D set of projections $g(t, \theta)$. This relationship is defined for every t and θ by

$$g(t, \theta) = \int_{R^2} f(x, y)\, \delta(t - w \cdot [x, y]^T)\, dxdy \tag{2.18}$$

and

$$w = [\cos(\theta), \sin(\theta)]^T \tag{2.19}$$

The function $g(t, \theta)$ is termed the *2D Radon Transform of f*. Further, the function $g(t, \theta)$ at a particular angle θ defines the *projection of f* at angle θ denoted by $g_\theta(t)$. Figure 2-1 depicts the relationship between the function $f(x, y)$ and its projection $g_\theta(t)$ at angle θ. The *sinogram* of f is the image formed from the values of $g(t, \theta)$ plotted as intensity values over the range $-\infty \leq t \leq +\infty$, $0° \leq \theta < 180°$.

An interesting property of the Radon transform is the fact that projections taken 180° apart are related by

$$g(t, \theta) = g(-t, \theta + 180°) \tag{2.20}$$

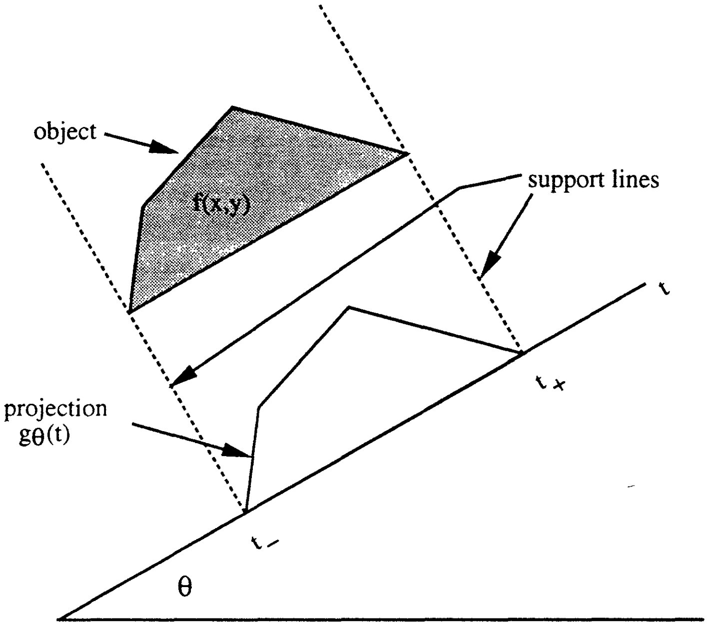

Figure 2-1: Relationship of the Radon Transform Parameters.

Therefore, to reconstruct an object from projections, data is only required over a 180° angular range because of the consistency requirement of Equation (2.20).

The line integral of f at a particular t and θ along the line $L(t, \theta)$ is depicted in Figure 2-2. Further, the line integrals that form the projection $g_\theta(t)$ are taken along a set of parallel lines that are denoted by $L_\theta(t)$. The ***support values*** are the two extrema, t_- and t_+, that define the set $L_\theta(t)$. The lines corresponding to these two values of t are termed ***support lines*** (refer to Figure 2-1).

Although Equation (2.18) is theoretically invertible for a large class of functions $f(x, y)$, in most practical applications an exact inversion is impossible because of the physical constraints of the system. For example, in medical imaging factors such as patient dosage and sensor availability limit both the number of line integrals per projection and the number of projection angles. Thus, practical tomography problems have both a finite number of angles and a finite number of line integrals. The primary focus of computerized tomography is to find an approximate or pseudo-inverse for the Radon Transform when projections are given for finite t and θ.

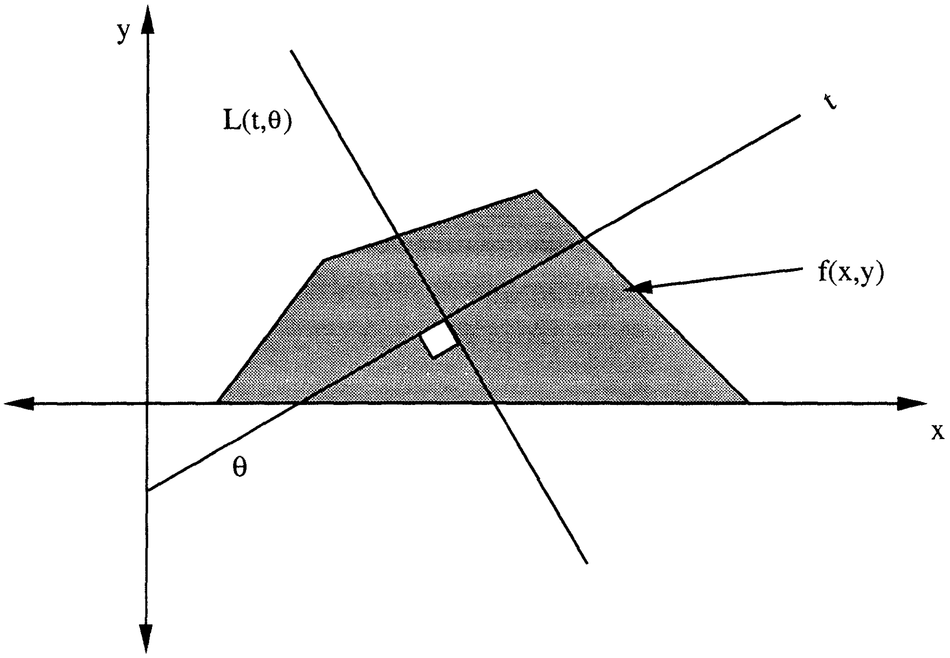

Figure 2-2: The Line $L(t, \theta)$.

2.1.5 Knot Locations and Chord Lengths

The projection of a binary polygonal object is a piecewise-linear spline function (see Figure 2-1). Each *knot location*, or position of abrupt change in the slope of the spline function [8], in the projection corresponds to the location of the projection of one of the vertices of the object. Except in degenerate cases where two vertices lie along the same line perpendicular to the projection, in each projection of the polygon the number of knots is equal to the number of vertices of the object. Figure 2-3 shows the relationship between the knot locations and the vertices of a polygonal object. From this diagram, it is apparent that the outer knot locations ($z1$ and $z4$ in the figure) also define the support values described previously. The *chord length* of a binary object is the thickness of the object along the line $L(t, \theta)$. In the projection of a binary object, the magnitude of the projection at a particular location $g_{\theta_k}(t_i)$, is equal to the chord length of the object along the line $L(t_i, \theta_k)$. Therefore, the projection of a binary object is simply a collection of chord lengths of the object. From the above discussion, it is apparent that chord length measurements and knot locations are closely related geometrically in both Radon and object space. Figure 2-3 shows the relationship between a 4-gon $f(x, y)$, its 1D projection $g_\theta(t)$, the knot locations, and chord lengths. Further, as Figure 2-3 demonstrates, the chord lengths at the support values are

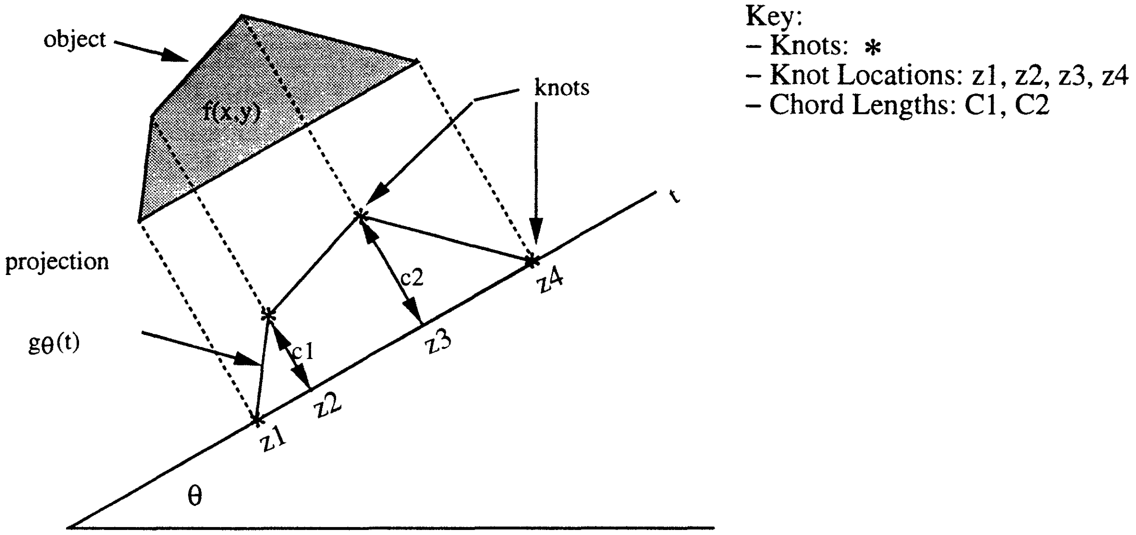

Figure 2-3: Relationship of Object, Projection, Knots, and Chords.

always zero.

2.2 Problem Statement

This thesis formulates reconstruction algorithms from sparse, noisy projection data. Specifically, the algorithms reconstruct binary polygonal objects using measurements of knot locations from the projection data[3]. Additionally, chord length measurements are incorporated to provide a more robust estimate of the object.

The problem of reconstructing a polygonal object from knot location data is one of *data association.* That is, given a set of ordered knot locations $z_1(\theta) \leq z_2(\theta) \ldots \leq z_n(\theta)$ from a projection, $g_\theta(t)$, the correspondence of the knots to the vertices of the object is unknown without knowledge of the object. In the first view, we arbitrarily assume that $z_1(\theta_1)$ corresponds to the unknown object vertex a, $z_2(\theta_1)$ corresponds to the unknown object vertex b, etc. Thus, the knot to vertex correspondence in the first view of a 4-gon can be summarized by *abcd.* As the angular view changes, so does this association of the ordered knot locations to the object's vertices. For example, in view 1 of Figure 2-4 the

---

[3] The idea of reconstructing polygonal objects from knot locations is a natural extension of the work of Prince [12] who used support lines to reconstruct the convex hull of an object. Chapter 3 details previous work in this and other related areas.

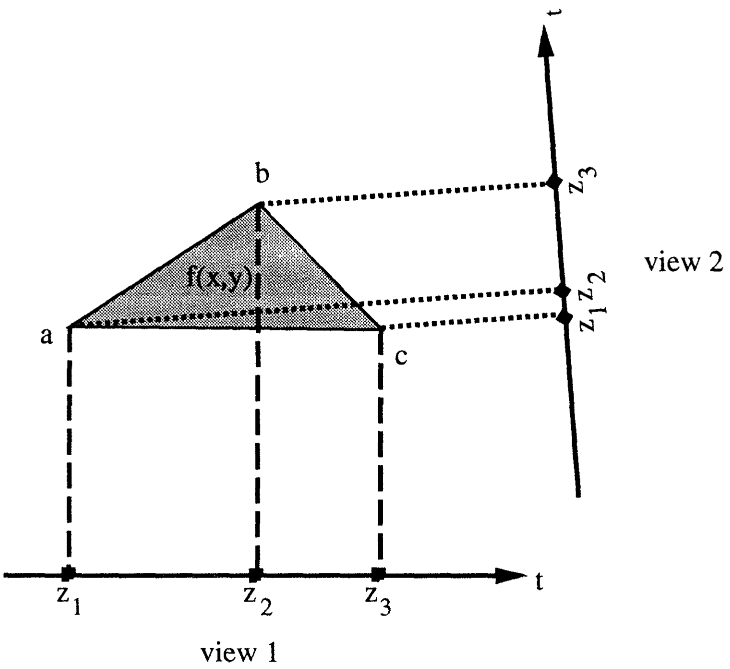

Figure 2-4: Example of Data Association of Knot Locations to the Vertices of an Object.

association of $z_1(\theta_1)$, $z_2(\theta_1)$, and $z_3(\theta_1)$ to object vertices would be labeled abc (consistent with the notation defined above). Therefore, the first knot, $z_1(\theta_1)$, corresponds to vertex a in the object. Referring again to Figure 2-4 we note that in the second view, however, the first knot, $z_1(\theta_2)$ corresponds to vertex c in the object. Thus, the knot locations have effectively "switched" their correspondences with the object. The knot location to vertex association in the second view can be summarized by cab because $z_1(\theta_2)$ corresponds to vertex c in the object, $z_2(\theta_2)$ corresponds to vertex a in the object, and $z_3(\theta_2)$ corresponds to vertex b in the object. We define a *switch* in two knots to occur when two vertices in the object lie along the same line, perpendicular to the projection. Prior to two knots switching, the knot locations more closer together with increasing θ. Finally, when the two vertices of the object are aligned along the same line, the knot locations for the two vertices are the same and the knot associations switch when θ increases. In a reconstruction algorithm from knot locations the knot to vertex data association must be determined at each projection.

Figure 2-5 shows the geometry of a noise-free reconstruction scenario. In the absence of noise, the knot locations of only three angular projections are needed to exactly triangulate the views and determine the vertices of the object. However, one misplaced knot can have tremendous effects on the overall reconstruction. Figure 2-6 demonstrates that movement of

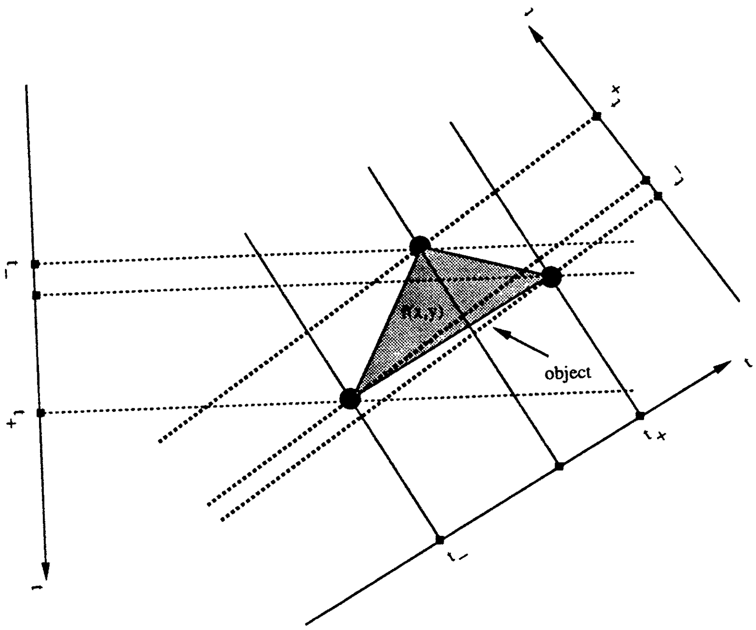

Figure 2-5: Object Reconstruction from Knot Location Data From 3 Noise-Free Views.

one knot location produces an inconsistent triangulation. The triangulation is inconsistent because there are only two points where three knot positions intersect to define a vertex location (instead of the three required to determine the vertices of the triangle).

The problem posed in this thesis is the following: Given noisy knot locations from sparse and noisy projection data, it is not possible to exactly triangulate the knot locations to determine the vertices of the polygon. Thus, the association of the knot location data to the vertices of the object from view-to-view is unknown. The objective of this investigation is to formulate an estimation algorithm that will incorporate knot locations, chord length data, noise statistics and prior geometric information to produce a robust estimate of the vertices of the object. This work seeks to reduce the effects of noise by using a hypothesis testing algorithm to robustly estimate the vertices of the polygon from noisy, inconsistent data. Each possible data association (within the constraints outlined by the algorithm) is considered a possible hypothesis and the objective is to determine the most likely hypothesis (using ML techniques) and determine the estimate of the object's vertices based on this hypothesis. The inclusion of chord length into the reconstruction algorithm enhances

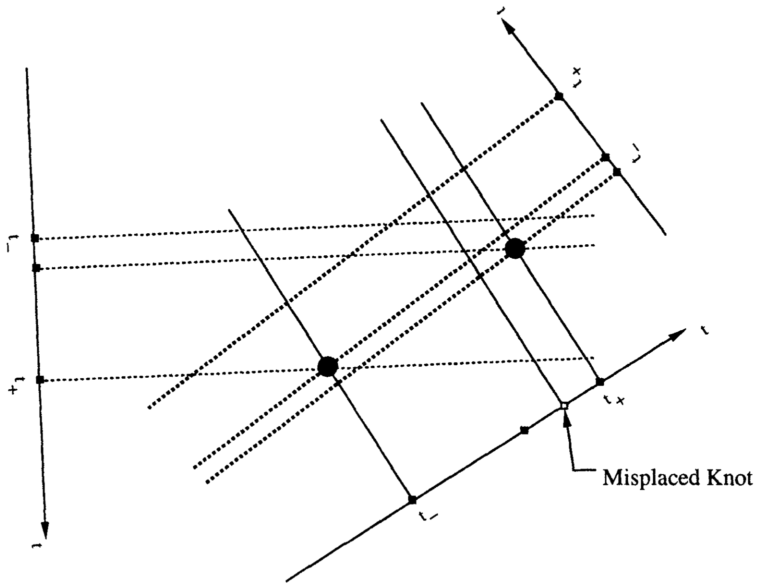

Figure 2-6: Inconsistent Triangulation of an Object Due to a Single Misplaced Knot Measurement.

the performance by using these data measurements to help determine the optimal data association hypothesis.

Chapter 3

Problem Formulation

Having defined the relevant terminology and problem statement in Chapter 2, we now move on to discuss related research and provide a more detailed explanation of the problem formulation. We begin in Section 3.1.1 by reviewing some of the more recent methods to detect object boundaries from tomographic data and describing related research in stochastic estimation of objects from tomographic data. Section 3.1.2 continues the review of relevant previous work by relating the tomographic reconstruction problem considered in this thesis to the multitarget tracking problem of radar. The second half of this chapter discusses some of the common issues related to the algorithms contained in this thesis. These include: assumptions (Section 3.2.1), triangulation geometry (Section 3.2.2), noise models (Section 3.2.3), initial view (Section 3.2.4), and performance measures (Section 3.2.5).

3.1 Previous Work

3.1.1 Related Research in Computerized Tomography

Recent Work in Boundary Detection Algorithms

In this section we briefly review some of the recent work used to detect object boundaries in tomographic imaging. Because the boundary detection of objects is the ultimate goal of many computerized tomography applications, a number of algorithms have been developed to detect edges and boundaries from tomographic data. Two different approaches can be used to detect edges of an object. The first is a two step procedure that reconstructs the object and then detects the boundary from the reconstructed image. In the second method, the boundaries of the object are obtained directly from the projection data. In the presence of noisy or sparse projection data, the second method is often preferred because the artifacts produced by classical algorithms during the reconstruction process render the application of standard edge detection methods on the image ineffective. In addition, even if the noise on the projection data is white, the reconstruction process produces an image with colored noise. Because of this, much of the recent work in computerized tomography has focused on this second approach to boundary detection.

Srinivasa *et al* [16] propose a method of edge detection of internal and external object boundaries without first reconstructing the object. They base their analysis on the use of the Marr-Hildreth edge detection operator, $\Delta^2 G(x,y)$, which uses the Gaussian function for filtering and the Laplacian for differentiation and detects edges by zero crossings in the filtered data. Instead of applying this edge detector to the reconstructed image, Srinivasa *et al* show that by taking the Radon transform of the Marr-Hildreth operator at each view angle, $p_\theta(t) = \mathcal{R}(\Delta^2 G(x,y))$, they can convert the 2D convolution step required to detect edges in an image to a 1D convolution in Radon space. They then combine this 1D function with the reconstruction filter, $r(t)$, and apply this "new" filter to the projection data in each view. After filtering each view, the filtered projections are then backprojected to obtain the function, $\tilde{f}''(x,y)$. Edges are then detected by determining the zero crossings in the resulting image. This method requires a large number of projections and Srinivasa *et al* consider only the noise free case.

Similarly, Thirion [17] proposes internal and external object boundary detection without image reconstruction. In his work, Thirion develops a technique which he terms *geometric tomography* to reconstruct object boundaries. His methodology is to reconstruct object boundaries directly from the sinogram data. Basically, he performs an edge enhancement on the sinogram, detects the "oriented dual curves" corresponding to inner and outer boundaries of the object, and reconstructs the object's boundaries from the dual curves. Thus, the object boundaries are reconstructed directly from the sinogram data. In his work, Thirion considers only the noise free case with dense sinogram data. In the presence of even small amounts of noise it may be difficult in many cases to extract the dual curves necessary for his edge detection algorithm.

The primary departure of this thesis from the work mentioned above is that sparse *and* noisy projection data is assumed. Further, it is also assumed that noise statistics and possibly some *a priori* geometric information are available. Unlike Thirion, the approach will be to reconstruct a polygon with a finite number of vertices. Because the sinogram of each vertex is a sinusoid, it should be easier to detect these known functions in the noisy sinogram data than it would be to detect the unknown "oriented dual curves" that Thirion enhances in his noise-free algorithm.

Related Research in Stochastic Estimation of Objects

Geometric estimation of objects from noisy tomographic data have been approached with a number of methods. Rossi and Willsky [15] used a finite parametric approach to characterize the size, shape, and orientation of objects. The primary methodology in this approach is a ML parameter estimation formulation. The accuracy of the estimate was evaluated using the Cramer-Rao lower bound on the variance of the estimates. The results of this study showed that size and orientation can be estimated more accurately than elongation and that estimates of orientation require a minimum elongation that is inversely related to the measurement SNR. Bresler and Macovski [2] extended this work to 3D reconstructions.

The primary focus of this thesis is an extension of the work of Prince and Willsky [11,12]. Instead of a finite parameter approach, they propose a number of reconstruction algorithms to characterize the size and shape of the object based on support line measurements of noisy projection data. Given the support lines at known angles equally spaced from $(0°, 360°]$, noise statistics, and *a priori* information such as relative smoothness, Prince and Willsky formulate the reconstruction as a constrained optimization problem. The algorithms use varying degrees of *a priori* geometric information. The formulations range from a ML estimate to a mini-max algorithm that maximizes the minimum discrete radius of curvature of the object. Further, Prince [12] proposes a knot location method which models each projection as a continuous piecewise-linear waveform. Using Kalman filtering techniques, he formulates a method to detect the two extreme knot locations, t_- and t_+, which correspond to the location of the support lines. This thesis will continue the work of Prince by using the additional information of the internal knot locations and chord length information to formulate a robust estimate of the size and shape of the object.

3.1.2 Relationship to Radar Multitarget Tracking Algorithms

In addition to the aforementioned work in tomography, the reconstruction of the vertices of a binary polygonal object from knot location measurements is very closely linked to the multitarget tracking problem of radar. The geometry of the tomography problem can be viewed in two ways: (1) as a stationary object with measurements taken at different angular positions or (2) as a rotating object with measurements taken at a single, fixed view. Although both geometries are equivalent (see [5]), the former is the more accepted framework for tomography problems as this is the method generally used to collect the data. By viewing the vertex reconstruction problem using the latter geometry, however, the close association to the single-sensor multitarget tracking problem becomes apparent. In this framework, the association of knots to object vertices is analogous to the association of radar reports to targets. Further, the vertex locations of the tomographic reconstruction problem correspond to the target locations of the radar problem.

In a single-sensor multitarget tracking problem, the objective is to collect radar data, associate the reports to targets, and estimate the locations of the targets. The problem consists of two distinct stages: data association and target estimation. In the knot reconstruction problem formulated in this thesis, the association of the knot location data to the vertices of the object from view-to-view is unknown. Like the multiple target tracking problem, the problem is one of data association; unlike the radar problem, however, there are knot correspondences that are impossible because of the geometric constraints imposed on the behavior of the vertices of a polygon. Further, the tomography problem is simplified because issues of "target maneuverablility" and dynamics are not necessary when we view the object as stationary with views taken at different angular positions; while in the radar context, the targets are free to move in and out of the field of view at different speeds relative to each other.

Still, these two problems are very similar and many of the estimation techniques developed in the radar context are applicable to this problem. Many algorithms, both optimal (Bayesian) and sub-optimal, have been formulated to solve the multitarget tracking problem. A number of these algorithms are based on an adapted form of a hybrid state estimation problem (see [1,6,7,13]) which proposes a simultaneous solution to a discrete-state estimation problem (data associations) and a continuous-state estimation problem (target locations). The reconstruction algorithms developed in this thesis are very similar to these hybrid state algorithms. The vertex reconstruction problem is viewed as a discrete-state data association problem where each hypothesis corresponds to a possible association of knot measurements to vertices. Similarly, it is also an estimation problem to determine the vertex locations. Thus, much of the methodology of these multitarget tracking problems can be applied to the vertex reconstruction problem.

In addition to formulating the problem as the simultaneous estimation of the discrete valued hypotheses (data associations) and continuous valued "targets" (object vertices), the algorithms used in this thesis also include the *target tree* approach described by Kurien *et al* [7] to represent the possible data associations. Further, if the number of vertices in the object is unknown *a priori*, the knot location data may contain "false alarms" as formulated in the radar problem. Thus, the object reconstruction can be extended to include the global hypothesis techniques used in the Track-Oriented approach of multitarget tracking [7].

Because of the huge number of possible hypotheses, the screening and pruning techniques developed in the multitarget approaches are also implemented in the reconstruction algorithms in this thesis. More specifically, the screening technique called *gating*, which limits the measurement space within which the reports for targets are expected to lie, can be used to limit the data associations for the object vertices in views that are taken in close angular proximity because the knot switches are constrained by the object's geometry (see Chapter 4). Also, pruning techniques such as the $N-scan$ approximation described by Kurien [6] can be used to limit the total number of hypotheses.

3.2 Problem Development

3.2.1 Assumptions

In the reconstruction algorithms developed in this thesis there are a number of underlying assumptions concerning both the *a priori* object information and the acquisition of input data.

In terms of *a priori* geometric information, we assume that the object to be reconstructed is a binary, convex polygonal object with a known number of sides. Although convexity is not required to reconstruct the vertices, it is necessary when determining the

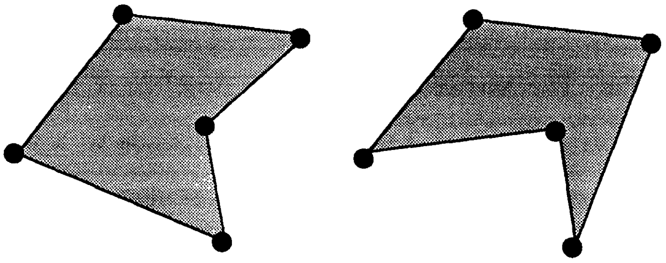

(a) (b)

Figure 3-1: Non-Unique Connection of the Vertices of a Non-Convex Object.

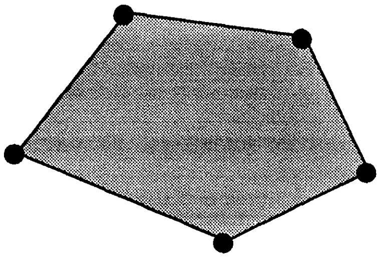

Figure 3-2: Unique Connection of the Vertices of a Convex Object.

unique connection of the vertices of the object. Figure 3-1 shows how 5 vertices can be connected to form two different non-convex objects. The connection of the vertices of a convex object, such as the object in Figure 3-2, is unique. The chord length measurement data can be effectively used to disambiguate the possible vertex connections but this involves a more complex reconstruction analysis. After the vertices have been estimated, all possible connections of the vertices would have to be evaluated using the chord length data to determine the optimal connection. Thus, in the interests of reducing model complexity we consider only convex objects.

The second set of assumptions concern the acquisition of the measurement data. First, we assume that the tomographic projections are taken at known angular views over the range $[0°, 180°)$. Second, we assume that the knot location measurements have already been extracted from the projection data independently of the algorithms developed in this thesis, for example using a procedure similar to the work of Mier-Muth and Willsky [8]. In an actual tomographic reconstruction scenario, the system would be similar to that of Figure 3-3. The noisy projection data (or chord length measurements) would be input into a knot extraction algorithm which would produce knot location measurements and their corresponding noise statistics. These data, along with the original projection data,

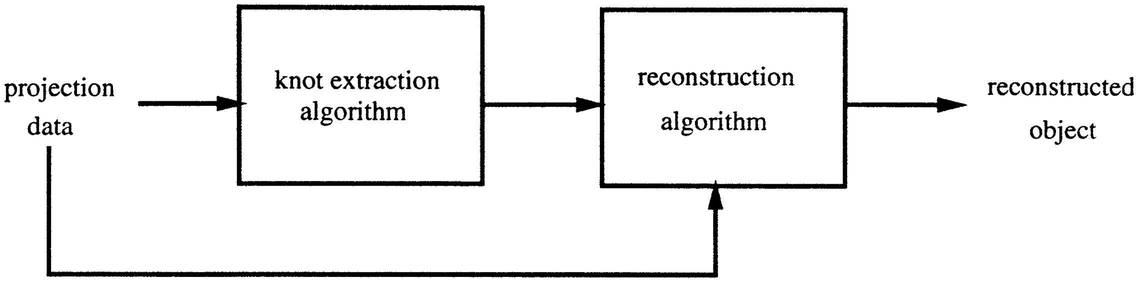

Figure 3-3: Application of Reconstruction Algorithms to Tomographic Imaging.

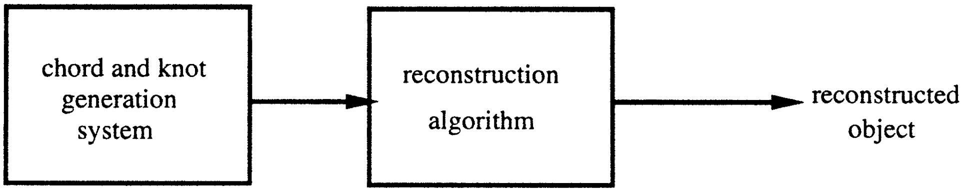

Figure 3-4: Reconstruction Scenario Considered in this Work.

would then be used as the input to one of the reconstruction algorithms developed in this thesis. We assume in this thesis that the pre-processing (knot extraction) has already occurred. Thus, the inputs to our system are noisy knot and chord measurements. The block diagram of the system modeled in this thesis is shown in Figure 3-4. The knot and chord measurements are generated from the underlying object and perturbed with additive white Gaussian noise (see Section 3.2.3 for a discussion of the noise models applied to the data). This simulated data set is then used as the input to the reconstruction algorithms developed in this thesis.

3.2.2 Triangulation Geometry

In the noise free case, *any* three views can be used to triangulate the knots and reconstruct the vertices of a polygonal object (see Section 2.2). When noise is present in the knot location measurements, however, the views that are chosen have a huge impact on the quality of the reconstruction. As an example of this, Figure 3-5(a) shows the triangulation (noise-free) of views taken in close angular proximity while Figure 3-5(b) shows the triangulation of widely separated angular views of the same object.

From these figures, it is apparent that if the same amount of noise were added to both reconstruction scenarios, the error would be much more significant in the case of the views taken in close proximity. Thus, the triangulation geometry is improved by taking widely

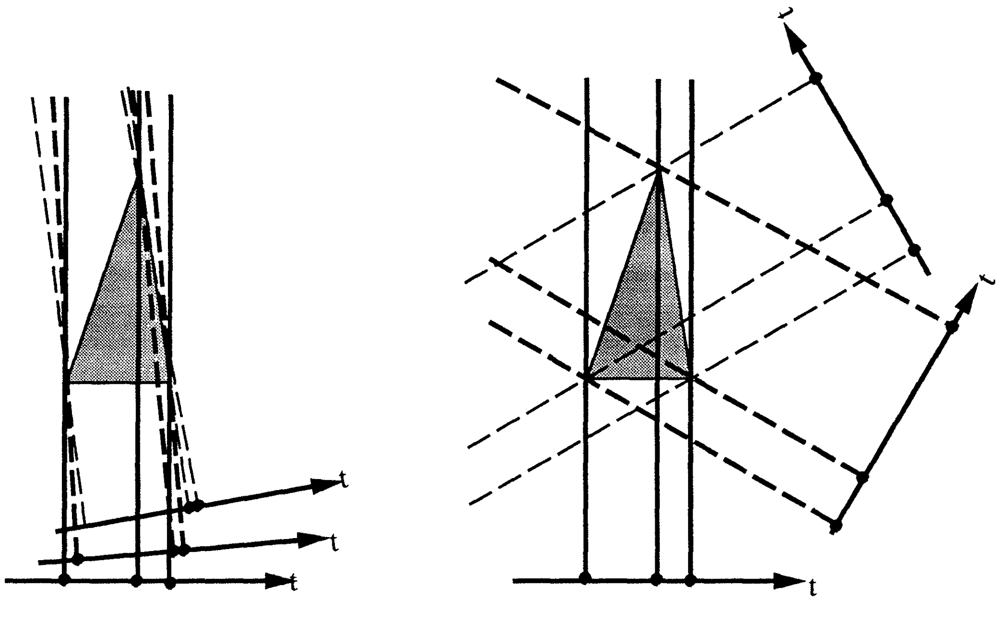

(a) (b)

Figure 3-5: Example of Triangulation Geometry of Views in Close Proximity Versus Views Taken at Wide Angular Spacing.

spaced angular views. From an algorithmic viewpoint, however, the advantage in considering views in close proximity is that the number of knot location switches that can occur for view-to-view is greatly reduced (i.e. the *gating* techniques described in Section 3.1.2); therefore, the number of hypothesized data associations is reduced and the overall complexity of the reconstruction algorithm is reduced. Because both geometries have advantages and disadvantages, Chapter 4 presents an algorithm that incorporates closely spaced views sequentially (poor triangulation geometry, reduced model complexity) while Chapter 5 presents an algorithm that incorporates views non-sequentially (better triangulation geometry, increased model complexity).

As a final note, another advantage of the algorithms developed in this thesis is that they allow reconstruction over any angular range with any angular spacing (i.e. non-uniform spacing is allowed). Most classical reconstruction techniques require views over a full angular range with uniform spacing. This advantage is important for applications such as non-destructive testing where views are sparse and angular separations are generally non-uniform.

3.2.3 Noise Models

In this section we resolve some of the issues concerning the assumptions of noise models used as input to the algorithms developed in this thesis.

Consistent with previous stochastic models of sinogram data [9, 12, 14], we assume throughout this thesis that the available sinogram data, or equivalently chord length data, are discrete samples of $g(t, \theta)$ corrupted by additive white Gaussian noise of known intensity. Specifically, for each chord our observation is given by

$$m_j(\theta_i) = g(t_j, \theta_i) + v(t_j, \theta_i), \qquad 1 \leq i \leq k, \;\; 1 \leq j \leq p \tag{3.1}$$

where k is the total number of projections, p is the number of chord measurements per view, and $v(t_j, \theta_i)$ are independent, identically distributed (i.i.d.) Gaussian random variables with known variance σ_m^2.

One difficulty in characterizing the amount of noise added to the chord data is that every chord in every view has a different magnitude. Thus, it is not possible to characterize the signal-to-noise ratio (SNR) easily. In this thesis we will use an approach similar to those used by Prince [12] and Milanfar [9][1]. We define the SNR per sample on the chord length data as

$$\text{SNR} = 10 \log \frac{\sum_{i,j} g^2(t_j, \theta_i)/d}{\sigma_m^2} \tag{3.2}$$

where $d = k \times p$ is the total number of chords length measurements and σ_m^2 is the variance of the i.i.d. noise v in the observations.

In order to completely characterize the noise on the knot locations, knowledge of the method used to extract the knots from the projection data is necessary. Assuming that this data extraction is complete, the knot extraction algorithm would provide the noise statistics of the knot measurements in an ideal reconstruction scenario. Because this algorithm would extract the knots directly from the projection data there would undoubtedly be a correlation between the noise on the knot locations and the noise on the chord length measurements.

---

[1] The only departure from the formulation used by Milanfar is that we use the traditional base 10 logarithm instead of the natural logarithm used by Milanfar

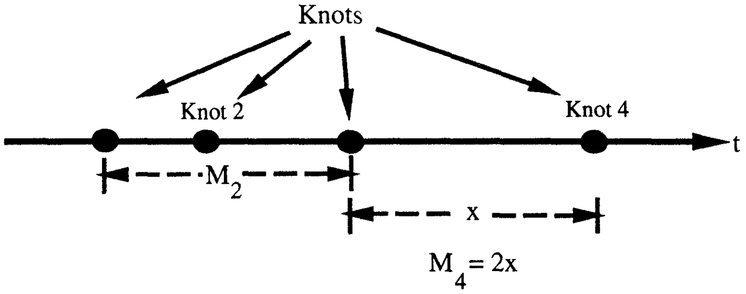

Figure 3-6: Example of the Values of M in the calculation of $\sigma^2_{z_j}(\theta_i) = \frac{K\sigma^2_m}{|\Delta s_j|^2 M_j}$.

In the absence of a specific knot extraction algorithm, we assume in this work that the noise on the knot measurements is additive Gaussian noise (for simplicity). To set the level of noise on each knot we use the following intuitive formulation

$$\sigma^2_{z_j}(\theta_i) = \frac{K\sigma^2_m}{|\Delta s_j|^2 M_j} \qquad 1 \leq i \leq k, \;\; 1 \leq j \leq n \tag{3.3}$$

where

$\sigma^2_{z_j}(\theta_i)$ = noise variance on knot location measurement $z_j(\theta_i)$
K = scaling constant
σ^2_m = noise variance on chord length measurements
$|\Delta s_j|$ = absolute change in slope at knot location measurement $z_j(\theta_i)$
M_j = distance measure between adjacent knot locations

and k and n are the total number of projections and number of knot locations per view, respectively. This equation can be interpreted as follows. The noise variance on knot j is proportional to the variance of the noise on the chords, denoted σ^2_m, to capture the idea that higher errors in the chord measurements will translate to higher errors in the knot location measurements during the knot extraction process. It is inversely proportional to the magnitude of the change in slope, $|\Delta s_j|$, at the knot location because a large change in slope should make it easier to extract the knot. Further, it is also inversely proportional to the distance M_j between the knot being extracted and the adjacent knots (for example, see M_2 in Figure 3-6). For the case of external knots, M_j is defined as twice the distance from the nearest knot (for example, see M_4 in Figure 3-6). The implication is that the greater the separation between adjacent knots, the easier it is to extract the knots.

To determine the value of the constant K we somewhat arbitrarily assign a value such that the average standard deviation of knot location noise terms is 1.5 times greater than the standard deviation of the noise on the chords

$$K(\frac{1}{d}\sum_{i,j}\sigma_{z_j}(\theta_i)) = 1.5\sigma_m \tag{3.4}$$

or

$$K = \frac{1.5\sigma_m d}{\sum_{i,j}\sigma_{z_j}(\theta_i)} \tag{3.5}$$

where $d = k \times n$ is the total number of knot location measurements. Although we have attempted to base this noise model on estimation principles, this model is admittedly somewhat arbitrary. Some adjustment to the knot location noise model may be required with the application of a specific knot extraction algorithm. The idea here is to get a rough approximation of the relative performance of the algorithms.

The noise models presented here are used throughout this thesis for the generation of data. Reconstructions are performed for scenarios where these applied noise statistics are assumed known and also for cases where the noise model for the chords is known and the noise on the knots is assumed i.i.d. with the variance equal to $(1.5\sigma_m)^2$. The i.i.d. model is a non-realistic model in that it assumes no information is available from the knot extraction algorithm about the relative "goodness" of one knot measurement compared to another. From the statistically-based GLR knot-location method used by Prince [12] to detect the support knots from projection data, we know that in reality more precise statistical information would be available from the knot extraction algorithm. Therefore, we consider this i.i.d. model a worst case assessment of the algorithm's performance.

3.2.4 Initial View - A Max-Min Approach

A major concern for the algorithms developed in this thesis is the correct assessment of the data association of the ordered knot locations measurements to the object's vertices. If the algorithm wrongly assigns this association, the wrong hypotheses will be retained and the optimal solution will never be reached. In order to provide an initial view into the algorithm with a favorable triangulation geometry, the algorithm begins by searching through all possible views and starts the algorithm with the view that maximizes the minimum (Max-Min) separation between adjacent knot location measurements. Using the property of the Radon transform given by Equation (2.20), the views that occur prior to the Max-Min view are appended to the original data such that their new angular position is $\theta_{new} = \theta_{old} + 180^\circ$ and the knot location measurements and chord length measurements are altered to be consistent with the requirements of Equation (2.20). The hope in this approach is that the Max-Min view and the views in close proximity to it will provide a good triangulation geometry and the algorithm will start off by retaining the correct hypotheses. In addition, we note that if all hypotheses are retained at each step of the algorithm, the initial view is unimportant because the optimal hypothesis is never deleted from the possible hypotheses retained at each step of the algorithm.

3.2.5 Performance Measures

Two measures are used to evaluate the performance of the reconstruction algorithms. The first is the Hausdorff distance, a metric on the set of all convex sets, $\mathcal{K}$. The second is the Symmetric Difference measure, a geometrically based measure of the difference between two convex sets. Both offer insight into the quality of the reconstruction; yet, both have deficiencies in their ability to quantify the "closeness" between the true object and its reconstruction. In this section each method is presented and its merits and deficiencies are noted.

In the algorithms that follow, the reconstructions are performed by minimizing the L^2 norm on measurements in the Radon Space. Regardless, however, the minimization of reconstruction error in the object space is of direct interest in quantifying the quality of the reconstructed object. Although it sometimes produces ambiguous results, the Hausdorff distance is appealing because it is a metric on $\mathcal{K}$. The Symmetric Difference measure is attractive because it is a geometrically intuitive measure of the difference between two convex sets. The one disadvantage of the Symmetric Difference measure is that it is not a "metric" (or a true measure of the distance between two sets). Because both measures have benefits and deficiencies, we will incorporate both measures into our analysis of the reconstruction results thus allowing a more comprehensive analysis of the algorithms.

Hausdorff Distance

The Hausdorff distance is a *true* metric on the set of all convex sets $\mathcal{K}$ in the metric space X; it is attractive because it offers a mathematically based definition of the distance between two convex sets.

To describe the Hausdorff distance we first need to define a few terms of metric spaces. First, the notion of distance between two sets in a metric space X is defined as

$$d(F,G) = \inf\{\ d(x,y) \mid x \in F,\ y \in G\} \tag{3.6}$$

where F and G are subsets of the metric space. Further, a *ball* around a set is defined

$$B(F,\rho) = \{\ x \in X \mid d(x,F) \leq \rho\} \tag{3.7}$$

where $\rho \geq 0$. Figure 3-7(a) shows an example of a ball, $B(F, \rho_F)$, around the set F.

Using these definitions, the Hausdorff distance is defined

$$\Delta^H(F,G) = \inf\{\ \rho \mid F \subset B(G,\rho)\ \ and\ \ G \subset B(F,\rho)\} \tag{3.8}$$

An intuitive explanation of the Hausdorff distance can been seen in Figure 3-7. Figure 3-7(a) shows the minimum amount, ρ_F, that F would have to be "uniformly expanded" so that G could be contained in it. Similarly, Figure 3-7(b) shows the minimum amount, ρ_G, that G would have to be expanded to contain F. The Hausdorff distance is the maximum of these two "minimum inflation factors".

The primary deficiency of this metric is that it can produce non-intuitive results for some convex sets. An example of such a result is shown in Figure 3-8. Figure 3-8(a) shows the Hausdorff distance between the two sets F and H. Similarly, Figure 3-8(b) shows the Hausdorff distance between the same set F and a new set G. Although G and H are significantly different to the eye, both have the same Hausdorff distance when compared with the set F. Thus, this metric can be somewhat misleading in quantifying the quality of a reconstruction. Therefore, a second measure of distance is introduced to provide a more geometrically intuitive performance measure.

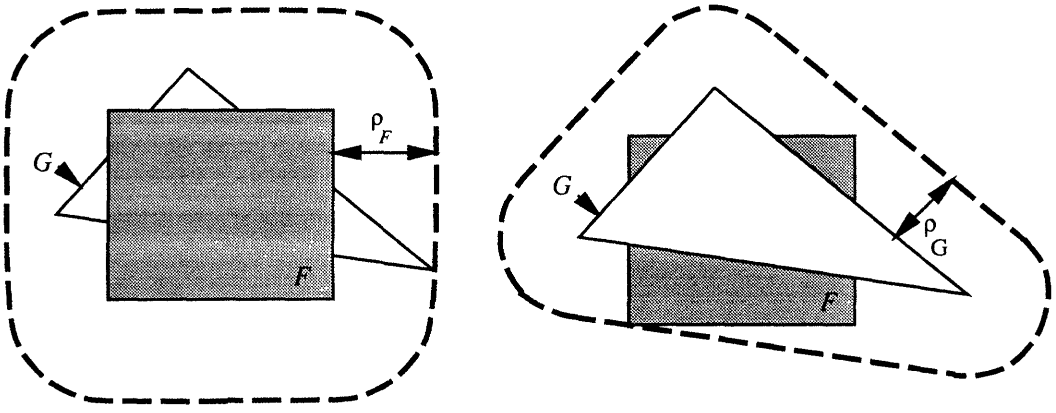

(a) (b)

Figure 3-7: Example of the Hausdorff Distance.

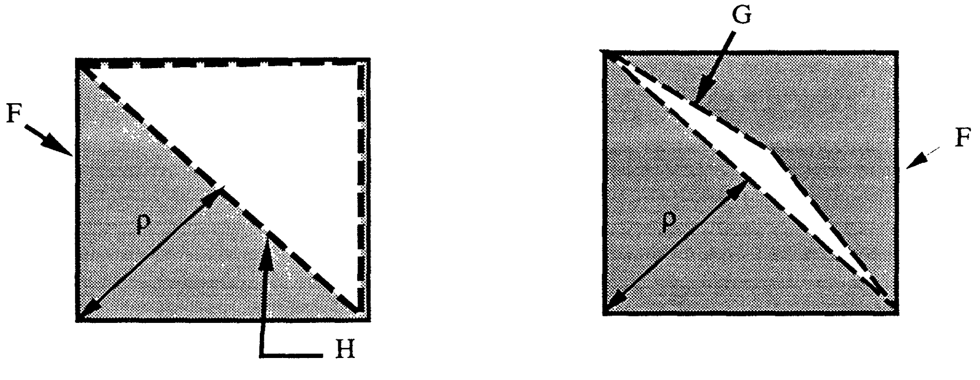

(a) (b)

Figure 3-8: Hausdorff Distance Inconsistency.

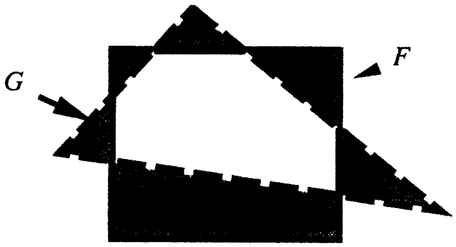

Figure 3-9: Example of the Symmetric Difference Measure.

Symmetric Difference Measure

The second method used to quantify the performance of the reconstruction algorithms is the Symmetric Difference measure. Unlike the Hausdorff distance which is a metric on $\mathcal{K}$, the set of all convex sets, the Symmetric Difference measure has no simple distance interpretation. Instead, it is a geometrically based measure on the difference between two convex sets. This measure is defined

$$\Delta^{S}(F, G) = \mu(F \Delta G) = \mu[(F \cup G)\ (F \cap G)] \tag{3.9}$$

where F and G denote convex sets and the operator $\mu(\cdot)$ denotes the area of the argument. Figure 3-9 shows two sets F and G and their corresponding Symmetric Difference measure.

A final note on the performance analysis is the manner in which we will use these distance measures. In order to make reconstruction results comparable, the measurements are generated as a percent. For the Hausdorff distance the percent Hausdorff error is defined

$$\%\ \text{Hausdorff error} = \frac{\Delta^{H}(\mathcal{S}, \widehat{\mathcal{S}})}{\Delta^{H}(\mathcal{S}, \emptyset)} \times 100\% \tag{3.10}$$

where $\mathcal{S}$ corresponds to the true object, $\widehat{\mathcal{S}}$ corresponds to the reconstruction, and $\emptyset$ denotes the set composed of the single point at the origin. Similarly, the Symmetric Difference error is defined

$$\%\ \text{Symmetric Difference error} = \frac{\Delta^{S}(\mathcal{S}, \widehat{\mathcal{S}})}{\mu(\mathcal{S})} \times 100\% \tag{3.11}$$

These error measures provide a means to quantify the quality of the reconstructions for the algorithms developed in Chapters 4, 5, 6, and 7.

Chapter 4

Sequential Reconstruction Algorithm

With this chapter we begin a sequence of three chapters dedicated to the reconstruction of the vertices of two dimensional (2D) binary polygonal objects from knot location and chord length data. We start in this chapter with the most structured of the three algorithms: an estimation algorithm that incorporates knot location measurements sequentially, with increasing angular position. Chapter 5 extends the algorithm to incorporate views non-sequentially. Finally, in Chapter 6 a full nonlinear estimation algorithm that reconstructs polygonal objects from knot locations and chord length measurements is developed. In all three of these chapters the number of vertices is assumed known and each knot location measurement corresponds to exactly one vertex in the object. Finally, Chapter 7 deals with the problems of inconsistent data measurements (missed knot detections).

This chapter begins in Section 4.1 with the presentation of the methodology and overview of the Sequential Reconstruction Algorithm. The next three sections detail each of the three distinct stages of the algorithm. Specifically, we present data association in Section 4.2, vertex estimation in Section 4.3, and data processing in Section 4.4. Finally, sample reconstructions and Monte-Carlo simulations are presented in Section 4.5 and conclusions in Section 4.6.

4.1 Overview of the Algorithm

This chapter is concerned with the reconstruction of binary polygonal objects from measurements of the knot locations of an n-gon (assuming that the number of sides is known). As outlined in Section 3.2.1, we assume that the projections are taken at known angular views and that the knot locations have been extracted from the projection data independently of this algorithm using a procedure such as that of Mier-Muth and Willsky [8]. Also, for simplicity we assume that all objects are binary and convex as discussed in Section 3.2.1. Throughout this chapter we will make extensive use of the definitions and terminology presented in Chapter 2 and problem formulation of Section 3.2.

The algorithm developed here incorporates the measurements of the knots at each view sequentially, with increasing angular position. The geometry of this problem can be viewed as a fixed (rigid) object with tomographic measurements taken at different views; equivalently, the problem can be viewed as measurements taken at a single, known viewpoint of an object rigidly rotating at a fixed velocity and viewed at equal time intervals (see [5]). By viewing the vertices with a fixed viewpoint geometry, the problem can be related to the single-sensor multitarget tracking problem (see section 3.1.2). In this framework, the association of knots to object vertices is analogous to the association of radar reports to targets. Further, the vertex locations of the tomographic reconstruction problem correspond to the target locations of the radar problem. Thus, like the radar multitarget tracking problem, this reconstruction problem is also developed by estimating both discrete and continuous values. Unlike the radar problem, there are a number of simplifying assumptions because of the geometry of the problem.

Figure 4-1 depicts the block diagram of the methodology used in this algorithm. Basically, the algorithm is a three stage process. In the first stage, all possible associations of knot location measurements to vertex locations are enumerated to form discrete-valued states or hypothesis. Once all of the possible data associations are determined, the second stage of the algorithm estimates the vertex locations for each discrete hypothesis under the assumption that it is the correct hypothesis. Finally, in the third step, the algorithm determines the likelihood that each hypothesis is correct given measurements of the knot locations *and* chord lengths. The likelihood ratios of these hypotheses are then used to prune the set of all possible discrete hypothesis for the next step of the algorithm and obtain the

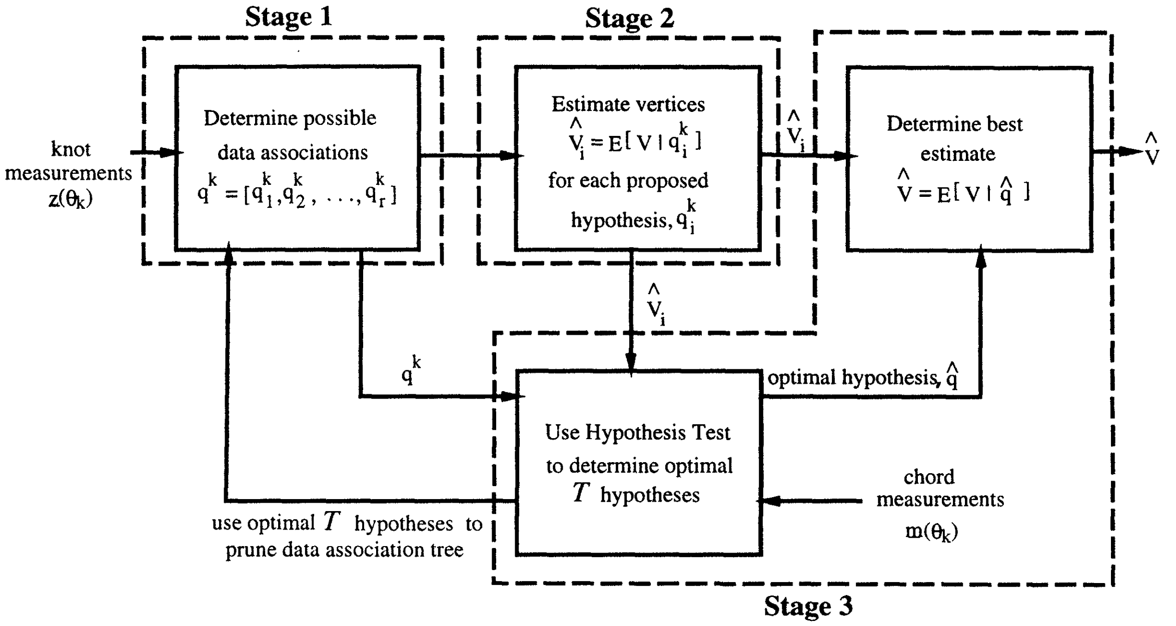

Figure 4-1: Block Diagram of Sequential Reconstruction Algorithm.

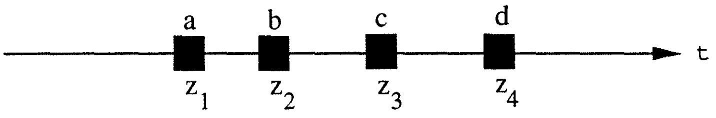

Figure 4-2: Knot Location Data.

optimal estimates of the vertices.

4.2 Part I : Data Association

In the problem of object reconstruction from knot locations, each knot corresponds to a vertex in the object. Prior to the reconstruction, however, this correspondence is unknown and the problem is one of data association. If the knot locations in the first view ($\theta_1 = 0°$) are associated to vertices a, b, c, and d as shown in Figure 4-2 (and denoted $abcd$), the goal is to determine the possible data association switches that can occur in these knots as the views are varied from $\theta_1 = 0°$ to $\theta_m = 180° - \Delta$ (assuming m projections are available and $\Delta = \frac{180°}{m}$).

Because of the geometry of the problem, there are a number of constraints in the allow-

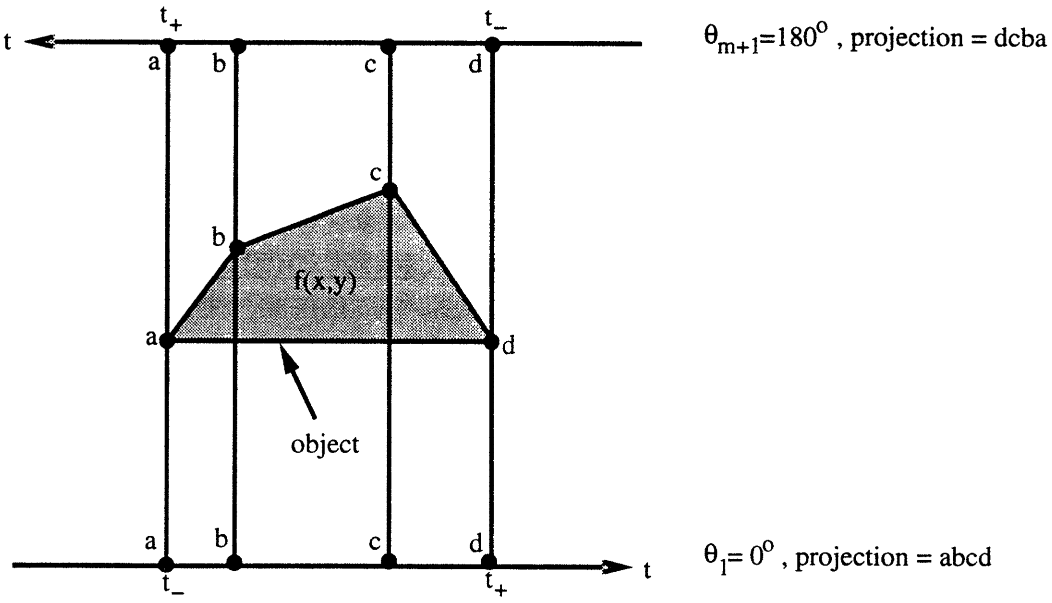

Figure 4-3: Relationship of Knot Location Data at $\theta_1 = 0°$ and $\theta = 180°$.

able knot switches that can occur if the knots are viewed as a continuous function of θ[1]:

A. As shown in Figure 4-3, if the knots in the projection at $\theta_1 = 0°$ are labeled *a*, *b*, *c*, and *d*, these knots occur in the reverse order in the final projection at $\theta_{m+1} = 180°$. Thus, for the case of a 4-gon a knot association hypothesis that starts with *abcd* must end with *dcba*.

B. A switch in two knots occurs when two vertices lie along the same line, perpendicular to the projection. Therefore, for an n-sided object, exactly $n(n-1)$ switches occur over a complete and continuous 360° range of projections. Further, because each switch of vertices occurs exactly 180° apart (because the vertices must lie along the same line), $\frac{n(n-1)}{2}$ switches occur in the range from $\theta_1 = 0°$ to $\theta_{m+1} = 180°$. Therefore, in a 4-gon exactly 6 knot location switches will occur.

C. Because two vertices have to be aligned on a line perpendicular to the projection for the knots to switch, it is not possible for two knots to switch twice in a row (i.e. you can't *undo* what was done in the previous step) with increasing θ.

---

[1] Note that the examples are given for the case of a 4-gon but can be naturally extended to the general case of an n-gon

Using these geometric constraints and the fact that the views input into this algorithm lie in close angular proximity, we assume for the Sequential Algorithm:

1. The following *gating* approach is used in this algorithm: from view-to-view the angular increments are small enough so that only adjacent knots can switch.
2. We have to account for the possibility that from one view to the next, no knots will change position.

One final note, for notational convenience we denote a switch in the associations of knots in positions j and $j+1$ by $(j, j+1)$. For example, the hypothesis $abcd \rightarrow bacd$ would be denoted $(1,2)$ because the knot associations in positions 1 and 2 have switched. Also, no change in the knot associations from one view to the next is denoted NC.

Given the above constraints, the enumeration of all possible knot association hypotheses can be shown in a tree. The root of the tree corresponds to the initial association of the knots to the vertices in the unknown object. We label the first measurement $abcd$ which denotes that knot location measurements $z_1 - z_4$ correspond to the unknown vertices $a - d$, respectively (for a 4-gon). Each branch of the tree corresponds to the inclusion of a new set of knot location measurements and gives the possible way in which the knot to vertex associations may have changed since the last measurement. An example of the hypothesis tree for a 4-gon is shown in Figure 4-4. We can see from this tree that the number of possible hypotheses grow significantly with added views, even given the geometric constraints of this problem. Further, the complexity of this tree increases substantially as the number of vertices in the object increases. A pruning algorithm to reduce the number of hypotheses is introduced in Section 4.4.

Using the notation of the *Multitracker* algorithm developed by Kurien [6] we let q_i^k denote one particular association or hypothesis of the first k views of knot location measurements to the vertices of the object. Each potential hypothesis for the object, q_i^k, is represented on the hypothesis tree by a trace of successive branches from the root of the tree to a branch representing the incorporation of the kth measurement. For example, the path traced out in dashed lines in Figure 4-4 represents the hypothesis q_1^3 (assuming that the $r = 17$ potential hypotheses are labeled from the top down), which corresponds to the knot location hypothesis $abcd \rightarrow badc \rightarrow badc$. Therefore, each sequence q_i^k gives one possible

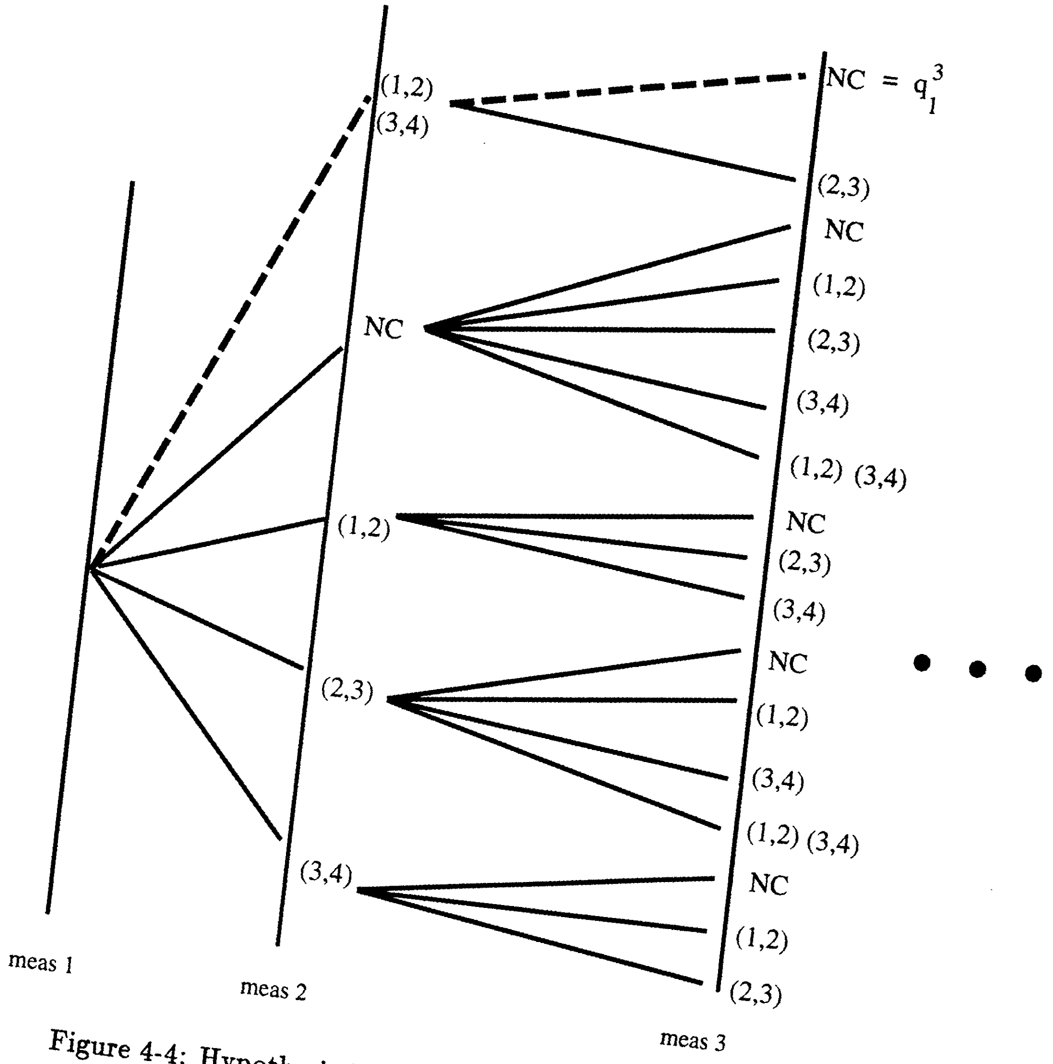

Figure 4-4: Hypothesis Tree for Sequential Reconstruction Algorithm.

interpretation of the association of the knot location measurements to the vertices of the object for the first k views input into the reconstruction algorithm. Further, we denote the set of *all* possible data associations up to measurement k by the set $q^k = [q_1^k,\ q_2^k,\ \ldots,\ q_r^k]$.

4.3 Part II: Estimation of the Vertices

The objective of this algorithm is to estimate the vertices of an object from measurements of its projections at known angular views. We denote the vertex locations, V, by the following matrix

$$V = \begin{bmatrix} x_1 & x_2 & \cdots & x_n \\ y_1 & y_2 & \cdots & y_n \end{bmatrix}$$

For each view θ_j, the assumed model relating the vertices to the knot location data is

$$\underline{z}(\theta_j) = \underline{h}(V, \theta_j) + \underline{n}(\theta_j) \tag{4.1}$$

where

$$\begin{aligned} \underline{z}(\theta_j) &= [z_1(\theta_j)\ z_2(\theta_j)\ \ldots\ z_n(\theta_j)]\,, \qquad z_1(\theta_j) \leq z_2(\theta_j) \leq\ \ldots\ \leq z_n(\theta_j) \\ &= \text{row vector of noisy, ordered knot location measurements} \\ V &= \text{vertices of the object(defined above)} \\ \underline{h}(V, \theta_j) &= \text{nonlinear function of the vertices} \end{aligned}$$

and $\underline{n}(\theta_j) \sim N(\underline{0}, \Lambda_j)$ is a vector independent Gaussian noise process with the variance on each knot as defined in Section 3.2.3 so that Λ_j is a diagonal matrix. The vector $\underline{h}(V, \theta_j)$ can be represented by

$$\underline{h}(V, \theta_j) = \text{sort}(\underline{c}(\theta_j) V) \tag{4.2}$$

where $\underline{c}(\theta_j) = [\cos(\theta_j),\ \sin(\theta_j)]$ is the projection matrix at angle θ_j, $\underline{c}(\theta_j)V$ is the projection of the vertices, and sort$(\cdot)$ is an operation that sorts the vertex projections in order of increasing value. Thus, the problem is a nonlinear function of the vertex locations.

If the measurements at each view are combined (stacked), the overall measurement equation becomes

$$\begin{aligned} Z^k &= [\underline{z}^T(\theta_1) \mid \underline{z}^T(\theta_2) \mid \ldots \mid \underline{z}^T(\theta_k)]^T \\ &= H(V) + \mathcal{N} \end{aligned} \tag{4.3}$$

where

$$\begin{aligned} H(V) &= [\underline{h}^T(V, \theta_1) \mid \underline{h}^T(V, \theta_2) \mid \ldots \mid \underline{h}^T(V, \theta_k)]^T \\ \mathcal{N} &= [\underline{n}^T(\theta_1) \mid \underline{n}^T(\theta_2) \mid \ldots \mid \underline{n}^T(\theta_k)]^T \end{aligned}$$

The objective of this algorithm is to estimate V given Z^k and the noise statistics. Because H is a nonlinear function of V, this problem is inherently nonlinear.

The formulation proposed in this chapter is a hypothesis testing algorithm which will linearize the estimation problem with respect to V at each stage by assuming that the association of the knot locations to the vertices is known. Later, in Chapter 6, we return to the full nonlinear problem and consider a nonlinear approach to estimating the vertices from Equation (4.3).

If we now assume that the knot associations q_i^k are known, Equation (4.1) can be written as:

$$\underline{z}(\theta_j) = \underline{c}(\theta_j) V P_j(q_i^k) + \underline{n}(\theta_j) \tag{4.4}$$

where

$$\begin{aligned} \underline{c}(\theta_j) &= [\cos(\theta_j),\ \sin(\theta_j)] \\ P_j(q_i^k) &= \text{permutation matrix dependent on the data associations} \end{aligned}$$

The term $\underline{c}(\theta_j)V$ is the location of each vertex in the projection (i.e. the knot locations) while $P_j(q_i^k)$ is a permutation matrix that orders the knot locations based on the hypothesized set of data associations given by q_i^k. The correspondence of the knot measurements, $\underline{z}(\theta_j)$, to the vertices of the object is lost without knowledge of P_j. This equation is now linear with respect to the vertices, V, and an estimate of V can be obtained using ML estimation techniques.

First, we post-multiply each measurement equation by the appropriate permutation matrix

$$\underline{z}(\theta_j) P_j^{-1}(q_i^k) = \underline{c}(\theta_j) V + \underline{n}(\theta_j) P_j^{-1}(q_i^k) \tag{4.5}$$

or

$$\underline{\tilde{z}}(\theta_j) = \underline{c}(\theta_j) V + \underline{\tilde{n}}(\theta_j) \tag{4.6}$$

where

$$\begin{aligned}
\tilde{\underline{z}}(\theta_j) = \underline{z}(\theta_j)P_j^{-1}(q_i^k) &= [\tilde{z}_1(\theta_j),\ \tilde{z}_2(\theta_j), \ldots,\ \tilde{z}_n(\theta_j)] \\
\tilde{\underline{n}}(\theta_j) = \underline{n}(\theta_j)P_j^{-1}(q_i^k) &= [\tilde{n}_1(\theta_j),\ \tilde{n}_2(\theta_j), \ldots,\ \tilde{n}_n(\theta_j)]
\end{aligned}$$

Note that P_j is always invertible because it is a permutation matrix. Recalling that each column of V is a vertex of the object, we can separate the problem into n separate estimation problems (one for each vertex) and estimate each vertex, $\underline{V}_m$, separately using the ML estimate equation defined in Section 2.1.2. Therefore, the measurement equation for <u>each vertex</u> becomes

$$\underline{Z}_m = C\underline{V}_m + \underline{\mathcal{N}}_m \tag{4.7}$$

where

$$\begin{aligned}
\underline{Z}_m &= [\tilde{z}_m(\theta_1),\ \tilde{z}_m(\theta_2), \ldots,\ \tilde{z}_m(\theta_k)]^T \\
C &= [\underline{c}(\theta_1)^T,\ \underline{c}(\theta_2)^T, \ldots,\ \underline{c}(\theta_k)^T]^T \\
\underline{V}_m &= [x_m\ y_m]^T \\
\underline{\mathcal{N}}_m &= [\tilde{n}_m(\theta_1),\ \tilde{n}_m(\theta_2), \ldots,\ \tilde{n}_m(\theta_k)]^T
\end{aligned}$$

and the ML estimate of $\underline{V}_m$ is calculated using Equation (2.14) from Section 2.1.2

$$\hat{\underline{V}}_m(\underline{Z}_m) = (C^T R_m^{-1} C)^{-1} C^T R_m^{-1} \underline{Z}_m \tag{4.8}$$

where the covariance matrix R_m is obtained by taking the appropriate terms from each diagonal Λ_j matrix (since the noise is uncorrelated and the matrix P_j simply changes the ordering of the knot measurements).

4.4 Part III: Data Processing

Because of the enormous number of possible hypotheses that occur as views are added, it is necessary for computational purposes to *prune* the number of possible hypotheses as new data are added. Further, we also require a procedure to determine the best data association hypothesis and thus the optimal estimate. For these reasons, a method of evaluating each hypothesis is incorporated into the algorithm. The quality of each proposed hypothesis is evaluated using the ordered knot location measurements Z^k, the ML estimate $\widehat{V}$ from Part II, and chord length measurements collected over the same angular range.

4.4.1 Modeling of Chord Length Measurements

The chord length measurements as a function of the vertices, V, for each view θ_j are modeled by:

$$\underline{m}(\theta_j) = \underline{f}(V, \theta_j) + \underline{v}(\theta_j) \tag{4.9}$$

where $\underline{m}(\theta_j)$ is the vector of chord length measurements at angle θ_j, $\underline{f}(V, \theta_j)$ is a nonlinear function relating the chord measurements and the vertices of the object, and $\underline{v}(\theta_j) \sim N(\underline{0}, \sigma_m^2 I)$ is independent Gaussian noise as defined in Section 3.2.3.

The measurements at each view are combined (stacked) to form a single chord length measurement equation

$$M^k = F(V) + \Upsilon \tag{4.10}$$

where

$$\begin{aligned} F(V) &= [\underline{f}^T(V,\theta_1) \mid \underline{f}^T(V,\theta_2) \mid \ldots \mid \underline{f}^T(V,\theta_k)]^T \\ \Upsilon &= [\underline{v}^T(\theta_1) \mid \underline{v}^T(\theta_2) \mid \ldots \mid \underline{v}^T(\theta_k)]^T \end{aligned}$$

Unlike the knot location data, the matrix $\underline{f}(V, \theta_j)$ does not become a linear relationship in V once the data associations are known[2]. Thus, it is difficult to estimate the vertices V even given the data associations (unlike the knot locations). However, the chord length data can be used effectively with the knot locations to help prune the possible hypotheses as described in the following section.

---

[2] The form of this function is rather complicated to describe. Rather than give an explicit formula, we include an explanation and the code used to calculate the chord measurements as a function of the vertex locations in Appendix A.

4.4.2 Hypothesis Test

If the measurements of the knot locations and chord lengths are combined to form the vector $X^k = [(Z^k)^T \mid (M^k)^T]^T$, the hypotheses can be evaluated using a form of the M-ary ML Generalized Likelihood Decision Rule (refer to Section 2.1.3)

$$\begin{aligned} d(X) &= q_i^k \text{ if for all } q_i^k \neq q_j^k \\ \log[p_{x|\,q_i^k,\widehat{V}}(X^k|\,q_i^k,\widehat{V}(Z^k|\,q_i^k))] &\geq \log[p_{x|\,q_j^k,\widehat{V}}(X^k|\,q_j^k,\widehat{V}(Z^k|\,q_j^k))] \end{aligned} \qquad (4.11)$$

where $\widehat{V}(Z^k|q_i^k)$ denotes the ML estimate of the vertices of the object given hypothesis q_i^k and knot measurements Z^k. The primary departure of this decision rule from the standard Generalized Likelihood Ratio Test is that the unknown non-random quantity, V, is estimated from a subset of the total measurements (i.e. from the knot location measurements but not from the chord measurements).

After evaluating the hypotheses using the above decision rule the top T hypotheses are retained and the remainder are discarded. In the next step of the algorithm the possible data associations for the new data are only included for these T hypotheses and the other hypotheses are effectively *pruned* from the tree. In the final step of the algorithm, the optimal hypothesis $\hat{q}$, as determined by the above decision rule, is retained as the final estimate of the true data associations. Similarly, the optimal estimate of V is the ML estimate given that $\hat{q}$ is the true set of data associations. Experimental results of this algorithm follow in the next section.

4.5 Experimental Results

4.5.1 Test Objects

Throughout the analysis of the reconstruction algorithms developed in this thesis, two binary test objects are used to characterize the performance of the algorithms. Both objects are 4-gons that are contained within the region defined in Cartesian coordinates by the four points $(1,1), (1,-1), (-1,-1)$, and $(-1,1)$ (or in polar coordinates by a *ball* of radius $\rho = \sqrt{2}$ centered at the origin). The first object is the "kite" shown in Figure 4-5. The kite was chosen because of its contrasting projection width as the projection angle is varied over the

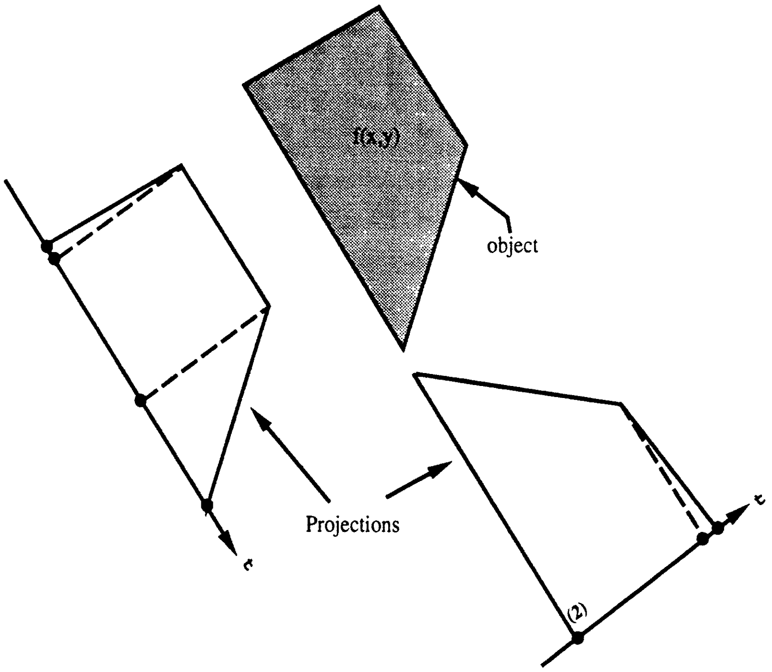

Figure 4-5: Kite Test Object.

$[0^\circ, 180^\circ)$ allowable range. When a projection is taken along a line parallel to its longest side, the distance between the two support knots is roughly twice that of the distance between the support knots taken of a projection taken at 90° rotation (shown in Figure 4-5). In contrast, for our second object we chose the square shown in Figure 4-6. At its minimum, the distance between the support knot locations is .701 times the distance at its maximum (refer to Figure 4-6). Thus, although the internal knots are allowed to move arbitrarily close together as the view changes, the width of the field containing the knots does not vary greatly (unlike the kite). By exploring sample reconstructions of both objects, we hope to reveal both strengths and weaknesses in each reconstruction algorithm.

4.5.2 Sample Reconstructions

Figure 4-7 shows sample reconstructions of the two test objects from 27 uniformly spaced angular views over the range $[0^\circ, 180^\circ)$. Each view contains 5 chord measurements with the SNR levels on the chord measurements set at 10dB. The chord measurements are uniformly sampled over the range defined by $t = [-1, 1]$. Note that this results in views where some of the chord measurements are zero (i.e. the object is not in the field of view) and conversely, in views where the chord measurements are confined to the interior of the object (for the

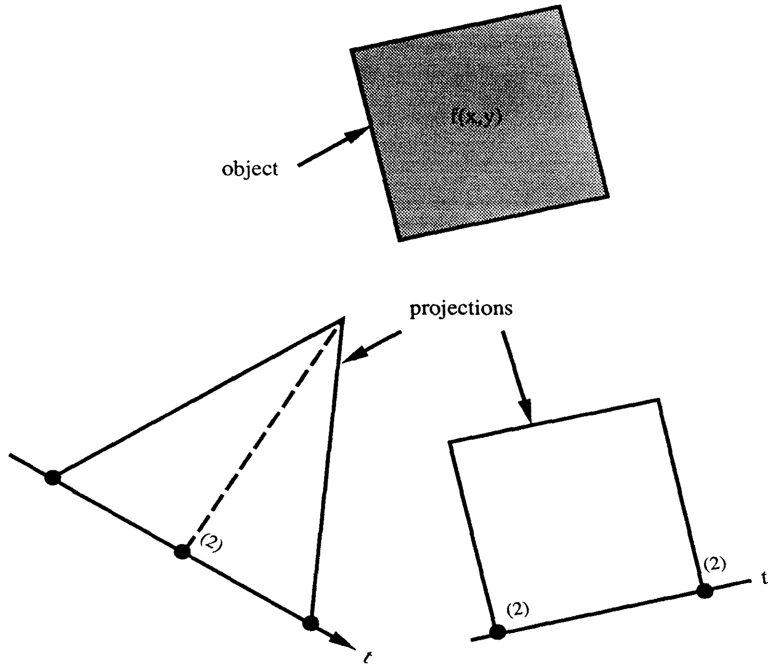

Figure 4-6: Square Test Object.

Error Measures	Kite Test Object		Square Test Object	
	True Model	i.i.d. Model	True Model	i.i.d. Model
Hausdorff Error (%)	16.22	36.67	32.16	44.61
Symm. Diff. Error (%)	19.35	38.42	17.71	33.62

Table 4.1: Comparison of Errors for Sample Reconstructions Using Sequential Reconstruction Algorithm.

kite object). Note also that T=10 hypotheses were retained in each step of the algorithm. Given the above description, there are a total of 108 knot location measurements and 135 chord measurements used in each reconstruction. In both Figure 4-7(a) and Figure 4-7(b), there are two reconstructions. The first one, denoted by the dashed lines is a reconstruction in which the correct applied noise model was assumed known. In the second reconstruction, denoted by dotted lines, the assumed model for the knot variances is an i.i.d. noise model with $\Lambda_Z = (1.5\sigma_m)^2 I$ where σ_m^2 is the noise variance on the chord length measurements[3]. Table 4.1 shows a comparison of both the Hausdorff and Symmetric Difference errors of these two reconstructions.

---

[3] Refer to Section 3.2.3 for a complete description of the noise models used in this thesis.

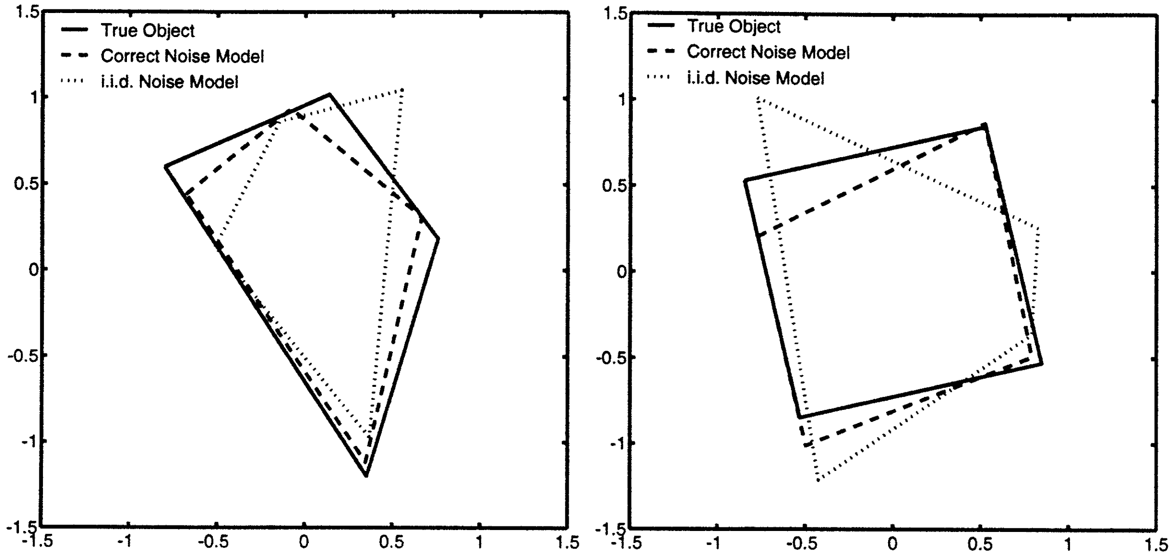

(a) Kite Reconstruction (b) Square Reconstruction

Figure 4-7: Sample Reconstructions of Sequential Algorithm: 27 views, 10 hypothesis pruning, 10dB SNR on chords.

The relative performance of both reconstructions was consistent. In both cases, the reconstruction that assumed the correct applied noise model for the knot locations outperformed the reconstruction that assumed the i.i.d. noise model. As discussed in Section 3.2.3, the i.i.d. model is viewed as a worst case performance assessment and therefore the higher errors are expected. We note that the Non-Sequential Algorithm presented in Chapter 5 deals more robustly with this noise model and the Nonlinear Algorithm presented in Chapter 6 significantly reduces the error due to the i.i.d. model assumptions.

Also, it is interesting to note that the Symmetric Difference errors were larger than the corresponding Hausdorff error in the reconstructions of the kite; while in the reconstructions of the square, the Hausdorff errors were larger than the corresponding Symmetric Difference errors. This occurs because of the different criteria used to quantify the distance between two convex sets employed by these two error measures (refer to Section 3.2.5 for a complete description of these error measures). Although the reconstruction errors of the kite were less than those of the square for most of the error measures in Table 4.1, this result is only for one possible sample path. We note this behavior but defer any comparative analysis of the algorithm in reconstructing these two objects to the Monte-Carlo analysis of Section 4.5.3.

4.5.3 Monte-Carlo simulations

In the following sections Monte-Carlo simulations of reconstructions are used to test the performance of the Sequential Reconstruction Algorithm. With the exception of the last section which compares the performance of reconstructions of both the kite and square test objects, all of the simulations are performed using the kite test object. The Monte-Carlo simulations consist of 100 independent reconstructions for each reconstruction scenario. Additionally, in each plot error bars denote the 95% confidence intervals of the sample mean values that result from the 100 runs of the algorithm. Again, as defined previously, all chord length data for each projection were uniformly sampled over the region $t = [-1, 1]$. *Unless otherwise stated*, $T = 15$ hypotheses were retained at each step of the algorithm, 5 chord measurements were taken per view, and the SNR was set to 10dB on the chord length measurements. Reconstructions are done for cases that assume the correct applied model for the noise on the knot location data and also for cases that assume an i.i.d. noise model such that $\Lambda_z = (1.5\sigma_m)^2 I$ where σ_m^2 is the noise on the chord measurements (refer to Section 3.2.3 for a detailed description of the noise models used in this thesis). Finally, in all of the simulations the projection data used to perform the reconstructions are uniformly spaced samples over the angular region $[0°, 180°)$.

Reconstruction Error as a Function of Retained Hypotheses

Figure 4-8 is a plot of the Hausdorff and Symmetric Difference errors of the kite reconstruction as a function of the number of hypotheses, T, retained in each step of the algorithm. Each reconstruction consists of 18 views and reconstructions are performed for $T = 5$, 10, 15, and 20 retained hypotheses. As described previously, reconstructions are done assuming both the correct applied noise model and the i.i.d. model.

As expected, as the number of retained hypotheses, T, is increased, the reconstruction results improve. If the number of retained hypotheses is small, the algorithm cannot correct itself after a sequence of poor projection data because the true hypothesis will already have been pruned from the hypothesis tree. Conversely, if the algorithm where to retain all hypotheses, the optimal estimate (as defined by the cost criterion used to prune the hypotheses) would always be obtained. We also note that the reconstruction error drops most dramatically when the number of hypotheses is increased from 5 to 10. This error

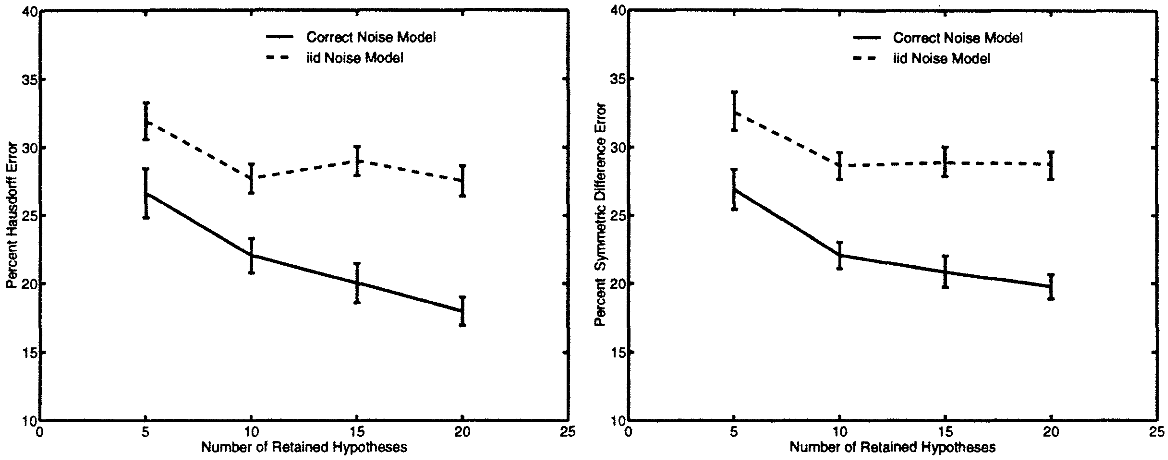

(a) Hausdorff Error

(b) Symmetric Difference Error

Figure 4-8: Performance as a Function of Retained Hypotheses Used in Sequential Reconstruction Algorithm.

reduction is greater than subsequent decreases (as the number of retained hypotheses is increased) because each time that T is increased the probability that the optimal hypotheses or one close to it will be included in the retained hypotheses also increases (i.e. this is a case of diminishing returns).

As expected, reconstructions based on the correct applied knot noise model outperform the reconstructions that assume an i.i.d. knot noise model. Another interesting implication of Figure 4-8 is that the i.i.d. model seems to level off much more quickly as a function of the retained hypotheses; implying that retaining more hypotheses will not obtain a better estimate. This behavior can be attributed to the uniform weighting, via the inverse of the covariance matrix of the assumed noise model, of all knot data in the cost function regardless of the amount of noise actually present in each knot measurement(i.e. a least squares cost criterion versus a weighted least squares cost criterion).

Reconstruction Error as a Function of Chord Measurements

In this section we analyze the effect of the number of chord measurements used to prune the hypotheses in the Part III of the Sequential Algorithm. Each reconstruction is based on 18 uniformly spaced views or a total of 72 knot location measurements. We set the variance on the chord length measurements to a constant value throughout this experiment. Thus,

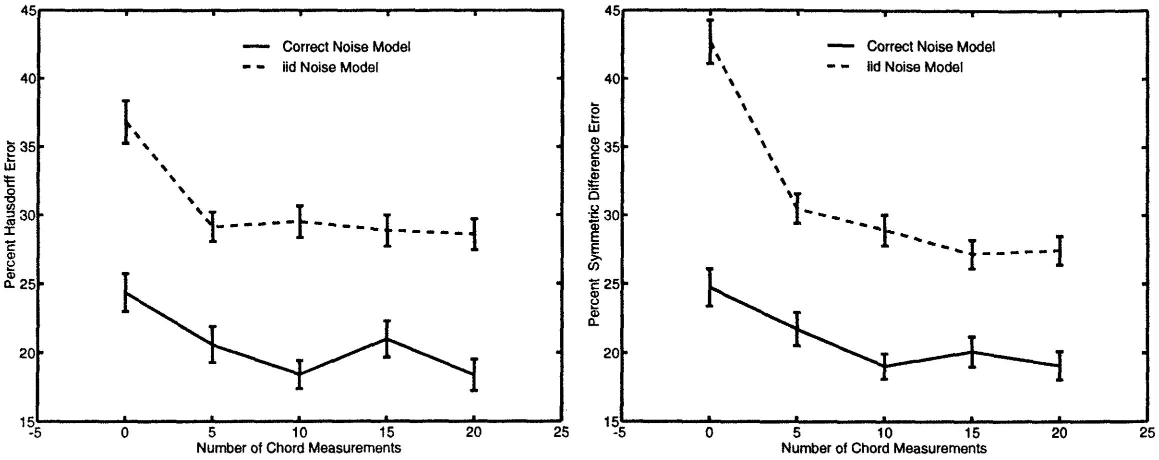

(a) Hausdorff Error

(b) Symmetric Difference Error

Figure 4-9: Performance as a Function of Number of Chords Used in the Sequential Reconstruction Algorithm.

as the number of chord length measurements per view is varied, the SNR (as defined in Section 3.2.3) will also change. The constant variance on the chord lengths was chosen for the case of 10 chord measurements per view at a SNR of 10dB and the noise on the knot locations was set with Equation (3.3) using this constant chord variance. Noisy chord data were generated for 0, 5, 10, 15, and 20 chord measurements per view. The simulation results are shown in Figure 4-9. Figure 4-9(a) presents the Hausdorff error results while Figure 4-9(b) presents the Symmetric Difference error results.

The overall behavior of the curves in Figure 4-9 is as expected. The reconstruction error for both measures decreases as the number of chords used in the algorithm is increased. Further, the most dramatic reduction in error occurs as the number of chords is increased from 0 to 5 measurements per view. This lends credibility to the fact that the geometric information contained in the chord data is an essential part of this reconstruction algorithm. As expected, the reconstructions that assume the correct applied noise model (shown in the solid lines) consistently have a lower error. An interesting observation regarding the Symmetric Difference error results for the i.i.d. noise model is that the Symmetric Difference error is significantly larger than the Hausdorff error when the number of chords measurements is small (refer to Figure 4-9). This implies that without an accurate knot location noise model (as in the case of the i.i.d. model), the chord measurements are extremely important in reducing the Symmetric Difference error between the reconstruction and the true object.

Reconstruction Error as a Function of SNR and Number of Views

This section presents the results of a series of independent Monte-Carlo simulations performed at various SNRs for different numbers of views (refer to Figure 4-10) to characterize the effect of these parameters on the Sequential Algorithm. Reconstructions are performed for SNRs of 0dB, 10dB, 20dB, and 30dB and numbers of views equal to 9, 18, 27, and 54. Figure 4-10(a) and (b) present the Hausdorff and Symmetric Difference errors, respectively, for reconstructions assuming the correct applied knot location noise model. Similarly, Figure 4-10(c) and (d) present the same error measures for reconstructions assuming an i.i.d. noise model for the knot location data.

There are a number of observations consistent with all of the reconstruction results presented in Figure 4-10. First, as the SNR increases the percent error decreases (as expected). The most dramatic decrease in error occurs when the SNR is increased from 0dB to 10dB. This result is attributed to the algorithm's inability to resolve the knot-to-vertex data associations in a high noise scenario. Another result that is common to all reconstructions is that the relative shape of the resulting curves at each SNR is the same. Thus, at a given SNR all of the reconstruction errors behave similarly as a function of the number of views even if the relative level is different. In addition, although the errors for both measures decrease initially as the number of views is increased, in all cases (except at SNR=0dB) the error increases as the number of views becomes large. Although this result may seem counter-intuitive at first, it can be explained by the method used to reconstruct the objects in this algorithm. As the number of views increases, the angular separation between views becomes small. From one data set to the next, the views change very little. When the knot locations are close in a particular view and noise is added to the system, the algorithm may retain the wrong hypotheses. If the next view has the same data quality (which is expected if the views are taken very closely together), the wrong hypotheses continue to be retained. If enough "bad" views are added in a sequence, the optimal hypothesis may be pruned and discarded, lost to the algorithm forever. Thus, the algorithm returns a sub-optimal estimate such as the results of Figure 4-10 show for large numbers of views. Figure 4-11 demonstrates

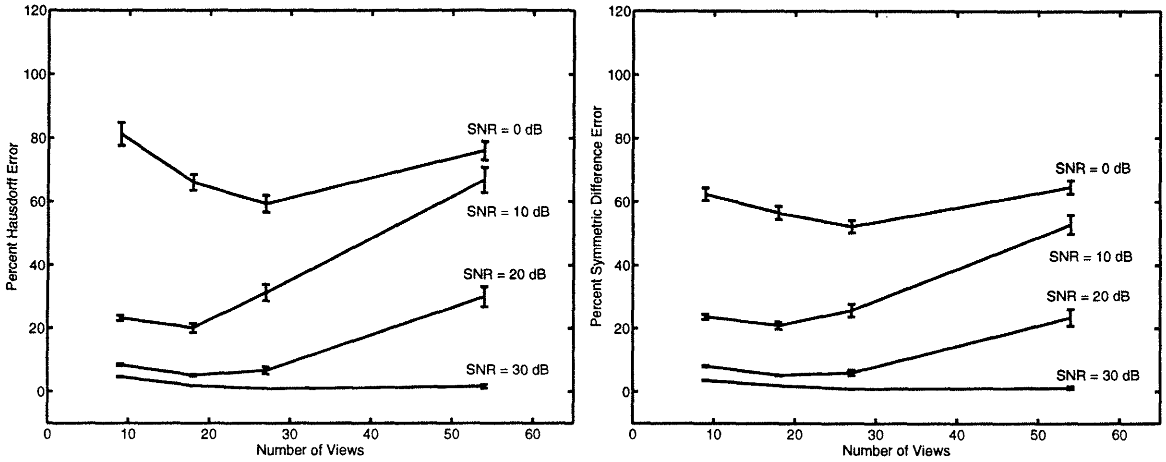

(a) Hausdorff Error, Correct Model (b) Symm. Diff. Error, Correct Model

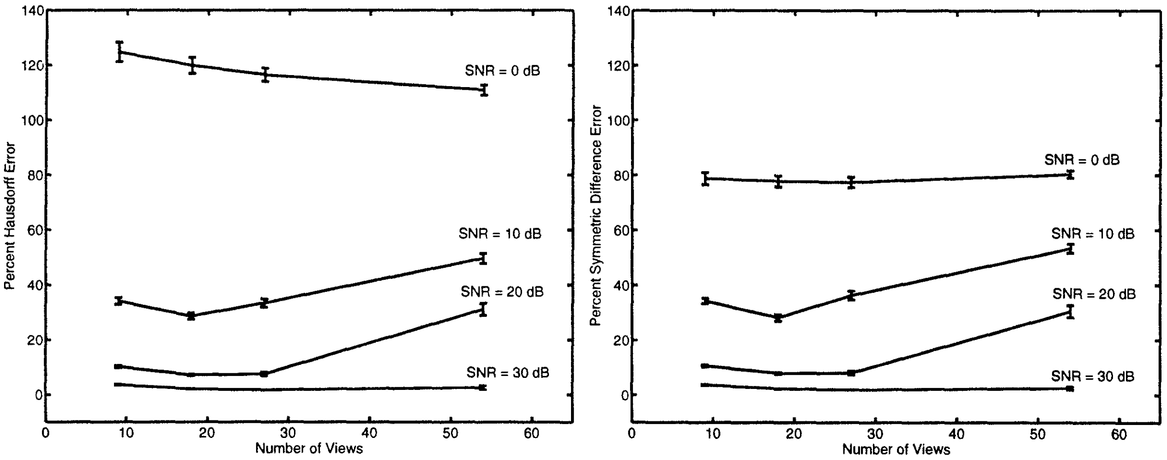

(c) Hausdorff Error, i.i.d. Model (d) Symm. Diff. Error, i.i.d. Model

Figure 4-10: Performance as a Function of Number of Views Used in the Sequential Reconstruction Algorithm. (a),(b) Correct Noise Model; (c),(d) i.i.d. Noise Model

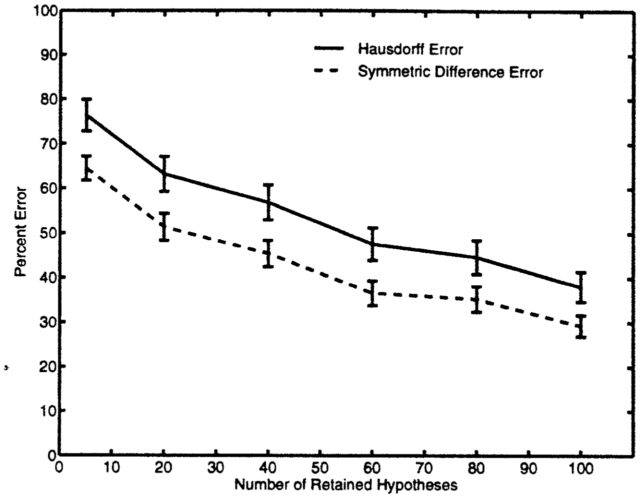

Figure 4-11: Sequential Algorithm as a Function of Retained Hypotheses: 54 views, SNR=10dB, 5 Chords, T=15 Hypotheses.

that if the number of retained hypotheses is increased the reconstruction error for scenarios of 54 views assuming the correct noise model is reduced dramatically. As a final note, we stress that the Non-Sequential Algorithm of Chapter 5 corrects the problem of a sequence of poor views by reordering the angular views before they are introduced into the algorithm (at the expense of model complexity).

Finally, we note some of the differences between the respective results of Figure 4-10. Consistent with previous sample reconstructions of Section 4.5.2, the i.i.d. noise model reconstructions produce a higher percent error for both the Hausdorff and Symmetric Difference measures. Note also that the results of the i.i.d. reconstructions tended to be flatter (i.e. did not vary as much) as a function of the number of views. Although the Hausdorff error and Symmetric Difference error curves seem to have the same relative shape in most cases, the Symmetric Difference error values where smaller. In addition, there are instances where the two measures produce conflicting results. For example in Figure 4-10(c) and (d) the Hausdorff error decreases as the number of views is increased from 27 to 54 views at SNR=0dB while the Symmetric Difference error increased slightly under the same conditions. These discrepancies can be attributed to the fact that each measure evaluates the difference between two convex sets using a different criterion (see Section 3.2.5). For example, in the sample reconstructions of Section 4.5.2, the i.i.d. kite reconstruction (in dotted lines in Figure 4-7(a) had a smaller Hausdorff error than the corresponding reconstruction error for the square, while the i.i.d. square reconstruction (in dotted lines in Figure 4-7(b)) had a smaller Symmetric Difference error. Although it is arguable which of these two reconstructions is "better", in most cases the Symmetric Difference error is a more intuitive measure.

Reconstruction Comparison of Two Test Objects

Having completed the Monte-Carlo analysis of the Sequential Algorithm for a particular object, the kite, we now turn to a comparative analysis. Figure 4-12 shows the Monte-Carlo reconstruction errors for both the kite and the square test objects as a function of the number of views used in the reconstruction. Reconstructions were performed for 9, 18, 27, and 54 views. Figure 4-12 (a) and (b) are Hausdorff and Symmetric Difference errors for reconstructions assuming the correct applied knot location noise model while Figure 4-12 (c) and (d) are the corresponding error measures under the assumption of the i.i.d. noise model.

With the exception of the Hausdorff error for the i.i.d. noise model (Figure 4-12(c)), the reconstruction of the kite seemed to slightly outperform that of the square for small numbers for views while the opposite was true for larger numbers of views. This result may be attributable to the sensitivity of some of the views of the kite to noise. Because the kite has views where the distance between the support knots is relatively small (refer to Section 4.5.1), this may result in a sequence of "poor" measurements that produce error in the reconstruction at large numbers of views. In the case of the Hausdorff error of the i.i.d. model reconstruction of Figure 4-12(c), the kite's reconstruction errors are less than the square reconstruction errors for all numbers of views. These results conflict slightly with the corresponding Symmetric Difference error analysis of Figure 4-12(d). Again, these discrepancies are attributed to the different performance criteria of the two measures. Despite these small differences, however, all of the curves have the same general shape. In addition, the reconstruction errors of both objects are on the same order of magnitude. Therefore, the algorithm does not seem to produce significantly better reconstructions for either object.

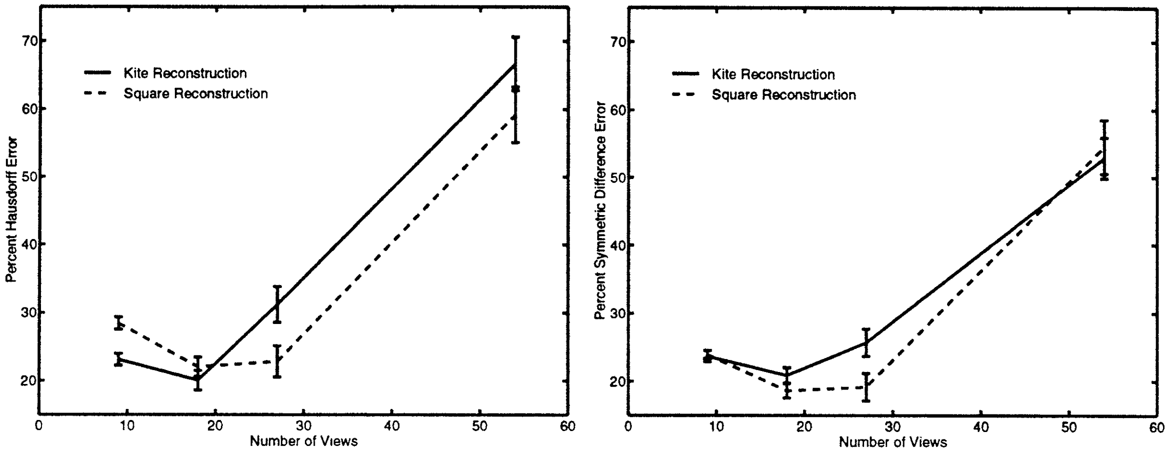

(a) Hausdorff Error, Correct Model (b) Symm. Diff. Error, Correct Model

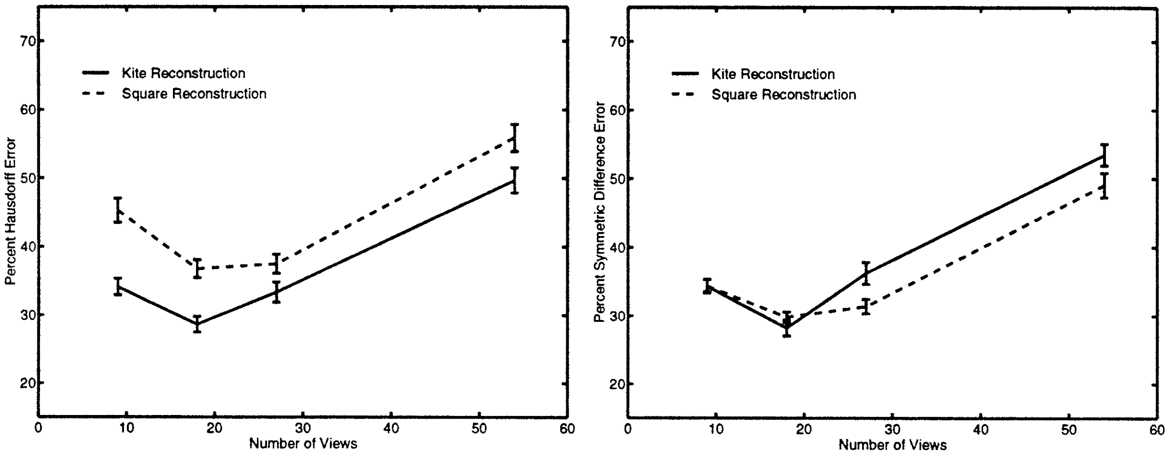

(c) Hausdorff Error, i.i.d. Model (d) Symm. Diff. Error, i.i.d. Model

Figure 4-12: Comparison of Kite and Square Test Objects for Sequential Reconstruction Algorithm: SNR=10dB, T=15 retained hypotheses, 5 chord measurements per view. (a),(b) Correct Noise Model; (c),(d) i.i.d. Noise Model.

4.5.4 Limited Angle Tomography and Non-Uniform Projection Angles

In this final section of experimental results we perform reconstructions using the Sequential Algorithm over limited angular regions and with non-uniformly spaced projection data. This is an important feature of this algorithm as many classical reconstruction techniques require uniformly spaced angular projections over the complete range $[0°, 180°)$.

Reconstruction Error as a Function of Limited Angle Projection Data

In this section we evaluate the reconstruction performance of the Sequential Algorithm over limited angular ranges. In each reconstruction, we generate a full set of projection data (i.e. chord measurements and knot location measurements) over the $[0°, 180°)$ angular range with 5 chord measurements per view at a SNR of 10dB. Like all previous experiments, the chord length measurements are uniformly sampled over the range $t = [-1, 1]$. Next, we perform reconstructions on *subsets* of the projection data over limited angular ranges using the Sequential Reconstruction Algorithm. Specifically, reconstructions were performed over the following angular ranges: $[0°, 45°)$, $[0°, 70°)$, and $[0°, 90°)$. Each subset of projection data contains 18 uniformly spaced views over the specified angular range. For each angular range, 100 independent Monte-Carlo reconstructions were performed with $T{=}15$ hypotheses retained at each step of the Sequential Algorithm.

Figure 4-13 shows the errors resulting from reconstructions of the kite as a function of the angular range. As the results from both the Hausdorff and Symmetric Difference errors in Figure 4-13(a) and Figure 4-13(b) demonstrate, the algorithm was able to produce reasonable reconstructions over each angular range. As expected, the reconstruction errors of both measures decreased as the angular range is increased. This is attributed to the improvement in triangulation geometry obtained by spacing the views over a wider angular region coupled with the problem that the Sequential Algorithm has with retaining the wrong hypotheses if the views are "too" close (refer to the discussion in Section 4.5.3 under the performance as a function of SNR). Although these results show that this algorithm is capable of reconstructing over limited angles, the success is somewhat limited because of the relatively high reconstruction errors (although this would decrease as the SNR is increased). However, looking ahead, we note that the results of the Non-Sequential Algorithm of Chapter 5 demonstrate greatly improved overall performance, especially for small

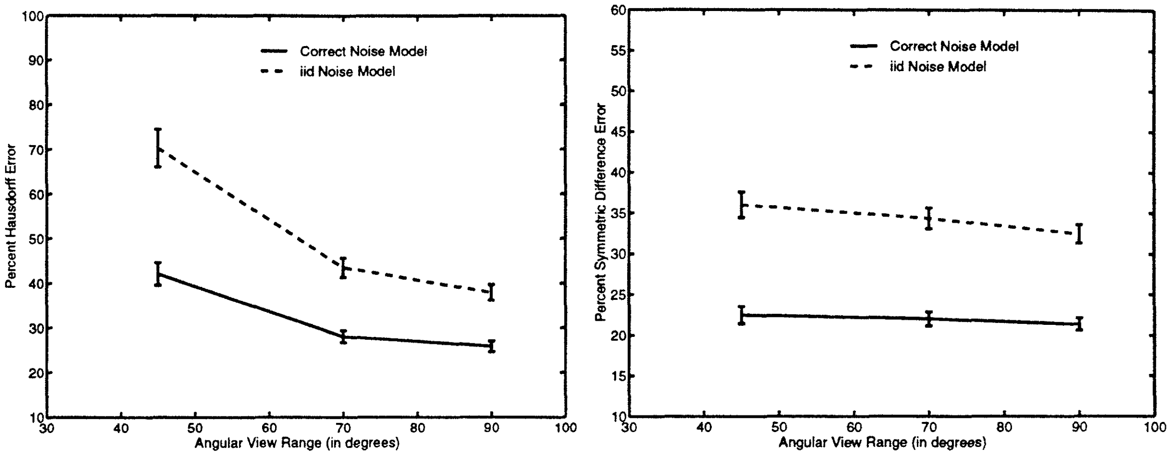

(a) Hausdorff Error

(b) Symmetric Difference Error

Figure 4-13: Performance as a Function of the Angular Range Using the Sequential Algorithm.

Error Measures	True Model	i.i.d. Model
Hausdorff Error (%)	10.75	14.53
Symm. Diff. Error (%)	19.88	17.05

Table 4.2: Sample Limited Angle Reconstruction Over the Range $[0°, 90°)$ Using the Sequential Algorithm.

angular ranges.

In addition to the reconstruction errors resulting from the Monte-Carlo reconstruction results, Figure 4-14 shows a sample reconstruction. This reconstruction was generated under the conditions listed above with projection data over the angular range $[0°, 90°)$. As the resulting plot shows, the algorithm was able to produce a good quality reconstruction in the face of limitations on both the angular range and the number of projections (18 views) in a low SNR situation (10dB). Additionally, Table 4.2 summarizes the reconstruction error for this particular sample.

Reconstruction Using Non-Uniform Angular Spacing

In addition to the ability to reconstruct over limited angular ranges, the Sequential Algorithm also has the ability to perform reconstructions on non-uniformly spaced projection data. Figure 4-15 shows a sample reconstruction from 27 non-uniformly spaced views over

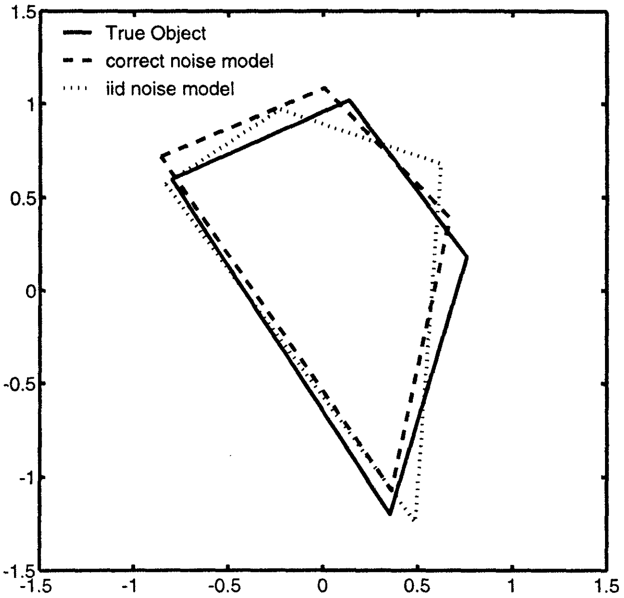

Figure 4-14: Sample Reconstruction Over Limited Angular Range ($[0^\circ, 90^\circ)$) Using the Sequential Algorithm.

the $[0^\circ, 180^\circ)$ angular range. In this reconstruction, each projection contained 5 chord measurements at a SNR of 10dB with $T = 15$ hypotheses retained at each step of the algorithm. To determine the non-uniform projection angles, we started with the initial projection at $\theta_1 = 0^\circ$. The second projection was at $\theta_2 = 3.5^\circ$ and all subsequent angles were generated with the following equation

$$\theta_i = \frac{7}{6}\theta_{i-1} < 180^\circ \tag{4.12}$$

The solid line in Figure 4-15 represents the true object and the dashed and dotted lines are the reconstructions assuming the correct knot location noise model and the i.i.d noise model respectively. In addition, Table 4.3 summarizes the Hausdorff and Symmetric Difference reconstruction errors for these results. Although only for one sample path, these results show the performance of this algorithm is not limited to a uniformly spaced projection data set.

4.6 Conclusions

In this chapter, we have studied a finite parameter reconstruction algorithm to reconstruct the vertices of a binary polygonal object from sequential measurements of knot location and

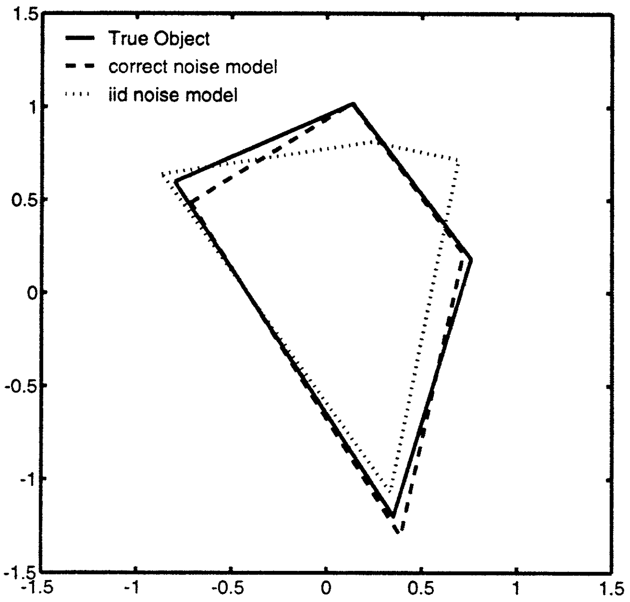

Figure 4-15: Sample Reconstruction with Non-Uniformly Spaced Projection Angles Using the Sequential Algorithm.

Error Measures	True Model	i.i.d. Model
Hausdorff Error (%)	11.19	20.84
Symm. Diff. Error (%)	7.23	22.71

Table 4.3: Comparison of Errors for Non-Uniform Angular Spacing Using the Sequential Algorithm.

chord length data. The reconstruction was posed as a combined hypothesis test-estimation problem; and the proposed solution was a combination of a Generalized Likelihood Ratio Test to solve the discrete hypothesis testing problem and Maximum Likelihood estimation technique to solve the estimation problem. The results of this algorithm were mixed. For relatively small amounts of data (as compared to traditional reconstruction algorithms like filtered backprojection [4] and even the statistically based ML reconstructions of Milanfar [9]), the results were encouraging. Although the algorithm was unable to disambiguate the data association problem for low SNRs (0dB), it was robust for SNRs at or above 10dB. The analysis also confirmed that the inclusion of chord length data to prune the possible hypotheses was a useful method of reducing the error in the reconstructed object. Further, given the reconstructions of the two test objects, the algorithm appears robust to different types of objects. Finally, we showed that the Sequential Algorithm is capable of reconstructing objects in situations of limited angular measurements and non-uniformly spaced projection data.

The primary deficiency was the degradation of the reconstruction results as the number of views increased for a fixed value of retained hypotheses. This is the result of the sub-optimal pruning algorithm implemented to reduce the total number of discrete data association hypotheses. We showed that one possible method of combating this problem is simply to retain more hypotheses. In the next chapter, we approach this problem with a second method: reordering the views so that they are introduced into the algorithm *non-sequentially*. Additionally, the performance of the Sequential Algorithm when the i.i.d. knot location noise model was assumed produced significantly larger reconstruction errors than the reconstructions that assume the correct applied model. Although we reiterate that this model is a worst case performance assessment (refer to Section 3.2.3), these errors are improved in the Non-Sequential Algorithm of Chapter 5 and significantly attenuated in the Nonlinear Reconstruction Algorithm of Chapter 6 which uses *both* knot location and chord length measurements to estimate the vertices of the object.

4-A Generating Chord Lengths From Vertex Locations

Because the relationship relating the chord lengths to the vertices of an object is complex, rather than giving an explicit formula to describe this relationship in this section we present the method and the C code used in this thesis to determine the chord lengths from the vertex locations of a binary convex object. The approach is to triangulate the object (i.e. divide the convex object into non-overlapping triangles), take the projections of each triangle at t_0, the position in the projection $g_\theta(t_0)$ where the chord length is desired, and sum the chord lengths of all of the triangles along this position to determine the chord length of the original n-gon. The reason for choosing this approach is twofold. First, determining the chord length of a triangle from its vertices is trivial; second, there is a simple method to triangluate a binary, convex n-gon into $n+1$ non-overlapping triangles.

As described in Section 2.1.5 the projection of a binary polygonal object is a piecewise-linear spline function. For a triangle, that simply means that each projection is itself a triangle as demonstrated in Figure 4-16. Therefore, to completely specify the projection of a triangle, we only need to determine the knot locations z_1, z_2, and z_3, and the chord length h at the internal knot location (refer to Figure 4-16). From these values, we can determine the equations of the two lines on either side of the internal knot location and thus determine the value of the projection at any value of t in the projection $g_\theta(t)$. All of these values can be easily obtained from the vertices of the triangle. Referring again to Figure 4-16, we note that if projections are taken of the three vertices of the triangle two of the values correspond to the support, which define the base b, and the third projected vertex corresponds to the internal knot location z_2. To completely specify the projection, we need only determine h. Using the fact that the area of the object and the area under the projection are the same (this is a well known property of the Radon transform), we can determine h or the value of the projection at the internal knot by using the equation of the area of a triangle

$$A = \frac{1}{2}bh \tag{4.13}$$

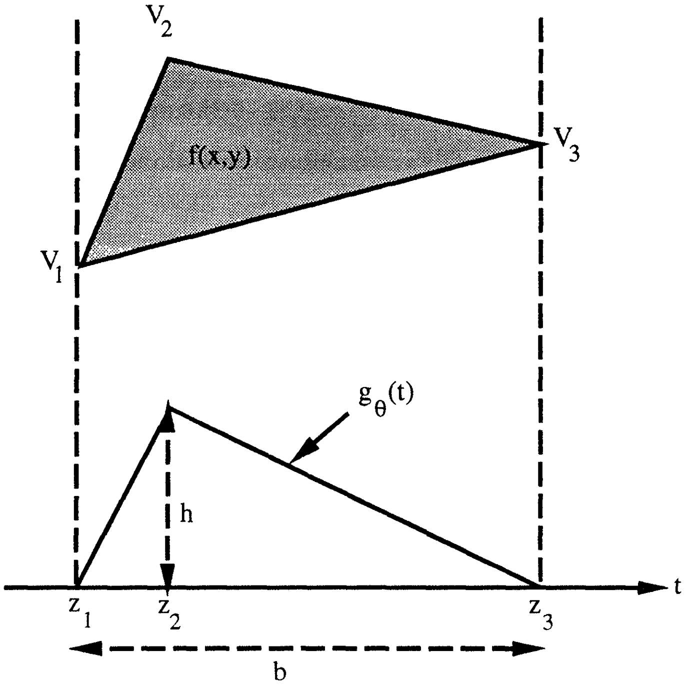

Figure 4-16: The Projection of a Triangle.

or

$$h = \frac{2A}{b} \tag{4.14}$$

where A is the area of the triangle (easily obtained from the vertices of the object), b is the base of the triangle (defined by the support knots), and h is the chord length or value of the projection at the internal knot location z_2. With these values, the entire projection of the triangle is defined and to calculate any desired chord length value we simply determine on which side of the internal knot the chord measurement lies and evaluate the correct equation of the line. It is obvious for an object with more sides (think of a square for example) this approach will not work because the determination of the chord lengths at the internal knot location is not as simple as the case of the triangle.

The second reason for choosing this approach is that the triangulation of a convex n-gon is trivial (procedure below). Once we have the non-overlapping triangles, the projections of the triangles at the desired position t_0 are taken and the sum of the $n+1$ chord lengths of the triangles determine the chord length of the overall object. This summation is possible because the Radon transform is a linear operation.

In summary, the overall approach is the following:

1. A vertex's angular position can be defined by the angular position of a vector from the origin to the vertex (assuming that the origin is located inside of the polygon). For a convex object, ordering the vertices in increasing angular order defines the connection of the vertices to form the object. We begin by ordering the vertices of the convex n-gon in increasing angular order to define the connection of the vertices.
2. The x and y coordinates of all of the vertices are then averaged to define the center of mass of the object. For a convex object, the center of mass in guaranteed to lie within the object.
3. The object is then triangulated by taking the ordered vertices of the object in pairs and the center of mass to form $n + 1$ non-overlapping triangles from the object. An example of this procedure is shown in Figure 4-17.
4. To determine the chord length at a particular location t_0 in the projection $g_\theta(t)$, the projection of each triangle is taken at the required position and the sum of the projections of the $n + 1$ non-overlapping triangles at this position forms the chord length of the overall object at t_0.
5. Step 4 is repeated for each chord length calculation.

The code used to generate these results is the following:

```
#include <math.h>
#include <mex.h>
#include <stdlib.h>
/* --------------------------------------------
function "c_chordgen"
-----------------
Determines the chord lengths and knot locations of a binary
  polygonal object.
Inputs:
  1. "num_ver"  : number of vertices in the object.
  2. "views"    : number of projection angles.
  3. "theta"    : array of desired projection angles.
  4. "ver"      : matrix of vertices (2*num_ver) [stored by ROWS].
```

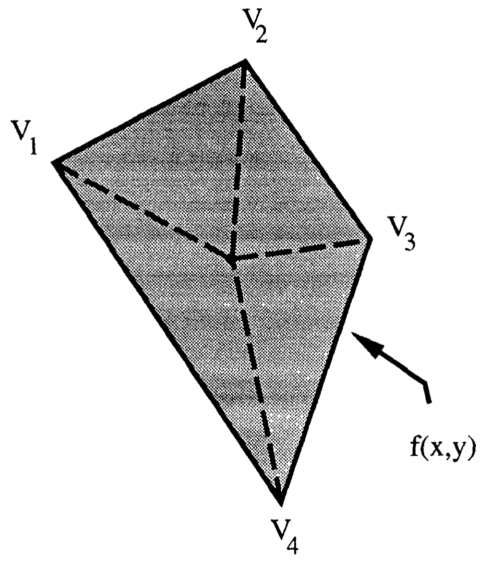

Figure 4-17: Triangulation of Kite Object.

```
     5. "num_chords": number of desired chord measurements.
    Outputs:
     1. "y1"        : matrix (num_ver*views) of knot locations, each column
                      corresponding to a projection [stored by ROWS].
     2. "chords1"   : matrix (num_chords*views) of chord lengths, each column
                      corresponding to a projection [stored by ROWS].
-------------------------------------------------------*/
c_chordgen(y1,chords1,num_ver,views,theta,ver,num_chords)
double y1[],chords1[];
double *num_ver, *views;
double theta[],ver[],*num_chords;
{
double *P, *temp1;
double *arr,*tri,drr[3],crr[3];
double sumx,sumy,fix,cent[2];
double *brr,s[9],*A1,*t1,*pos;
double *v, *Angle,*t,de,*t2,*t3,delta;
double temptri[6],chord;
double m,b,temp[3],h;
int p;
int j,nrt2,k,nrt,nct,num,a,nviews,nchords,l;
fix=*num_chords;
nchords=(int)(fix);
```

```
fix= *views;
nviews=(int)(fix);
fix = *num_ver;
num=(int)(fix);
sumx=0.0;
sumy=0.0;

v=(double *)malloc(sizeof(double)*(2*num));
tri=(double *)malloc(sizeof(double)*(num*2*3));
P = (double *)malloc(sizeof(double)*(num * num));
temp1 = (double *)malloc(sizeof(double)*(num));
brr = (double *)malloc(sizeof(double)*(num));
arr = (double *)malloc(sizeof(double)*(num));
Angle = (double *)malloc(sizeof(double)*(2*nviews));
t = (double *)malloc(sizeof(double)*(2*nviews*num));
t1 = (double *)malloc(sizeof(double)*(3*nviews*num));
t2 = (double *)malloc(sizeof(double)*(3*nviews*2));
t3 = (double *)malloc(sizeof(double)*(3*nviews*num*2));
A1 = (double *)malloc(sizeof(double)*(num));
pos= (double *)malloc(sizeof(double)*nchords);

nrt2=2;

/* determine the center of mass of the object */
for(j=0;j<num;j++) {
  sumx=ver[j]+sumx;
  sumy=ver[j+num]+sumy;
}
cent[0]=sumx/fix;
cent[1]=sumy/fix;

/* order the vertices */
c_convex_order(v,ver,num_ver);

matzero(P,&num,&num);
for(j=0;j<num;j++) P[num*j+(num-j-1)] = 1.0;
nrt2=2;
nrt=num*2;
```

```
nct=3;
matzero(tri,&nrt,&nct);

/* triangulate the object */
for(j=0;j<num;j++)
{
 for(k=0;k<2;k++)
   {
    if(j==num-1 && k==1)
    {
         tri[((num-1)*6)+1]=v[0];
     tri[((num-1)*6)+4]=v[num];
    }
    else
    {
         tri[(j*6)+k]=v[j+k];
     tri[(j*6)+k+3]=v[j+k+num];
     }
   }
 tri[(j*6)+2]=cent[0];
 tri[(j*6)+5]=cent[1];
}

/* determine the projection matrix from the array of angles */
for(j=0;j<nviews;j++) {
  Angle[(j*2)]=cos(theta[j]);
  Angle[(j*2)+1]=sin(theta[j]);
}

/* calculate the knot locations from the vertices */
matmult(Angle,&nviews,&nrt2,ver,&nrt2,&num,t,&nviews,&num);
for(j=0;j<num;j++) brr[j]=1;
for(j=0;j<nviews;j++) {
  for(k=0;k<num;k++) arr[k]=t[(j*num)+k];
  piksr2(num,arr,brr);
  for(k=0;k<num;k++) y1[(k*nviews)+j]=arr[k];
}

/* determine the area of each of the n+1 triangles */
```

```
for(j=0;j<3;j++) s[j]=1.0;
for(j=0;j<3;j++) drr[j]=j+1.0;
nrt=1;
for(j=0;j<num;j++) {
    for(k=0;k<2;k++) {
     for(l=0;l<3;l++) {
          s[(k*3)+l+3]=tri[(j*6)+(k*3)+l];
          temptri[(k*3)+l]=tri[(j*6)+(k*3)+l];
     }
   }
   de= ((-1)*s[3]*(s[8]-s[7]))+(s[4]*(s[8]-s[6]))-(s[5]*(s[7]-s[6]));
   if (de < 0) de=de*(-1);
   A1[j]=de;
   matmult(Angle,&nviews,&nrt2,temptri,&nrt2,&nct,t2,&nviews,&nct);
   for(l=0;l<nviews;l++){
     for(k=0;k<3;k++) crr[k]=t2[(l*3)+k];
     piksr2(nct,crr,drr);
     for(k=0;k<3;k++) t3[(j*nviews*3)+(3*l)+k]=crr[k];
   }
}

/* calculate the chord lengths over the range t=[-1,1] (uniformly sampled) */
if (nchords==0)
    matzero(chords1,&nchords,&nviews);
else {
  if(nchords==1){
    pos[0]=0.0;
  }
  else {
    delta=2.0/(nchords-1);
    pos[0]=-1.0;
    for(j=1;j<nchords;j++) pos[j]=pos[j-1]+delta;
  }

  for(j=0;j<nchords;j++) {
    for(k=0;k<nviews;k++) {
      chord=0.0;
      for(l=0;l<num;l++){
            for(p=0;p<3;p++) temp[p]=t3[(k*3)+(l*nviews*3)+p];
```

```
          h=A1[l]/(temp[2]-temp[0]);
          if((temp[1] >= (pos[j]-1e-3)) && (temp[1] <= (pos[j]+1e-3)))
              chord=chord+h;
          else if((temp[1] > pos[j]) && (pos[j] > temp[0])) {
            if ((temp[1]-temp[0]) >0 )
              m=(h/(temp[1]-temp[0]));
            else
              m=-(h/(temp[1]-temp[0]));
            b=-m*temp[0];
            chord=chord+(m*pos[j])+b;
          }
        else if((temp[1]<pos[j]) && (pos[j]<temp[2])) {
            if ((temp[2]-temp[1]) >0 )
              m=-(h/(temp[2]-temp[1]));
            else
              m=(h/(temp[2]-temp[1]));
            b=-m*temp[2];
            chord=chord+(m*pos[j])+b;
          }
      }
      chords1[(j*nviews)+k]=chord;

    }
  }
}

free(v); free(tri); free(P); free(templ); free(brr); free(arr); free(Angle);
free(t); free(t1); free(t2); free(t3); free(A1); free(pos);
}

#include <math.h>
#include <mex.h>
/* ------------------------------------------------------------
 function "c_convex_order"
 -------------------
 Orders the vertices of a convex polygonal object.
 Inputs:
    1. "Ver"      : matrix of vertices (2*num_ver) [stored by ROWS].
    2. "num_ver"  : number of vertices in the object.
```

```
 Outputs:
     1. "v"        : matrix of vertices (2*num_ver) [stored by ROWS].
                     v=[x1 x2 x3 ... xn
                              y1 y2 y3 ... yn]
--------------------------------------------------------*/
c_convex_order(v,Ver,num_ver)
double v[],Ver[];
double *num_ver;
{
double *P,*temp1,*arr,fix,*brr,*brr1;
int j,nr,num,a,nr2;

fix = *num_ver;
num=(int)(fix);

P = (double *)malloc(sizeof(double)*(num * num));
temp1 = (double *)malloc(sizeof(double)*(num));
brr = (double *)malloc(sizeof(double)*(num));
brr1 = (double *)malloc(2*sizeof(double)*(num));

for(j=0;j<num;j++) brr[j]=j+1.0;
nr=1;
nr2=2;
matzero(P,&num,&num);
for(j=0;j<num;j++) P[num*j+(num-j-1)] = 1.0;

/* determine the angle of each vertex */
for(j=0;j<num;j++)
  {
    temp1[j]=atan2(Ver[j+num],Ver[j]);
  }

/* sort the vertices */
piksr2(num,temp1,brr);

/* order the vertices */
matmult(brr,&nr,&num,P,&num,&num,brr1,&nr,&num);

for(j=0;j<num;j++)
```

```
{
  a=(int)(brr1[j]);
  v[j]=Ver[a-1];
  v[j+num]=Ver[a-1+num];
}
free(P); free(temp1); free(brr); free(brr1);
}

#include <stdio.h>
/* ---------------------------------------------
 function "matzero"
 ------------------
 Zeros out a matrix.
 Inputs:
       1. "matrix1": The first matrix.
       2. "nr1"    : Number of rows in above.
       3. "nc1"    : Number of columns in above.
-----------------------------------------------*/
matzero (matrix1, nr1, nc1)
double *matrix1;
int *nr1, *nc1;
{
  int i,j;
  for ( i = 0 ; i <= (*nr1)-1 ; i++ ) {
    for ( j = 0 ; j <= (*nc1)-1 ; j++ ) {
      matrix1[ (*nc1)*i + j ] = 0.0;
    }
  }
  return;
}

#include <stdio.h>
/* ---------------------------------------------
 function "matmult"
 ------------------
 Multiplies two (double) matrices.
 Inputs:
       1. "matrix1": The first matrix.
       2. "nr1"    : Number of rows in above.
```

```
    3. "nc1"     : Number of columns in above.
    4. "matrix2": The second matrix.
    5. "nr2"     : Number of rows in above.
    6. "nc2"     : Number of columns in above.
 Outputs:
    1. "matrixr": The result of matrix1*matrix2.
    2. "nrr"     : Number of rows in above.
    3. "ncr"     : Number of columns in above.
------------------------------------------------------*/
matmult (matrix1, nr1, nc1, matrix2, nr2, nc2, matrixr, nrr, ncr )
int *nr1, *nc1, *nr2, *nc2, *nrr, *ncr;
double *matrix1, *matrix2, *matrixr;
{
 int i, j, k;
 double sum;
/* Check for size compatibility.*/
 if ( *nc1 != *nr2 ) {
  fprintf( stderr, "ERROR in MATMULT: Matrix size mismatch\n" );
  exit(1);
 }
/* The size of the multiplied matrix. */
 *nrr = *nr1;
 *ncr = *nc2;
/* Multiply the matrices. */
 for ( i = 0; i <= (*nr1)-1; i++ ) {
  for ( j = 0; j <= (*nc2)-1; j++ ) {
   sum = 0.0;
   for ( k = 0; k <= (*nc1)-1; k++ ) {
      sum += matrix1[ (*nc1)*i + k ] * matrix2[ (*nc2)*k + j ];
   }
   matrixr[ (*nc2)*i + j ] = sum;
  }
 }
 return;
}
/* END function "matmult" ----------------------*/


/*------------------------------------------------
```

function "piksr2"

Sorts 2 arrays. Sorts array arr[1...n] into ascending numerical order by straight insertion, while making the corresponding rearrangement of the array brr[1 ...n]

Inputs:

1. *"n" size of arrays arr and brr*
2. *"arr" array to be sorted*
3. *"brr" array to be sorted*

Outputs:

1. *"arr" sorted version*
2. *"brr" sorted version*

```
________________________________*/
piksr2(n,arr,brr)
int n;
double arr[];
double brr[];
{
        int i,j;
        double a;
        double b;
        for (j=1;j<n;j++) {
                a=arr[j];
                b=brr[j];
                i=j−1;
                while (i >= 0 && arr[i] > a) {
                        arr[i+1]=arr[i];
                        brr[i+1]=brr[i];
                        i−−;
                }
                arr[i+1]=a;
                brr[i+1]=b;
        }
}
```

Chapter 5

Non-Sequential Reconstruction Algorithm

In this chapter we develop a second algorithm, which we term the Non-Sequential Reconstruction Algorithm, to reconstruct binary polygonal objects from knot location data and chord length measurements. Unlike the reconstruction algorithm developed in Chapter 4, the algorithm developed in this chapter allows views of knot location data to be introduced into the reconstruction algorithm in any angular order. In Section 5.1 we outline the general framework of the Non-Sequential Reconstruction Algorithm and how its methodology differs from the Sequential Algorithm developed previously. Next, Section 5.2 describes the criteria used to generate the possible hypotheses at each stage of the reconstruction and presents a general overview of the algorithmic procedure. Finally, Section 5.3 presents sample reconstructions and Monte-Carlo simulations of this algorithm and Section 5.4 summarizes our conclusions on the results of this algorithm.

5.1 Overview of the Non-Sequential Algorithm

Like the previous chapter, this chapter is concerned with the reconstruction of binary polygonal objects from measurements of the knot locations of an n-gon (assuming that the number of sides is known); unlike the previous chapter, the views are introduced into the algorithm non-sequentially. As outlined in Section 3.2.1, we assume that the projections are taken at known angular views and that the knot locations have been extracted from the projection data independently of this algorithm using a procedure such as that of Mier-Muth and Willsky [8]. Also, for simplicity we assume that all objects are binary and convex as discussed in Section 3.2.1. Throughout this chapter we will make extensive use of the definitions and terminology presented in Chapter 2 and problem formulation of Section 3.2. In addition, many of the definitions and procedures introduced in Chapter 4 for the Sequential Algorithm are also used for the Non-Sequential Algorithm.

Essentially, the methodology of the Non-Sequential Algorithm is the same as that of the Sequential Algorithm developed in Chapter 4. The general framework is a multi-stage estimation algorithm that simultaneously estimates discrete-valued states, or data association hypotheses, and continuous-valued estimates of the vertices of the object based on these discrete hypotheses (refer to the block diagram of Figure 4-1). In fact, the only difference between these two algorithms is the method used in Stage 1 (see Section 4.2) to evaluate the possible data associations and form the set of discrete hypotheses.

Instead of appending each new view right after the last projection (i.e. in close angular proximity), this algorithm allows new projection data to be inserted at any angular location over the $[0°, 180°)$ allowable angular range. The primary advantage of incorporating views in any order is that the triangulation geometry (see Section 3.2.2) of this algorithm is greatly improved by introducing views into the algorithm at wide angular spacing. Unlike the Sequential Reconstruction Algorithm which uses a gating approach to limit the allowable knot switches by assuming that only adjacent knots can switch their data association with the vertices of the object from one view to the next, the Non-Sequential Algorithm allows *any* switch as long as it is possible given the geometry of the tomography problem. Because of its ability to accept any angular separation, this algorithm can also be used for limited angle reconstructions or reconstructions with few but widely spaced views. As a result of the increased flexibility, however, the number of possible hypotheses is increased and the evaluation of the possible hypotheses is slightly more complex.

5.2 Generation of Data Association Hypotheses

In the Sequential Reconstruction Algorithm developed in the previous chapter, the number of possible knot-to-vertex associations (and ultimately the number of proposed hypotheses) was severely limited by the geometric constraints of reconstructing a binary polygonal object coupled with the assumption that the projection angles were closely spaced. Because views were assumed to be added into the algorithm in close angular proximity, a gating approach was used to constrain the number of hypotheses by allowing only adjacent knot-to-vertex data associations to switch from one view to the next. Although the Non-Sequential Algorithm cannot use this gating technique (because there can be large angular separations between adjacent views), the number of possible data associations can still be greatly reduced from the $n!$ vertex permutations that can occur in an n-sided object. If the knots in the projection data are viewed as a continuous function of θ, the following geometric constraints that were first described in Chapter 4 (and repeated here for convenience) are still valid[1]:

A. If the knots in the projection at $\theta_1 = 0°$ are labeled a, b, c, and d, these knots occur in the reverse order in the final projection at $\theta_{m+1} = 180°$. Thus, for the case of a 4-gon a knot association hypothesis that starts with *abcd* must end with *dcba*.

B. A switch in two knots occurs when two vertices lie along the same line, perpendicular to the projection. Therefore, for an n-sided object, exactly $n(n-1)$ switches occur over a complete and continuous 360° range of projections. Further, because each switch of vertices occurs exactly 180° apart (because the vertices must lie along the same line), $\frac{n(n-1)}{2}$ switches occur in the range from $\theta_1 = 0°$ to $\theta_{m+1} = 180°$. Therefore, in a 4-gon exactly 6 knot location switches will occur.

C. Because two vertices have to be aligned on a line perpendicular to the projection for the knots to switch, it is not possible for two knots to switch twice in a row (i.e. you can't *undo* what was done in the previous step) with increasing θ.

In the absence of a gating approach such as that used in Chapter 4, the above criteria are used to develop the following "rules" to determine the allowable knot orderings and possible hypotheses at each insertion of data[2]:

1. *If a projection is inserted between two views that have the same knot-to-vertex association, the inserted view must also have that ordering.*

   This follows from a combination of constraints A, B, and C. If two views have the

---

[1] Note that the examples are given for the case of a 4-gon but can be naturally extended to the general case of an n-gon

[2] Again, these examples are given for the case of a 4-gon but can be naturally extended to the general case of an n-gon

same set of knot-to-vertex associations and a view is inserted between them, the only way that the inserted association could be different is if the knots associations switch and then switch back again. Over the 360° angular range two vertices of an object are aligned, and therefore switch associations, exactly twice (i.e. only when they lie along the same line). Further, because they must be along the same line we know that the knot switches in the two vertices of the object will occur exactly 180° apart. Thus, in order for the knots to switch and then switch back again, the angular range would have to be greater than 180°. Because we consider only the range $[0°, 180°)$ this is not possible. Therefore, a view inserted between two views with the same knot-to-vertex association must also have that set of associations.

2. *Once a knot-to-vertex association has been used somewhere in a hypothesis, it can never occur again after the knots switch from that position.*
   This follows from the analysis under rule 1. An angular separation of greater than 180° is required for the knot-to-vertex associations to switch and then switch back.

3. *If a knot-to-vertex association is used somewhere in the proposed hypothesis path, its reciprocal cannot be used anywhere in the hypothesis (with the exception of the initial and final associations).* From the consistency requirements of the Radon transform (refer to Equation (2.20)) we know that the knot locations must adhere to the following property:

$$z(\theta_i) = [z_1(\theta_i) \ldots z_n(\theta_i)] = [-z_n(\theta_i + 180°) \ldots - z_1(\theta_i + 180°)] \tag{5.1}$$

   and from the geometry of this problem (refer to Figure 4-3) this corresponds to an exact reversal of the knot-to-vertex associations. Thus, if the knot associations switch at exactly θ_i, the switch to the reciprocal of the knot-to-vertex associations occurs at exactly $\theta_i + 180°$. Because we consider a range of less that 180° the reciprocal ordering cannot be a valid association over the range that we consider. The one exception to this rule is the first set of data associations, *abcd*. Because the knots switch to form this set of associations prior to the initial view at $\theta = 0°$ the reciprocal set of associations *dcba* will occur before θ reaches 180°.

4. Each hypothesis must start with abcd and end with dcba.

   This rule follows directly from constraint A.

5. Only $\frac{n(n-1)}{2}$ knot switches are allowed to occur in each hypothesis path. After this, only the data associations adjacent to the inserted knot projection are allowed.

   This rule follows directly from constraint B.

These rules are only a subset of the possible rules that can be derived to determine the allowable data associations given the geometry of this problem. They are meant only to provide a simple and efficient means of reducing the total number of hypotheses that must be considered at each stage of the algorithm. A precise statement of the space defining the set of valid data associations is still an open issue.

As an example of an application of rules *1-5* consider the 3 hypotheses given in the first column of Table 5.1. The data associations for each hypothesis are listed in increasing angular order in the first column and the arrow represents the point at which we wish to insert a new set of projection data. Without any regard for the geometry of this problem, there would be 4! or 24 possible permutations of $a, b, c,$ and d that could be used to represent the data association of the 4 vertices to the 4 knot location measurements that could be added to each hypothesis. The center columns of Table 5.1 show which views can be *eliminated* with the implementation of the 5 rules. In fact, looking at the last column, which gives the allowable knot location data associations for this point of insertion into each of the hypotheses we see that the number of potential associations has been reduced from 24 to 18, 1, and 2 for the respective 3 hypotheses. Note also, that as the assumed hypothesis becomes larger and more defined that the number of possible data associations that can be inserted greatly diminishes. In fact, after the $\frac{n(n-1)}{2}$ switches are chosen (such as the third hypothesis given in the table), only the 2 data associations adjacent to any inserted view are allowed.

Using the above "rules", the possible data associations are enumerated in a hypothesis tree similar to the one used in the Sequential Algorithm of Chapter 4. The primary difference is that the views are not in increasing angular order as you move from the root to a branch of the tree. Instead, as you move through the tree, the data associations are listed in the order that they are introduced into the algorithm (which doesn't necessarily correspond to increasing angular order). Each branch in the tree must keep track of the angle at which

	Associations Eliminated by Respective Rule					
Hypothesis	Rule 1	Rule 2	Rule 3	Rule 4	Rule 5	Possible Data Assoc.
abcd ▭⇐ *bacd* *bdac* *dbca* *dcba*		*dbca* *dcba* *bdac*	*cadb* *acbd* *dcab*			*cdab cabd dacb* *adcb acdb dabc* *adbc abdc abcd* *cdba cbda cbad* *bdca dcda bcad* *dbac badc bacd*
abcd *bdac* ▭⇐ *bdac* *dcba*	All except *bdac*					*bdac*
abcd ▭⇐ *bacd* *badc* *bdac* *dbac* *dbca* *dcba*					All except *abcd* and *bacd*	*abcd* *bacd*

Table 5.1: Example Evaluations of Possible Knot Location Data Associations.

the projection data is inserted. Note also that because at least three projections are needed to triangulate the knot location data and form an object, the algorithm must retain *all* possible permutations of the knot location data until three views are obtained. Therefore, for the second view, $n!$ permutations of the knot location data must be considered regardless of the angular position. An example of a hypothesis tree for a 4-gon is shown in Figure 5-1. As with the Sequential Algorithm of the previous chapter, the initial knot location data is arbitrarily assigned to the unknown vertex locations as *abcd*. Note that this tree is significantly more complex than the corresponding tree given for the sequential algorithm (see Figure 4-4). Each branch of the tree gives the association of the knots chosen for that particular angle. Like the hypothesis tree developed in Chapter 4, a path from the root to a branch of the tree represents one possible hypothesized set of data associations, q_i^k, while $q^k = [q_1^k, q_2^k, \ldots, q_r^k]$ represents the set of all r hypotheses up to the kth measurement. Unlike the previous chapter, however, these data associations must be reordered to put them in increasing angular position. For example, the hypothesis q_2^4 which is represented by the the set of data associations traced out in dashed lines in Figure 5-1, represents the hypothesis $abcd \rightarrow abcd \rightarrow abcd \rightarrow bacd$ when reordered.

As mentioned previously, apart from the difference in the generation of the data association hypotheses, the format of the Non-Sequential Algorithm is the same as that of the Sequential Algorithm. Thus, the algorithm consists of three distinct stages: generation of data association hypotheses, estimation of vertex locations, and evaluation of data association hypotheses. The block diagram of Figure 4-1 summarizes these stages. Using the above data association rules with each set of knot location data, the algorithm evaluates the set of possible associations for each existing hypothesis q_i^{k-1}, $1 \leq i \leq T$ in the hypothesis tree and determines what knot associations are possible for the kth measurement given the geometric constraints outlined above (and demonstrated in Table 5.1). This set of valid associations forms the set of hypotheses q^k which corresponds to Stage 1 of the reconstruction algorithm. Stages 2 and 3 of the algorithm, which evaluate and prune the hypotheses, are performed using the same method as the Sequential Algorithm (see Sections 4.3 and 4.4). Again, chord length data are used in addition to the knot location data to prune the possible hypotheses by using the likelihood ratio tests defined by Equation (4.11). All of the calculations and evaluation procedures are identical to those of the Sequential

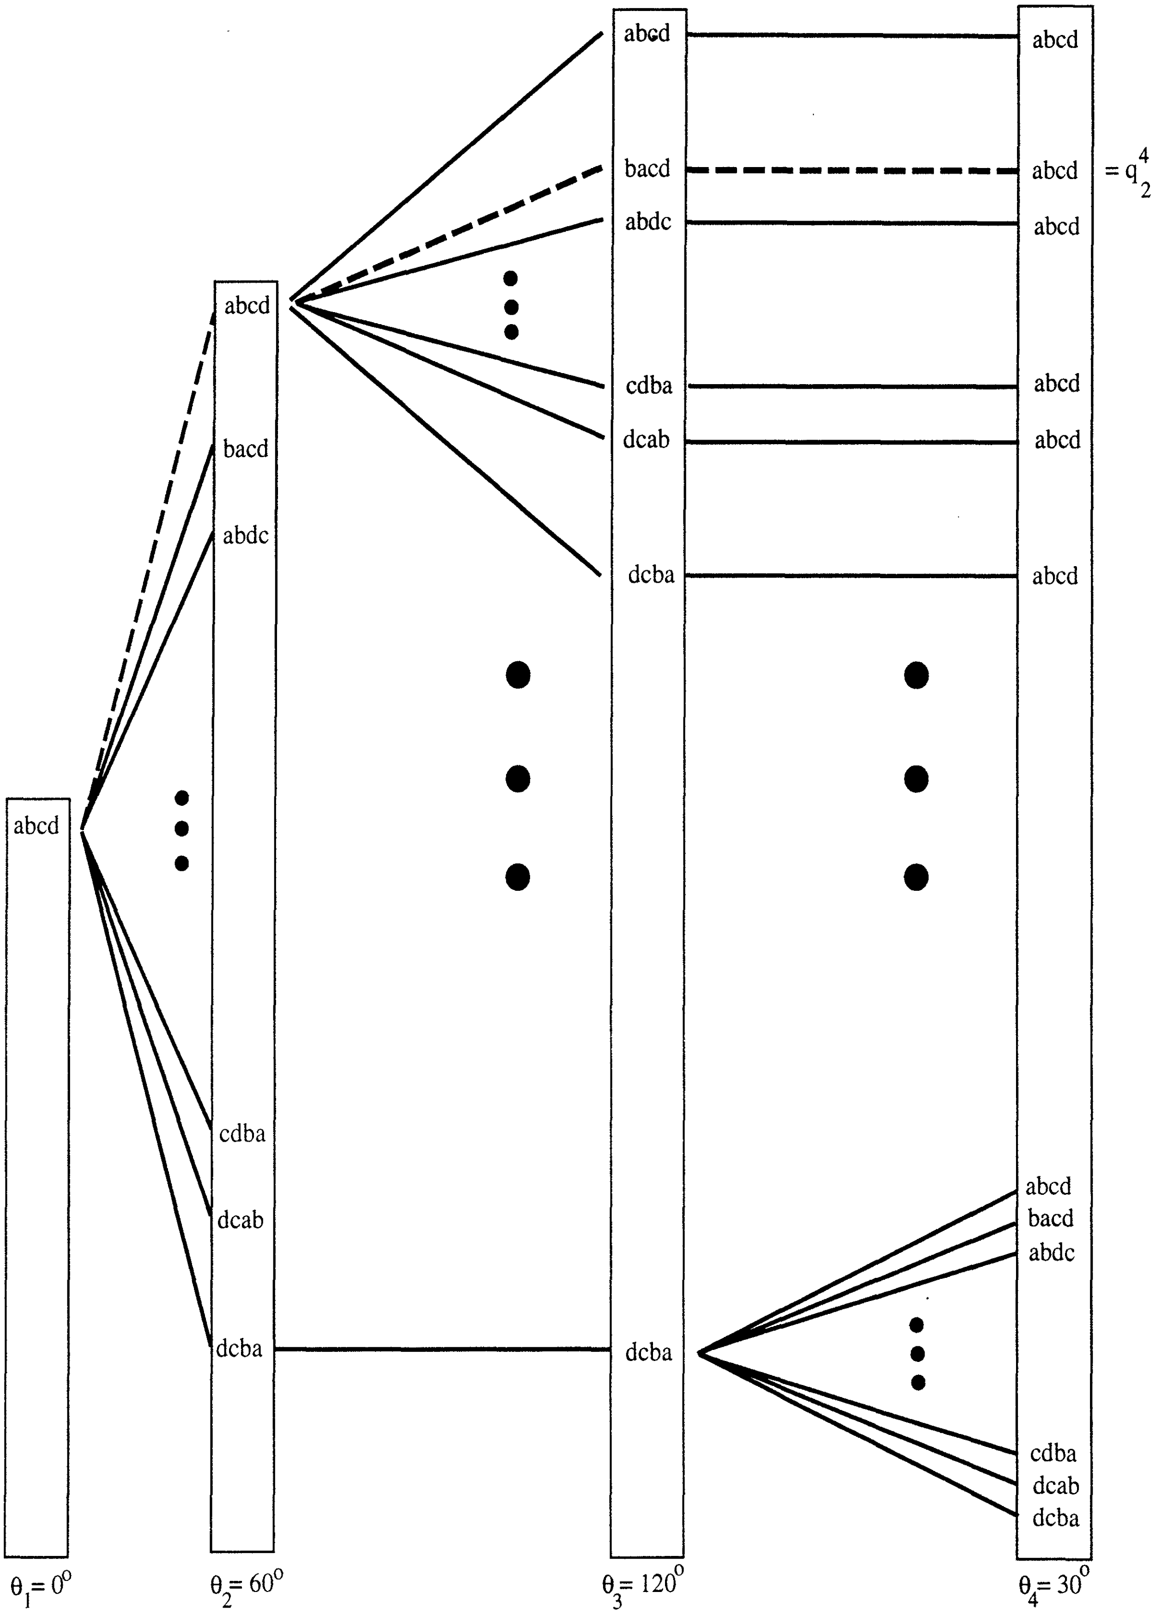

Figure 5-1: Hypothesis Tree for Non-Sequential Reconstruction Algorithm.

Algorithm with the exception of the determination of the discrete hypotheses, as described above. Experimental Results of this algorithm follow in the next section.

5.3 Experimental Results

In this section, we present some sample reconstructions and Monte-Carlo simulations of the proposed Non-Sequential Reconstruction Algorithm. As with the Sequential Algorithm, we perform our analysis on two test objects: a "kite" and a square (refer to Section 4.5.1, Figures 4-5 and 4-6). There are a number of specifications that are common to all of the reconstructions performed in the following sections.

First, at the start of each reconstruction the initial view is determined by the Max-Min Approach defined in Section 3.2.4. If the Max-Min view is labeled $\tilde{\theta}_1 = 0^\circ$, the views are then divided by angular range into three sets with roughly the same number of views (i.e. $[0^\circ, 60^\circ)$, $[60^\circ, 120^\circ)$, and $[120^\circ, 180^\circ)$)[3]. The views are taken from each set one by one and added into the algorithm such that the separation between the new data and the views input into the algorithm at the last pass through the three sets of views is maximized. For example, the first three views are approximately at $\tilde{\theta}_1 = 0^\circ$, $\tilde{\theta}_2 = 60^\circ$, and $\tilde{\theta}_3 = 120^\circ$. For the next view, we return to set 1 and pick the view that is as close to 30° as possible (i.e. midway between $\tilde{\theta}_1$ and $\tilde{\theta}_2$). Similarly, for the 5th view the projection angle is chosen from set 2 such that the view is midway between the views at $\tilde{\theta}_2$ and $\tilde{\theta}_3$. This continues until all of the views are input into the reconstruction algorithm.

Second, as with the Sequential Algorithm, the chord measurements are uniformly sampled over the region defined in Radon space by $t = [-1, 1]$. Again, as stated in the previous chapter, this results in views where some of the chord measurements are zero (i.e. the object is not in the field of view) and conversely, in views where the chord measurements are confined to the interior of the object (for the kite object).

The noise on the chord measurements is i.i.d. Gaussian noise with variance σ_m^2 and the corresponding noise on each knot location measurement is given by Equation (3.3). Reconstructions are performed for cases where the correct applied knot location noise model

---

[3] These angular ranges are approximate and depend on the angular separation of the views in the projection data.

is assumed known and also for cases where the assumed model is an i.i.d. model with covariance $\Lambda_z = (1.5\sigma_m)^2 I$ where σ_m^2 is the variance of the noise on the chord length measurements[4]. Finally, as described in Section 3.2.5 the percent Hausdorff error and percent Symmetric Difference error are used for the performance analysis.

5.3.1 Sample Reconstructions

This section contains results of sample reconstructions using the Non-Sequential Algorithm. Note that for all of these results the true object is depicted by solid lines while the reconstructions are represented by dashed or dotted lines. The sample reconstructions of this section use the same data set used to generate the sample reconstructions of the Sequential Algorithm (refer to Section 4.5.2). For both objects, the square and the kite, noisy chord length and knot location data were generated so that the SNR on the chord length data was set to 10dB. Additionally, projection data was taken at 27 equally spaced projections over the interval $[0^\circ, 180^\circ)$ with 5 chord measurements in each view. Thus, each object had a total of 108 knot location measurements (used to estimate the vertices and prune the hypotheses) and 135 chord measurements (used to prune the hypotheses).

Figure 5-2(a) shows the sample reconstruction of the kite. Two reconstructions are shown: the dashed line represents the reconstruction generated assuming the correct applied noise model on the knot measurements while the dotted line represents the reconstruction assuming the i.i.d. noise model. Similarly, Figure 5-2(b) shows a sample reconstruction of the square. Like the kite reconstruction, the dashed line represents the reconstruction assuming the correct applied noise model while the dotted line represents the i.i.d. model reconstruction. Table 5.2 summarizes the Hausdorff error and Symmetric Difference error for both sample reconstructions.

As in the Sequential Algorithm, the reconstruction errors assuming the correct applied noise model were significantly less that those assuming the i.i.d. model for all of the cases. This result is expected as the i.i.d. model weighs all knot location data equally in the cost criterion used to determine the optimal hypothesis regardless of the amount of noise actually applied to the knot. As a result, poor measurements are considered to be as important as

---

[4] Refer to Section 3.2.3 for complete descriptions of the noise models and SNR definition used in this thesis.

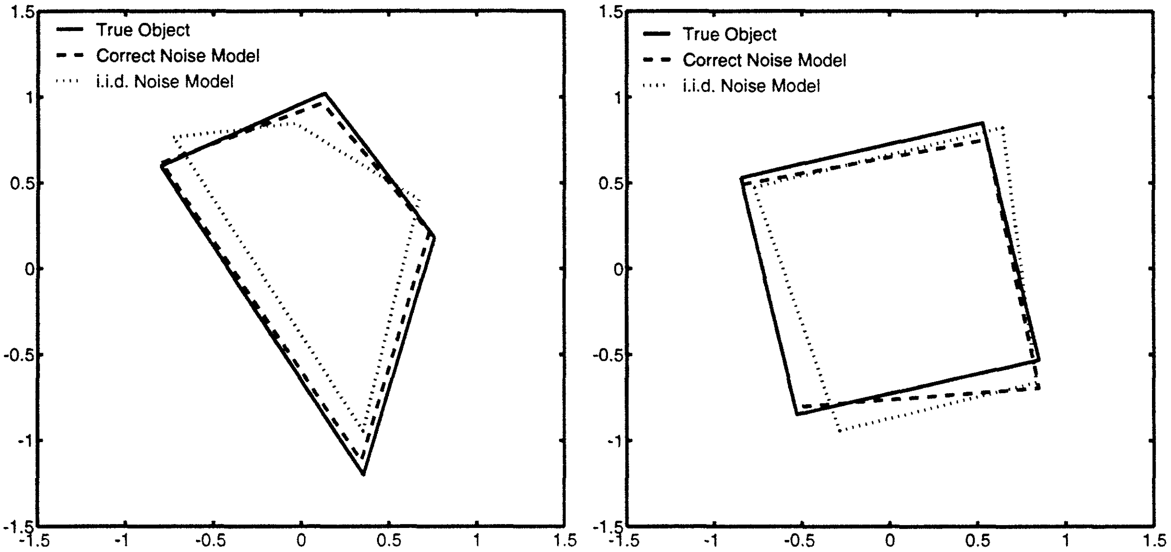

(a) Kite Reconstruction

(b) Square Reconstruction

Figure 5-2: Sample Reconstructions of Non-Sequential Algorithm, 27 views, 10 hypothesis pruning, 10dB SNR on chords.

Error Measures	Kite Object		Square Object	
	True Model	i.i.d. Model	True Model	i.i.d. Model
Hausdorff Error (%)	6.57	20.41	15.79	19.22
Symm. Diff. Error (%)	8.08	32.66	10.67	24.19

Table 5.2: Comparison of Errors for Sample Reconstructions Using Non-Sequential Algorithm.

measurements with little noise which results in a degradation in the reconstruction quality in most cases. It is interesting to note that for the Non-Sequential Algorithm, the kite reconstruction has lower error measures for the correct assumed model while the square has lower error measures for the i.i.d. assumed model. As in the previous chapter, we acknowledge this result but defer any comparative analysis of the performance of the algorithm in reconstructing these two objects to the Monte-Carlo analysis of Section 5.3.2. In addition, we note that the overall performance of this set of reconstructions was far superior to those of the Sequential Algorithm shown in Table 4.1 (the same input data was used in both reconstruction algorithms). In all error measures, the Non-Sequential Algorithm produced smaller reconstruction errors. In particular, both the Hausdorff and Symmetric Difference errors in the correct assumed model reconstructions were nearly halved in the Non-Sequential reconstruction results.

5.3.2 Monte-Carlo Simulations

In the following sections Monte-Carlo simulations of reconstructions are used to test the performance of the Non-Sequential Reconstruction Algorithm. As in the previous chapter, all simulations are performed using the kite test object (with the exception of the last section which is a comparison of reconstructions of the kite and square). The Monte-Carlo simulations consist of 100 independent reconstructions for each reconstruction scenario. In each plot error bars denote the 95% confidence intervals of the sample mean values that result from the 100 runs of the algorithm. As defined at the beginning of this section, the chord length data from each projection are uniform samples over the region $t = [-1, 1]$. Further, reconstructions are done for cases that assume the correct applied noise model for the knot location data and also for cases that assume an i.i.d. model [5]. The projection data for all simulations are uniformly spaced samples over the angular region $[0°, 180°)$. Finally, *unless otherwise stated*, $T = 15$ hypotheses are retained in each step of the algorithm, 5 chord measurements are taken per view, and the SNR on the chord length measurements is set to 10dB.

---

[5] Refer to Section 3.2.3 for a complete discussion of the noise models used in this thesis.

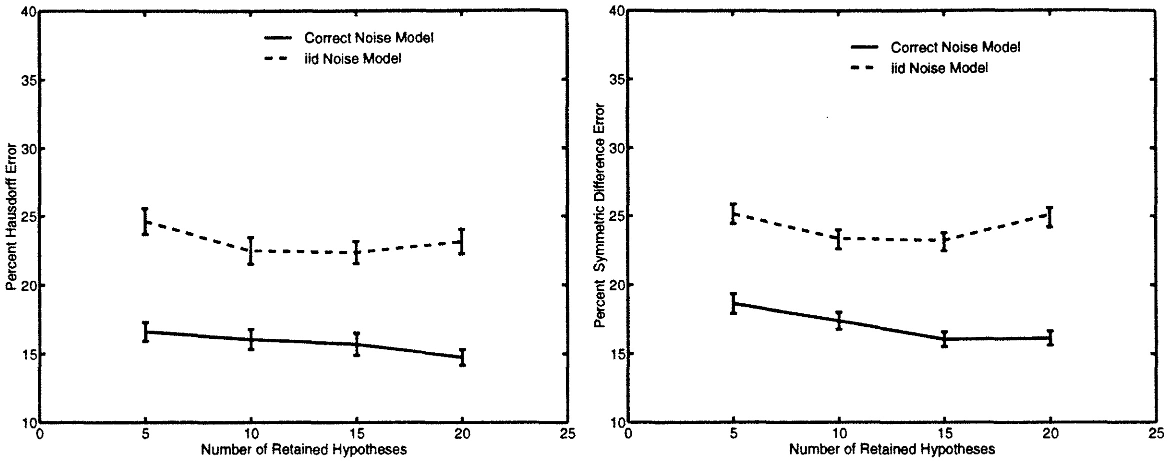

(a) Hausdorff Error

(b) Symmetric Difference Error

Figure 5-3: Performance as a Function of Number of Retained Hypotheses at Each Step of the Non-Sequential Reconstruction Algorithm.

Reconstruction Error as a Function of Retained Hypotheses

Figure 5-3 shows the performance of the Non-Sequential Algorithm as a function of the number of retained hypotheses, T, at each step of the algorithm. The projection data for each reconstruction consists of 18 views (for a total of 72 knot measurements and 90 chord measurements). Reconstructions were performed for $T=$ 5, 10, 15, and 20 retained hypotheses. Figure 5-3(a) is a plot of the Hausdorff error while Figure 5-3(b) is a plot of the Symmetric Difference error. In both plots, the solid line corresponds to the reconstruction assuming the correct applied model while the dashed line corresponds to the reconstruction assuming the i.i.d. model.

As expected, the number of retained hypotheses is inversely related to the reconstruction error. Unlike the Sequential Algorithm, these results do not display a huge reduction in the error as the number of hypotheses is increased. Instead, all of the error curves are somewhat flat and decrease relatively linearly as T is increased over the range from 5 to 20. This result is attributed to the improved triangulation geometry introduced by altering the order in which the views are entered into the algorithm. Basically, the algorithm is able to recognize the optimal hypothesis even when only a few hypotheses are retained at each step of the algorithm (unlike the Sequential Algorithm).

Consistent with the results of the previous section, the reconstructions that assume the correct applied noise model clearly outperform the reconstructions that assume the i.i.d. model. Further, despite the differences in the relative error, all of the curves have the same shape as a function of the number of retained hypotheses.

In addition, we also note that the level of error in the Non-Sequential Algorithm reconstructions is considerably less than those of the Sequential Algorithm for all of the performance measures. For example, in Figure 5-3(a) the percent Hausdorff error for the reconstruction assuming the correct applied noise model (solid line plot) ranges from 16.62% error with 5 retained hypotheses to 14.75% error with 20 retained hypotheses. The corresponding result for the Sequential Algorithm (shown by the solid line in Figure 4-8(a)) ranges from 26.64% error with 5 retained hypotheses to 17.99% error with 20 retained hypotheses. These results are representative of the performance comparisons of the reconstruction results of Figure 4-8 and Figure 5-3. This supports the claim that the improved triangulation geometry of the Non-Sequential Algorithm results in better reconstructions given the same input data.

Reconstruction Error as a Function of Chord Measurements

In this section we analyze the performance of the Non-Sequential Algorithm as a function of the number of chord measurements used in the reconstruction algorithm. Each reconstruction is based on 18 uniformly spaced views or a total of 72 knot location measurements. We set the variance on the chord length measurements to a constant value throughout this experiment. Thus, as the number of chord length measurements per view is varied, the SNR (as defined in Section 3.2.3) will also change. The constant variance on the chord lengths was chosen for the case of 10 chord measurements per view at a SNR of 10dB and the noise on the knot locations was set with Equation (3.3) using this constant chord variance. Noisy chord data were generated for 0, 5, 10, 15, and 20 chord measurements per view and the resulting Hausdorff and Symmetric Difference reconstruction errors are plotted in Figure 5-4 (a) and (b), respectively. As in the previous section, the solid line represents reconstructions assuming the correct applied knot location noise model while the dashed line represents reconstructions assuming the i.i.d. noise model.

The overall behavior of the curves in Figure 5-4 is as expected. The reconstruction error for both measures decreases as the number of chord measurements used in the algorithm is

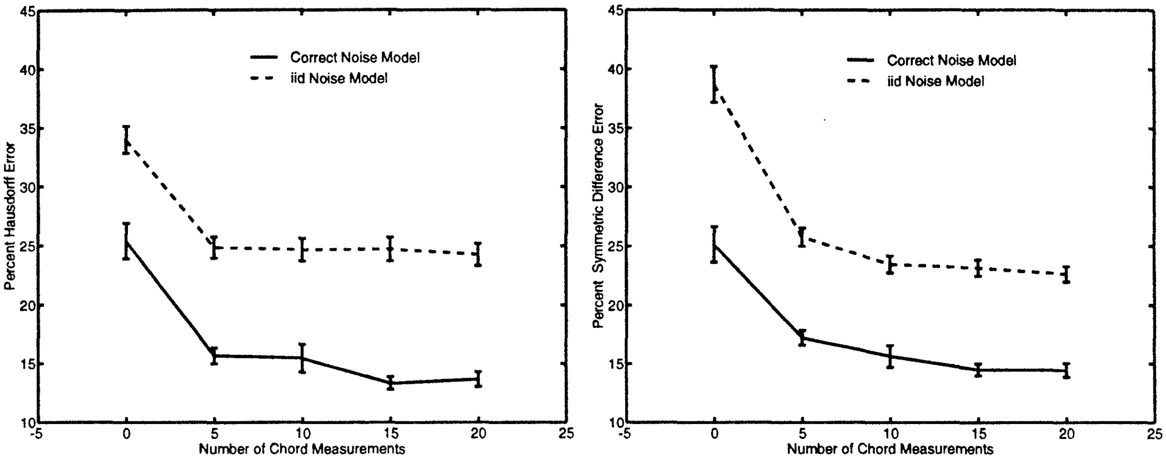

(a) Hausdorff Error

(b) Symmetric Difference Error

Figure 5-4: Performance as a Function of Number of Chords Used in the Non-Sequential Reconstruction Algorithm.

increased. As with the Sequential Algorithm, the most significant drop in error occurs as the number of chords is increased for 0 to 5 samples per view. Again this lends credibility to the importance of the chord measurements in determining the optimal hypothesis. As with the performance analysis of the Sequential Algorithm, all of the error curves behave similarly as a function of the number of chord measurements and the reconstructions that assume the correct applied knot location noise model outperform the reconstructions that assume the i.i.d. model. Further, as described in Section 4.5.3 of the previous chapter, the Symmetric Difference error is much more sensitive to the number of chord measurements, especially in the i.i.d. model reconstructions.

There are a number of significant differences between the reconstruction results presented here and the corresponding results of the Sequential Algorithm of Chapter 4. First, the relative level of error for all of the reconstruction scenarios is significantly less for the Non-Sequential Algorithm. Again, this is attributed to the improved triangulation geometry of the Non-Sequential Algorithm. The second difference is that the reconstruction errors assuming the correct knot location noise model are a much stronger function of the number of chord length measurements per view in the Non-Sequential Algorithm. This implies that the combination of the improved triangulation geometry coupled with the chord measurements help the algorithm to recognize the optimal hypothesis. This claim is particularly supported by comparing the reduction in both the Hausdorff error and the Symmetric Difference error as the number of chords is increased from 0 to 5 in the results assuming the correct noise model of Figure 4-9 and Figure 5-4. As these results demonstrate, the error reduction (for both measures) is much greater for the Non-Sequential Algorithm.

Reconstruction Error as a Function of SNR and Number of Views

This section presents the results of a series of independent Monte-Carlo simulations performed at various SNRs and numbers of views (refer to Figure 5-5) to characterize the effect of these parameters on the Non-Sequential Algorithm. Reconstructions are performed for SNRs of 0dB, 10dB, 20dB, and 30dB and numbers of views equal to 9, 18, 27, and 54. Figure 5-5(a) and (b) present the Hausdorff and Symmetric Difference errors, respectively, for reconstructions assuming the correct applied knot location noise model. Similarly, Figure 5-5(c) and (d) present the same error measures for reconstructions assuming an i.i.d. noise model for the knot location data.

There are a number of observations that are consistent with all of the results presented in Figure 5-5. The first observation is that the most dramatic decrease in error occurs when the SNR is increased from 0dB to 10dB. As with the Sequential Algorithm, the Non-Sequential Algorithm is still unable to resolve the knot-to-vertex data associations in a high noise scenario. The second observation common to all reconstructions is that all of the error curves behave similarly as a function of the number of views. Each error curve decreases initially and then flattens out as a function of the number of views for each SNR. Thus, at a given SNR all of the reconstructions behave similarly as a function of the number of views even though the relative error levels vary as a function of the SNR.

There are also a number of differences in the reconstruction results of Figure 5-5. The first is that the results that assume an i.i.d. knot location noise model are slightly larger than those that assume the correct applied noise model (as expected). The second difference is that at a SNR=0dB the Hausdorff error is higher than the Symmetric Difference error in both reconstruction cases. This difference is attributed to the different method of characterizing the distance between two convex sets employed by the two error measures (refer to Section 3.2.5).

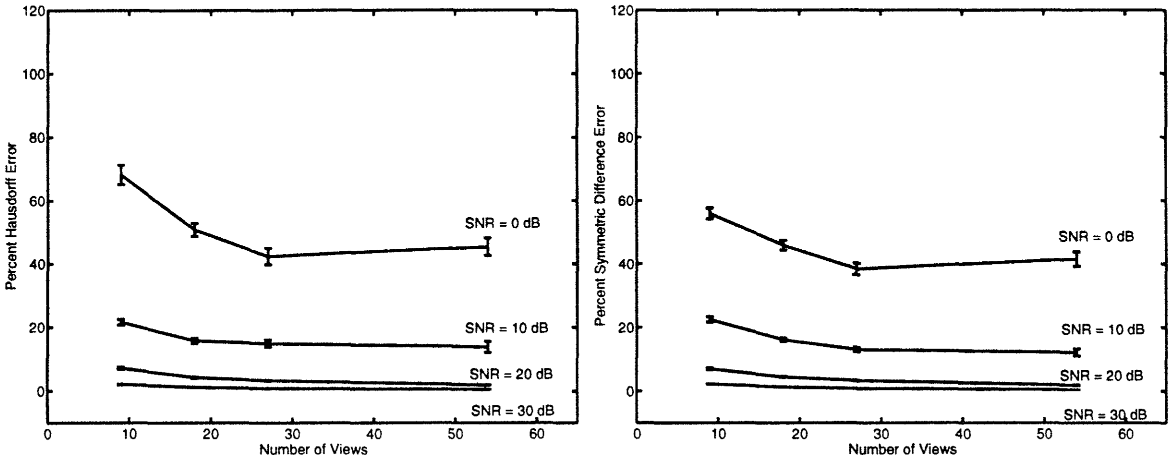

(a) Hausdorff Error, Correct Model (b) Symm. Diff. Error, Correct Model

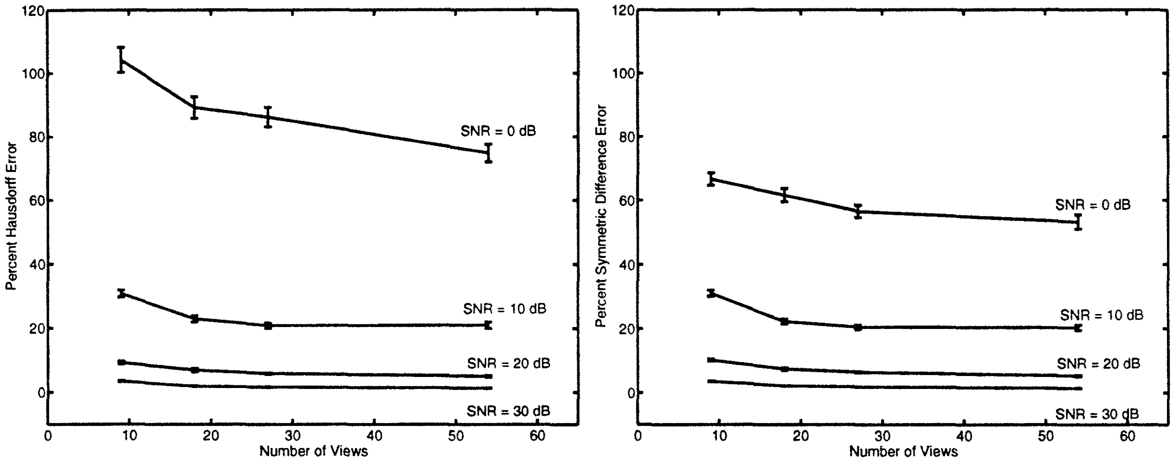

(c) Hausdorff Error, i.i.d. Model (d) Symm. Diff. Error, i.i.d. Model

Figure 5-5: Performance as a Function of Number of Views Used in the Non-Sequential Reconstruction Algorithm. (a),(b) Correct Noise Model; (c),(d) i.i.d. Noise Model

Finally, we comment on the differences between the reconstruction errors of the Non-Sequential Algorithm presented here and those of the Sequential Algorithm presented in Figure 4-10 in the previous chapter. The most obvious difference between these results is that the Non-Sequential Algorithm results do not degrade as the number of views becomes large. This difference occurs for two reasons. First, because of the improved triangulation geometry of the Non-Sequential Algorithm obtained by the non-sequential reordering of the views. Second, unlike the Sequential Algorithm, the Non-Sequential Algorithm does not have the problem of retaining the wrong hypotheses due to a series of poor views in close angular proximity that the Sequential Algorithm has (refer to Section 4.5.3). Additionally, we note that the relative level of error in all of the Non-Sequential reconstructions is less for all numbers of views and SNRs. Consistent with that observation, for the same number of experiments, the error bars of the Non-Sequential Algorithm are smaller implying that the variance of the estimates is also smaller.

Reconstruction Comparison of Two Test Objects

Having completed the Monte-Carlo analysis of the Non-Sequential Algorithm for a particular object, the kite, we now turn to a comparative analysis. Figure 5-6 shows the Monte-Carlo reconstruction errors for both the kite and the square test objects as a function of the number of views used in the reconstruction. Figure 5-6 (a) and (b) are Hausdorff and Symmetric Difference errors for reconstructions assuming the correct applied knot location noise model while Figure 5-6 (c) and (d) are the corresponding error measures under the assumption of an i.i.d. noise model.

The first observation that is common to all plots of Figure 5-6 is that the reconstruction of the kite outperformed that of the square for both the Hausdorff and Symmetric Difference error under both reconstruction scenarios (i.e. assuming the correct applied knot location noise model and assuming the i.i.d. noise model). This is attributed to the fact that three of the vertices of the kite are relatively close together. If the algorithm is off slightly, the reconstruction errors do not suffer as much as in the reconstruction of the square where the vertices are widely spaced. The difference in error seems to be more prevalent in the Hausdorff error measures. This is probably because a slight rotation and elongation of the square can cause large errors in the percent Hausdorff error. Regardless of these differences,

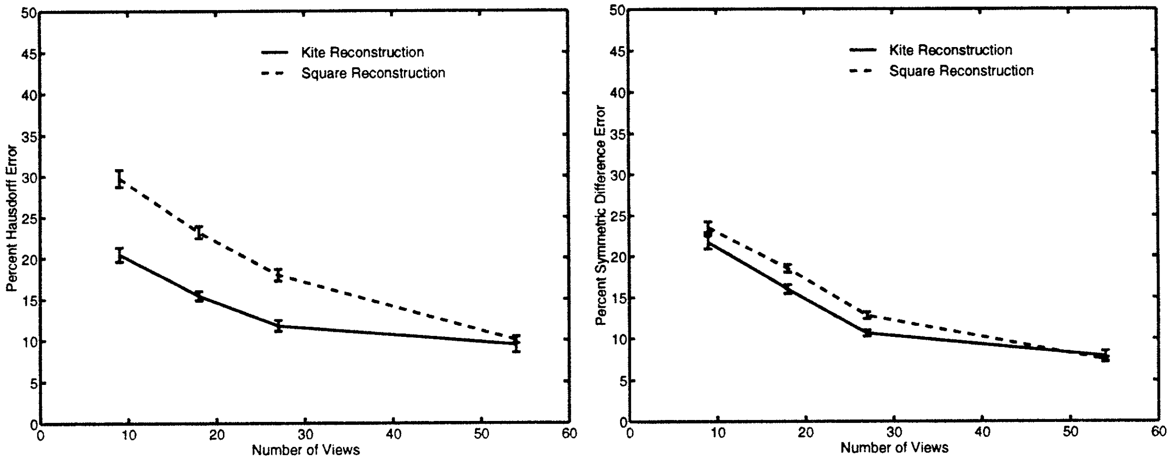

(a) Hausdorff Error, Correct Model

(b) Symm. Diff. Error, Correct Model

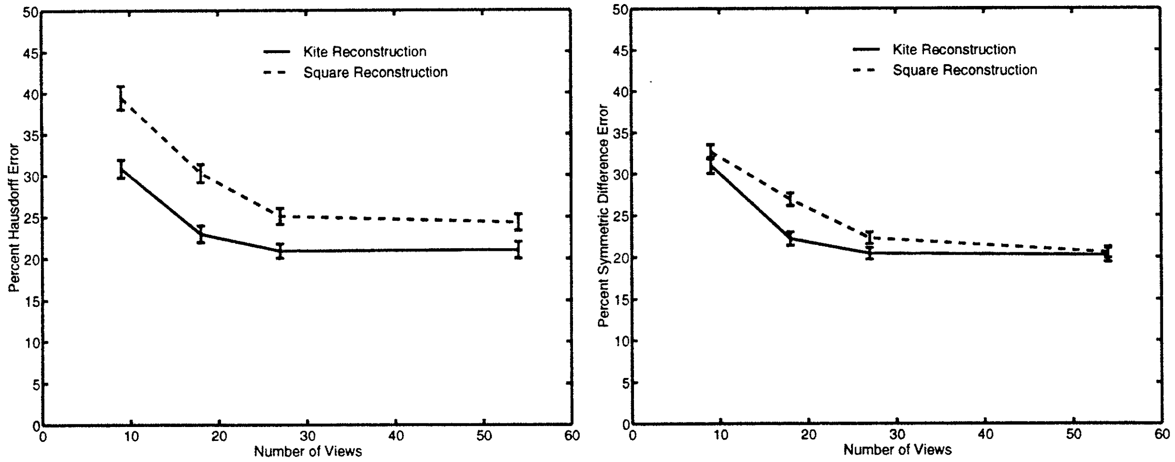

(c) Hausdorff Error, i.i.d. Model

(d) Symm. Diff. Error, i.i.d. Model

Figure 5-6: Comparison of Kite and Square Test Objects for Non-Sequential Reconstruction Algorithm: SNR=10dB, T=15 retained hypotheses, 5 chord measurements per view. (a),(b) Correct Noise Model; (c),(d) i.i.d. Noise Model.

however, the reconstruction errors of both objects seem to converge as the number of views is increased for both the Hausdorff and Symmetric Difference errors.

Another interesting observation is that the performance of both objects degrades by approximately the same amount as the assumed model is switched from the correct applied knot location noise model to the i.i.d. model. Thus, the noise modeling errors do not seem to effect one object more than the other.

The primary difference between the reconstructions presented in this section and the corresponding section of Chapter 4 is that the error decreases for large numbers of views in the Non-Sequential Algorithm. Again, this is attributed to the non-sequential ordering of views. An additional result of this improved triangulation geometry is that the performance of the kite is consistently better than that of the square for the Non-Sequential Algorithm. Thus, poor triangulation geometry seems to have a greater effect on the elongated kite object. In addition, consistent with the results of the previous section, for the same number of experiments the error bars of the Non-Sequential Algorithm are smaller than those of the Sequential Algorithm, implying that the variance of the estimates is smaller.

Despite the differences in the reconstruction errors of these two objects, the error plots of the kite and the square have the same general shape as a function of the number of views. Further, the reconstruction errors of both objects have the same order of magnitude. Therefore, as in the case of the Sequential Algorithm, the Non-Sequential Algorithm does not produce significantly better reconstructions for one object over the other.

5.3.3 Limited Angle Tomography and Non-Uniform Projection Angles

In this final section of experimental results we perform reconstructions using the Non-Sequential Algorithm over limited angular regions and with non-uniformly spaced projection data. As stated in the previous chapter on the Sequential Algorithm, this is an important feature of this algorithm as many classical reconstruction techniques require uniformly spaced angular projections over the complete range $[0°, 180°)$.

Reconstruction Error as a Function of Limited Angle Projection Data

In this section we evaluate the reconstruction performance of the Non-Sequential Algorithm over limited angular ranges. In each reconstruction, we generate a full set of projection data (i.e. chord measurements and knot location measurements) over the $[0°, 180°)$ angular range with 5 chord measurements per view at a SNR of 10dB. Like all previous experiments, the chord length measurements are uniformly sampled over the range $t = [-1, 1]$. Next, we perform reconstructions on ***subsets*** of the projection data over limited angular ranges using the Non-Sequential Reconstruction Algorithm. Specifically, reconstructions were performed over the following angular ranges: $[0°, 20°)$, $[0°, 35°)$, $[0°, 45°)$, $[0°, 70°)$, and $[0°, 90°)$. Each subset of projection data contains 18 uniformly spaced views over the specified angular range. For each angular range, 100 independent Monte-Carlo reconstructions were performed with T=15 hypotheses retained at each step of the Non-Sequential Algorithm.

Figure 5-7 shows the reconstruction error results of reconstructions of the kite over limited angular ranges. As the results from both the Hausdorff and Symmetric Difference errors in Figure 5-7(a) and Figure 5-7(b) demonstrate, the reconstruction errors decrease significantly as the maximum angular view is increased, with the greatest drop occurring for the increase from 20° to 30°. This result is attributed to poor triangulation geometry over the smaller angular ranges. In addition, we note that when the view range reaches 45° both the Hausdorff and the Symmetric Difference error measures appear to level off as a function of the angular range. This result occurs for the reconstruction assuming the correct applied noise model and also for reconstruction assuming the i.i.d. noise model. Thus, the algorithm appears to obtain a robust estimate for the kite (for the case of 18 views) if the view range is greater than 45°. Also, although the relative level of the i.i.d. reconstruction is larger for both measures, we note that it behaves exactly as the correct model as a function of the view angle range.

In addition to the reconstruction errors resulting from the Monte-Carlo reconstruction results, Figure 5-8 shows a sample reconstruction resulting from a single sample path. This reconstruction was generated under the conditions listed above with projection data over the angular range $[0°, 90°)$. As the resulting plot shows, the algorithm was able to produce a good quality reconstruction in the face of limitations on both the angular range and the number of projections (18 views) in a low SNR situation (10dB). Additionally, Table 5.3

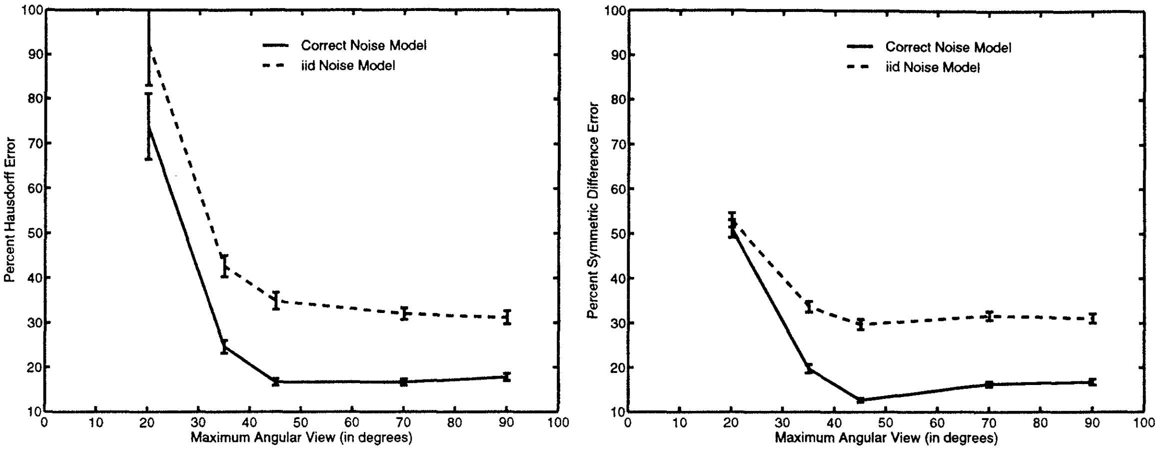

(a) Hausdorff Error

(b) Symmetric Difference Error

Figure 5-7: Performance as a Function of the Angular Range Using the Non-Sequential Algorithm.

Error Measures	True Model	i.i.d. Model
Hausdorff Error (%)	7.32	16.95
Symm. Diff. Error (%)	12.97	24.20

Table 5.3: Sample Limited Angle Reconstruction Over the Range $[0°, 90°)$ Using the Non-Sequential Algorithm.

summarizes the reconstruction error for this sample path.

In addition to the above experiment, we performed a second set of Monte-Carlo simulations over the range $[0°, 45°)$. For this case, 100 reconstructions were performed for cases of 4, 5, 9, 18, and 27 views. The remainder of the parameters (chord length, SNR, etc.) were set exactly as in the previous experiment.

The results of this experiment, shown in Figure 5-9, show that the limited angle reconstruction errors decrease as a function of the number of views used to perform the reconstruction for the Non-Sequential Algorithm. As the number of views were increased while holding the angular view range at 45°, both the Hausdorff and Symmetric Difference errors decreased significantly. In addition, although the reconstruction errors assuming the i.i.d. noise model are higher (as expected), the behavior as a function of the number of views is consistent with the reconstructions assuming the correct applied noise model. Overall, the results of Figure 5-9 are significant in that they show that the Non-Sequential Algorithm

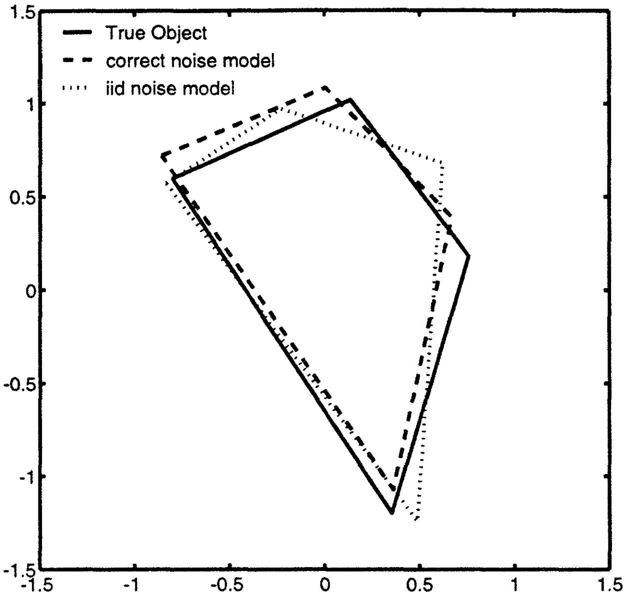

Figure 5-8: Sample Reconstruction Over Limited Angular Range ($[0°, 90°)$) Using the Non-Sequential Algorithm.

behaves robustly under limited angle reconstruction scenarios.

Reconstruction Using Non-Uniform Angular Spacing

In addition to the ability to reconstruct over limited angular ranges, the Sequential Algorithm also has the ability to perform reconstructions on non-uniformly spaced projection data. Figure 5-10 shows a sample reconstruction from 27 non-uniformly spaced views over the $[0°, 180°)$ angular range. In this reconstruction, each projection contained 5 chord measurements at a SNR of 10dB with $T = 15$ hypotheses retained at each step of the algorithm. To determine the projection angles, we started with the initial projection at $\theta_1 = 0°$. The second projection was at $\theta_2 = 3.5°$ and all subsequent angles were generated with the following equation

$$\theta_i \;=\; \frac{7}{6}\theta_{i-1} \;<\; 180° \tag{5.2}$$

The solid line in Figure 5-10 represents the true object and the dashed and dotted lines are the reconstructions assuming the correct knot location noise model and the i.i.d noise model, respectively. In addition, Table 5.4 summarizes the Hausdorff and Symmetric Difference reconstruction errors for these results. Although only for one sample path, these results

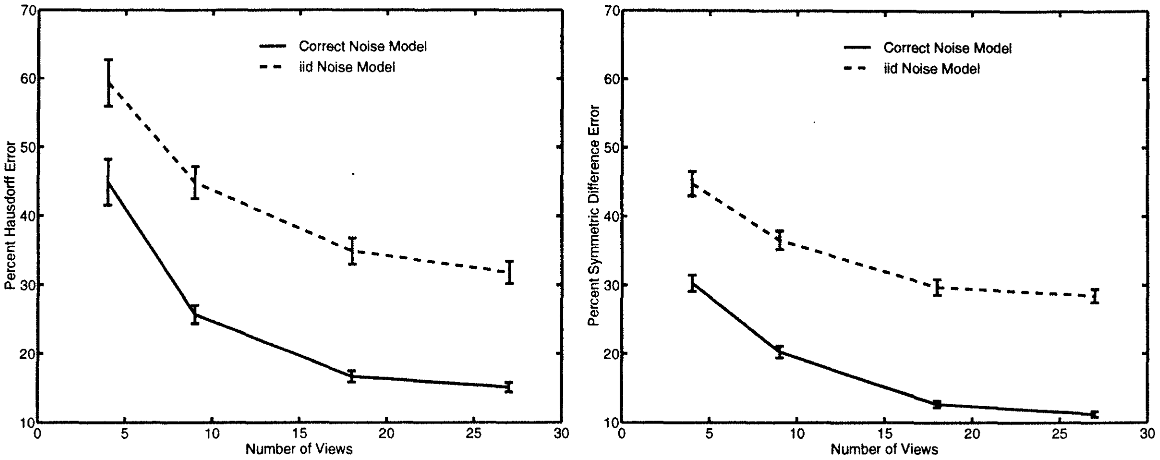

(a) Hausdorff Error

(b) Symmetric Difference Error

Figure 5-9: Performance as a Function of Number of Views with Projection Angle Range of 45° Using the Non-Sequential Algorithm.

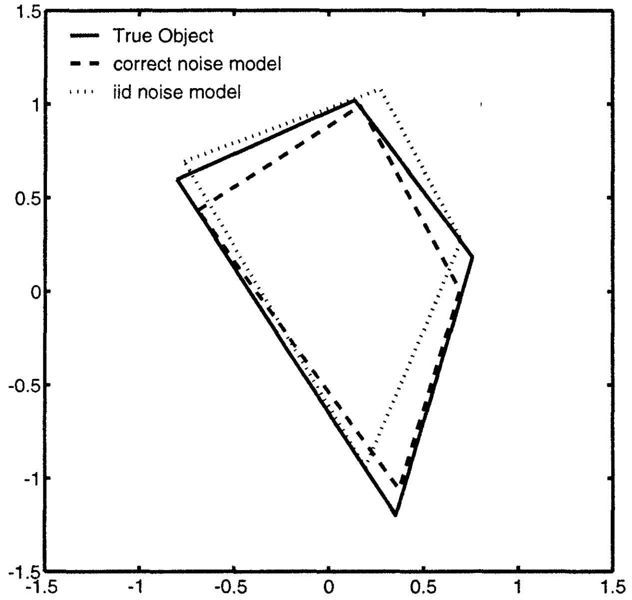

Figure 5-10: Sample Reconstruction with Non-Uniformly Spaced Projection Angles Using the Non-Sequential Algorithm.

Error Measures	True Model	i.i.d. Model
Hausdorff Error (%)	15.91	24.21
Symm. Diff. Error (%)	17.75	22.76

Table 5.4: Comparison of Errors for Non-Uniform Angular Spacing Using the Non-Sequential Algorithm.

show the performance of this algorithm is not limited to a uniformly spaced projection data set.

5.4 Conclusions

In this chapter, we extended the finite parameter reconstruction algorithm of the previous chapter to allow views to be added into the algorithm non-sequentially. At the expense of a slightly more complex method of evaluating the discrete hypotheses, the new algorithm resulted in both increased flexibility (views can be input into the algorithm over any angular range in any angular order) and improved reconstruction performance (as measured by the Hausdorff and Symmetric Difference percent errors).

In every reconstruction scenario the Non-Sequential Algorithm results outperformed those of the Sequential Algorithm of the previous chapter. The overall performance of this algorithm was consistent and produced robust object estimates for relatively small amounts of projection data. In addition to producing better reconstructions, the Non-Sequential Algorithm also supported the results of the previous chapter in the assessment of the importance of chord length measurements in improving the hypothesis selection performance. Unlike the Sequential Algorithm, the results of the Non-Sequential Algorithm did not degrade as the number of views increased for a fixed number of retained hypotheses. The Non-Sequential Algorithm was also less sensitive to the number of retained hypotheses at each step of the algorithm. Additionally, the reconstruction error for both the Hausdorff and Symmetric Difference errors approached a limiting value as each of the parameters (number of hypotheses, number of chords, number of views) was increased. This implies a fixed level of performance for a given amount of noise in the measurements. Finally, the Non-Sequential Algorithm showed a marked improvement over the Sequential Algorithm in its ability to perform robust estimates over limited angular regions. In the next chapter, we seek to overcome this limiting threshold by using the output of the Sequential Algorithm or Non-Sequential Algorithm as input to a nonlinear reconstruction algorithm that reconstructs the object's vertices from knot location *and* chord length measurements.

Chapter 6

Nonlinear Reconstruction Algorithm

This chapter presents the third and final reconstruction algorithm developed in this thesis. Unlike the previous two algorithms, this algorithm is a fully nonlinear reconstruction algorithm that uses both knot location *and* chord length data to estimate the vertex locations of a convex, binary polygonal object. Section 6.1 presents an overview of this algorithm including initialization, cost criterion, and a block diagram of the overall reconstruction scenario. Sample reconstructions and Monte-Carlo simulations are presented in Section 6.2 and conclusions regarding the results of this algorithm are found in Section 6.3.

6.1 Overview of Nonlinear Algorithm

In this chapter we develop a nonlinear estimation algorithm to estimate the vertices of a binary polygonal object from noisy measurements of knot locations and chord lengths. Like the Sequential and Non-Sequential Algorithms developed in the previous two chapters, the estimate is obtained by using ML estimation techniques. Specifically, we seek the value of the vertices, V, which makes the observed data most likely

$$\widehat{V}_{ML} = \arg\max_{V} \; p_{x|V}(\underline{X} \mid V) \tag{6.1}$$

where $p_{x|V}(\underline{X} \mid V)$ denotes the conditional probability density of the observations, $\underline{X}$, given the vector V of vertices of the object. Because the measurements are corrupted by independent, additive Gaussian noise, this formulation reduces to (refer to Section 2.1.2)

$$\hat{V}_{ML} = \arg\min_{V} \quad (\underline{X} - \bar{X})^T \Lambda_{x|V}^{-1} (\underline{X} - \bar{X}) \tag{6.2}$$

where $\bar{X}$ and $\Lambda_{x|v}$ denote the mean and covariance of the conditional density $p_{x|V}(\underline{X}|\ V)$.

In the previous algorithms developed in this thesis, vertex estimates were based solely on the knot location data and a set of discrete data association hypotheses. Under each hypothesis, we found that the knot-to-vertex estimation problem was linear and Equation (4.8) was used to determine the ML estimate of the vertices based on the knot location data. The chord length data, in addition to the knot location data, were used to prune the set of discrete data association hypotheses.

Unlike the previous method, however, the Nonlinear Algorithm developed in this chapter obtains the estimate defined by the cost criterion of Equation (6.2) from measurements of knot locations *and* chord lengths. The measurement equations relating the knot location measurements, Z, and the chord length measurements, M, to the vertices of the object were first defined in Section 4.3 and Section 4.4. We repeat them here for convenience

$$\begin{aligned} Z &= H(V) + \mathcal{N} \\ M &= F(V) + \Upsilon \end{aligned}$$

where H and F are system matrices that rely nonlinearly on the vertices V, and, $\mathcal{N}$ and Υ are independent Gaussian noise processes with covariances Λ_z and Λ_m, respectively. If these two equations are *stacked* to form a single measurement equation, the resulting equation becomes

$$\underline{X} = \begin{bmatrix} Z \\ M \end{bmatrix} = \begin{bmatrix} H(V) + \mathcal{N} \\ F(V) + \Upsilon \end{bmatrix} \tag{6.3}$$

and the mean and covariance matrix corresponding to the conditional probability density function $p_{x|V}(\underline{X}|V)$ are

$$\bar{X} = \begin{bmatrix} H(V) \\ F(V) \end{bmatrix} \tag{6.4}$$

$$\Lambda_{x|V} = \begin{bmatrix} \Lambda_z & 0 \\ 0 & \Lambda_m \end{bmatrix} \tag{6.5}$$

Note that Λ_z, Λ_m, and hence $\Lambda_{x|V}$ are diagonal matrices because of the assumptions of independent noise on the chord length and knot location measurements[1]. If we substitute the above values into the ML estimation equation given by Equation (6.2), the ML estimate of the vertices can be written

$$\begin{aligned} \widehat{V}_{ML} &= \arg\min_{V} \left(\begin{bmatrix} Z \\ M \end{bmatrix} - \begin{bmatrix} H(V) \\ F(V) \end{bmatrix} \right)^T \begin{bmatrix} \Lambda_z & 0 \\ 0 & \Lambda_m \end{bmatrix}^{-1} \left(\begin{bmatrix} Z \\ M \end{bmatrix} - \begin{bmatrix} H(V) \\ F(V) \end{bmatrix} \right) \\ &= \arg\min_{V} \; (Z - H(V))^T \Lambda_z^{-1} (Z - H(V)) + (M - F(V))^T \Lambda_m^{-1} (M - F(V)) \quad (6.6) \end{aligned}$$

which is a Weighted Nonlinear Least Squares Error (WNLSE) formulation. The terms of the matrix $\Lambda_{x|V}^{-1}$, which is a diagonal matrix because of the assumption of uncorrelated noise, comprise the weights in the least squares formulation. Because the inverse covariance matrix is often considered a measure of "information", the interpretation of $\Lambda_{x|V}^{-1}$ as weighting matrix makes intuitive sense.

Because of the nonlinear dependence of the system matrices H and F on the vertices of the object, this problem is a highly nonlinear minimization problem. Our approach to solving Equation (6.6) given measurements of chord lengths and knot locations is via a nonlinear optimization routine that will minimize Equation (6.6) with respect to V given measurement data, noise statistics, and an initial guess, $\widehat{V}_o$. Because the search space is $2n$ dimensional for an n-gon, the cost criterion given by Equation (6.6) is a complicated cost function. To avoid problems of differentiability of the cost function, we use the Nelder-Mead nonlinear optimization algorithm[2] to minimize Equation (6.6) with respect to V. This algorithm was chosen because of its ability to deal robustly with highly discontinuous cost criteria. Further, as a nonlinear function of a multi-dimensional vector, we expect the cost criterion given by Equation (6.6) to be plagued by numerous extraneous minima. Thus, to determine the global minima, a good initial guess is essential. We use the output of either the Sequential Algorithm of Chapter 4 or the Non-Sequential Algorithm of Chapter 5 to

---

[1] Refer to Section 3.2.3 for a complete description of the noise models used throughout this thesis.

[2] This algorithm was implemented via the Matlab function FMINS.

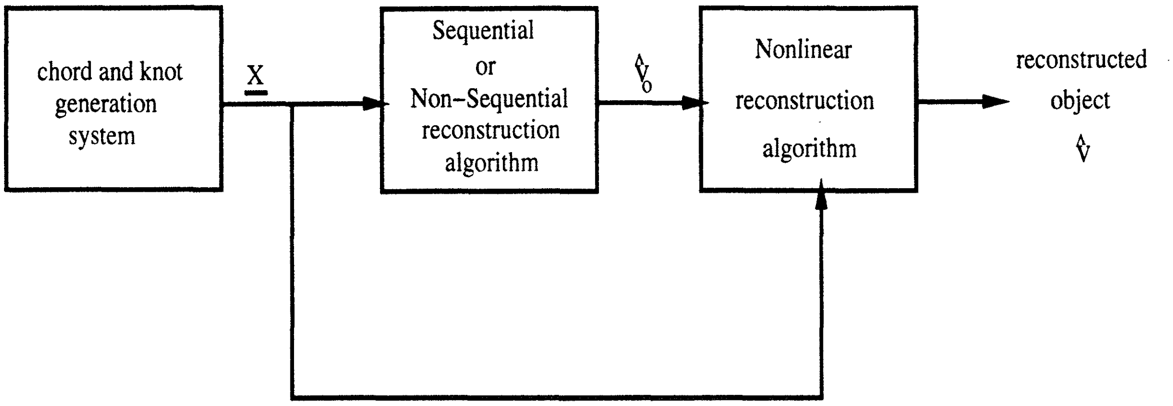

Figure 6-1: Block Diagram of Nonlinear Reconstruction Scenario.

provide the initial estimate of V to used in this Nonlinear Algorithm.

A block diagram of the reconstruction scenario is shown in Figure 6-1. As described in Section 3.2.1, chord length and knot location data are obtained from the underlying object and perturbed with additive, independent Gaussian noise (hence the knot and chord generation block in Figure 6-1). The Sequential Algorithm or Non-Sequential Algorithm estimates the vertices of the object via the simultaneous solution of a discrete-valued hypotheses testing problem and a continuous-valued ML estimation problem. As outlined in Section 4.3, by conditioning on the discrete data association hypotheses the knot location measurement equation is a linear function of the vertices of the object. Thus, the ML estimate of each vertex for a given hypothesis can be obtained using Equation (4.8). Both knot location and chord length data are used to prune the set of discrete data association hypotheses at each step of the algorithm with the likelihood ratio test given by Equation (4.11). The estimate of V corresponding to the optimal hypothesis (and based on knot location data alone) is then used as the initial estimate for the Nonlinear Algorithm. This initial guess, $\widehat{V}_o$, and the original noisy knot location data and the chord length data are then input to the Nonlinear Reconstruction Algorithm, which finds the vertex estimates by minimizing the ML cost criterion of Equation (6.6) based on measurements of *both* knot location data and chord length data and produces the final estimate $\widehat{V}$. Experimental Results of this algorithm follow in the next section.

6.2 Experimental Results

As with the previous two reconstruction algorithms, the performance of the Nonlinear Reconstruction Algorithm is tested with a combination of sample reconstructions and Monte-Carlo simulations. There are a number of specifications that are common to all of the reconstructions performed in this section. As before, the kite and square test objects first described in Section 4.5.1 are used to test the reconstruction algorithm. The initial estimates of the vertices of the object that are input to the Nonlinear Algorithm are obtained by using either the Sequential Reconstruction Algorithm of Chapter 4 or the Non-Sequential Reconstruction Algorithm of Chapter 5. The i.i.d. knot location noise model (presented in Section 3.2.3) is assumed for *all* of the reconstruction of this chapter. In addition, a number of assumptions used in the previous two chapters also hold for the this chapter; namely, the Max-Min Approach (described in Section 3.2.4) is used to determine the initial view, the chord measurements are uniformly sampled over the region $t = [-1, 1]$, and the percent Hausdorff error and percent Symmetric Difference error (described in Section 3.2.5) are used to characterize reconstruction performance. Additionally, for the Non-Sequential Algorithm, the reordering of views outlined in Section 5.3 is used throughout this chapter when the Non-Sequential Algorithm is used to calculate the initial guess.

6.2.1 Sample Reconstructions

This section contains results of sample reconstructions using the Nonlinear Reconstruction Algorithm. Note that for all of these results the true object is depicted by solid lines, the initial guess is represented by the dash-dotted lines, and the final nonlinear reconstruction is denoted by dotted lines. The sample reconstructions of this section use the same data set used to generate the sample reconstructions of the Sequential Algorithm (refer to Section 4.5.2) and the Non-Sequential Algorithm (refer to Section 5.3.1). For both objects, the square and the kite, noisy chord length and knot location data were generated so that the SNR on the chord length data was set to 10dB and in each reconstruction the assumed knot location noise model was the i.i.d. model described in Section 3.2.3. Projection data was taken at 27 equally spaced angles over the interval $[0°, 180°)$ with 5 equally spaced chord measurements over the range $t = [-1, 1]$ in each view. Thus, each data set had a total of 108 knot location measurements and 135 chord measurements.

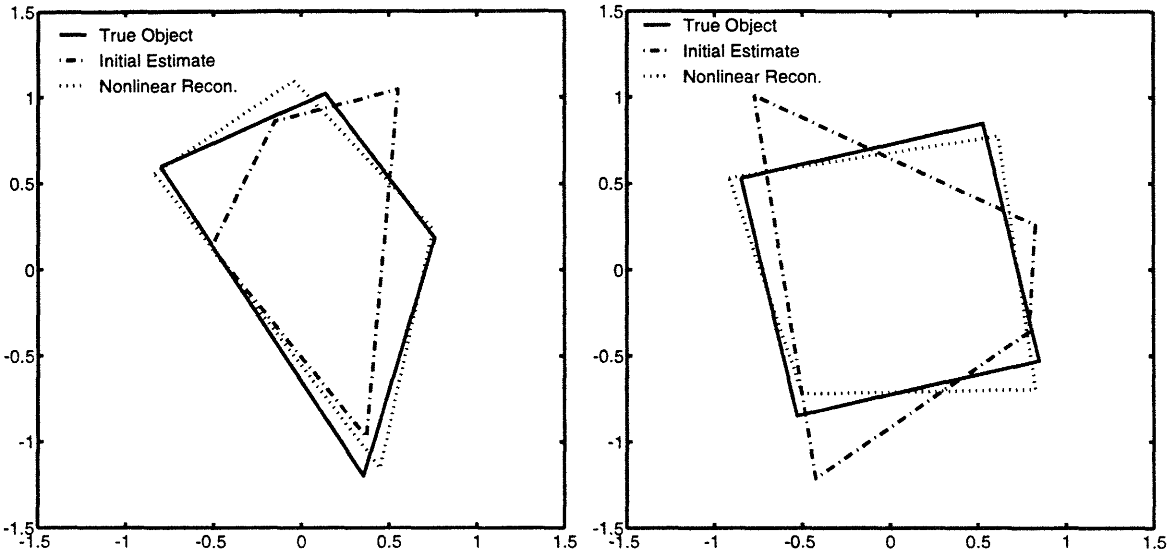

(a) Kite Reconstruction

(b) Square Reconstruction

Figure 6-2: Sample Reconstructions of Nonlinear Algorithm (Seq. Alg. initial guess), 27 views, 10 hypothesis pruning, 10dB SNR on chords.

Error Measures	Kite Test Object		Square Test Object	
	$\hat{V}_o$	$\hat{V}$	$\hat{V}_o$	$\hat{V}$
Hausdorff Error (%)	36.67	10.75	44.61	15.66
Symm. Diff. Error (%)	38.42	11.51	33.62	11.66

Table 6.1: Comparison of Errors for Sample Nonlinear Reconstructions - Sequential Algorithm Initial Guess.

Figure 6-2 (a) and (b) show sample reconstructions of the kite and square using the Sequential Reconstruction Algorithm to generate the initial guess. Table 6.1 summarizes the Hausdorff error and Symmetric Difference error for the sample reconstructions of Figure 6-2. Similarly, Figure 6-3 (a) and (b) show sample reconstructions of the kite and square using the Non-Sequential Algorithm to generate the initial guess. Table 6.2 summarizes the Hausdorff error and Symmetric Difference error for the sample reconstructions of Figure 6-3.

As the results in Table 6.1 and Table 6.2 summarize, the Nonlinear Reconstruction Algorithm significantly improves the reconstruction results of both objects from initial estimates provided by either the Sequential or Non-Sequential Algorithm. The reconstruction errors of the initial estimates were halved in nearly all of the cases. Additionally, the non-

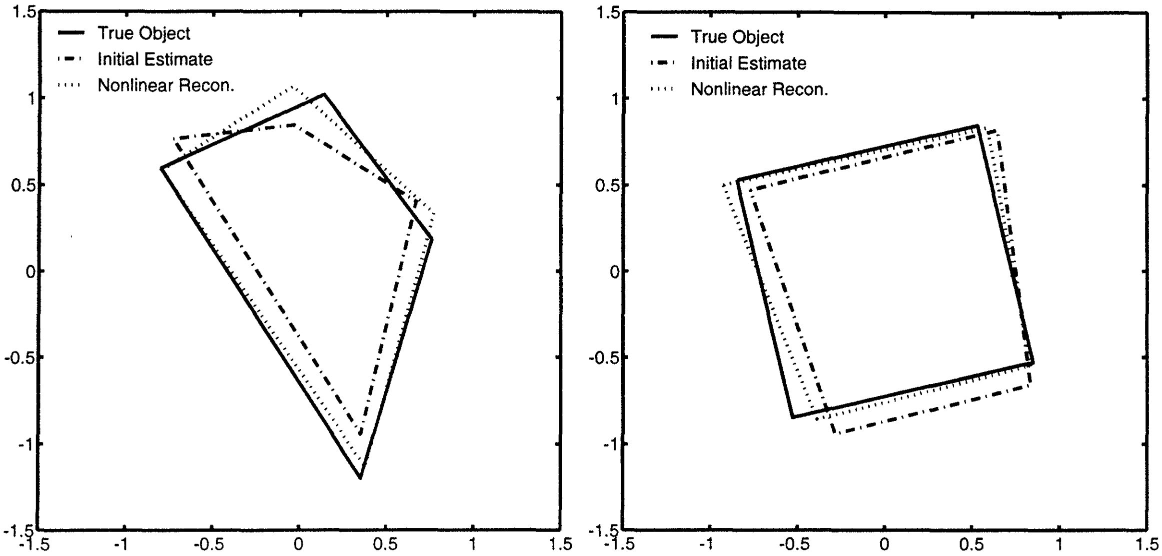

(a) Kite Reconstruction

(b) Square Reconstruction

Figure 6-3: Sample Reconstructions of Nonlinear Algorithm (Non-Seq. Alg. initial guess), 27 views, 10 hypothesis pruning, 10dB SNR on chords.

Error Measures	Kite Test Object		Square Test Object	
	$\widehat{V}_o$	$\widehat{V}$	$\widehat{V}_o$	$\widehat{V}$
Hausdorff Error (%)	20.41	8.88	19.22	12.60
Symm. Diff. Error (%)	32.66	9.52	24.19	12.61

Table 6.2: Comparison of Errors for Sample Nonlinear Reconstructions - Non-Sequential Algorithm Initial Guess.

linear reconstructions for each object have nearly the same errors, regardless of whether the Sequential and Non-Sequential Algorithm was used to generate the initial estimate of the objects vertices (although the Non-Sequential Algorithm does perform *slightly* better in all of the cases). Yet, despite the similar levels of error, as a comparison of the reconstructed objects of Figure 6-2 and Figure 6-3 show, the reconstructions for each object are slightly different for different initial estimates of the vertex locations. This implies that the cost function of Equation (6.6) is indeed a complicated function with numerous local minima. For example, referring to the square reconstructions of Figure 6-2(b) and Figure 6-3(b) we note that although both reconstructions had the same measurement data, they had different initial vertex estimates and ultimately converged to different final estimates of the vertices. Despite this, both reconstructions produce reconstruction errors that are very close for both the Hausdorff and Symmetric Difference errors.

The reconstruction errors of the kite were less than those of the square for this sample path in all of the Nonlinear Reconstruction results. Again, we acknowledge this result but defer comparative analysis of the performance of the algorithm in reconstructing these two objects to the Monte-Carlo results of Section 6.2.2.

6.2.2 Monte-Carlo Simulations

In the following sections Monte-Carlo simulations of reconstructions are used to test the performance of the Nonlinear Reconstruction Algorithm. As in the previous chapter, all simulations are performed using the kite test object (with the exception of the last section which is a comparison of reconstructions of the kite and square). The Monte-Carlo simulations consist of 50 independent reconstructions for each reconstruction scenario. In each plot error bars denote the 95% confidence intervals of the sample mean values that result from the 50 runs of the algorithm. The initial estimates of the vertices of the object are generated using the Non-Sequential Algorithm with the i.i.d. assumed knot location noise model[3]. The projection data for all simulations are uniformly spaced samples over the angular region $[0^\circ, 180^\circ)$. Finally, *unless otherwise stated*, $T = 15$ hypotheses are retained in each step of the Non-Sequential Algorithm, 5 uniformly spaced chord measurements over the range $t = [-1, 1]$ are taken per view, and the SNR on the chord length measurements

---

[3] Refer to Section 3.2.3 for a complete discussion of the noise models used in this thesis.

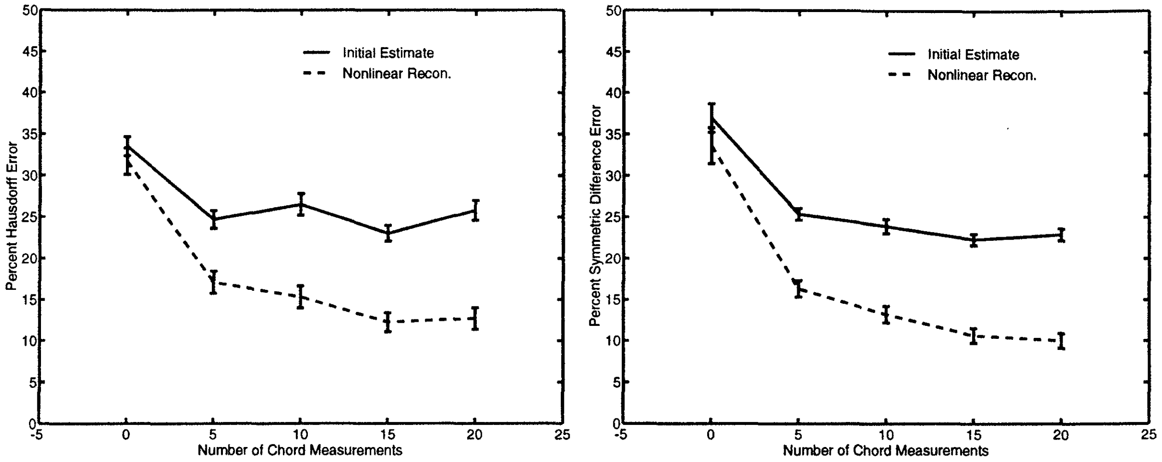

(a) Hausdorff Error

(b) Symmetric Difference Error

Figure 6-4: Performance as a Function of Number of Chords Used in the Nonlinear Reconstruction Algorithm.

is set to 10dB.

Reconstruction Error as a Function of Chord Measurements

In this section we analyze the performance of the Nonlinear Reconstruction Algorithm as a function of the number of chord measurements used in the reconstruction algorithm. Each reconstruction is based on 18 uniformly spaced views or a total of 72 knot location measurements. We set the variance on the chord length measurements to a constant value throughout this experiment. Thus, as the number of chord length measurements per view is varied, the SNR (as defined in Section 3.2.3) will also change. The constant variance on the chord lengths was chosen for the case of 10 chord measurements per view at a SNR of 10dB and the noise on the knot locations was set with Equation (3.3) using this constant chord variance. Noisy chord data were generated for 0, 5, 10, 15, and 20 chord measurements per view and the resulting Hausdorff and Symmetric Difference reconstruction errors are plotted in Figure 6-4 (a) and (b), respectively. The solid line represents reconstruction errors of the initial estimate generated from the Non-Sequential Algorithm while the dashed line represents reconstruction errors from the Nonlinear Reconstruction Algorithm.

The error measures of Figure 6-4 behave as expected as a function of the number of chord measurements. The Nonlinear Algorithm reduces the overall reconstruction error for both the Hausdorff error and Symmetric Difference error. Further, as expected, as the number of chord length measurements increases, the difference between the reconstruction errors of the initial estimate and those of the final nonlinear reconstruction become more pronounced. This occurs because the Nonlinear Algorithm uses the information in the chord measurements to obtain the estimate of the vertices while the Non-Sequential Algorithm (used to generate the initial estimate) only uses chord length to prune the possible discrete data association hypotheses. Thus, with more chord data, the Nonlinear Algorithm obtains a better estimate while the initial estimate from the Non-Sequential Algorithm levels off as the number of chords increases because it has obtained the best estimate that it can reach from the set of discrete hypotheses.

As with the reconstruction results of the first two algorithms, the greatest reduction in error (for both error measures) occurs when the number of chord measurements is increased from 0 to 5 samples per view. This makes sense in the Nonlinear Algorithm for two reasons. First, the Non-Sequential Algorithm produces a better initial estimate for the Nonlinear Algorithm (as shown in Figure 6-4); second, because the Nonlinear Algorithm uses chord length data to estimate the vertex locations (unlike the algorithms developed previously). Therefore, at the first inclusion of chord measurements we expect that the estimate would improve significantly. As the error curves of Figure 6-4 demonstrate, the Hausdorff error and Symmetric Difference error from the Nonlinear Algorithm are halved as the number of chords is increased from 0 to 5 measurements per view.

Reconstruction Error as a Function of SNR and Number of Views

This section presents the results of a series of independent Monte-Carlo simulations performed at various SNRs and numbers of views (refer to Figure 6-5) to characterize the effect of these parameters on the Nonlinear Algorithm. Reconstructions are performed for SNRs of 0dB, 10dB, 20dB, and 30dB and numbers of views equal to 9, 18, 27, and 54. Figure 6-5 presents the Hausdorff and Symmetric Difference errors for both the initial estimate (generated from the Non-Sequential Algorithm) and the nonlinear reconstruction for four cases: (a) SNR = 0 dB, (b) SNR = 10dB, (c) SNR = 20dB, and (d) SNR = 30dB.

There are a number of interesting observations regarding the results presented in Figure 6-5. The first observation is that all of the error curves, whether from the initial estimate

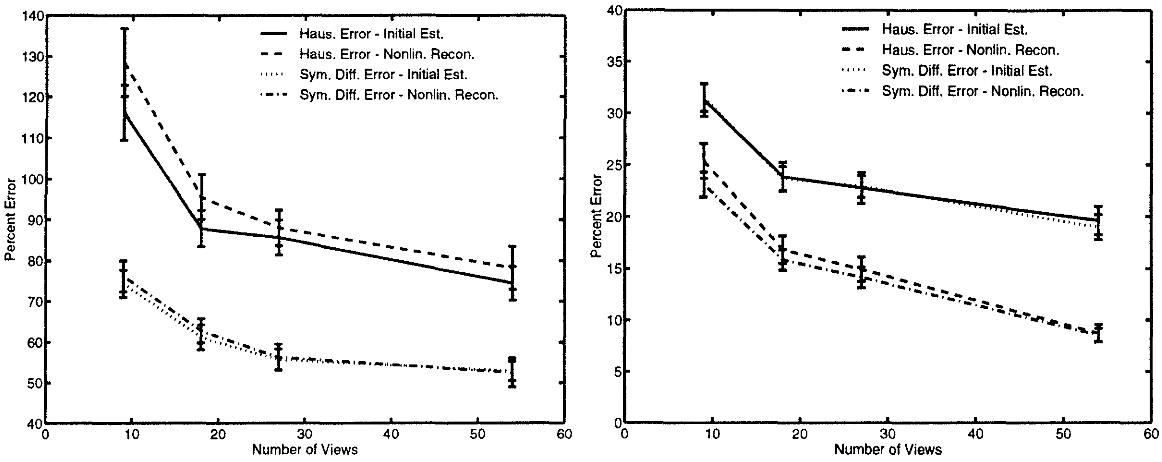

(a) SNR = 0dB

(b) SNR = 10dB

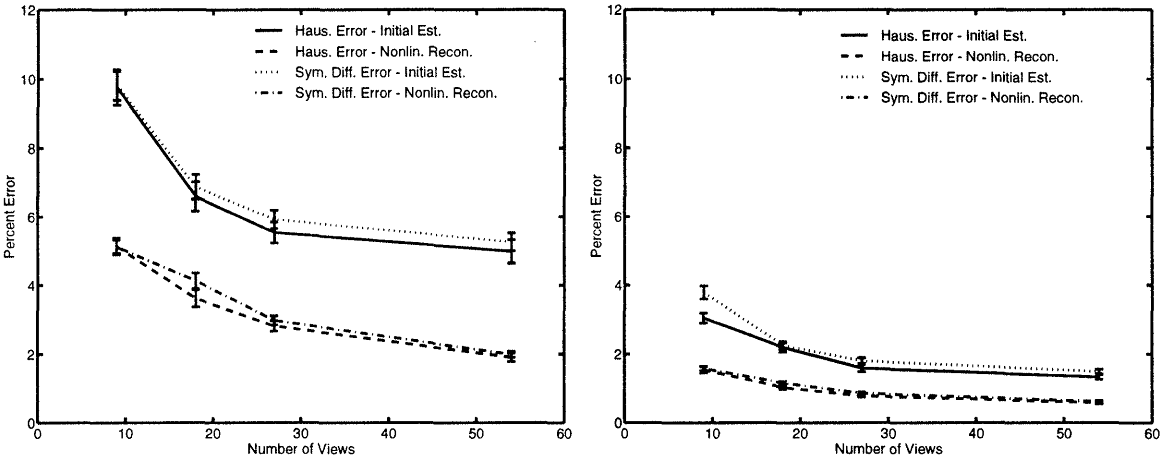

(c) SNR = 20dB

(d) SNR = 30dB

Figure 6-5: Performance as a Function of Number of Views Used in the Nonlinear Reconstruction Algorithm at different SNRs.

or the nonlinear reconstruction, behave similarly as a function of the number of views. As expected, each error curve decreases as the number of views is increased. Additionally, the reconstruction error decreased as a function of increasing SNR for all plots (as expected). We also note that the Nonlinear Reconstruction Algorithm produces significantly better reconstructions as measured by the Hausdorff error and Symmetric Difference error for all cases except those at an SNR equal to 0dB. Consistent with the results of the previous section, we see that the nonlinear reconstruction results halve the errors for cases of 18, 27, or 54 views (with a slightly smaller error reduction in the case of 9 views). Thus, the sample reconstructions of Figure 6-2 and Figure 6-3 are representative of the kind of error reduction possible with the nonlinear algorithm (as noted in the previous section, the reconstruction errors are nearly halved in the sample reconstructions of Section 6.2.1). In addition, for the same number of experiments the error bars of the Nonlinear Algorithm's results are smaller, implying that the variance of the estimates is smaller. As for the reconstruction results at a SNR of 0dB, the errors are also consistent with the results of the previous two chapters. The initial estimate provided to the nonlinear optimization algorithm is too far away from the true vertices of the object, supporting the claim that a good initial guess is important for the nonlinear optimization routine. Additionally, the knot location and chord length measurements may be too noisy to resolve this problem. As a result, the algorithm produces results that have reconstruction errors on the same order of magnitude as the initial estimate.

Nonlinear Reconstruction Without Knot Location Measurements

In this section we present Monte-Carlo simulations of the Nonlinear Reconstruction Algorithm that use the Non-Sequential Algorithm to generate the initial estimate but use *only* chord length data in the nonlinear optimization routine to minimize the cost criterion given by Equation (6.6). This approach was taken because the i.i.d. knot location noise model is a poor characterization of the noise actually added to the system.

Figure 6-6 (a) and (b) show the resulting percent Hausdorff and Symmetric Difference errors resulting from this experiment. Each reconstruction contained 5 chord measurements in each view with the SNR set at 10dB. As shown in the figure, reconstructions were performed with 9, 18, and 27 views. In both plots the solid line is the error in the initial estimate which

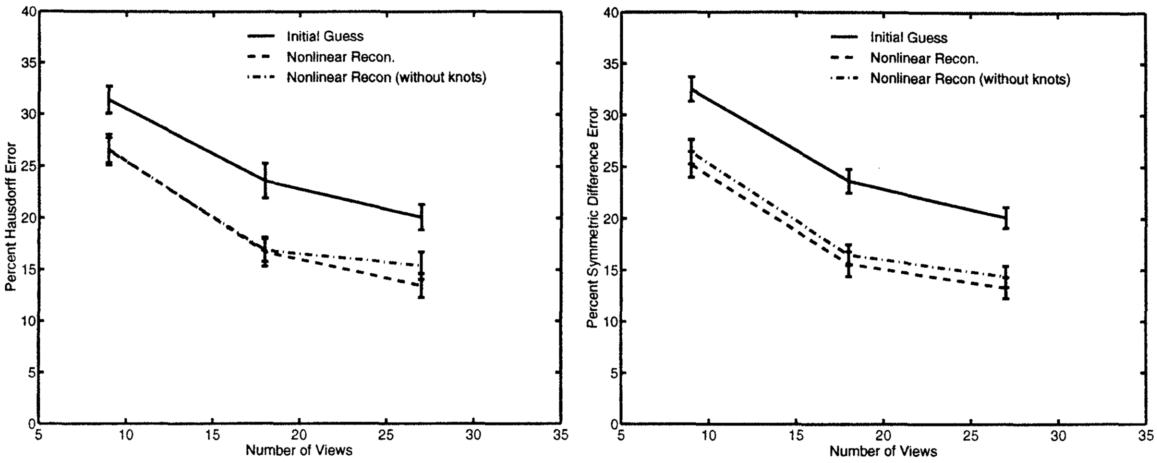

(a) Hausdorff Error

(b) Symmetric Difference Error

Figure 6-6: Nonlinear Reconstruction Without Knot Location Measurements.

is generated from the Non-Sequential Algorithm, the dashed line represents reconstructions using both knot location and chord length data in the nonlinear reconstruction, and finally the dash-dotted line is the reconstruction errors for the nonlinear reconstruction using only chords in the reconstruction. As the Hausdorff error and Symmetric Difference error of Figure 6-6 demonstrate, both of the nonlinear reconstructions produce better estimates than the initial estimate generated by the Non-Sequential Algorithm. In fact, for small numbers of views, the results were almost identical. As the number of views increased, however, the reconstructions *without* knot location data produced higher reconstruction errors for both measures. This shows that despite the badly mismatched model for the knot location noise, the knot location measurements produce valuable geometric information and add to the overall nonlinear reconstruction results.

Nonlinear Reconstruction With Top T Hypotheses

We alter the format of the Nonlinear Reconstruction Algorithm in this section slightly. Instead of initializing the nonlinear optimization with the estimate corresponding to the top data association hypothesis in the Non-Sequential Algorithm, we perform T nonlinear reconstructions initialized with the estimates corresponding to the top T data association hypotheses that result from the Non-Sequential Algorithm. After all of the nonlinear reconstructions are complete, the estimate that minimizes the cost criterion of Equation (6.6) is

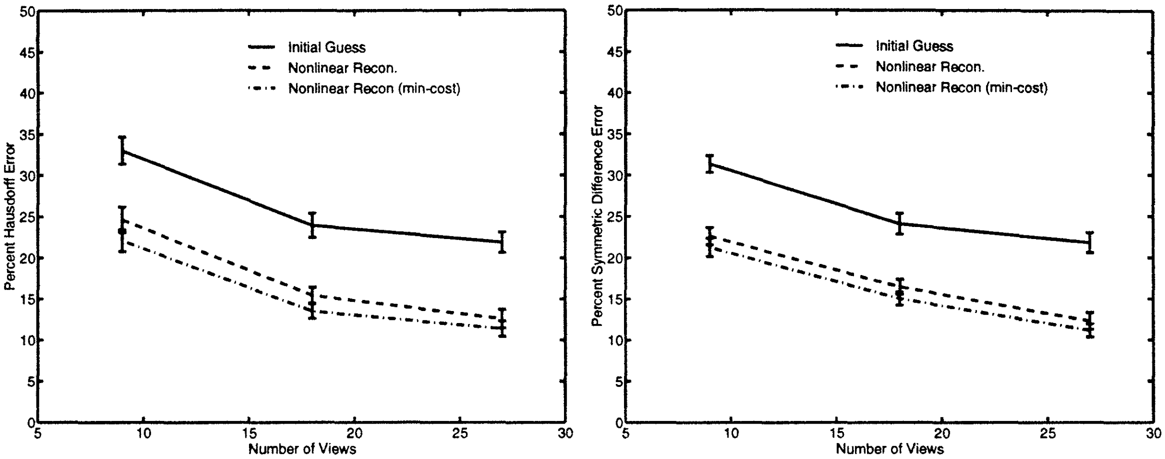

(a) Hausdorff Error

(b) Symmetric Difference Error

Figure 6-7: Nonlinear Reconstruction Without Top T Hypotheses.

chosen as the final estimate. The approach is taken so that the addition of chord length data in the nonlinear estimate can have even greater impact on choosing the optimal estimate.

Each reconstruction contained 5 chord measurements per view with the SNR on the chord length measurements set at 10dB. This corresponds to a total of 72 knot location measurements and 90 chord length measurements. Further, 10 hypotheses were retained at each step of the Non-Sequential Algorithm. The estimates corresponding to the final 10 hypotheses of the Non-Sequential Algorithm are used as the initial estimates to the Nonlinear Algorithm, and the reconstruction with the smallest resulting error is kept as the final estimate. Figure 6-7 (a) and (b) are plots of the Hausdorff and Symmetric Difference errors, respectively, that resulted from this analysis. In both plots, the solid line represents the reconstruction error of the top estimate provided by the Non-Sequential Algorithm, the dashed line represents the results from the Nonlinear Algorithm using only the top hypothesis from the Non-Sequential Algorithm, and finally the dash-dotted line represents the results from performing nonlinear reconstructions on all 10 hypotheses and choosing the one with the smallest error. As shown in Figure 6-7, this simulation was run for 9, 18, and 27 views.

As the results of Figure 6-7 demonstrate, the reconstruction obtained by choosing the reconstruction that minimizes the cost function of Equation (6.6) from the top 10 hypotheses outperformed the reconstruction using only the top initial estimate from the Non-Sequential

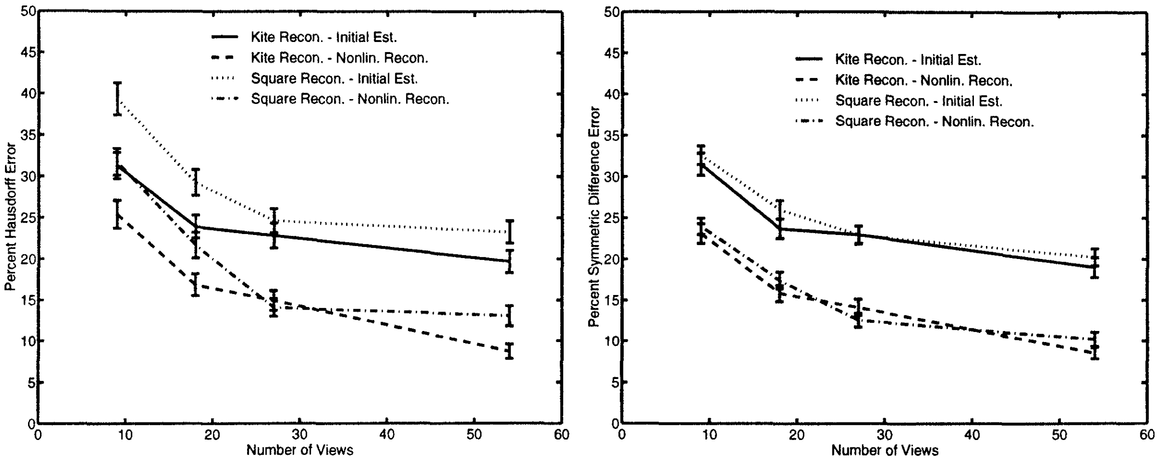

(a) Hausdorff Error

(b) Symmetric Difference Error

Figure 6-8: Comparison of Kite and Square Test Objects for Nonlinear Reconstruction Algorithm: SNR=10dB, T=15 retained hypotheses, 5 chord measurements per view.

Algorithm (as expected). It is also interesting to note that for both error measures, the relative decrease in reconstruction error was the same amount, independent of the number of views. Again, this result demonstrates the importance of the initial estimate in the Nonlinear Algorithm. For each nonlinear estimate generated from the initial estimate corresponding to one of the top 10 hypotheses of the Non-Sequential Algorithm the measurement data was identical. The only difference was in the initial estimate. Yet, as the results of Figure 6-7 show, this initial estimate is important when trying to minimize the cost function of Equation (6.6).

Reconstruction Comparison of Two Test Objects

Having completed the Monte-Carlo analysis of the Nonlinear Reconstruction Algorithm for a particular object, the kite, we now turn to a comparative analysis. The reconstructions were performed with 5 chord measurements per view and the SNR on the chord measurements set at 10dB. As stated previously, the initial estimates were obtained from the Non-Sequential Algorithm assuming an i.i.d. noise model on the knot location measurements. Reconstructions were performed for 9, 18, 27, and 54 projections. Figure 6-8 shows the Monte-Carlo reconstruction errors for both the kite and the square test objects as a function of the number of views used in the reconstruction.

Consistent with the results of the previous chapter, the initial estimates generated from the Non-Sequential Algorithm for the kite object have slightly smaller reconstruction errors than those of the square. The nonlinear reconstruction curves, shown in dashed lines for the kite and dash-dotted lines for the square in Figure 6-8 are basically shifted versions of the initial estimate error curves. Therefore, the overall improvement in the reconstruction of each object over the initial estimate is approximately the same. It is interesting to note that there are greater discrepancies between the reconstruction errors for the kite and square in the Hausdorff errors of Figure 6-8(a) than the corresponding Symmetric Difference errors of Figure 6-8(b). This is attributed to the fact that a slight elongation and rotation of the square can have a significantly large effect on the Hausdorff error while the change in the Symmetric Difference error is less severe. Overall, however, the initial reconstructions of both the kite and the square have the same relative amounts of error (as discussed in Chapter 5) and the improvement over this initial estimate due to the Nonlinear Reconstruction Algorithm is approximately the same for both objects.

6.3 Conclusions

In this chapter we extended our reconstruction approach to a full nonlinear estimation algorithm from knot location *and* chord length data. The approach was to use the output of either the Sequential Algorithm of Chapter 4 or the Non-Sequential Algorithm of Chapter 5 as the initial estimate to a nonlinear optimization routine which minimizes the cost function given by Equation (6.6) using the noisy knot location and chord length measurements and the corresponding noise statistics for each data set. Note that despite the assumption throughout this experiment of the i.i.d. noise model on the knot location data (which we consider to be a worst case assessment of the algorithm's performance), the results showed a vast improvement in the reconstruction quality as compared to both the Sequential Reconstruction Algorithm and the Non-Sequential Algorithm (in terms of the percent Hausdorff and Symmetric Difference errors). This improvement is attributed to the incorporation of the chord length measurements (which we model with the correct applied noise model), which have a significant impact on the estimate of the vertex locations of the object.

Overall, the Nonlinear Algorithm significantly reduced the reconstruction errors as measured by both the Hausdorff error and Symmetric Difference error. For the sample reconstructions of Section 6.2.1 the Nonlinear Algorithm halved the percent error measures for both the Hausdorff and Symmetric Difference errors in nearly all of the cases. As expected, the Monte-Carlo reconstruction experiments also showed a vast improvement over the corresponding errors obtained using the Sequential or Non-Sequential Algorithms alone. The Nonlinear Algorithm was able to significantly reduce the reconstruction error for all experiments except the cases where the SNR was 0dB. For this 0dB case, the algorithm was unable to recover from the poor initial estimate (here the resulting errors were essentially the same as the errors in the initial estimates). In summary, the percent reconstruction error decreased as the number of chords increased, number of retained hypotheses increased, number of views increased, and the SNR increased (as expected in all cases). Additionally, the performance of the algorithm had a strong dependence on the chord measurements when the number of chord measurements per view was small. In addition to containing important information to regularize the reconstruction problem, the chord measurements are extremely important because of the poor assumed noise model on the knot location measurements. The algorithm incorrectly weights the importance of the knot measurements and thus make the correctly weighted chord measurements extremely important in determining the optimal estimate of the vertices. Despite the poor modeling of the knot location data, we found that the knot location data still had a positive impact on the overall reconstruction quality. We also note the importance of the initial estimate on the performance of the reconstruction algorithm. As results of the Monte-Carlo analysis retaining the top T hypotheses from the Non-Sequential Algorithm showed, lower reconstruction errors can be obtained from the same data set solely by varying the initial estimate input to the Nonlinear Reconstruction Algorithm. Finally, in the comparison of the two test objects, the Nonlinear Algorithm reduced the reconstruction error for both objects by approximately the same amount. Basically the Nonlinear reconstruction error curves were shifted (lower) versions of the error curves of the initial estimates. Thus, consistent with the results presented in the previous chapter on the Non-Sequential Algorithm, the result of the comparison of the reconstructions of the two test objects showed that the Nonlinear Algorithm did not produce significantly better reconstructions for either object.

Chapter 7

Resolving Data Inconsistencies - Missing Knot Measurements

This chapter returns to the Non-Sequential Reconstruction Algorithm of Chapter 5 and presents an adapted version of this algorithm, which we term the *Extended* Non-Sequential Algorithm. The Extended Non-Sequential Algorithm is capable of dealing with *missed detections*, or missing knot location measurements in the projection data. This chapter is organized as follows. Section 7.1 describes the underlying assumptions of this problem. Section 7.2 outlines the general framework of the Extended Non-Sequential Reconstruction Algorithm and how it differs for the original Non-Sequential Algorithm, describes the criteria used to generate the possible hypotheses at each step of the reconstruction, and presents a general overview of the algorithmic procedure. Section 7.3 explains the criteria used to remove knot location measurements from the projection data during the data generation process. Section 7.4 presents sample reconstructions, Monte-Carlo simulations of this algorithm, and an implementation of the Nonlinear Algorithm of Chapter 6 with the Extended Non-Sequential Reconstruction Algorithm used to generate the initial estimate. Finally, Section 7.5 summarizes our conclusions on the results of this algorithm.

7.1 Initial Assumptions

Like the algorithms developed in the previous three chapters, the algorithm developed in this chapter is concerned with the reconstruction of binary polygonal objects from measurements of knot locations and chord lengths. Unlike the previous chapters, however, the Extended Non-Sequential Algorithm incorporates the additional complexity of missing knot location measurements. As in all previous chapters, we make use of many of the definitions and terminology presented in Chapter 2 and the problem formulation of Section 3.2. Additionally, because the algorithm developed in this chapter is an extension of the Non-Sequential Algorithm presented in Chapter 5, we make use of many of the definitions and procedures introduced in that chapter.

The underlying assumptions for this chapter are essentially the same as in all other chapters of this thesis. To review, we assume that we wish to reconstruct the vertices of a binary polygonal object with a known number of sides from measurements of knot location and chord length data. Also, for simplicity, we consider only convex objects as discussed in Section 3.2.1. We assume that the projections are taken at known angular views and that the knots have been extracted from the projection data independently of this algorithm using a procedure such as that of Mier-Muth and Willsky [8]. In addition, we assume that all objects are convex n-gons and that the support knots will not be missed detections. This assumption is based on the work of Prince [12] who was able to successfully detect the support knot locations using Kalman filtering methods. Thus, missed detections are only assumed to occur at internal knot locations. Again, since it is outside the scope of this thesis to actually implement a knot extraction procedure, we take the more simplistic approach of directly generating noisy knot observations with missing data measurements. Thus, as outlined in Section 3.2.1, our reconstruction scenario consists of two stages: the generation/simulation of the noisy knot location and chord length data and the implementation of the reconstruction algorithm. The chord lengths and knot locations are generated from the underlying object and perturbed with independent Gaussian noise using the noise models specified in Section 3.2.3 and discussed later in this chapter. A block diagram of the reconstruction scenario is shown in Figure 3-4.

We begin in the next section with a brief review of the Non-Sequential Algorithm and explain how it must be altered to accommodate measurements with missing knots. Because we directly generate the data, we then present the methodology used to generate knot location data with missing measurements followed by experimental results from reconstructions using the Extended Non-Sequential Algorithm.

7.2 Overview of the Extended Non-Sequential Reconstruction Algorithm

Recall from Chapter 5 that the Non-Sequential Reconstruction Algorithm is a multi-stage estimation algorithm that simultaneously estimates discrete-valued states, or data association hypotheses, and continuous-valued estimates of the vertices of a polygonal object based on these discrete hypotheses. There are three distinct stages in this algorithm. The first stage enumerates all of the possible associations of knot location measurements to vertex locations. Once the possible data associations are determined, the second stage of the algorithm estimates the vertex locations for each discrete hypothesis under the assumption that it is the correct hypothesis. Finally, the third stage uses a generalized likelihood ratio test to determine the top T hypotheses for the next step of the algorithm. The optimal hypothesis is the most likely hypothesis given by the ratio test and the estimate based on this hypothesis is the optimal estimate of the vertex locations of the object. The block diagram of this algorithm can be found in Figure 4-1. Recall also that the only difference between the Sequential Reconstruction Algorithm of Chapter 4 and the Non-Sequential Algorithm is that the Non-Sequential Algorithm allows views to be input into the algorithm in any order. This simply amounts to a different method to determine the possible data associations at each step of the algorithm.

The only difference between the original Non-Sequential Algorithm and the algorithm developed in this chapter is that we now have incomplete knot location measurement data. In terms of the algorithm, this translates into keeping track of both the association of the hypothesized underlying knot locations to the vertices (as before) *and* the association of the measurements to these ordered knot locations $z_1 \leq z_2 \leq \cdots \leq z_n$. This means that each hypothesis must now clearly define which knot locations are missing measurements. Thus, a missing knot location measurement is analogous to a missed detection in the multitarget radar tracking problem [6] and simply adds another level of complexity to the discrete data association hypotheses.

In the Non-Sequential Algorithm, at the kth step of the algorithm, the possible knot-to-vertex data associations for the insertion of the kth view into an existing hypothesis are evaluated with five "rules" (which we repeat here for convenience):

1. *If a projection is inserted between two views that have the same knot-to-vertex association, the inserted view must also have that ordering.*

2. *Once a knot-to-vertex association has been used somewhere in a hypothesis, it can never occur again after the knots switch from that position.*

3. *If a knot-to-vertex association is used somewhere in the proposed hypothesis path, its reciprocal cannot be used anywhere in the hypothesis (with the exception of the initial and final associations).*

4. *Each hypothesis must start with abcd and end with dcba.*

5. *Only* $\frac{n(n-1)}{2}$ *knot switches are allowed to occur in each hypothesis path. After this, only the data associations adjacent to the inserted knot projection are allowed.*

Using the above data association rules with each new view of knot location data, the algorithm evaluates the set of possible associations for each existing hypothesis q_i^{k-1}, $1 \leq i \leq T$ and determines what knot associations are possible for the kth measurement given the geometric constraints outlined above. This set of valid associations forms the set of hypotheses q^k which corresponds to Stage 1 of the reconstruction algorithm. Further, we demonstrated that these associations can be enumerated in a tree (refer to Section 5.2, Figure 5-1). A trace of successive branches from the root to a branch of the tree denotes one possible data association hypothesis q_i^k and the set of all possible hypotheses is denoted by $q^k = [q_1^k, q_2^k, \cdots, q_r^k]$.

To develop the Extended Non-Sequential Algorithm, we simply add another step to the hypothesis generation stage. At the kth measurement, after we have determined the possible knot-to-vertex associations using the above rules (which may be done in the absence of data), we check to see if the kth measurement is incomplete (it has fewer measurements than vertices). If there are fewer measurements than knot locations at that measurement, then we also must associate with each possible knot-to-vertex association a hypothesized measurement-to-knot association. This procedure can best be explained in terms of an example.

Referring to the first column of Table 7.1 we have a possible data association hypothesis for the first 7 views. We now wish to insert the 8th view into the position denoted by the

Hypothesis	Possible Data Assoc. Non-Seq. Algorithm	Final Data Assoc. Extended Non-Seq. Algorithm
abcd	*abcd*	*abcd*
▭ ⇐	*bacd*	1101
bacd		*abcd*
badc		1011
bdac		*bacd*
dbac		1101
dbca		*bacd*
dcba		1011

Table 7.1: Example Evaluation of Possible Knot Location Data Associations for Extended Non-Sequential Algorithm.

arrow. Using the five rules given by the Non-Sequential Algorithm we narrow the possible knot-to-vertex data associations to those listed in the second column. If this view contains only three measurements, we must now associate with each of the knot-to-vertex associations in column 2, a corresponding measurement-to-knot location association. Keeping in mind that we have assumed *only* the loss of internal knot location data and the fact that both the measurements and knot locations are ordered in increasing order (i.e. we don't have to consider all permutations of the measurements because they must be ordered in increasing order), we need only specify which knots have measurements and which do not to uniquely specify this association. For notational simplicity, we denote a knot location with a measurement by a 1 and one without a measurement by a 0. For example, 1001 corresponds to measurements only at the two support knots while 1111 corresponds to a full data set. Using this notation, the last column of Table 7.1 gives the revised list of data association hypotheses.

In addition, the set of all hypotheses can still be enumerated in a slightly altered version of the hypotheses tree given in Chapter 5. Instead of just the knot-to-vertex associations at each view, we must give *both* knot-to-vertex associations *and* measurement-to-knot associations. For example, Figure 7-1 presents an example hypothesis tree for a 4-gon. We have assumed in this example that the first three views expanded in this tree are full sets of knot projection data but that the fourth only has three measurements. If we compare this tree to the corresponding tree for the Non-Sequential Algorithm (Figure 5-1) we find that the

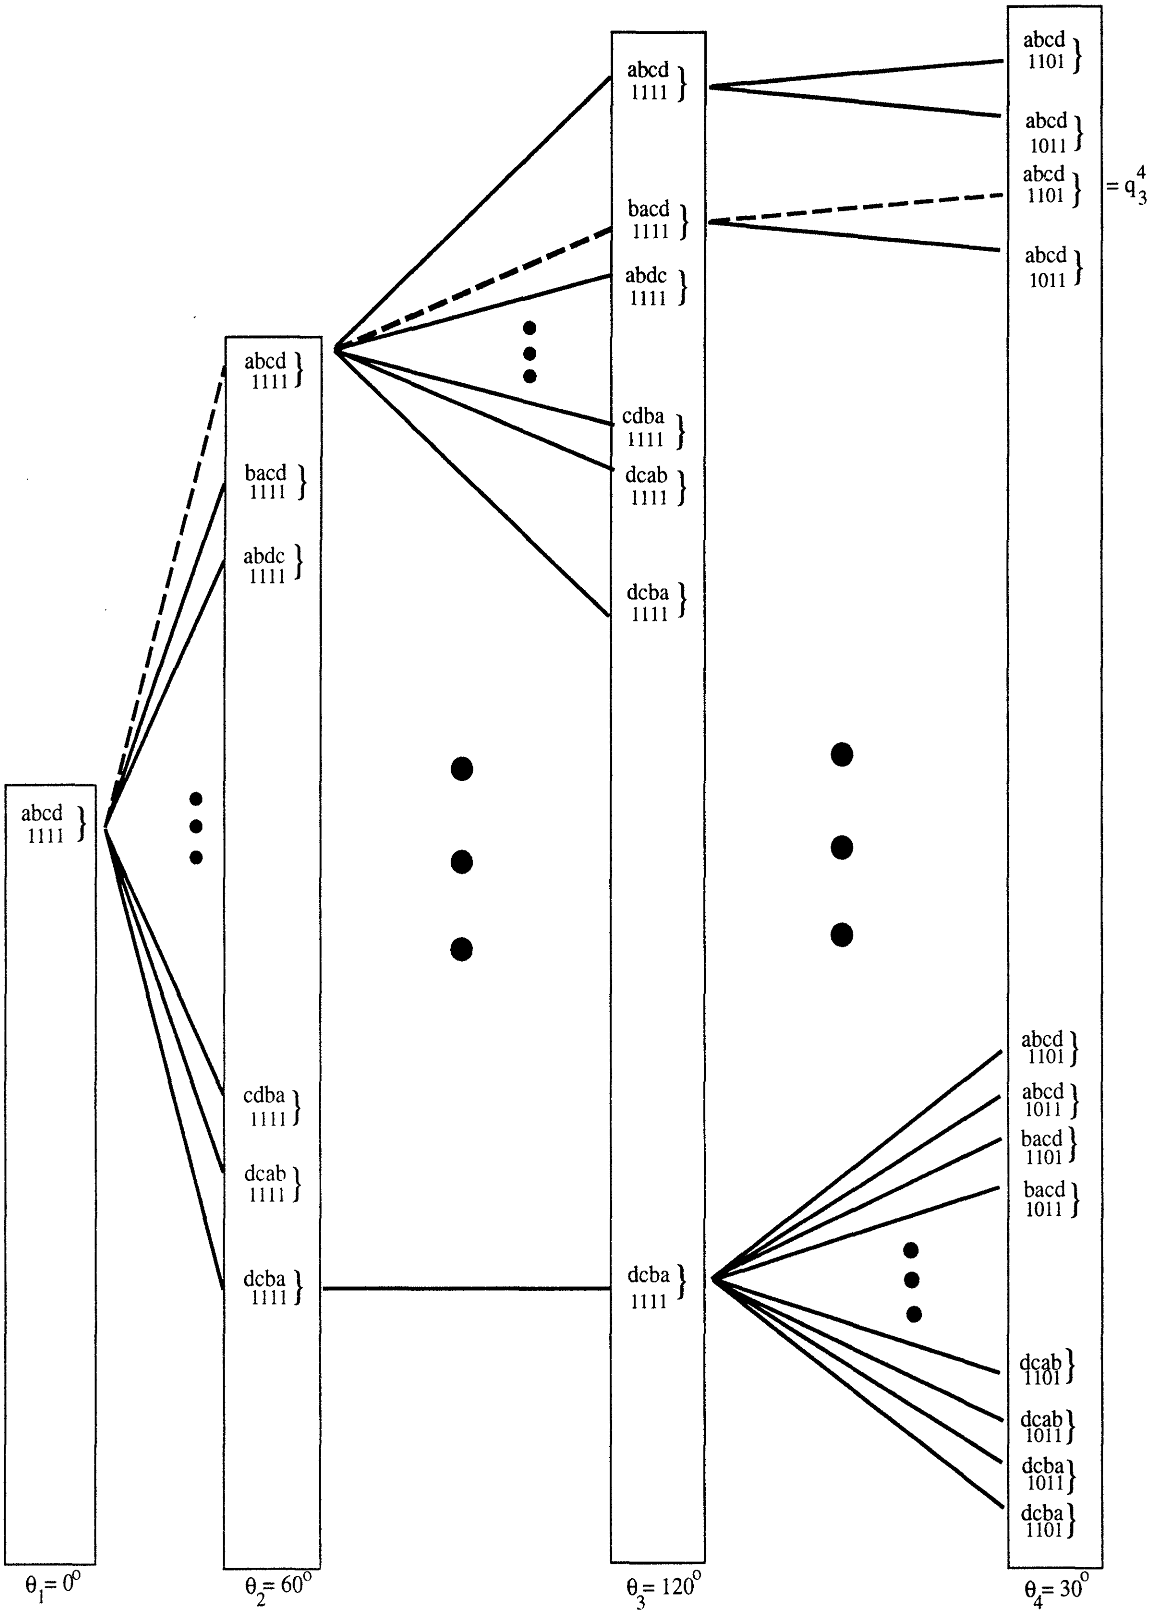

Figure 7-1: Hypothesis Tree for the Extended Non-Sequential Algorithm.

number of hypotheses at the fourth measurement is doubled because we have to account for the possibility that the missing knot location measurement is in either the second or third position for each of the possible knot-to-vertex associations. Yet, a single hypothesis is still represented by a trace of successive branches from the root to a branch of the tree. For example, the hypothesis q_3^4 shown in dashed lines in the figure represents the data association hypothesis (when reordered in increasing angular order):

$$\left.\begin{matrix} abcd \\ 1111 \end{matrix}\right\} \Rightarrow \left.\begin{matrix} abcd \\ 1101 \end{matrix}\right\} \Rightarrow \left.\begin{matrix} abcd \\ 1111 \end{matrix}\right\} \Rightarrow \left.\begin{matrix} bacd \\ 1111 \end{matrix}\right\}$$

Thus, the problem is essentially the same as in the original Non-Sequential Algorithm with one key difference: each hypothesis must identify both the association of the knot locations to the vertices of the object (as before), and, the association of the measurements to these knot locations (i.e. hypothesizing which knots are missing measurements). The remainder of the algorithm is essentially the same as in Chapter 4 and Chapter 5. Stages 2 and 3 of the algorithm, which evaluate and prune the hypotheses, are performed using the same method as the Sequential Algorithm (see Sections 4.3 and 4.4). Again, chord length data are used in addition to the knot location data to prune the possible hypotheses by using the likelihood ratio tests defined by Equation (4.11) on the observed data. All of the calculations and evaluation procedures are identical to those of the Sequential Algorithm and Non-Sequential Algorithm with the exception of the determination of the discrete hypotheses, as described above. The next section develops the method we used to remove knots from our simulated knot location data in order to test the Extended Non-Sequential Algorithm.

7.3 Data Generation

A *missed detection* is a knot location that the knot extraction algorithm is unable to detect from the noisy projection data. In a complete reconstruction scheme a specific measure for the likelihood of missing a detection is dependent upon the particular method used to extract the knots from the projection data. Since we directly generate our data, however, we simply use two parameters to quantify two geometrically intuitive criteria for missing detections in the noisy projection data: the absolute value of the change in slope at the knot location and the minimum distance to adjacent knots[1].

The knot location data generation process used in this chapter is composed of three steps:

1. Generation of the noise-free knot locations from the projection of the vertices of the underlying object at the desired angular locations.
2. Removal of some knot locations from the full set of projection data via the criteria that are developed in this section.
3. Addition of independent Gaussian noise to the remaining knot location data via the noise model outlined in this section.

We begin by developing the criteria to remove knots from the noise free projection data.

7.3.1 Knot Removal Criterion

We define the absolute change in slope at a particular knot location z_i as

$$|\Delta s_i| = |s_{i-1} - s_i| \tag{7.1}$$

where s_{i-1} is the slope of the line segment of the projection just before the knot location and s_i is the slope of the line segment of the projection just after the knot location. The absolute value of the change in slope is an important parameter in the ability to detect a knot. A knot that has a small change in slope would be difficult to detect in noisy projection data (think of the two limiting cases: first, the absolute change in slope, $|\Delta s|$, as it approaches zero and second, the absolute change in slope, $|\Delta s|$, as it approaches infinity). For example, Figure 7-2 is the projection of a binary polygonal object. From the projection, is obvious that detecting the knot at location z_3 would be much easier than the knot at z_2, especially with the addition of noise. Consistent with this intuitive assessment, the absolute change in slope at z_3, $|\Delta s_3|$, is greater than the absolute change in slope at z_2, $|\Delta s_2|$.

The absolute change in slope is not the only factor that affects the detectability of knots in the projection data however. For example, the second knot location, z_2, in Figure 7-3(a)

---

[1] These two parameters are very similar to the values in the denominator of Equation (3.3) which is used to set the variance on the knot locations. Refer to Section 3.2.3 for the noise models used in this thesis.

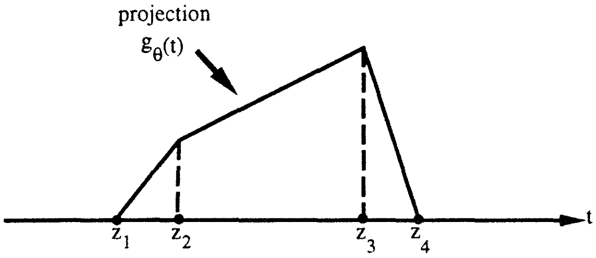

Figure 7-2: Effect of Slope on Knot Extraction.

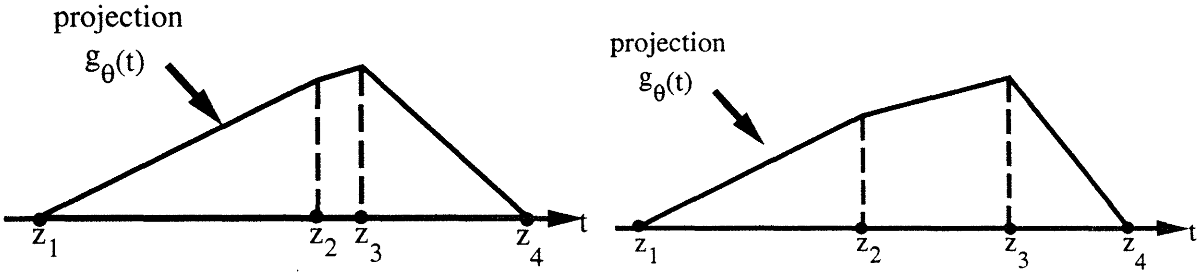

Figure 7-3: Effect of Minimum Separation on Knot Extraction.

and Figure 7-3(b) both have the same slope change, yet, because of the smaller distance between the adjacent knots in Figure 7-3(a), we expect that these knots would be more easily misidentified as a single knot instead of two. Thus, we use the minimum separation between adjacent knot locations as our second criterion. We define the minimum separation of knot location z_i to be

$$\widetilde{M}_i = \min[(z_{i+1} - z_i), (z_i - z_{i-1})] \tag{7.2}$$

Because we only remove internal knot locations, there are always adjacent knot locations on either side of z_i.

We incorporate both the absolute change in slope and the minimum separation between adjacent knots into a single threshold test for the generation of missed knots in the data:

- Threshold Test 1

  If z_i is an internal (noiseless) knot location and

  $$|\Delta s_i| \widetilde{M}_i \leq t_1 \tag{7.3}$$

  knot z_i is removed for the data set; if not, z_i is retained.

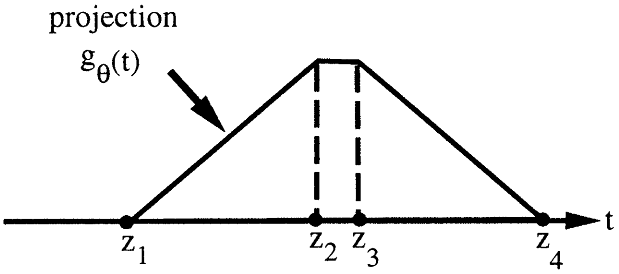

Figure 7-4: Effect of Symmetric Projections on Knot Extraction.

In Equation (7.3), $|\Delta s_i|$ is the absolute change in slope at knot z_i, $\widetilde{M}_i$ is the minimum separation defined above, and t_1 is an arbitrary threshold used on the set of all knot location projection data to test the knots for removal. In this chapter, we vary the parameter t_1 over the set of values 0, .75, 1.03, and 1.075 to test the performance of our algorithm.

In addition to the above threshold test, we also consider a second test to remove knot location data. If two internal knots are close together and *both* fall below the threshold t_1 using the test given by Equation (7.3), it doesn't make sense intuitively to remove both of the knots. For example, in the projection of Figure 7-4, if the threshold t_1 were large enough, both of these knot would be removed based solely on the criterion given in Threshold Test 1 (they both have the same absolute change in slope and minimum separation), implying that the knot extraction algorithm was unable to detect either knot location. A more reasonable assumption would be to replace both knot locations with a single aggregate knot location. Consistent with this logic, if two adjacent internal knot locations, z_i and z_{i+1}, fail test 1 above, a second test is performed

- Threshold Test 2

  If z_i and z_{i+1} are adjacent internal knot locations with

  $$|\Delta s_i|\widetilde{M}_i \leq t_1 \quad \textit{and} \quad |\Delta s_{i+1}|\widetilde{M}_{i+1} \leq t_1$$

  and

  $$(z_{i+1} - z_i) \quad \leq \quad t_2 \tag{7.4}$$

replace both knot locations with the new knot location:

$$\tilde{z} \;=\; (z_{i+1} + z_i)/2 \tag{7.5}$$

In this test, a second threshold, t_2, is introduced. This threshold defines the maximum allowable distance between two adjacent knots such that they are replaced by the average knot $\tilde{z}$ if they both fail Threshold Test 1. In the simulations of this thesis, we assume a constant threshold of $t_2 = .2$. This value was chosen based on the fact that in most views the width of the support of our test objects[2] was ~ 1. Thus, the minimum distance is on the order of 20% of the overall support width. Figure 7-4 is an example of a projection where the average knot would replace the actual knot measurements if threshold t_1 is set sufficiently large enough (and similarly t_2 set to an appropriate level such that it is greater than the distance between z_2 and z_3).

In summary, we present the knot removal procedure based on the above criteria. We consider the case of a 4-gon but this procedure is easily extended for an object with more vertices (and thus more knot locations). For a given set of noise free knot location data, the internal knots of a 4-gon are z_2 and z_3 as shown in Figure 7-4 and the knot removal procedure is the following:

1. *Perform Threshold Test 1.*
    - *If* $|\Delta s_2|\widetilde{M}_2 > t_1$ *and* $|\Delta s_3|\widetilde{M}_3 > t_1$, *no knots are removed from the view.*
    - *If* $|\Delta s_2|\widetilde{M}_2 \leq t_1$ *and* $|\Delta s_3|\widetilde{M}_3 > t_1$, z_2 *is removed from the view.*
    - *If* $|\Delta s_2|\widetilde{M}_2 > t_1$ *and* $|\Delta s_3|\widetilde{M}_3 \leq t_1$, z_3 *removed from the view.*
    - *If* $|\Delta s_2|\widetilde{M}_2 \leq t_1$ *and* $|\Delta s_3|\widetilde{M}_3 \leq t_1$, z_2 *and* z_3 *are removed from the view and step 2 is performed.*
2. *Perform Threshold Test 2.*
   If both z_2 *and* z_3 *are being removed according to Threshold Test 1 and if* $(z_3 - z_2) \leq t_2$, *the two knots are replaced by the "averaged" knot* $\tilde{z}$ *such that*

$$\tilde{z} = (z_3 + z_2)/2$$

---

[2] Refer to Section 4.5.1 for a discussion of the test objects used in this thesis.

In addition to the procedure and assumptions outlined above, for simplicity in our algorithm we also assume that the maximum threshold, t_1, for a given set of projection data is chosen such that two full views of knot location data remain. This is required because the initial estimate of the n vertices of the object occurs with three views and the two full views insure that there will be at least two measurements to estimate the two parameters, x_m and y_m, of each vertex (i.e. to avoid an underdetermined set of equations). Although three full views are necessary to uniquely define the vertices in a noise free reconstruction scenario we do not require this because our algorithm hypothesizes all possible data associations. In addition, we note that with increased algorithmic complexity we could do without the requirement of two full views by waiting until enough measurements were obtained to estimate each vertex (but this entails retaining all possible hypotheses until there are enough measurements to generate the first complete estimate of all vertex locations).

Based on the above criteria, Table 7.2 and Table 7.3 summarize the various thresholds, t_1, considered in this chapter and how the procedure outlined above affects the kite and square test objects respectively with t_2 set at 0.2. The tables detail for a given threshold and number of views (assuming that the views are uniformly spaced over the range $[0°, 180°)$) how many knot will be removed in both absolute numbers and as a percentage of the total number of knots in the projection data. An X in any position in either of the two tables denotes that the threshold t_1 is too high for the projection data to retain the required two full views. We note that for each object at a given threshold t_1, the percentage of knots removed basically remains constant as the number of views is varied. In addition, the kite test object is much more sensitive to knot removal at lower threshold values.

7.3.2 Noise Models

In addition to altering the knot location projection data by removing knots from the projections, we must also alter the noise model developed in Section 3.2.3 to accommodate the new "averaged" knot locations. Basically the approach taken is as follows:

1. The variance of every original noise-free knot remains given by Equation (3.3).

Views	Meas. on Data Set	$t_1 = 0.0$	$t_1 = .75$	$t_1 = 1.03$	$t_1 = 1.075$
9	Number Removed	0	3	7	X
	Percent Removed (%)	0	8.33	19.44	X
18	Number Removed	0	6	15	18
	Percent Removed (%)	0	8.33	20.83	25.00
27	Number of Removed	0	9	22	29
	Percent Removed (%)	0	8.33	20.37	26.85
54	Number Removed	0	19	46	56
	Percent Removed (%)	0	8.80	21.30	25.92

Table 7.2: Table of Conversions for Kite Test Object - Thresholds, Number of Missing Knot Location Measurements, Percent Missing Knot Measurements.

Views	Meas. on Data Set	$t_1 = 0.0$	$t_1 = .75$	$t_1 = 1.03$	$t_1 = 1.075$
9	Number Removed	0	3	5	5
	Percent Removed (%)	0	8.33	13.89	13.89
18	Number Removed	0	6	10	10
	Percent Removed (%)	0	8.33	13.89	13.89
27	Number of Removed	0	8	12	12
	Percent Removed (%)	0	7.40	11.11	11.11
54	Number Removed	0	16	24	24
	Percent Removed (%)	0	7.40	11.11	11.11

Table 7.3: Table of Conversions for Square Test Object - Thresholds, Number of Missing Knot Location Measurements, Percent Missing Knot Measurements.

2. If two internal knots z_i and z_{i+1} are replaced by the averaged knot $\tilde{z}$, the variance of this new knot is given by:

$$\sigma_{\tilde{z}}^2 = \max(\sigma_{z_{i+1}}^2, \sigma_{z_i}^2) \tag{7.6}$$

It is arguable what the actual variance of this "new" knot would be as it depends on the specific knot extraction algorithm, but we choose the variance defined in Equation (7.6) as a conservative assessment of this value. After the knot extraction is complete, Gaussian noise is added to the remaining projection data using the knot location noise model described above and knot location and chord length measurements are input to the Extended Non-Sequential Algorithm detailed in Section 7.2.

7.4 Experimental Results

In this section, we present some sample reconstructions and Monte-Carlo simulations of the proposed Extended Non-Sequential Reconstruction Algorithm. As in previous chapters, we perform our analysis on two test objects: a "kite" and a square (refer to Section 4.5.1, Figures 4-5 and 4-6). There are a number of specifications that are common to all of the reconstructions performed in the following sections.

First, consistent with the requirement of at least two full views of knot location data to obtain the initial estimate of the vertices of the object in the first step of the algorithm (refer to Section 7.3), we require that the first two views (three if possible) entered into the algorithm contain full sets of knot location data. Thus, at the start of each reconstruction the initial view is determined by the Max-Min Approach defined in Section 3.2.4 with the stipulation that it contain a full set of knot location data. The algorithm then reorders the remaining views using the procedure for the Non-Sequential Algorithm outlined in Section 5.3 with one small deviation. After the views are divided into three angular ranges, we move our choice for the second view position (and third view position, if possible) the minimum distance such that it (they) will contain a full set of projection data. The remaining views are reordered exactly as outlined in Section 5.3.

As in previous chapters, the chord measurements are uniformly sampled over the region defined in Radon space by $t = [-1, 1]$. Again, as stated previously, this results in views where some of the chord measurements are zero (i.e. the object is not in the field of view) and conversely, in views where the chord measurements are confined to the interior of the object (for the kite object).

The noise on the chord measurements is i.i.d. Gaussian noise with variance σ_m^2 and the corresponding noise on each knot location measurement is given by the formulation outlined in the previous section. Reconstructions are performed for cases where the correct applied knot location noise model is assumed known and also for cases where the assumed model is an i.i.d. model with covariance $\Lambda_z = (1.5\sigma_m)^2 I$ where σ_m^2 is the variance of the noise on the chord length measurements. Additionally, as described in Section 3.2.5 the percent Hausdorff error and percent Symmetric Difference error are used for the performance analysis.

Finally, for all of the simulations that follow, we vary the knot removal threshold, t_1, over the set of values $t_1 = 0, .75, 1.03,$ and 1.075. The second threshold, t_2, which defines the maximum distance at which two adjacent knot locations are replaced by their average value, is set at the constant value of 0.2 (refer to Section 7.3 for a complete discussion of these parameters). Table 7.2 and Table 7.3 relate the values of t_1 to the number of knots removed and the percentage of knots removed for projection data containing 9, 18, 27, or 54 views for each test object. These conversions provide a more intuitive understanding of the relationship between a given value of t_1 and its affect on the knot location projection data for a given object.

7.4.1 Sample Reconstructions

This section contains results of sample reconstructions using the Extended Non-Sequential Reconstruction Algorithm. Note that for all of these results the true object is depicted by solid lines while the reconstructions are represented by dashed or dotted lines. For both objects, the square and the kite, noisy chord length and knot location data were generated so that the SNR on the chord length data was set to 10dB. Additionally, projection data was taken at 27 equally spaced points over the interval $[0°, 180°)$ with 5 chord measurements in each view. The reconstructions were performed for knot removal thresholds t_1 of 0, .75, 1.03, and 1.075.

Figure 7-5 shows sample reconstructions of the kite at various threshold values. Two

Thres.	Error Measures (t_1)	Kite Object		Square Object	
		True Model	i.i.d. Model	True Model	i.i.d. Model
$t_1 = 0$	Hausdorff (%)	14.87	25.60	15.54	23.08
	Symm. Diff. (%)	16.06	17.46	17.91	16.17
$t_1 = .75$	Hausdorff (%)	15.57	16.98	18.68	19.77
	Symm. Diff. (%)	11.27	16.55	11.31	25.42
$t_1 = 1.03$	Hausdorff (%)	21.96	43.17	18.93	18.94
	Symm. Diff. (%)	21.06	37.07	20.83	11.78
$t_1 = 1.075$	Hausdorff (%)	22.44	35.88	19.65	22.56
	Symm. Diff. (%)	20.74	34.86	14.23	20.65

Table 7.4: Comparison of Errors for Sample Reconstructions Using Non-Sequential Algorithm.

reconstructions are shown in each plot: the dashed line represents the reconstruction generated assuming the correct applied noise model on the knot location measurements while the dotted line represents the reconstruction assuming the i.i.d. noise model. Similarly, Figure 7-6 shows sample reconstructions of the square for various threshold levels. Like the kite reconstruction, the dashed line represents the reconstruction assuming the correct applied noise model while the dotted line represents the i.i.d. model reconstruction. Table 7.4 summarizes the Hausdorff error and Symmetric Difference error for both sample reconstructions. To insure that the sample path shown was indicative of the typical performance of the algorithm, at each threshold the sample reconstruction was chosen such that its Hausdorff error assuming the correct applied noise model was approximately equal to the average reconstruction error that resulted from 50 Monte-Carlo runs for that reconstruction scenario (refer to Section 7.4.2).

Although these results are only for one sample reconstruction for each of the above scenarios some interesting observations can be made from these results. First, as in all of the reconstructions from previous chapters, the reconstruction assuming the correct noise model consistently outperformed the reconstruction assuming the i.i.d. model. Referring to Table 7.4, we also note that the differences between the corresponding reconstruction errors assuming the correct noise model and assuming the i.i.d. model were smaller for the square. Further, consistent with these results the correct and i.i.d. reconstructions of Figure 7-6 seemed to be more closely associated while in some of the kite reconstructions shown in Figure 7-5 the reconstruction assuming the correct model was noticeably better than the

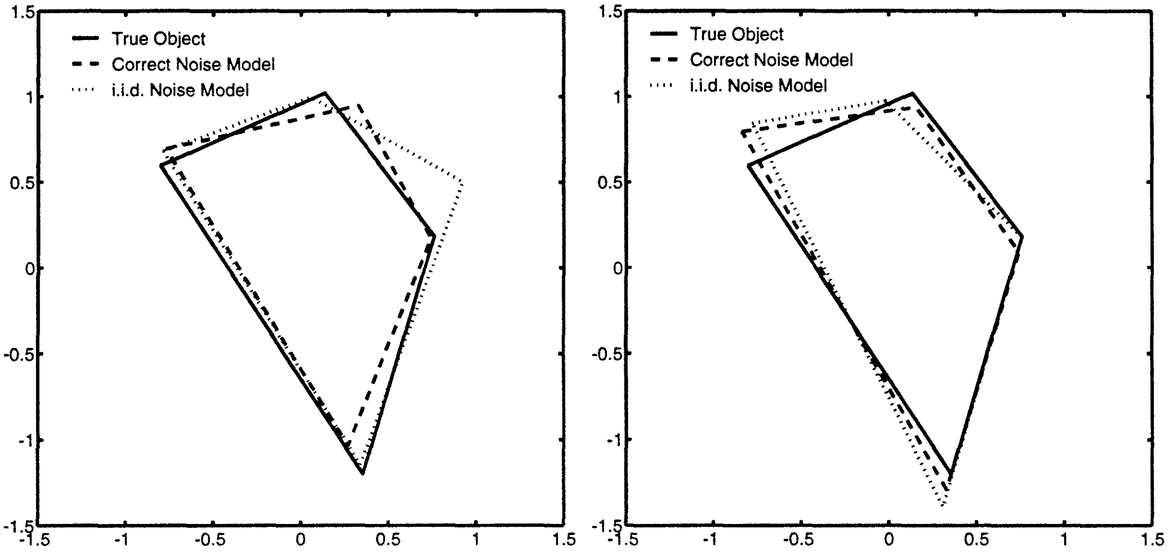

(a) t1 = 0

(b) t1 = .75

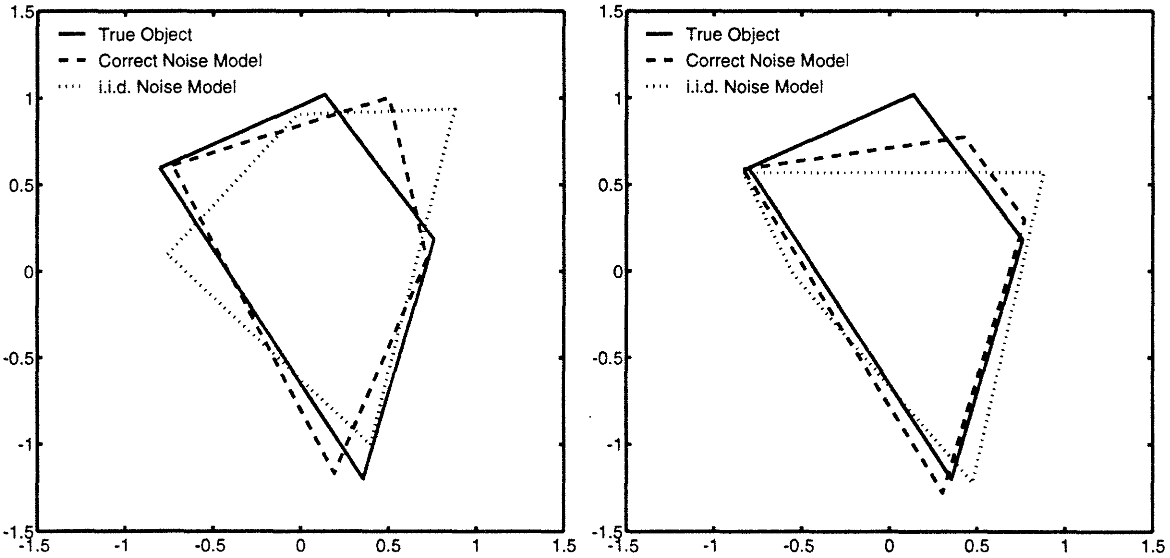

(c) t1 = 1.03

(d) t1 = 1.075

Figure 7-5: Sample Reconstructions for Kite Test Object Using Extended Non-Sequential Reconstruction Algorithm; (a) $t_1 = 0$, (b) $t_1 = .75$, (c) $t_1 = 1.03$, (d) $t_1 = 1.075$.

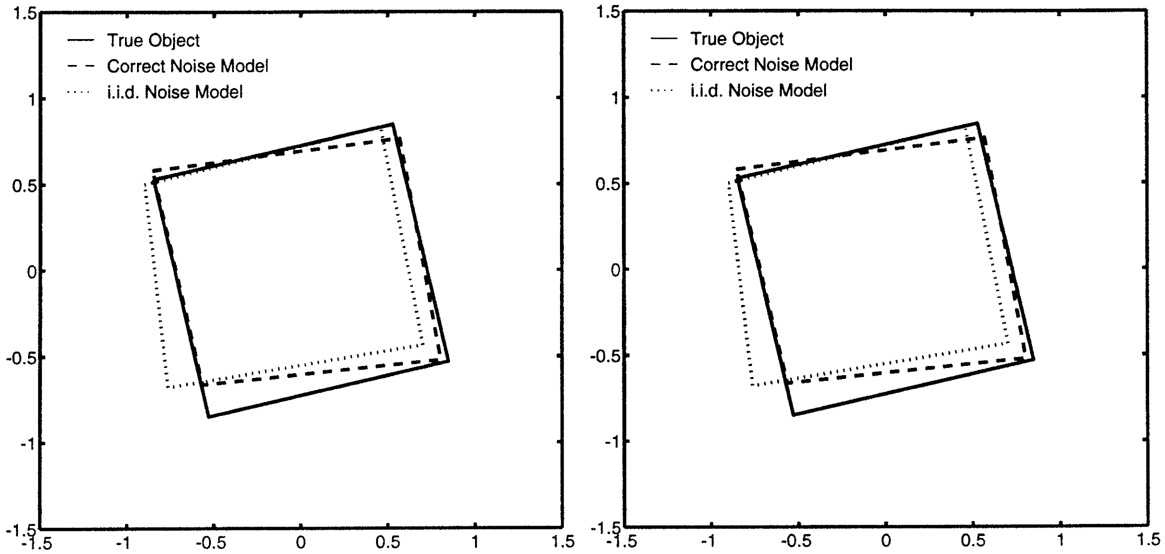

(a) t1 = 0

(b) t1 = .75

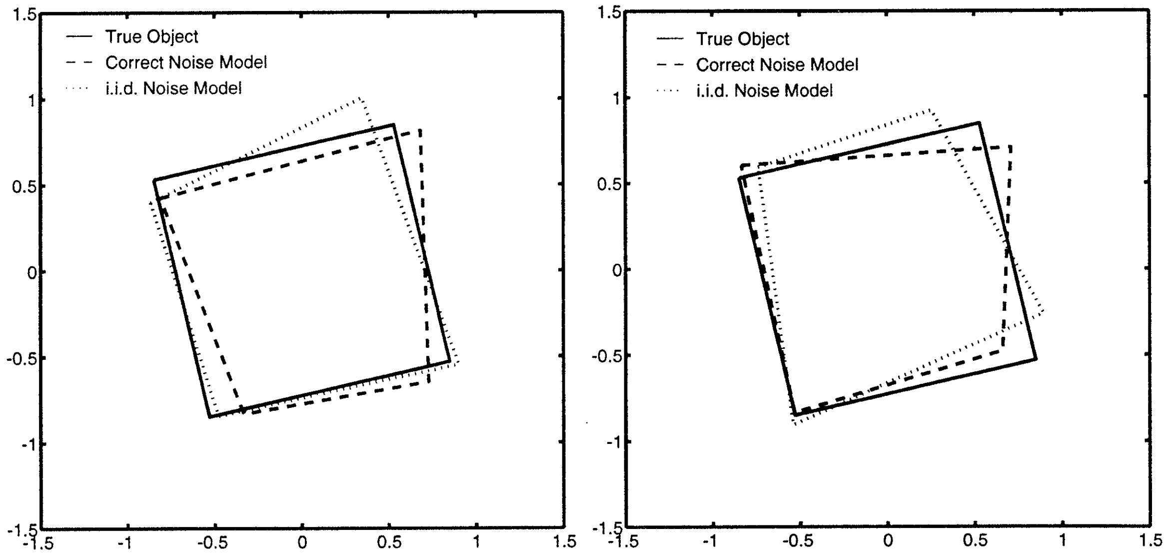

(c) t1 = 1.03

(d) t1 = 1.075

Figure 7-6: Sample Reconstructions for Square Test Object Using Extended Non-Sequential Reconstruction Algorithm; (a) $t_1 = 0$, (b) $t_1 = .75$, (c) $t_1 = 1.03$, (d) $t_1 = 1.075$.

i.i.d. model reconstruction (refer, for example, to Figure 7-5(c)). We also note that in the reconstructions of the kite, as the threshold increased and more knots were dropped or averaged, the reconstructions assuming both the correct applied noise model and the i.i.d. assumed noise model became more peaked and closer to a triangle. Regardless, however, both Figure 7-5 and Figure 7-6 demonstrate that viable reconstructions can be obtained from incomplete projection data.

7.4.2 Monte-Carlo simulations

In the following sections Monte-Carlo simulations of reconstructions are used to test the performance of the Extended Non-Sequential Reconstruction Algorithm. As in the previous chapters, all simulations are performed using the kite test object, with the exception of the last section which is a comparison of reconstructions of the kite and square. The Monte-Carlo simulations consist of 50 independent reconstructions for each reconstruction scenario (unless otherwise stated). In each plot error bars denote the 95% confidence intervals of the sample mean values that result from the 50 runs of the algorithm. As defined at the beginning of this section, the chord length data from each projection are uniform samples over the region $t = [-1, 1]$. Further, reconstructions are done for cases that assume the correct applied noise model for the knot location data and also for cases that assume an i.i.d. model. The projection data for all simulations are uniformly spaced samples over the angular region $[0°, 180°)$. Finally, *unless otherwise stated*, $T = 15$ hypotheses are retained in each step of the algorithm, 5 chord measurements are taken per view, and the SNR on the chord length measurements is set to 10dB. Additionally, each reconstruction scenario performs analysis for cases where the knot removal threshold, t_1, is 0, .75, 1.03, and 1.075 which corresponds to a $0 - 25\%$ (for the kite object) removal of knot location measurements from the projection data.

Reconstruction Error as a Function of Retained Hypotheses

Figure 7-7 shows the performance of the Extended Non-Sequential Algorithm as a function of the number of retained hypotheses, T, at each step of the algorithm. The projection data for each reconstruction consisted of 18 views. Reconstructions were performed for T= 5, 10, 15, and 20 retained hypotheses.

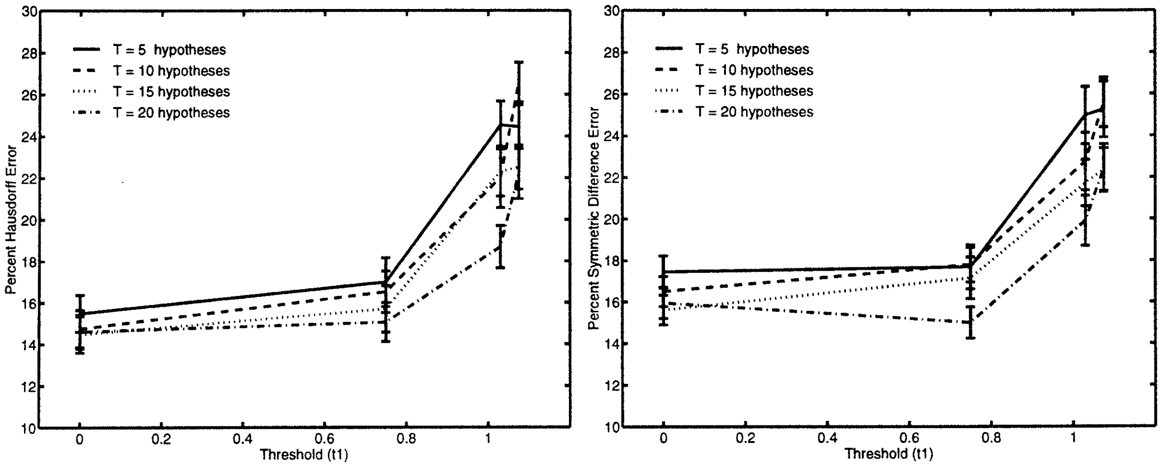

(a) Hausdorff Error, Correct Model (b) Symm. Diff. Error, Correct Model

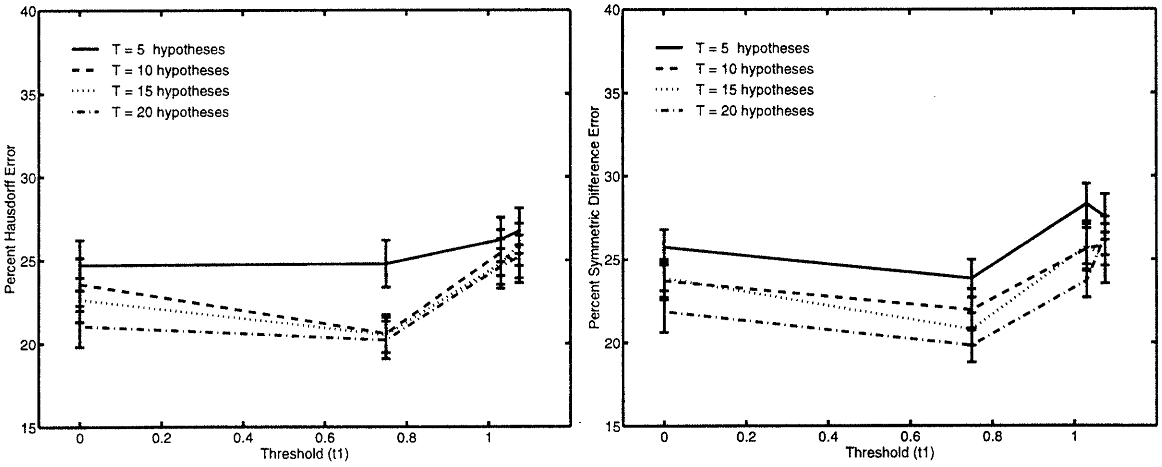

(c) Hausdorff Error, i.i.d. Model (d) Symm. Diff. Error, i.i.d. Model

Figure 7-7: Performance as a Function of Retained Hypotheses Used in the Extended Non-Sequential Reconstruction Algorithm.

As expected, in the plots of Figure 7-7 as the number of retained hypotheses, T, increases, the reconstruction error decreases for both the Hausdorff and Symmetric Difference errors over the range of the threshold values. In addition, in the cases where the correct noise model was assumed (Figure 7-7(a) and Figure 7-7(b)) the reconstruction error for all numbers of retained hypotheses increases as the threshold, t_1, increases. This is expected because even though the knots that are initially removed from the knot location projection data are the most noisy (according to both our noise model given by Equation (3.3) which is applied to the data and the knot removal procedure outlined in Section 7.3), by virtue of the fact that we know the noise statistics *exactly* these knots are still valuable information to the knot extraction process. Thus, missing knot location measurements in the reconstructions that use the correct applied noise model result in a degradation in the reconstruction quality (as denoted by the increasing error measures in the figure). We do note, however, that the increase in error is gradual at first and rises more rapidly once the threshold reaches $\sim .75$. Still, referring to Table 7.2 we see that at the largest considered threshold of $t_1 = 1.075$, 25% of the knot location data were missing from the original projection data yet the percent errors for both measures are well below 30%.

There are a couple of key differences between the reconstructions assuming the i.i.d. noise model and the reconstructions assuming the correct noise model (discussed above). The first is that the relative level of error of the i.i.d. noise model reconstructions is greater for all numbers of retained hypotheses and for all threshold levels (as expected). One of the interesting things to note from the results of Figure 7-7 (c) and (d) is that many of the reconstruction error curves decrease initially and although the curves increase as t_1 gets larger, the curves are much flatter than the corresponding results of Figure 7-7 (a) and (b). This too is expected. Because the i.i.d. model weights all knot location measurements equally, when the threshold t_1 is low the most noisy knots are missing from the projection data and the reconstruction can improve. Thus, the reconstruction improves initially and as t_1 continues to increase, the reconstruction errors start to increase as "good" knot location measurements are missing from the knot location data set.

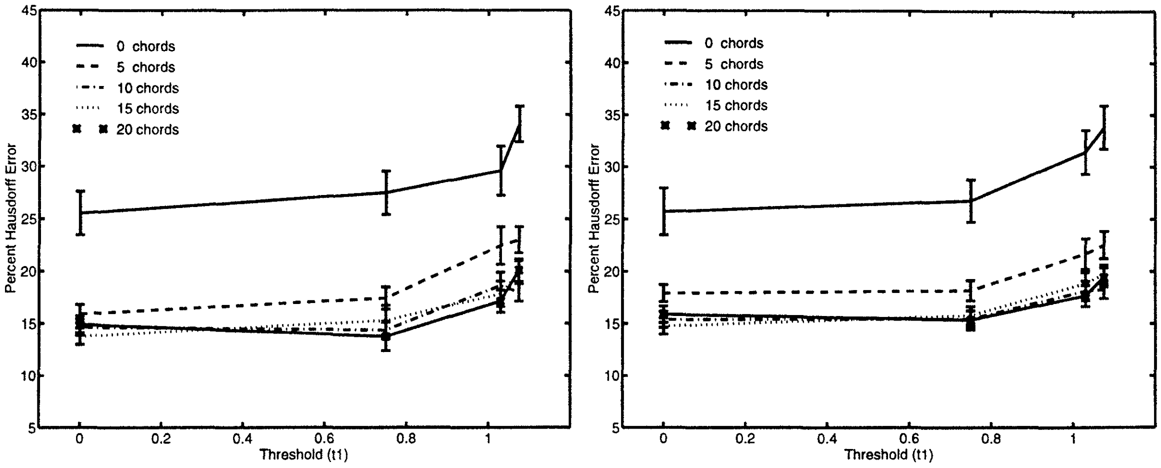

(a) Hausdorff Error (b) Symmetric Difference Error

Figure 7-8: Performance as a Function of Number of Chord Measurements Used in the Extended Non-Sequential Reconstruction Algorithm.

Reconstruction Error as a Function of Chord Measurements

In this section we analyze the performance of the Extended Non-Sequential Algorithm as a function of the number of chord measurements used in the reconstruction algorithm. We set the variance on the chord length measurements to a constant value throughout this experiment. Thus, as the number of chord length measurements per view is varied, the SNR (as defined in Section 3.2.3) will also change. The constant variance on the chord lengths is chosen for the case of 10 chord measurements per view at a SNR of 10dB and the noise on the knot locations was set with Equation (3.3) using this constant chord variance. Noisy chord data were generated for 0, 5, 10, 15, and 20 chord measurements per view. The resulting Hausdorff and Symmetric Difference reconstruction errors are shown in Figure 7-8 (a) and (b), respectively, for the case of reconstructions assuming the correct applied noise model.

As expected, the reconstruction errors in Figure 7-8 are inversely related to the number of chord measurements used in the reconstruction for all threshold values, t_1, considered in this experiment. Like the analysis of the previous section, the Symmetric Difference and Hausdorff errors increase as the value of the threshold increases. In addition, we see that the relative increase in error due to increasing the threshold values varied for different numbers of chord measurements. For example, the Hausdorff reconstruction error for the 0 chord measurement case increases 8.55% as t_1 is varied from 0 to 1.075, while the Hausdorff error for the 20 chord measurement case increases by only 5.18% over the same range of t_1 values. Thus, the inclusion of chord length is an effective method of reducing the error due to missing knot location data measurements. We also note that this is a case of diminishing returns as the error plots for 10, 15, and 20 chord measurements per view are very closely clumped together. This is attributed to the fact that although the chord measurements can have a huge effect on which hypotheses to retain as the most likely estimate of the knot-to-vertex association, ultimately it is only the knot location data in this algorithm that generate the estimate of the vertex locations. Finally, although we have only considered the case of the correct assumed noise model, it is reasonable to assume that the i.i.d. model would behave with characteristics similar to those shown in previous analysis (i.e. higher relative errors but the same general curve shape, slightly decreasing errors for low threshold values, etc.).

Reconstruction Error as a Function of SNR and Number of Views

This section presents the results of a series of independent Monte-Carlo simulations performed at various SNRs and numbers of views. First, we performed a set of reconstructions to characterize the effect of noise on the Extended Non-Sequential Algorithm (refer to Figure 7-9). Reconstructions are performed for SNRs of 0dB, 10dB, 20dB, and 30dB with 18 uniformly space views per reconstruction. Figure 7-9(a) and (b) present the Hausdorff and Symmetric Difference errors, respectively, for reconstructions assuming the correct applied knot location noise model. Similarly, Figure 7-9(c) and (d) present the same error measures for reconstructions assuming an i.i.d. noise model for the knot location data. Similar to the SNR experiment, Figure 7-10 demonstrates the results of a series of independent Monte-Carlo simulations performed with various numbers of views to characterize the effect of the number of angular views on the Extended Non-Sequential Algorithm. Reconstructions are performed for numbers of views equal to 9, 18, 27, and 54 with the SNR set at 10dB. Figure 7-10(a) and (b) present the Hausdorff and Symmetric Difference errors, respectively, for reconstructions assuming the correct applied knot location noise model. Similarly, Figure 7-10(c) and (d) present the same error measures for reconstructions assuming an i.i.d. noise model for the knot location data.

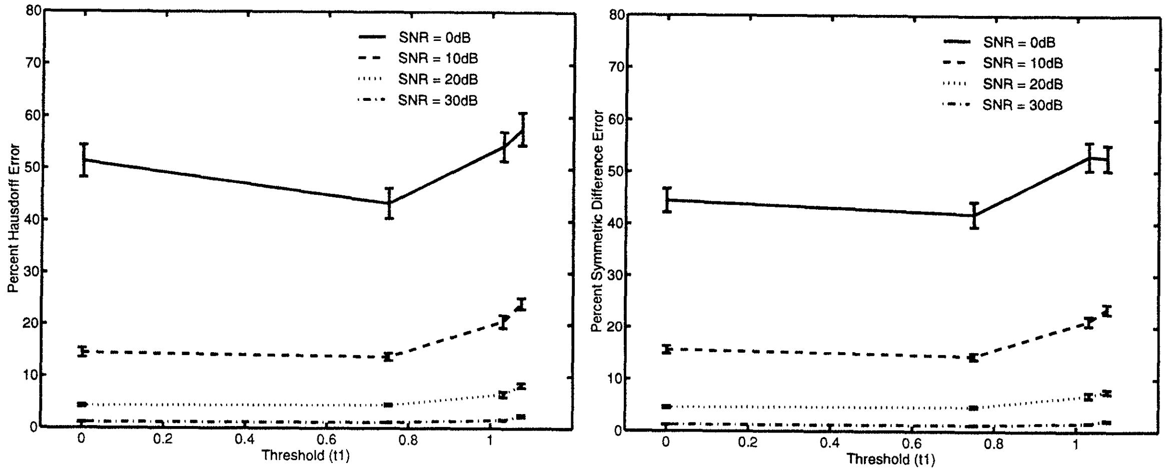

(a) Hausdorff Error, Correct Model (b) Symm. Diff. Error, Correct Model

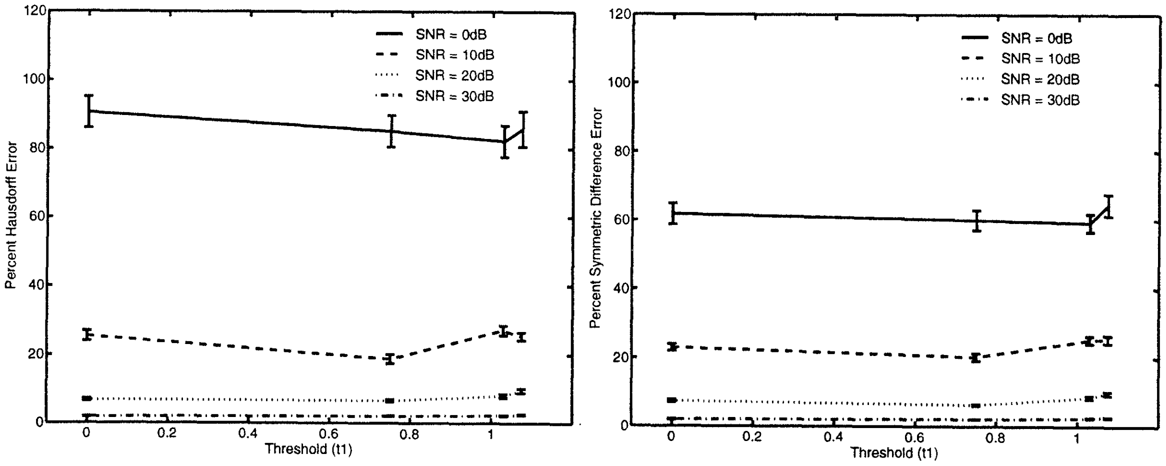

(c) Hausdorff Error, i.i.d. Model (d) Symm. Diff. Error, i.i.d. Model

Figure 7-9: Performance as a Function of SNR Used in the Extended Non-Sequential Reconstruction Algorithm. (a),(b) Correct Noise Model; (c),(d) i.i.d. Noise Model

As expected, for a fixed number of projections, both the Hausdorff and Symmetric Difference reconstruction errors decrease as the SNR increases at all threshold values. Further, consistent with the analysis as a function of retained hypotheses and chord length, the results of the correct noise model, shown in Figure 7-9(a) and (b), demonstrate an increase in the reconstruction error as the threshold increases for all SNRs. In addition, we see that the amount of the increase in error is dependent on the level of SNR. For example, the increase in the reconstruction error for the 0dB case is much more obvious than the increases at 20dB or 30dB. Overall, however, the error curves were relatively flat, showing that the reconstruction is robust even when the knot location projection data is missing as much as 25% of the original measurements. For the case of the i.i.d. model reconstructions (shown in Figure 7-9(c) and (d)), the errors decrease slightly for small threshold values (consistent with previous results) and remain relatively level until the threshold gets to the point where the percentage of knot removed becomes significant enough to adversely effect the reconstruction. Thus, the Extended Non-Sequential Algorithm continues to produce quality reconstructions for relatively large data losses. Even as the percentage of knots removed is increased to 25% of the overall knot location data, the resulting increase in error is not extreme (refer again to Figure 7-9). Finally, we note that both the i.i.d. and correct noise reconstructions seem to approach the same reconstruction error values for both the Hausdorff and Symmetric Difference errors as the threshold gets large.

In addition, the results of varying the number of views (refer to Figure 7-10) were consistent with both the expected results and those obtained in previous sections. In the reconstructions that assume the correct noise model the reconstruction errors increase as a function of the threshold; while for the i.i.d. noise model the reconstruction errors decreased initially and start to increase as the threshold gets large. In addition, the reconstruction errors for both of the measures decrease as the number of views increases (as expected).

Reconstruction Comparison of Two Test Objects

Having completed the Monte-Carlo analysis of the Non-Sequential Algorithm for a particular object, the kite, we now turn to a comparative analysis. Figure 7-11 shows the Monte-Carlo reconstruction errors for both the kite and the square test objects as a function of the number of views used in the reconstruction. Figure 7-11 (a) and (b) are Hausdorff and Symmetric

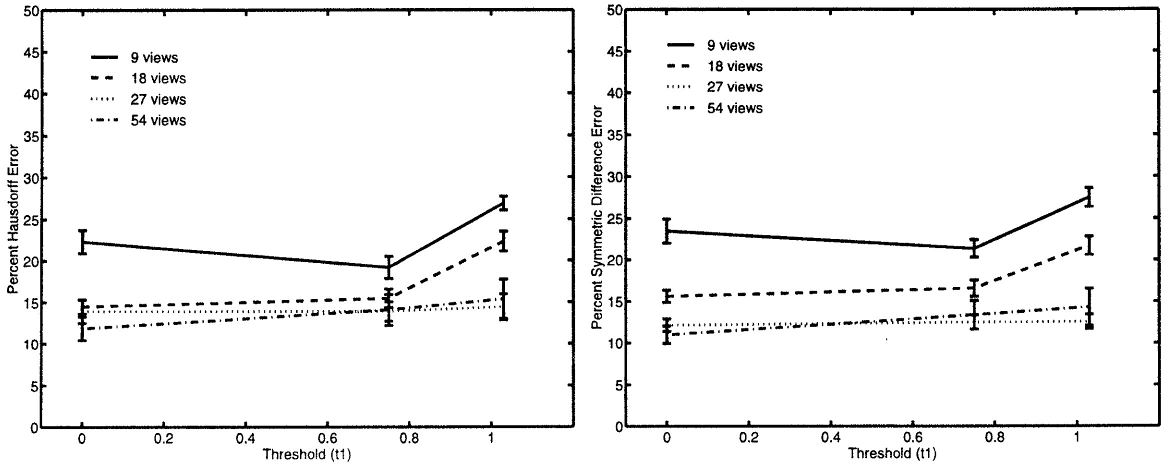

(a) Hausdorff Error, Correct Model (b) Symm. Diff. Error, Correct Model

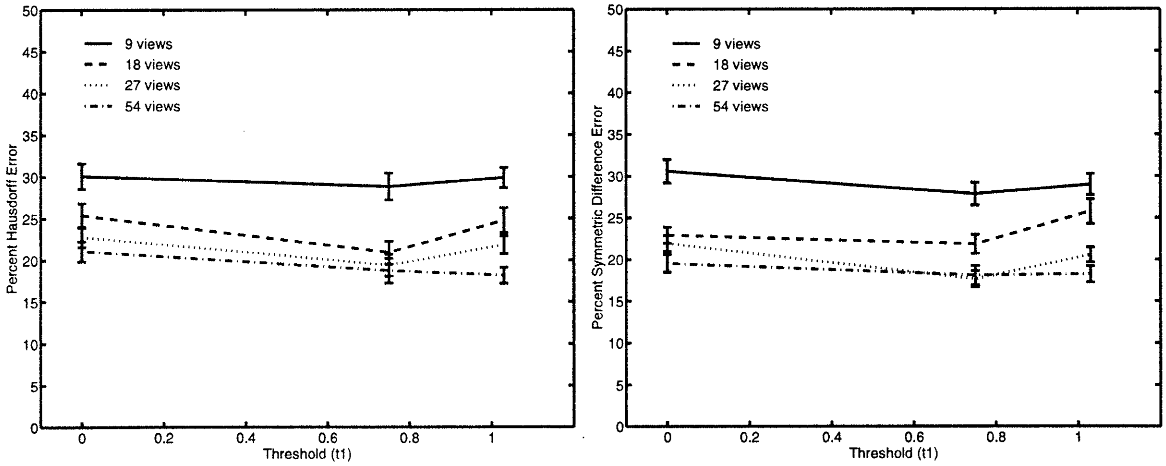

(c) Hausdorff Error, i.i.d. Model (d) Symm. Diff. Error, i.i.d. Model

Figure 7-10: Performance as a Function of Number of Views for the Extended Non-Sequential Reconstruction Algorithm. (a),(b) Correct Applied Noise Model; (c),(d) i.i.d. Noise Model.

Difference errors for reconstructions assuming the correct applied knot location noise model while Figure 7-11 (c) and (d) are the corresponding error measures under the assumption of an i.i.d. noise model.

From the results of the analysis assuming the correct applied noise model (Figure 7-11(a) and (b)), we see that the errors of the kite reconstruction remain relatively flat for small threshold values and then increase significantly for thresholds greater than ~0.75. Unlike the kite, the reconstruction errors of the square are relatively flat and increase only slightly over the entire range of threshold values. We note that for both the Hausdorff and Symmetric Difference errors the reconstruction error of the kite exceeds that of the square at the final threshold level of $t_1 = 1.075$. This result can be attributed to the amount of data that is missing from each object at the various threshold levels. Referring to Table 7.2 and Table 7.3 for the kite and square respectively we see that the maximum percent missing of knot locations for the kite at $t_1 = 1.075$ is 25% while for the square the percentage of missing data is only 13.89%.

In the i.i.d. reconstruction error results of Figure 7-11(c) and (d), we see that the reconstruction error of the square is significantly higher than that of the kite at smaller threshold levels. As the threshold increases, the reconstruction error for the square decreases for both the Hausdorff and Symmetric Difference errors. This decrease in reconstruction error occurs for the same reasons outlined previously. Basically, the i.i.d. model weights the really noisy knot location data equal to the "good" measurements. As the threshold starts to increase, the poor knot location measurements are the first values removed from the projection data and the reconstruction improves. When the threshold gets high enough, measurements that provide vital information to the reconstruction begin to be removed and the error starts to increase again (this can be seen in the case of the kite). Because the percentage of missing knots in the square is so low for the square, the reconstruction process has not reached the point where the errors have started to increase again. Finally, we note that the relative errors of the i.i.d. reconstructions are higher than the reconstructions assuming the correct noise model (as expected).

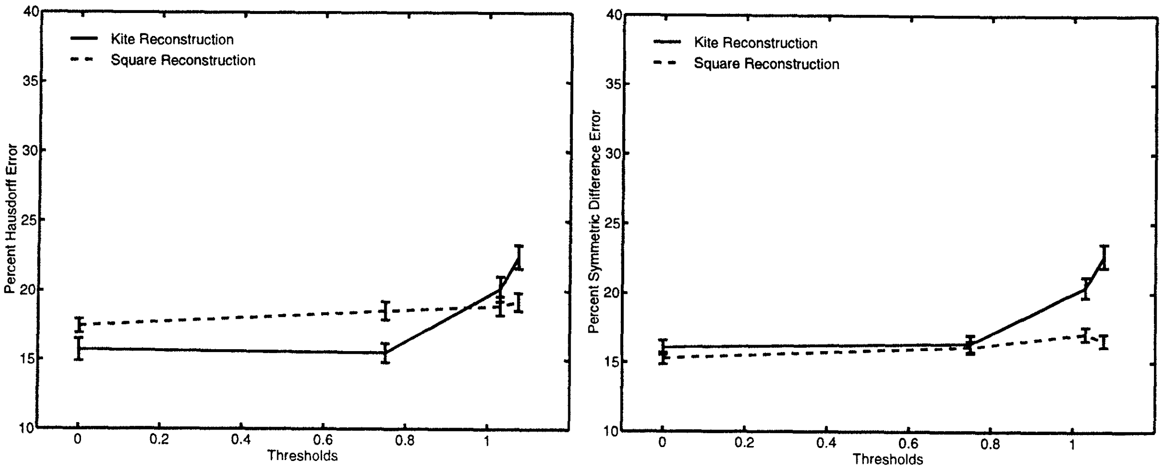

(a) Hausdorff Error, Correct Model (b) Symm. Diff. Error, Correct Model

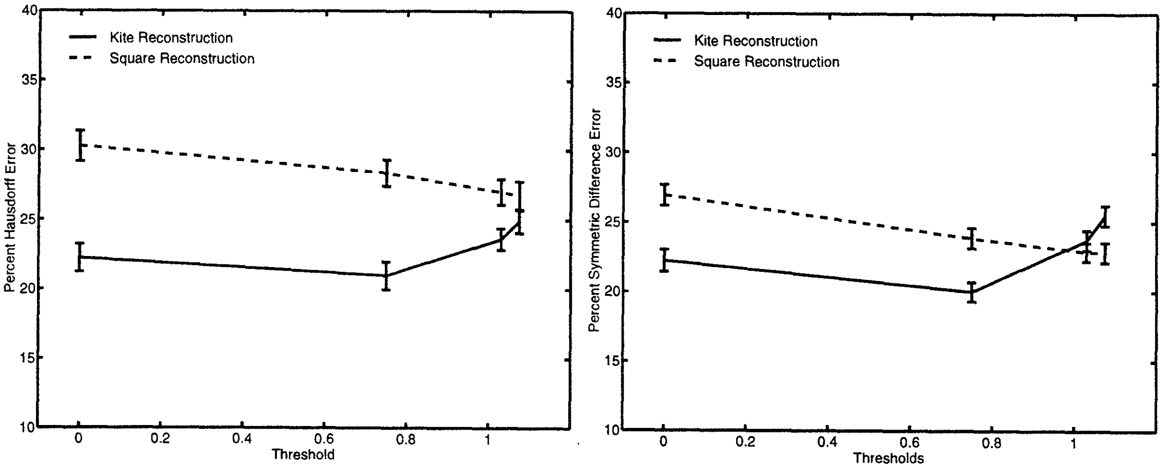

(c) Hausdorff Error, i.i.d. Model (d) Symm. Diff. Error, i.i.d. Model

Figure 7-11: Comparison of Kite and Square Test Objects for the Extended Non-Sequential Reconstruction Algorithm: SNR=10dB, T=15 retained hypotheses, 5 chord measurements per view. (a),(b) Correct Noise Model; (c),(d) i.i.d. Noise Model.

7.4.3 Nonlinear Reconstruction Algorithm With Extended Non-Sequential Reconstruction Algorithm Initial Estimate

In this section we use the Nonlinear Reconstruction Algorithm of Chapter 6 to reconstruct the vertices of the object using the estimate of the vertices generated from the Extended Non-Sequential Algorithm as the initial estimate and the chord length data and incomplete knot location data as the measurements. Reconstructions were performed with 18 uniformly spaced angular views over the range $[0°, 180°)$ with 5 chord measurements per view at a SNR of 10dB. As in all previous experiments, the chord measurements were uniformly spaced over the range $t = [-1, 1]$. In addition, we assume the i.i.d. noise model in our reconstructions (as was done in all experiments of Chapter 6). For each reconstruction scenario, 50 Monte-Carlo reconstructions were performed at each threshold t_1.

In addition, the one key difference between the implementation using the Extended Non-Sequential Algorithm as the initial guess to the Nonlinear Algorithm is that we must specify the hypothesized missing knot location measurements in the projection data to the Nonlinear Algorithm. Thus, to minimize the cost function of Equation (6.6) we must know the correspondence of the knot locations to the incomplete projection data. This information is available in the optimal discrete data association hypothesis at the output of the Extended Non-Sequential Reconstruction Algorithm and is passed into the Nonlinear Algorithm along with the incomplete measurement data.

Figure 7-12 shows the reconstruction results of the nonlinear reconstruction. The solid and dotted lines represent the Hausdorff error and Symmetric Difference error of the initial estimates generated from the Extended Non-Sequential Algorithm. The behavior of these error plots as a function of threshold level is consistent with the results of previous sections. The results of the nonlinear reconstruction, shown in the dashed and dash-dotted lines for the Hausdorff and Symmetric Difference errors respectively, demonstrate a marked improvement over the initial estimates. Both nonlinear reconstruction error curves behave in the same way as a function of the threshold as the initial estimates, but the relative error of each reconstruction has been significantly reduced. Thus, affirming the results of Chapter 6, the Nonlinear Reconstruction Algorithm resulted in a marked reconstruction improvement even when the threshold level was high and there were a large percentage of missing knot location measurements. Although this reduction in error is slightly larger for small thresholds,

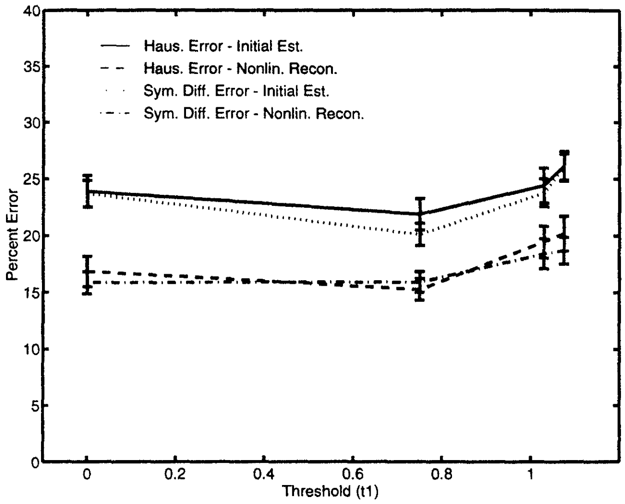

Figure 7-12: Performance of Nonlinear Reconstruction Algorithm Using Extended Non-Sequential Reconstruction Algorithm Initial Estimate.

there is still a noticeable improvement in high threshold reconstruction scenarios. Thus, the Nonlinear Reconstruction Algorithm again reduces the error significantly.

7.5 Conclusions

In this chapter, we extended the capability of the Non-Sequential Reconstruction Algorithm of Chapter 5 to include the possibility of missing knot location measurements in the projection data. The general reconstruction framework was still posed as a combination hypothesis test-estimation problem but each hypothesis consisted of two associations: the association of the knot locations to the vertices of the object and the association of the measurement data to the knot locations. As a result, the mechanics of the algorithm remained the same as in the original Non-Sequential Algorithm (i.e. the methods used to estimate the vertices, determine the optimal hypotheses, prune the hypotheses, etc.), but the number of possible hypotheses is increased significantly.

The results of the reconstructions in this chapter were encouraging. Even with losses of up to 25% of the original projection data, the algorithm was still able to produce quality reconstruction results. In addition, by retaining more hypotheses, using more chord data, and increasing the number of views (up to a point of diminishing returns), the reconstruction errors were decreased. We found that if the correct noise models are known, lower SNRs are much more sensitive to the loss of data than reconstruction scenarios where only i.i.d.

models were used. Finally, we note that the results of the nonlinear reconstruction showed that the reconstruction error can be significantly reduced even if a large percentage of knot location data is missing from the projection data. Thus, the Extended Non-Sequential Algorithm was able to deal remarkably well with incomplete knot location measurement data. This result is important in that a realistic implementation of this reconstruction algorithm would almost certainly have to deal with missing data measurements.

As a final note, we stress that the focus of this chapter has been to characterize the performance of the Extended Non-Sequential Algorithm when faced with the problem of missing knot location measurements. As such, we developed a simplistic approach for removing knots from the noise free knot location data. In reality, we expect that the number of knots that are missed during the knot extraction process to be directly related to the SNR of the projection data. One way of characterizing the missed detections as a function of SNR is with the following formulation:

- New Threshold Test 1

  If z_i is an internal (noiseless) knot location and

$$\widetilde{M}_i \leq 2\left(\frac{K^{.5}\sigma_m}{|\Delta s_i|(M_i)^{.5}}\right) \tag{7.7}$$

  knot z_i is removed for the data set; if not, z_i is retained.

This uses the standard deviation of the knot location noise model (refer to Section 3.2.3, Equation (3.3)) and the minimum separation to quantify a missed detection. As defined previously in Section 3.2.3, K is a scaling constant, σ_m^2 is the noise variance on chord length measurements, $|\Delta s_j|$ is the change in slope at knot location measurement $z_j(\theta_i)$, and M_j is the distance measure between adjacent knot locations. Basically, this formulation says that a knot is considered a missed detection if the probability of the knot switching positions with the closest adjacent knot position is greater than 2.3%. Using this formulation in conjunction with Threshold Test 2, Table 7.5 lists both the number of and percentage of knots that would be removed from the kite object based on the SNR values of 0dB, 10dB, 20dB, and 30dB used in this thesis (where the SNR is computed for the case of 5 chord measurements per view).

Views	Meas. on Data Set	SNR=0dB	SNR=10dB	SNR=20dB	SNR=30dB
9	Number Removed	17	15	2	1
	Percent Removed (%)	47.22	41.67	5.56	2.78
18	Number Removed	33	30	5	2
	Percent Removed (%)	45.83	41.67	6.94	2.78
27	Number of Removed	51	47	9	3
	Percent Removed (%)	47.22	43.52	8.33	2.78
54	Number Removed	102	95	19	7
	Percent Removed (%)	47.22	43.98	8.80	3.24

Table 7.5: New Table of Conversions for Kite Test Object - SNRs, Number of Missing Knot Location Measurements, Percent Missing Knot Measurements.

Chapter 8

Conclusions

8.1 Concluding Remarks

In this thesis, we have developed several algorithms to reconstruct convex polygonal objects from sparse and noisy measurements of projection data. Specifically, our techniques concentrated on reconstructions that use geometric-based information (chord lengths and knot locations) extracted directly for the projection data to estimate the vertices of the polygonal object. We have shown in these algorithms that information that has geometric importance in *both* object space and Radon space can be used to produce high quality reconstructions in the face of limited and noisy projection data.

The first two algorithms, the Sequential and Non-Sequential Algorithms, are based on techniques that are traditionally used in multitarget tracking problems to determine the position of targets. The methodology is a simultaneous solution to a discrete-valued hypothesis testing problem (to determine knot-to-vertex data associations) and a continuous-valued ML estimation problem (to determine the vertices of the polygon). By using knot location measurements, which are the locations of the projections of the vertices of the object in the projection, we were able to use the discrete data association hypotheses to arrive at a linear estimation problem for the vertices and the knot location measurements. This is significant because in general the vertices of a polygonal object are a nonlinear function of its projections (as demonstrated in the chord length analysis as a function of the vertices in the appendix of Chapter 4). Thus, in our framework, traditional ML estimation and M-ary hypothesis testing techniques are used to estimate the vertices of the object. In addition to the use of the knot location data, we also incorporated the chord length measurements into the algorithms to prune the set of possible data associations at each step of the algorithm and help determine the optimal data association hypothesis.

The Sequential Algorithm required views to be input into the algorithm in increasing angular order with relatively small angular increments. It used a gating approach to limit the number of discrete data association hypotheses by allowing only adjacent knot locations to switch associations with the vertices of the object from one view to the next. By relaxing the restrictions on the discrete data association hypotheses considered at each step of the algorithm, the Non-Sequential Algorithm allowed views to be incorporated into the algorithm in *any* angular order. Thus, the Non-Sequential Algorithm allowed greater flexibility at the expense of increased model complexity. In our analysis, we found that both algorithms were able to reconstruct convex, binary polygonal objects in sparse and noisy reconstruction scenarios. In addition, both algorithms were able to reconstruct objects from projections obtained over limited angular ranges and with non-uniform spacing. In all of these reconstruction scenarios, however, we found that the better triangulation geometry afforded by the Non-Sequential Algorithm (due to its flexibility in allowing views to be entered into the algorithm in any angular order) greatly reduced the reconstruction errors.

The third algorithm developed in this thesis, the Nonlinear Reconstruction Algorithm, directly uses *both* knot location and chord length data to estimate the vertices of the object by minimizing a ML-based cost criterion. The algorithm uses an initial estimate generated from either the Sequential or Non-Sequential Algorithms in addition to the measurement data to minimize the cost function and arrive at the estimate of the vertices of the object. Thus, the Nonlinear Algorithm is an extension of the work of the Sequential and Non-Sequential Algorithms. Although the cost function is a complex nonlinear function of the $2n$ parameters that define the vertices of the polygon, our analysis showed that for the case of a 4-gon the Nonlinear Algorithm was able to significantly reduce the reconstruction error (as measured by the Hausdorff and Symmetric Difference errors) for all numbers of views and SNRs above 10dB considered in this thesis. This result is attributed to the direct inclusion of both chord length and knot location measurements to arrive at the estimate of the vertex locations.

Finally, the last algorithm developed in this thesis, the Extended Non-Sequential Algorithm, expands the original Non-Sequential Algorithm to encompass situations of missing knot location measurements in the projection data. The approach was to expand the set of discrete data association hypotheses to consider all possible knot-to-vertex associations (as before) *and* all possible measurement-to-knot associations. This algorithm is an important extension of the original Non-Sequential Algorithm because it considers the realistic situation that a full set of knot location projection data may not always be available in an actual implementation of this algorithm. The reconstruction analysis of the Extended Non-Sequential Algorithm showed that the algorithm was able to robustly deal with losses of up to 25% of the original knot location data. Even with these large losses of data, the reconstruction errors only rose slightly. In addition, we showed that if the results of the Extended Non-Sequential Algorithm are used as the input to the Nonlinear Algorithm, the reconstruction error can be significantly reduced.

8.2 Future Work

Because the scope of this thesis was limited in a number of ways, there are several possible extensions of this work:

- *Knot Extraction Algorithm*

  Throughout the work of this thesis, we have generated both the chord length and knot location data directly from the underlying object and perturbed the noise-free data with additive, independent Gaussian noise based on noise models developed in this thesis. The measurements were then input into the reconstruction algorithms developed in this thesis. A natural extension of this work is to implement a knot extraction algorithm and provide a complete reconstruction system that would use the actual projection data at the input and produce estimates of the vertices of the object at the output.

- *Inclusion of False Alarms*

  Consistent with the analogy to the radar multitarget tracking problem, the algorithms developed in this thesis could be expanded to include false alarms, or erroneous measurements that do not correspond to actual vertices of the object, in the knot location measurements. With the added complexity of false alarms in conjunction with the possibility of missed detections, the algorithm could adaptively estimate the number of vertices of the object in addition to the estimates of their location.

- *Relaxation of Assumptions*
  To limit the scope of this thesis, a number of simplifying assumptions were made on the objects to be reconstructed and data used. For example, the assumption of convexity and the requirement of two full views of knot location data in the Extended Non-Sequential Reconstruction Algorithm can be removed at the expense of increasing the complexity of the algorithm.

Bibliography

[1] Henk A. P. Blom and Yaakov Bar-Shalom. The interacting multiple model algorithm for systems with markovian switching coefficients. *IEEE Transactions on Automatic Control*, 33(8):780 – 783, August 1988.

[2] Yoram Bresler and Albert Macovski. Three-dimensional reconstruction from projections with incomplete and noisy data by object estimation. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-35(8):1139–1152, August 1987.

[3] Rodney A. Brooks and Giovanni Di Chiro. Principles of computer assisted tomography (cat) in radiographic and radioisotropic imaging. *Phys. Med. Biol.*, 21(5):689 – 732, 1976.

[4] G.T. Herman. *Image Reconstruction from Projections.* Academic Press, 1980.

[5] W. C. Karl. *Reconstructing Objects From Projections.* PhD thesis, MIT, February 1991.

[6] T. Kurien. Issues in the design of practical multitarget tracking algorithms. In Y. Bar-Shalom, editor, *Multitarget-Multisensor Tracking: Applications and Advances*, chapter 3. Artech House, 1992.

[7] T. Kurien, A.L. Blitz, R.B. Washburn, and A.S. Willsky. Optimal manuever detection and estimation in multiobject tracking. In *Proceedings of the Sixth MIT/ONR Workshop on Command and Control*, pages 164 – 171, December 1983.

[8] A. M. Mier-Muth and Alan S. Willsky. A sequential method for spline approximation with variable knots. MIT Electronic Systems Laboratory ESL-P-759, MIT, 1977.

[9] Peyman Milanfar. *Geometric Estimation and Reconstruction From Tomographic Data.* PhD thesis, MIT, June 1993.

[10] Walter Munk and Carl Wunch. Ocean acoustic tomography: A sceme for larger scale monitoring. *Deep-Sea Research*, 26A:123–161, 1979.

[11] Jerry L. Prince and Alan S. Willsky. Reconstructing convex sets from support line measurements. Center for Intelligent Control Systems CICS-P-16, MIT, 1987.

[12] Jerry Ladd Prince. *Geometric Model-Based Estimation From Projections.* PhD thesis, MIT, January 1988.

[13] Donald B. Reid. An algorithm for tracking multiple targets. *IEEE Transactions on Automatic Control*, 24(6), December 1979.

[14] David J. Rossi. *Reconstruction from Projections Based on Detection and Estimation.* PhD thesis, MIT, October 1982.

[15] David J. Rossi and Alan S. Willsky. Object shape estimation from tomographic measurements - a performance analysis. Center for Intelligent Control Systems CICS-P-33, MIT, September 1987.

[16] N. Srinivasa, K. R. Ramakrishnan, and K. Rajgopal. Detection of edges from projections. *IEEE Transactions on Medical Imaging*, 11(1), March 1992.

[17] Jean-Philippe Thirion. Segmentation of tomographic data without image reconstruction. *IEEE Transactions on Medical Imaging*, 11(1), March 1992.

[18] Harry L. Van Trees. *Detection, Estimation, and Modulation Theory, Part I.* John Wiley and Sons, 1968.

[19] Alan S. Willsky. Stochastic processes, detection, and estimation. 6.432 Course Notes, MIT, 1991.